ROBOTICS IN EXTREME ENVIRONMENTS

EDITED BY: Chie Takahashi, Manuel Giuliani, Barry Lennox, William R. Hamel, Rustam Stolkin and Claudio Semini

PUBLISHED IN: Frontiers in Robotics and AI

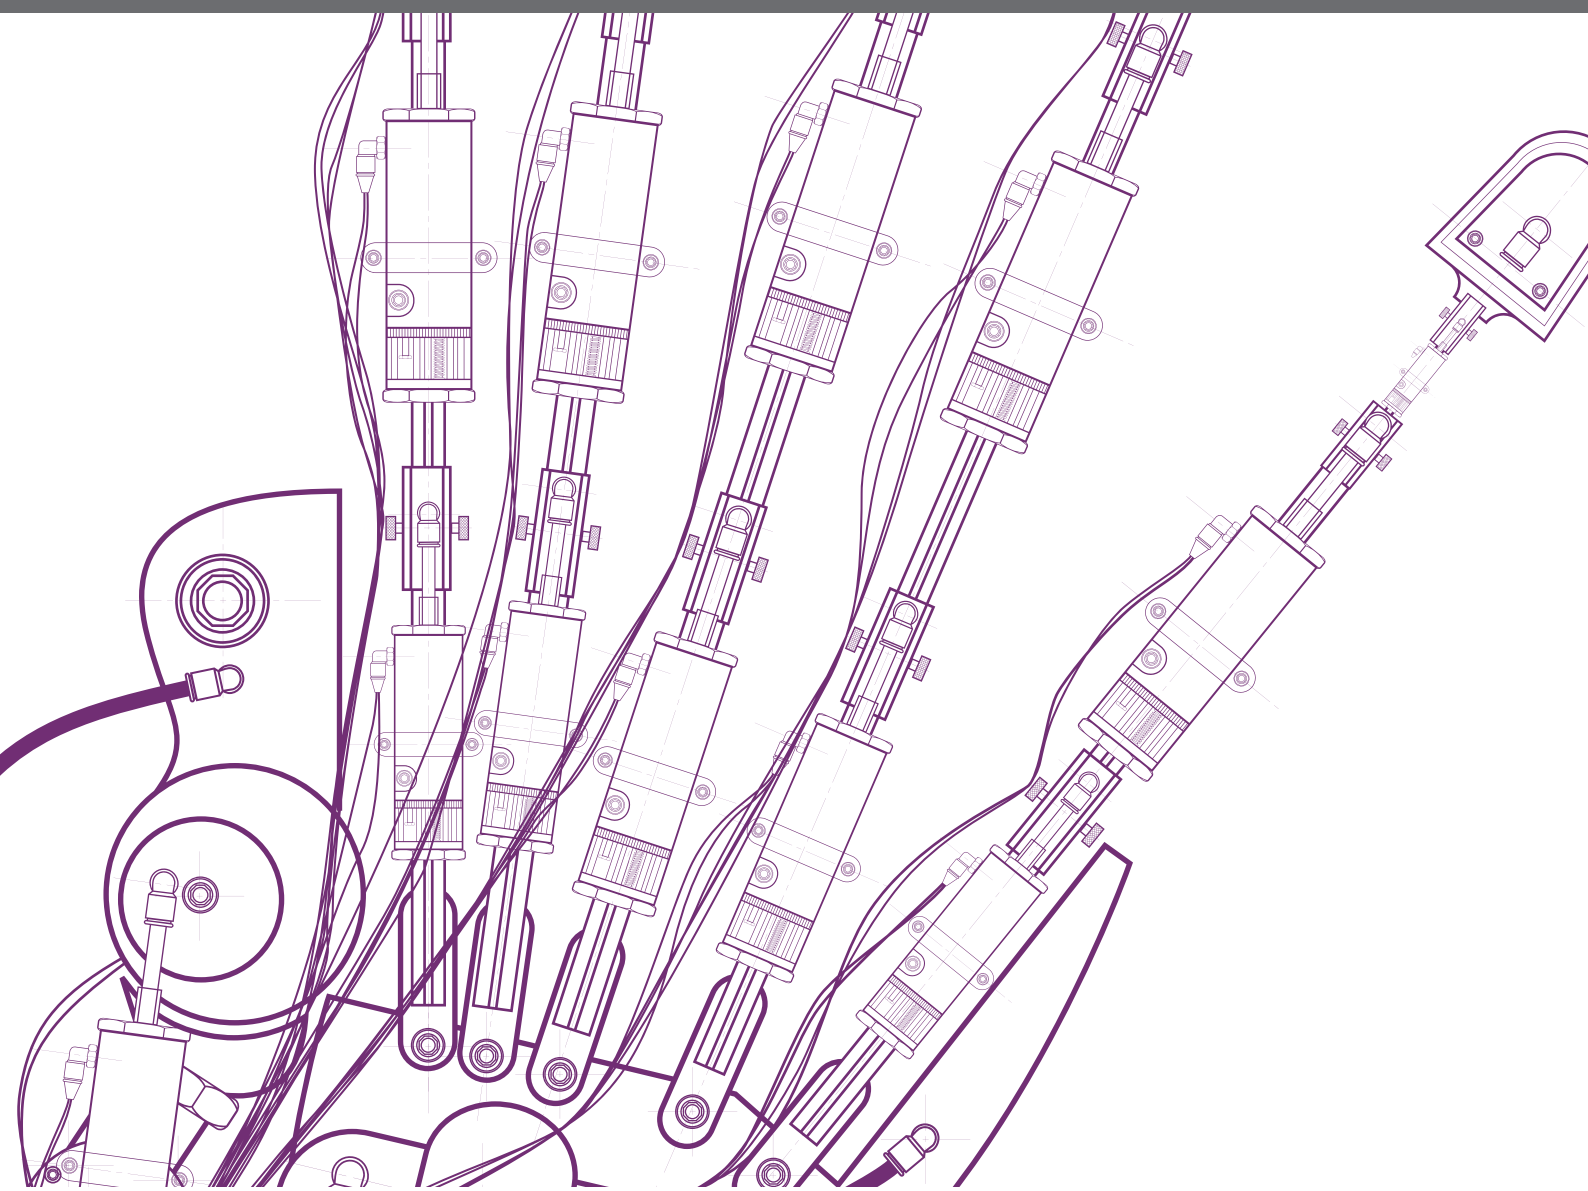

frontiers Research Topics

frontiers

ISSN 1664-8714
ISBN 978-2-88971-425-4
DOI 10.3389/978-2-88971-425-4

About Frontiers

Frontiers is more than just an open-access publisher of scholarly articles: it is a pioneering approach to the world of academia, radically improving the way scholarly research is managed. The grand vision of Frontiers is a world where all people have an equal opportunity to seek, share and generate knowledge. Frontiers provides immediate and permanent online open access to all its publications, but this alone is not enough to realize our grand goals.

Frontiers Journal Series

The Frontiers Journal Series is a multi-tier and interdisciplinary set of open-access, online journals, promising a paradigm shift from the current review, selection and dissemination processes in academic publishing. All Frontiers journals are driven by researchers for researchers; therefore, they constitute a service to the scholarly community. At the same time, the Frontiers Journal Series operates on a revolutionary invention, the tiered publishing system, initially addressing specific communities of scholars, and gradually climbing up to broader public understanding, thus serving the interests of the lay society, too.

Dedication to Quality

Each Frontiers article is a landmark of the highest quality, thanks to genuinely collaborative interactions between authors and review editors, who include some of the world's best academicians. Research must be certified by peers before entering a stream of knowledge that may eventually reach the public - and shape society; therefore, Frontiers only applies the most rigorous and unbiased reviews.
Frontiers revolutionizes research publishing by freely delivering the most outstanding research, evaluated with no bias from both the academic and social point of view. By applying the most advanced information technologies, Frontiers is catapulting scholarly publishing into a new generation.

What are Frontiers Research Topics?

Frontiers Research Topics are very popular trademarks of the Frontiers Journals Series: they are collections of at least ten articles, all centered on a particular subject. With their unique mix of varied contributions from Original Research to Review Articles, Frontiers Research Topics unify the most influential researchers, the latest key findings and historical advances in a hot research area! Find out more on how to host your own Frontiers Research Topic or contribute to one as an author by contacting the Frontiers Editorial Office: frontiersin.org/about/contact

ROBOTICS IN EXTREME ENVIRONMENTS

Topic Editors:
Chie Takahashi, University of Cambridge, United Kingdom
Manuel Giuliani, University of the West of England, United Kingdom
Barry Lennox, The University of Manchester, United Kingdom
William R. Hamel, The University of Tennessee, Knoxville, United States
Rustam Stolkin, University of Birmingham, United Kingdom
Claudio Semini, Italian Institute of Technology (IIT), Italy

Topic editor Rustam Stolkin is director of A.R.M Robotics Ltd. All other topic editors declare no competing interests with regards to the Research Topic subject.

Citation: Takahashi, C., Giuliani, M., Lennox, B., Hamel, W. R., Stolkin, R., Semini, C., eds. (2021). Robotics in Extreme Environments. Lausanne: Frontiers Media SA.
doi: 10.3389/978-2-88971-425-4

Table of Contents

frontiers in Robotics and AI

EDITORIAL
published: 19 August 2021
doi: 10.3389/frobt.2021.744092

Editorial: Robotics in Extreme Environments

Chie Takahashi[1], Manuel Giuliani[2], Barry Lennox[3], William R. Hamel[4], Rustam Stolkin[5] and Claudio Semini[6]*

[1]Department of Psychology, University of Cambridge, Cambridge, United Kingdom, [2]University of the West of England, Bristol, United Kingdom, [3]The University of Manchester, Manchester, United Kingdom, [4]The University of Tennessee, Knoxville, TN, United States, [5]University of Birmingham, Birmingham, United Kingdom, [6]Dynamic Legged Systems, Istituto Italiano di Tecnologia (IIT), Genova, Italy

Keywords: robotics, manipulation, autonomous systems, remote operation, human-robot interaction, sensing

Editorial on the Research Topic

Robotics in Extreme Environments

INTRODUCTION

The development and deployment of robotics technology in extreme environments, such as nuclear decommissioning, offshore maintenance, space exploration and deep mining has received considerable attention in recent years. In all of these areas, robots are required to reduce the risks associated with operations staff, typically by removing the requirement for people to enter the hazardous environments and to increase productivity in high consequence and cluttered facilities. There has been significant effort in robotics research in this area, for example to make robotic systems able to survive in the presence of high levels of radiation and toxic substances, to operate at extreme pressures and temperatures, and to complete tasks safely in unstructured environments. We have seen there are many cross-domain challenges that researchers in robotics in extreme environments are working on (**Figure 1**).

In the context of robotics in extreme environments we raised the research topic aiming to bring together the latest cutting-edge research in the field, to deepen the current understanding and to share research challenges. This e-book comprises a collection of eight articles, published by Frontiers in Robotics and AI on the topic.

OVERVIEW OF THE CONTENTS OF THE E-BOOK

Let's Push Things Forward: A Survey on Robot Pushing

Stüber et al. present an overview of the current state of the art in robot pushing. Pushing objects is a crucial skill that mobile robots require for a wide range of applications, including tasks in extreme environments. The authors compared more than 50 publications dividing them into six categories: purely analytical, hybrid, dynamic analysis, physics engine, data driven and deep learning. A special focus was given to the problem of motion prediction of the object to be pushed. The authors conclude the paper with final remarks including open problems.

OPEN ACCESS

Edited and reviewed by:
Antonio Manuel Pascoal, Universidade de Lisboa, Portugal

****Correspondence:***
Chie Takahashi
chie.k.takahashi@gmail.com

Specialty section:
This article was submitted to Robotic Control Systems, a section of the journal Frontiers in Robotics and AI

Received: *19 July 2021*
Accepted: *27 July 2021*
Published: *19 August 2021*

Citation:
Takahashi C, Giuliani M, Lennox B, Hamel WR, Stolkin R and Semini C (2021) Editorial: Robotics in Extreme Environments. Front. Robot. AI 8:744092. doi: 10.3389/frobt.2021.744092

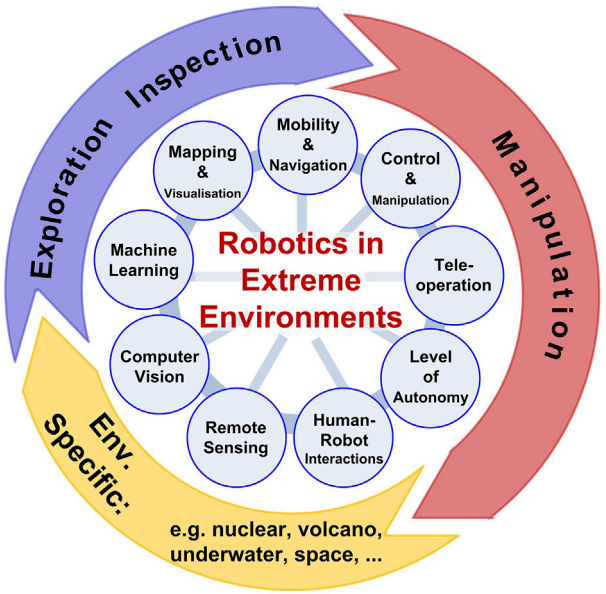

FIGURE 1 | Overview of research topics and challenges for Robotics in Extreme Environments.

Radiation Tolerance Testing Methodology of Robotic Manipulator Prior to Nuclear Waste Handling

Zhang, et al., present a radiation tolerance testing methodology for robotic manipulators, covering key aspects of the process from emulation of the radiation environment through to hot cell testing set up, data acquisition and analysis. The group from the Universities of Bristol and Manchester applied their methodology to a KUKA robotic arm and showed that it was surprisingly robust to radiation. These results suggest that the full spectrum of modern industrial robot products may be viable (with appropriate modifications) for practical use in future nuclear remote operations.

A Holistic Approach to Human-Supervised Humanoid Robot Operations in Extreme Environments

The paper (Wonsick et al.) written by researchers from Northeastern University and Irish Manufacturing Research gives attention to the physical interaction aspects of a humanoid using glove box ports in terms of the mechanics of surfaces, footing and extended-arm operations. The authors present their initial concepts and ideas about making humanoid operations human-supervised along the lines that will make sense for glove box operations and provide pragmatic insights into how humanoid robots will need to be extended to 1 day be used in extreme environments.

Radiation Mapping and Laser Profiling Using a Robotic Manipulator

Characterisation of nuclear materials, particularly in legacy facilities, is of great importance and in this article (White et al.), researchers from the University of Bristol, working with KUKA Systems, developed a robotic system, integrated with a radiation detector and time of flight sensor to provide measures of the distance to any object being scanned and the level of gamma activity at this location. By using the manipulator to raster over the surface of an object, the proposed technique is able to produce 3-dimensional radiation characterisation maps.

Radiological Mapping of Post-Disaster Nuclear Environments Using Fixed-Wing Unmanned Aerial Systems: A Study From Chernobyl

This article (Connor et al.) describes a study led by researchers from the University of Bristol into the use of a fixed-wing unmanned aerial system to map radioactive contamination across relatively large areas of land. The capabilities of the proposed system were demonstrated through deployment within the Chernobyl Exclusion Zone (CEZ), where it generated radiation dose-rate maps of large areas of land that were consistent with more expensive, manual surveys.

BVLOS Unoccupied Aerial Systems Operations in Highly-Turbulent Volcanic Plumes

Wood et al. provide insights into operating Unoccupied Aerial Systems (UAS) in highly-turbulent volcanic plumes. They present a detailed analysis of three missions in which the team flew a fixed wing UAS beyond visual line of sight into the plumes of Manam volcano, Papua New Guinea. The paper contains a detailed description of the used UAS and provides insights into plume sampling applications, with the authors giving recommendations for physical parameters and propulsion systems of aircraft used for taking measurements in turbulent volcanic plumes.

Automatic Fracture Characterization Using Tactile and Proximity Optical Sensing

Palermo et al. present a custom-designed integrated tactile and proximity sensor that can be used for automatic detection of surface cracks. This approach might be more suitable for operation in extreme environments where, for example, radiation may damage electronic components of commonly employed sensing devices. For the detection, the sensor slides across different surfaces and records data. Using machine

learning, the team can then classify fractures and other mechanical features with an average crack detection accuracy of ~94% and width classification accuracy of ~80%.

Simultaneous Material Segmentation and 3D Reconstruction in Industrial Scenarios

Zhao et al. present a novel transfer learning approach for material segmentation and categorization of RGB images, with a special focus on nuclear waste. The authors combine transfer learning with Recurrent Neural Networks to perform boundary-aware annotation and 3D semantic reconstruction. Additionally, the authors generated a new dataset that includes RGB image patches and fully pixel-wise annotated RGB images as a supplement to the public dataset *Materials in Context* (MINC).

CONCLUSIONS

The eight articles in this e-book show a wide range of state-of-the-art technologies and multidisciplinary approaches based on different use case scenarios. Each article shows the latest research progress and actively discusses the current technological problems. Through this article collection we can share the common issues and future perspectives. we expect this would contribute to make a breakthrough and promote optimal integration of systems from different fields of science and technologies.

AUTHOR CONTRIBUTIONS

All authors listed have contributed to the research topic and this editorial article substantially and intellectually. All agreed to the publication.

ACKNOWLEDGMENTS

We thank all authors and all reviewers for their valuable contributions.

Conflict of Interest: The author, RS, is a director of A.R.M Robotics Ltd.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

frontiers in Robotics and AI

ORIGINAL RESEARCH
published: 17 January 2020
doi: 10.3389/frobt.2019.00149

Radiological Mapping of Post-Disaster Nuclear Environments Using Fixed-Wing Unmanned Aerial Systems: A Study From Chornobyl

Dean T. Connor[1], Kieran Wood[2], Peter G. Martin[3], Sevda Goren[3], David Megson-Smith[1], Yannick Verbelen[3], Igor Chyzhevskyi[4], Serhii Kirieiev[4], Nick T. Smith[5], Tom Richardson[2] and Thomas B. Scott[3]*

[1] *Interface Analysis Centre, University of Bristol, Bristol, United Kingdom,* [2] *Aerospace Engineering, University of Bristol, Bristol, United Kingdom,* [3] *South West Nuclear Hub, University of Bristol, Bristol, United Kingdom,* [4] *SSE Ecocentre, Ministry of Ecology, Chornobyl, Ukraine,* [5] *National Nuclear Laboratory, Workington, United Kingdom*

In the immediate aftermath following a large-scale release of radioactive material into the environment, it is necessary to determine the spatial distribution of radioactivity quickly. At present, this is conducted by utilizing manned aircraft equipped with large-volume radiation detection systems. Whilst these are capable of mapping large areas quickly, they suffer from a low spatial resolution due to the operating altitude of the aircraft. They are also expensive to deploy and their manned nature means that the operators are still at risk of exposure to potentially harmful ionizing radiation. Previous studies have identified the feasibility of utilizing unmanned aerial systems (UASs) in monitoring radiation in post-disaster environments. However, the majority of these systems suffer from a limited range or are too heavy to be easily integrated into regulatory restrictions that exist on the deployment of UASs worldwide. This study presents a new radiation mapping UAS based on a lightweight (8 kg) fixed-wing unmanned aircraft and tests its suitability to mapping post-disaster radiation in the Chornobyl Exclusion Zone (CEZ). The system is capable of continuous flight for more than 1 h and can resolve small scale changes in dose-rate in high resolution (sub-20 m). It is envisaged that with some minor development, these systems could be utilized to map large areas of hazardous land without exposing a single operator to a harmful dose of ionizing radiation.

Keywords: radiation, Chornobyl, UAS (unmanned aircraft system), fixed-wing aerial surveys, post-disaster, cesium, nuclear, drones (UAV)

OPEN ACCESS

Edited by:
Chie Takahashi, University of Cambridge, United Kingdom

Reviewed by:
Jose Luis Sanchez-Lopez, University of Luxembourg, Luxembourg
Jonathan M. Aitken, University of Sheffield, United Kingdom

****Correspondence:***
Dean T. Connor
dean.connor@bristol.ac.uk

Specialty section:
This article was submitted to Robotic Control Systems, a section of the journal Frontiers in Robotics and AI

Received: *18 September 2019*
Accepted: *16 December 2019*
Published: *17 January 2020*

Citation:
Connor DT, Wood K, Martin PG, Goren S, Megson-Smith D, Verbelen Y, Chyzhevskyi I, Kirieiev S, Smith NT, Richardson T and Scott TB (2020) Radiological Mapping of Post-Disaster Nuclear Environments Using Fixed-Wing Unmanned Aerial Systems: A Study From Chornobyl. Front. Robot. AI 6:149. doi: 10.3389/frobt.2019.00149

1. INTRODUCTION

The large-scale release of radionuclides from the Chornobyl Nuclear Power Plant (ChNPP) remains the most significant nuclear accident in the history of civil nuclear power generation. During the 10 days of emissions from Reactor 4, approximately 11,780 PBq of radioactive material was released into the environment, including 1,700 PBq of 131I and 85 PBq of ^{137}Cs (Gudiksen et al., 1988; Cort et al., 1998; Smith and Beresford, 2005, p. 12). While the accident had far reaching environmental implications for a large area of Northern Europe, the area worst affected by the accident covers approximately 4,730 km^2 across modern day Ukraine (2,600 km^2) and Belarus (2,130 km^2). The area within Ukraine defines the Chornobyl Exclusion Zone (CEZ), which is an access controlled region, established in May 1986, designed to mitigate dose exposure to the public. Restrictions on access are still in place to this day, although tourism permits are currently available as part of official guided tours.

In the decades since the accident, there have been significant advancements in remote/automated characterization technologies that were not available to the responding forces in April 1986. Examples of these technologies include radiation-hard robotics systems, stand-off characterization systems and nuclear-focused unmanned aerial systems (UAS). In particular, UAS with radiation mapping capabilities have been shown on a number of occasions to provide excellent results when it comes to mapping radiation within post-disaster environments.

There are a number of formats of UASs that have been used to map radiation within the environment, these include helicopter-style systems (Towler et al., 2012; Furutani et al., 2013; Sanada and Torii, 2014), multi-rotor systems (MacFarlane et al., 2014; Martin et al., 2015, 2016; Burtniak et al., 2018; Connor et al., 2018a,b) and fixed-wing systems (Kurvinen et al., 2005; Pöllänen et al., 2009). Operating within the area affected by the Fukushima nuclear incident in 2011, Sanada and Torii (2014) demonstrated that UASs are capable of greatly improving the resolution of airborne radiation maps vs. the traditional method of using manned aircraft systems (MAS). Using a Yamaha-RMAX helicopter as a transport platform for a custom radiation mapping payload, several plumes of radioactivity were revealed that were not identifiable within the original manned aircraft survey. This result confirmed that UAS can offer greater detail within airborne radiometric surveys, albeit at the expense of absolute spatial coverage due to the increased time it takes to survey a comparable area.

As well as the heavier systems, like the RMAX presented in Sanada and Torii (2014) (94 kg), smaller UAS systems have been effectively demonstrated for a range of radiological mapping purposes. These include post-disaster environments (Martin et al., 2016; Burtniak et al., 2018; Connor et al., 2018a,b), indoor applications (Boudergui et al., 2011) and mapping naturally occurring radioactive material (NORM) (Martin et al., 2015; Šálek et al., 2018). A full review of UAS radiation mapping is provided within Connor et al. (2016). More recently, a lightweight, multi-rotor UAS was used to map radiation within the "Red Forest" region of the CEZ in especially high resolution (5 m pixel^{-1}). The system used within this study showed the capabilities of UASs to operate within even some of the most extreme radiological environments on earth with great effect. The ability of lightweight UASs to fly significantly lower than manned aircraft means that they can achieve improved spatial resolutions and sensitivities despite carrying much lighter (smaller-volume) payloads.

Terrestrial radiological measurements acquired from the air require a series of processing steps to convert the collected information into a true estimate of the activity of radioactive material present in the surveyed environment. Over approximately seven decades of operational experience of using manned aircraft, a defined series of processing steps have been established to accurately perform the correction of the raw measurements to a ground activity or dose-rate. This process is termed the "Spectral Windows" method and involves the segregation of the recorded gamma spectrum into a number of discrete energy windows that correspond to particular isotopes. The method is outlined in detail within Minty (1997) and can be used to map a range of radionuclides within the environment. However, these processing algorithms are optimized for large-volume (16–64 L) detectors that span gamma-ray energy ranges from 0 to 3 MeV, which are commonly used within manned aircraft surveys (International Atomic Energy Agency, 1991; Erdi-Krausz et al., 2003). The sensor packages used on unmanned aerial systems (UASs) are limited in terms of their volume (and hence detection efficiency) due to weight limits imposed by both physical and regulatory controls. This, in turn, limits certain key detection parameters that do not easily permit the use of all the steps within the standard processing algorithm encouraged by the IAEA. As increasing the payload capacity of UASs or reducing regulatory control is not often realistic or possible, alternative processing methods are required to overcome these hardware issues.

The major step within the workflow that small-volume systems struggle to overcome is in defining spectral stripping coefficients for removing scattered photons from erroneously appearing within lower energy spectral windows. This process occurs when an incident photon deposits only a portion of its energy to the detector medium before escaping the active volume without depositing further energy (Minty, 1997; Knoll, 2010). The occurrence is more pronounced in small-volume detectors due to the reduced interaction volume and causes the energy of the incident photon to be underestimated. The presence of high-energy emitters within the survey environment can therefore cause the intensity of lower-energy emitters to be overestimated. By using data acquired from well-characterized, infinite-yield sources (typically doped concrete pads) the magnitude of scattered photons falling within lower-energy windows can be determined through the definition of a series of simultaneous equations. These allow for the contribution to any energy window from one that exists at a higher energy to be removed, leaving only the true intensity behind. To correct for this, a signal within higher energy windows is required. This can be a problem for small-volume systems, as with limited detection volumes comes limited energy ranges. The contributions from abundant naturally occurring radioactive radionuclides (principally ^{238}U/^{235}U, ^{232}Th, and ^{40}K) need to be determined regardless of the primary target of the mapping procedure either for direct observation or for correction purposes. However, the characteristic spectral windows span energies in excess of 1.3 Mev (**Table 1**), which would require relatively large counting times or very intense sources to be able to accurately detect using small-volume detectors under aerial survey conditions. Therefore an alternative method for conducting this correction is required.

A further issue with utilizing small UASs into radiation mapping procedures is the limited battery life of the systems (Connor et al., 2016). This means that the absolute range of the aircraft is relatively low and therefore the potential area that can be covered in a single survey is small. This is highlighted in the research outlined above, as the investigations are limited to either relatively small sites, or present large data collection periods for more extensive areas. Increasing the range of radiation mapping UASs is therefore a necessary advancement for utilizing these platforms to monitor radiation over larger areas. At present,

TABLE 1 | Recommended spectral windows for radioelement mapping using manned aircraft.

Window	Nuclide	Energy range (MeV)
Total count	_	0.400–2.810
Cesium	^{137}Cs (0.662 MeV)	0.618–0.705
Potassium	^{40}K (1.460 MeV)	1.370–1.570
Uranium	^{214}Bi (1.765 MeV)	1.660–1.860
Thorium	^{208}Tl (2.614 MeV)	2.410–2.810

Modified from International Atomic Energy Agency (1991) and Erdi-Krausz et al. (2003).

manned-aircraft are the most utilized method of radiologically mapping large areas of land. However, there are a number of problems with using this method over utilizing UASs. Firstly, manned aircraft cannot operate at very low altitudes in the same way that UASs can due to regulatory restrictions. This dramatically reduces the spatial of the radiometric data collected by the platform (Connor et al., 2016). Furthermore, manned aicraft surveys are inherently more expensive to conduct than UAS surveys and still require the exposure of the pilot/crew to potentially harmful ionizing radiation. One of the methods to overcome the limitations presented by both multi-rotor UASs and manned aircraft is to utilize fixed-wing vehicles as the transport platform. Whilst designs and implementations of radiation mapping systems using these aircraft have been presented within the literature (Kurvinen et al., 2005; Pöllänen et al., 2009), they have never been deployed to map terrestrial radiation in non-controlled environments.

A potential reason for the lack of real-world implementations is that, from an operational perspective, using a fixed-wing UAS is more complicated than using an equivalent system incorporating a multi-rotor vehicle. This arises from the physical differences in the way each platform achieves flight and maneuvers in the air. Multi-rotor systems fundamentally act like helicopters with the ability to statically hover, maneuver along complex flight paths and take-off/land in confined spaces. Advances in autopilot technology effectively isolates the pilot from the vehicle behavior with commands in the form of simplified orthogonal motions. Conversely, fixed-wing systems require constant motion to stay aloft and typically require much larger open areas for take-off/landing. The pilot also needs to be more skilled and guide the aircraft using combined roll and pitch motions.

One of the advantages of the fixed-wing is its increased range and speed relative to multi-rotor systems. Take-off and landing zones do not need to be close to the proposed survey area due to the higher flight speeds resulting in short transit times. Assuming the survey area was well known to the operators (or sufficiently well-characterized by pre-flight and ground surveying methods), it would be entirely possible to operate a few kilometers away from a survey area. Being able to operate from a distance is a key parameter in responding to nuclear incidents, as it allows operators to be completely or partially removed from a potential hazard. At present, current safety regulations regarding flying beyond visual line of sight (BVLOS) make this type of operation difficult to implement routinely. However, it is possible to obtain exemptions for such flights. It is envisaged that the response to a nuclear incident would constitute a valid exemption to this restriction under a well planned safety case. The work presented within this study presents a new design for a lightweight fixed-wing radiation mapping UAS to monitor terrestrial radiation within post-disaster environments using a modified data processing procedure based on the "Spectral Windows" method. The study aims to demonstrate the advantages of deploying these systems to map relatively large areas of contaminated land in a short time frame within the CEZ (**Figure 1**) and discusses their suitability to the task vs. the traditional methods of mapping radiation in the environment. It is hypothesized that the utilization of these systems can bridge the gap in spatial resolution, aerial coverage and operating costs that exist between previous UAS surveys and surveys conducted using manned aircraft.

2. INSTRUMENTS AND METHODS

2.1. Fixed-Wing Mapping System

The UAS used within this investigation comprises a fixed-wing vehicle, an integrated radiation mapping payload, and associated ground support equipment. The vehicle was custom built at the University of Bristol based around the "Titan" airframe (Skywalker, Shenzhen, China) (**Figure 2**). The aircraft has a wingspan of 2.1 m and a take-off weight of 8.5 kg (1 kg payload). This particular system was advantageous because it could be hand-launched and recovered by parachute, therefore allowing deployment at any site with a reasonable clearing without a requirement for a runway. Power was provided by a 12.7 Ah, 6S 22.2 Lithium Polymer (LiPo) battery, giving an approximate flight duration of 50–70 min depending on the weather conditions during the flight.

The vehicle featured a full auto-pilot computer with supporting sensors [GNSS, barometric altitude, airspeed indicator (ASI), and IMU]. The autopilot was capable of navigating the aircraft along pre-planned waypoint missions. Three wireless links were used to interact with the vehicle during flight. The pilot safety link, operating on the 433 MHz frequency, was used for initializing the automatic flight and for manual control during the parachute landing. The second link was a bi-directional telemetry modem operating on the 868 MHz frequency and was used for monitoring of flight statistics (such as battery consumption) and to issue updated commands to the autopilot. The third link was a live FPV video stream from a camera in the nose of the aircraft operating on the 2.4 GHz frequency.

There are three internal cabins within the body of the "Titan". The first is the fuselage, which housed the control systems and batteries. The second is the payload bay, which contained the radiation sensor and its associated electronic control systems. The final cabin, located toward the tail of the aircraft, contained the parachute landing system and video transmitter.

The radiation sensing payload, supplied by Imitec (Bristol, UK), followed a similar design to the detection systems utilized within previous works by the authors (Martin et al., 2016, 2018; Connor et al., 2018a,b). However, modifications were made

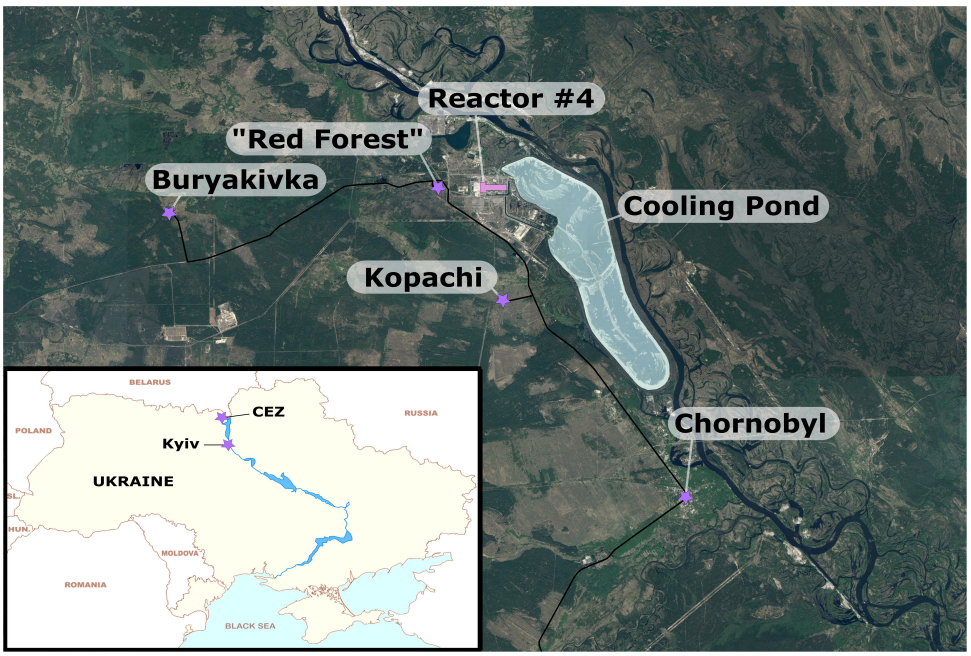

FIGURE 1 | Study locations within the exclusion zone. Buryakivka and Kopachi are abandoned settlements and the "Red Forest" is a natural area.

to the system design in anticipation of the new challenges posed by operating on a fixed-wing aircraft over a multi-rotor UAS. Whilst the previous system comprised of one solid state radiation detector, the updated system utilizes two SIGMA-50 CsI(Tl) scintillation detectors (Kromek Ltd., County Durham). This was implemented to increase the active detection volume of the system in order to offset the reduction in measured count-rates associated with operating at an increased altitude. An independent GNSS sensor allowed the radiation sensor payload to operate independently of the flying hardware with the exception of 12 V power supply. A secondary payload, comprising a higher power laser-range finder, was added by UOB to increase the fidelity of the height above ground measurements required for post-processing the radiation data. It must be noted that the data collected by the laser-range finder unit was not utilized within the processing in this study. This was due to practical considerations relating to the variable vegetation cover experienced across the surveyed area. The data from this system would have been used in an environment wherein vegetation cover does not extremely distort the apparent form of the land surface (e.g., an arid or urban environment).

2.2. Survey Locations

There are two separate zones in which the fixed-wing system was deployed within the CEZ. The first region lies between 12 and 13 km due west of the ChNPP [51.379 *N*, 29.916 *E*] (**Figure 3A**), near the former village of Buryakivka. This survey area covers 2.4 km^2 and bisects the main westward trending plume of radioactivity deposited from the accident. This area was used as an accessible initial test site to ensure the system behaved as expected within a lower-risk, but real-world, environment. The second region was significantly larger than the first, covering 12 km^2 of the land immediately to the west and south of the ChNPP Reactor 4 location. This area encompasses the "Red Forest" of Chornobyl [51.385 *N*, 30.051 *E*], as well as extending as far south as a former mechanical/farm yard near the settlement of Kopachi [51.345077 *N*, 30.111382 *E*] (**Figure 3B**).

2.3. Survey Methodology

2.3.1. Operational Restrictions

Operating UASs in real-world environments is subject to nationally varying sets of restrictions. These restrictions are in place to maintain airspace safety for all aerial users. Within the CEZ, all flights were conducted under standard visual line-of-sight (VLOS), meaning that direct and unaided visual contact with the aircraft was required at all times. It is not sufficient to be able to just physically see the aircraft within VLOS restrictions, instead the pilot must be able to ascertain the orientation of the vehicle and be able to avoid other air users and/or obstacles should they be present within the airspace. A number of take-off and land zones were therefore used to split the CEZ into several zones that could maximize coverage within the flight restrictions. No extra modifications were required to operate under these conditions. A minimum separation distance of 1 km from the New Safe Confinement Building (the old Reactor Four site) was also imposed in addition to the standard restrictions.

2.3.2. Flight Planning and Deployment

Each of the flights conducted using the fixed-wing system in the CEZ followed the same operational process. Once a target region was chosen, a small multi-rotor UAS (DJI Phantom 4 Pro) was deployed to investigate the area for obstacles and potential radio

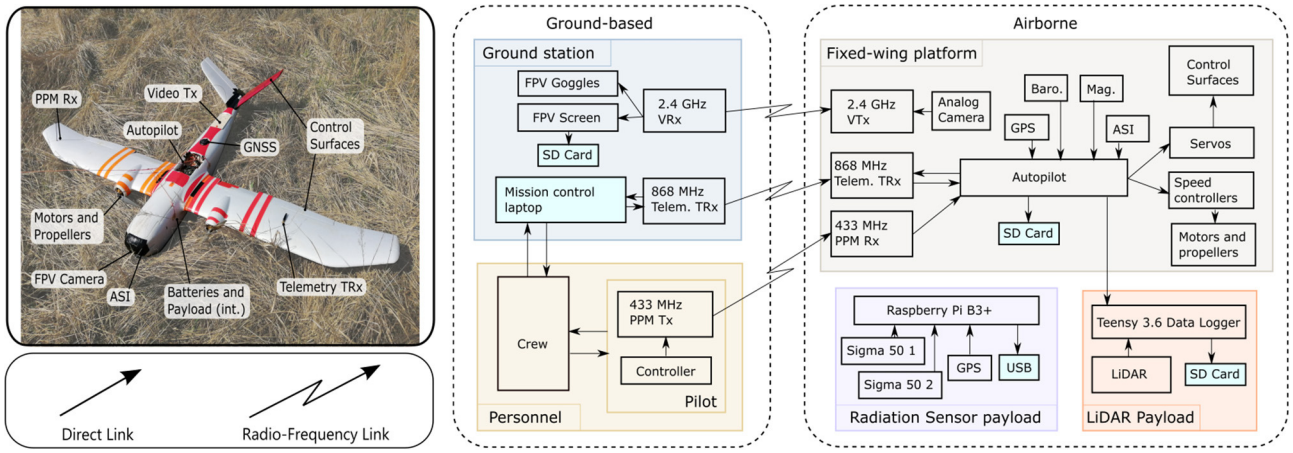

FIGURE 2 | The fixed-wing vehicle and sub-system interconnection diagram.

interference sources to determine the minimum altitude that the fixed-wing system can safely operate. The minimum altitude was set by adding a margin of 15 m to the measured altitude of the tallest structure present. Wherever possible, operating as low as reasonably practical was the target for every survey. Minor variations in the vegetation canopy height and the land surface throughout the zone resulted in survey altitudes of between 40 and 60 m above ground level (agl). Flight velocities were programmed to operate between 14 and 18 m s^{-1} airspeed, although local wind conditions created significant variations in the resulting ground-speed.

Once the flight altitude was decided, the flight paths were planned in a parallel-line raster scan format. The separation distance between each of the flight lines was equal to the survey altitude in all cases. At this flight line separation, the half-angle of the detectors produce a circle of investigation with a radius equal to the height above ground level of the aircraft (Martin, 2019, p. 77). Therefore, in the current flight planning configuration, the detector system is considerably oversampling the land; the measurement footprint from one flight line ends along the ground trace of the previous flight line. Whilst this is counter-intuitive to aiming to maximize spatial coverage, this compromise was necessary to demonstrate the range of the system whilst still maintaining VLOS with the aircraft. As the fixed-wing system has an inherent turning radius that is larger than the flight line separation, the parallel grid pattern was produced by flying in a series of laterally-offset loops rather than by flying each flight line in turn from one end of the survey to the other.

The surveys within the CEZ were completed semi-autonomously. The aircraft was hand-launched in automatic flight mode on take-off and remained in automatic mode until the landing phase of the flight. During the main flight phase, internal parameters (battery voltage, airspeed, ground-speed etc.) were monitored by the co-pilot using the ground station software. Whenever necessary, certain parameters were altered in-flight by the co-pilot under command from the pilot. Only the landing phase of the flight was conducted manually, with the pilot taking control of the aircraft upon its return to the take-off location. The automatic return to home function was programmed to bring the aircraft directly back to the take-off location whilst climbing to 70 m agl. Due to the lack of appropriate landing strips within the CEZ, the physical landing of the aircraft was accomplished through the deployment of a parachute.

2.4. Ground Investigation of Aerial Data

To further investigate the radiological measurements collected by the UAS, a series of ground-based measurements were collected within areas that were accessible to the ground team over the 6 days of data. These measurements were collected using a series of PED+ personal dosimeters (from Tracerco^{TM}) and ground-based versions of the fixed-wing mapping system presented in this study. The dosimeters were placed on the body of each of the team of six operators within the CEZ and continuously collected georeferenced dose-rate measurements every 2 s, allowing for aerial measurements to be correlated against ground measurements. The ground-based mapping systems were deployed as handheld devices that were carried by the operators throughout the target areas. Wherever safe and permissible, measurements were recorded by walking in a series of parallel grid lines to ensure optimal coverage throughout the area. More direct routes and quantification pathways were used if the dose-rates were considered high in order to minimize exposure to operators.

2.5. Calibrations and Data Processing

Using aerial platforms to collect measurements about the Earth's surface often requires a number of processing steps to correct the raw data to more appropriately reflect the original signal. This is because the collected signal may go through some elements of change during its travel between the source and the detector. The same is especially true for aerial gamma-ray measurements. The following procedure

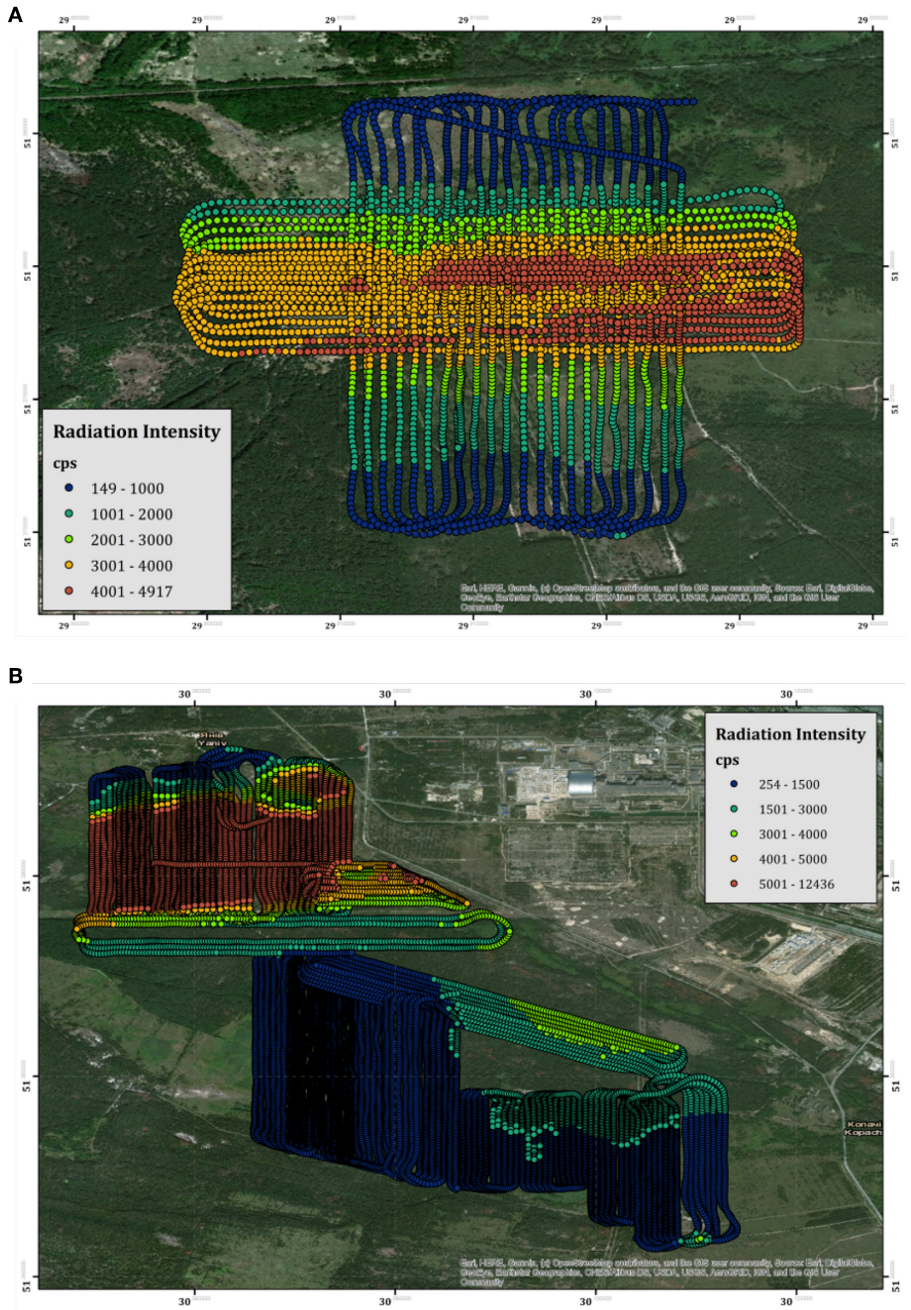

FIGURE 3 | A map showing all the measurements recorded by the fixed-wing system within **(A)** the Buryakivka testing region and **(B)** the region between the "Red Forest" and Kopachi. The points are color-scaled according to their unprocessed total count intensity.

outlines the process used within this study to convert the raw gamma-ray spectra obtained by the fixed-wing system into a cesium equivalent dose-rate (CED), in $\mu Sv\ h^{-1}$, at 1 m agl.

The raw 10 Hz measurements recorded by the system were first integrated into 1 Hz intervals and corrected for the dead time of the detector during the measurement. Then a correction for the cosmic, aircraft and radon background signal was applied to the data. This correction factor was determined by hovering the detection system at incrementally increasing altitudes over a wide meander bend of the River Pripyat, which is located <20 km from both the Buryakivka and Red Forest survey regions. The water within the river acts as an attenuating medium for the terrestrial signature of the total radiation flux within the surrounding area. With the terrestrial signal removed, the resultant recorded spectrum represents the contribution of all

other sources of radiation within the environment. This derived background signal was subtracted from every terrestrial spectrum recorded within the CEZ.

To keep integration times as short as possible and maximize the achievable spatial resolution of the surveys, a fast Fourier transform (FFT) was applied to the background-corrected spectrum to remove high-frequency noise. The appropriate filtering magnitudes/masks for this task were determined experimentally by exposing the detection systems to a small source of ^{137}Cs for intervals of a few seconds and visually inspecting the resultant unfiltered and filtered spectra. The energy range of the Sigma-50 detectors used within the payload is reported as 0 - 2 MeV. However the detection efficiency is significantly reduced beyond 1.2 MeV, meaning that observable signals within the potassium and uranium windows are unlikely even after FFT filtering. Hence, spectral stripping cannot be applied. Instead, a third-order polynomial estimation of the baseline of the region surrounding the ^{137}Cs spectral window was calculated and subtracted from the FFT spectrum, removing the Compton background below the peak and leaving only the direct measurements of ^{137}Cs.

Whilst the aircraft operates at a consistent altitude above the take-off location, the ground surface underneath the aircraft can vary significantly throughout the planned route. As a result, the separation distance between the detector and its target is also varying, which erroneously creates high or low artifacts within the recorded data set. To overcome this issue, the measurements were corrected to a consistent one meter agl using a correction curve derived from hovering the detection system at incrementally increasing altitude overland in the CEZ. This was conducted using a multi-rotor UAS over a fixed-point within a area of land presenting a uniform radiological signature throughout the field of view (FOV) of the detector at the highest altitude within the survey. Spectral information was collected for 90 s at while hovering at each of the 14 altitudes across the range 3 - 150m agl. The resulting spectra were split into four groups, delineated by the energy of the recorded photons (total counts, 0–0.2, 0.2–0.4, and 0.4–0.8 MeV), and the relationship between the altitude and energy-specific intensity was determined by fitting an exponential regression line in the form of Equation (1) (**Figure 4**).

$$I = I_0.e^{-kx} \quad (1)$$

In this relationship, I is the measured radiation intensity at altitude, x, agl, I_0 is the intensity at ground level ($x = 0$) and k is an experimentally derived constant encompassing the contributions of geometric dilution and attenuation by the atmosphere between the detector and the source.

The height agl of the system was obtained by subtracting the land surface height, obtained from an SRTM 30 m digital elevation model (NASA JPL, 2013) of the CEZ, away from the GPS altitude of the UAS. For this application of converting measurements to dose-rate at 1 m agl, correcting directly to ground level using the DEM was preferred over using the single-point laser range-finder. This is because the pulsed laser signal produced by the unit can be intercepted by surfaces (tall vegetation canopies in particular) before reaching the target ground surface, which can produce false altitude artifacts/inconsistent correction surfaces within the data. In an environment presenting less variability from vegetation cover, the inbuilt laser range-finder would have been used over the DEM-based correction.

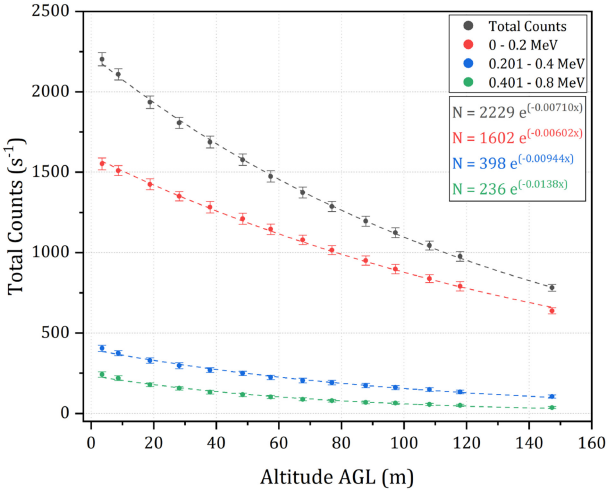

FIGURE 4 | A scatter plot depicting the results of the hover survey used to determine the altitude correction curves for the CEZ for all four energy groupings. The error bars for each measurement are equal to one standard deviation of the data.

Following the altitude correction of each measurement within the survey, the intensity of ^{137}Cs at 1 m agl are converted into a CED using a laboratory defined calibration. This conversion factor was determined by placing both the detection system and a PED+ personal dosimeter (from Tracerco^{TM}) at a range of distances away from several different sources of ^{137}Cs. In total, two sources were used (labeled RP5 and LRP10), varying in activity between 500 and 1,500 kBq. Background measurements were recorded for both sensors for 30 min and normalized with respect to live time. For the active measurements, both sensors were exposed to the source for 5 min at the same separation distance. The time series data recorded by the dosimeter and the net peak area of the ^{137}Cs peak were extracted and corrected for the background of the laboratory. The conversion factor between the CED and measured ^{137}Cs intensity is represented by the line of best fit of a scatter plot of the two variables (**Figure 5**).

3. RESULTS

3.1. Flight Characteristics and Duration

Over the 6 days of active fieldwork (17th April–24th April 2019), a total distance of 583.8 km was flown by the fixed-wing system, covering a total area of 14.8 km^2. **Table 2** shows the flight metrics for all flights conducted with the fixed-wing UAS. The system spent a total of 09h:17m:37s airborne, with an average flight time of 39m:50s. Buryakivka was the region in which the largest

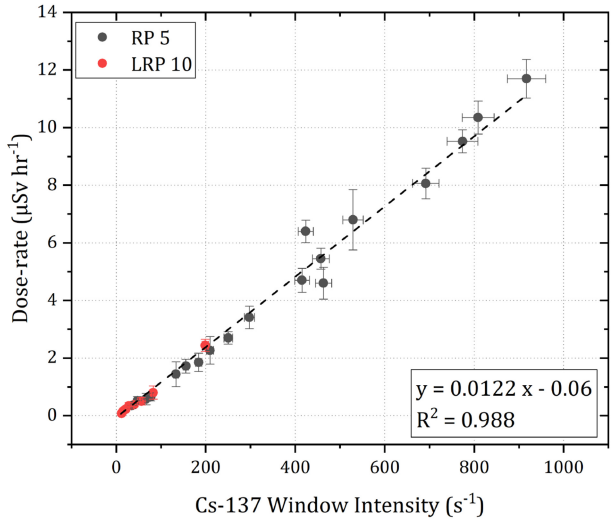

FIGURE 5 | A scatter plot detailing the relationship between the net peak counts of ^{137}Cs recorded by the detector and the dose-rate resulting from ^{137}Cs from two sources of different intensities.

TABLE 2 | Details of all of the flights conducted with the fixed-wing UAS within the CEZ.

Flight code	Date	Altitude (m)	Flight Length (km)	Duration (min:s)
Buryakivka Titan 1	17-04-2019	50	28.2	30:05
Buryakivka Titan 2	17-04-2019	45	28.0	31:51
Buryakivka Titan 3	17-04-2019	40	34.9	28:47
Buryakivka Titan 4	18-04-2019	40	42.7	39:02
Buryakivka Titan 5	18-04-2019	50	51.2	60:24
Buryakivka Titan 6	18-04-2019	50	51.7	48:42
Kopachi Titan 1	20-04-2019	50	52.9	52:23
Kopachi Titan 2	21-04-2019	50	60.1	65:58
Kopachi Titan 3	22-04-2019	50	39.7	30:11
Kopachi Titan 4	22-04-2019	50	39.1	30:36
Kopachi Titan 5	22-04-2019	60	39.8	30:31
RedForest Titan 1	23-04-2019	45	49.0	39:04
RedForest Titan 2	24-04-2019	45	44.6	39:23
RedForest Titan 3	24-04-2019	45	21.9	30:40

number of flights were conducted as this area was used as a semi-controlled testing zone. These test flights were used to fine tune all the aspects of the deployment procedure, including the pre-flight surveys, take-off and landing phases before operating closer to the ChNPP itself. The altitudes and flight velocities were varied slightly between each flight to reach the optimal level that the pilot was comfortable to operate subsequent surveys at. From these flights, it was determined that the optimal flight velocity was between 14 and 17 ms^{-1} airspeed, although in real-terms, the ground-speed value varied between individual flight paths within the same survey due to local wind conditions.

The largest average distance traveled in a single flight by the UAS was within the Kopachi region of the map [51.371 N, 30.065 E: 51.434 N, 30.114 E] with a value of 46.3 km. This region is dominated by open fields featuring small amounts of vegetation of <5 m in height, therefore presenting the most optimal conditions for all flight phases. The open fields provided space to take-off and land in any direction and permitted excellent visibility to the aircraft during the in-flight phase, allowing line-of-sight to be maintained easily. The longest distance covered was also in Kopachi, totalling 60.1 km in length. This value describes roughly the total distance that can be covered safely by the UAS given the current battery technology available to operators. If possible, all the flights used within the survey would be closer to this upper limit, but due to line-of-sight restrictions this was not possible in all parts of the surveyed areas.

3.2. Radiological Monitoring

3.2.1. Buryakivka

The results of the derived CED for the Buirakivka survey area are presented within **Figure 6**. The map within this figure is produced from three flights conducted at 40–45 m altitude agl, flights conducted at more elevated altitudes during the testing process have not been included within the map as many of these cover the same areas. An inverse distance weighting (IDW) interpolation algorithm has been applied to the data to produce the color-scaled CED overlay, which is presented at a pixel size of 20 × 20 m. This resolution was chosen as it is slightly coarser than the inline point spacing of the data set. The overall trend of the map follows the expected pattern from previous soil sampling investigations as presented within Kashparov et al. (2018), exhibiting a contaminant plume trending east to west, which drops off in intensity to the immediate north and south of the central line. The maximum CED measured within this area is 4.65 μSv h^{-1}, measured at 51.363198 N, 30.107020 E, which is more than 23 times greater than the average total background dose-rate of the UK (0.2 μSv h^{-1}).

3.2.2. Red Forest and Kopachi

The measured CED for the region surrounding the ChNPP is presented within **Figure 7**. The combined survey amalgamates the data from seven flights conducted over 4 days of deployment. Contrary to the data collected within the Buryakivka region (section 3.2.1), all the surveys conducted within this area are included within the presented data set (see **Table 2** for full flight details). The color-scaled overlay is once more presented at a pixel size of 20 × 20 m. As expected, the overall CED measured in the area surrounding the ChNPP is significantly larger than that measured in Buryakivka. The maximum CED successfully recorded by the fixed-wing system was 12.8 μSv h^{-1}, which is 2.8 times greater than the maximum CED recorded within the Buryakivka region. The map shows two main areas displaying elevated dose-rates. The first is a sharply delineated hot spot that extends immediately to the west of the ChNPP itself and covers the "Red Forest" area [51.379 N, 30.071 E]. The second is a much broader zone of elevated intensity, extending southwards from the plant toward the village of Kopachi [51.366 N, 30.100 E]. This overall trend is also depicted within the soil sampling

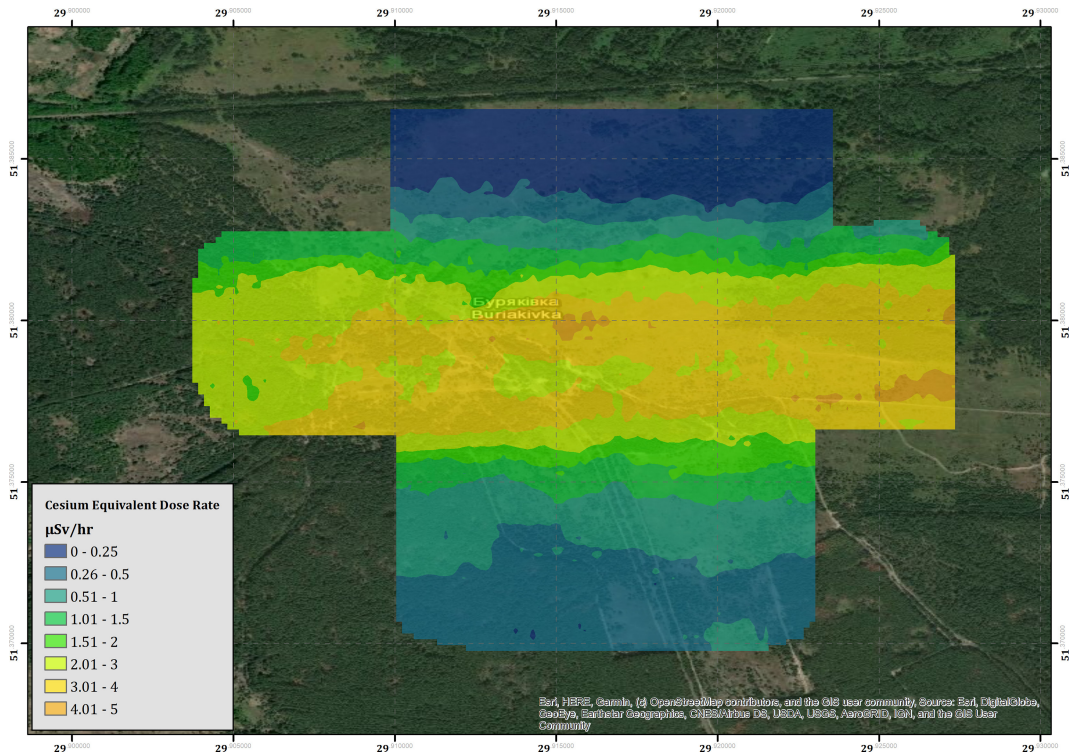

FIGURE 6 | The cesium equivalent dose-rate (CED) of the Buryakivka area.

investigations previously conducted by Kashparov et al. (2018), showing a general agreement between this dataset and previously published works from other institutions.

Located at the south-eastern corner of the area is a region of elevated dose-rate (3.3 μSv h^{-1}) that lies within an area of relatively low dose-rate (0.51–1.0 μSv h^{-1}). The hot spot [51.343843 N, 30.110399 E] manifests in an almost idealized point-source geometry when compared to the broad spreading of radioactivity evident within the measurements collected in the area surrounding it. The shape and location of the hot spot suggest that its presence is the result of anthropogenic concentration of radioactivity rather than the natural deposition following the accident. Dose-rate information could not be extracted from the cross-hatched area within **Figure 7** due to detector saturation issues.

4. DISCUSSION

4.1. Radiological Monitoring

Using the fixed-wing UAS, 15 km^2 of the CEZ was successfully mapped, detailing the dose-rate variation relating to the ^{137}Cs concentration in the ground and other surface features. Unfortunately, the high-intensity plume extending over the "Red Forest" could not be presented in its entirety as its significant radiological fingerprint caused an overloading (saturation) of the electronics of the detector. Whilst total count data could still be recorded, the measured gamma spectrum experienced significant degradation when exposed to total count rates of over 5,500 counts per second (cps). This reduction in signal quality manifests in the form of shifting the ^{137}Cs peak toward a lower gamma energy (i.e., shifting the peak to the left) and significantly increasing the full width at half maximum (FWHM) value. This leads to non-sensical values within the analysis, which have been redacted from the map presented within this study. Detailed analysis of the spectral data indicates that reliable information can be extracted from the detector up to around 5,250 cps, as the peak positions and FWHM values remain within the manufacturers tolerance. Any measurements with a total count rate of greater than 5,250 cps have been removed from the presented dataset.

Previous surveys have measured dose-rates within a small portion of the "Red Forest" area to be up to around 170 μSv h^{-1} (Burtniak et al., 2018). These surveys were conducted within the portion of the "Red Forest" that could not be mapped by our system at much lower altitudes (5 m) and much slower velocities that are typical of multi-rotor surveys. Despite being inherently unreliable, the total-count data recorded by the fixed-wing system (**Figure 3**) reported a maximum count-rate of 12,436 cps at 45 m altitude. Even though the measurements were saturated, using this count-rate as a minimum value for the radiological intensity within this area would produce an expected dose rate of at least 95 μSv h^{-1} (based upon the approximate ratio of the altitude corrected total intensity to cesium dose-rate). As the detector is facing an overload during these measurements, the real total counts value would most likely be greater, producing a larger CED.

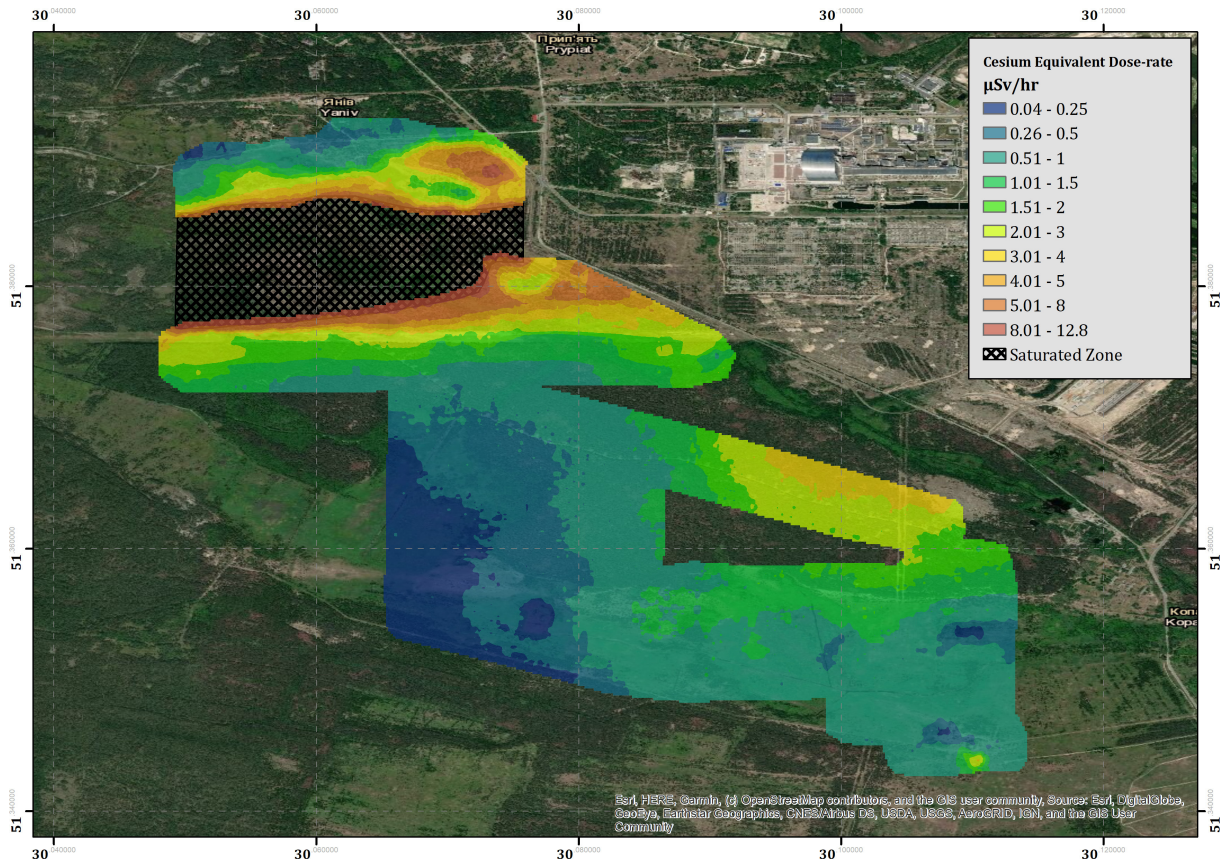

FIGURE 7 | The cesium equivalent dose-rate (CED) of the "Red Forest" area surrounding the ChNPP.

Overall, the general results from the UAS agrees well with previous datasets collected though other methods. The extensive soil sampling investigations published by Kashparov et al. (2018) provide excellent overall coverage and measurement accuracy throughout the CEZ, but fail to provide an easily repeatable method of monitoring radiation within post-disaster environments. The amount of labor-hours involved in conducting ground sampling surveys of this size are significant and the results are comparatively low-resolution when assessed against more mobile methods. The effects of this are best shown by the localized hot spot present near the south-eastern corner of the map [51.343843 N, 30.110399 E], which was previously unreported in literature until this study was conducted.

Following the identification of the hot spot from the raw data collected by the aircraft, a ground-based team was deployed to investigate the area covered by the elevated intensity region. Upon arrival, this team used SIGMA-50 detectors, Geiger-Muller (GM) tubes and PED+ personal dosimeters to monitor the radioactive output of this region. The source of the radioactivity was determined to be a series of funnel-shaped metal structures that seem to have been used to mechanically sort through material in an attempt to reduce the overall volume of contaminated material following the accident (**Figure 8**). These structures will be referred to as "hoppers" for the remainder of this document. The residual radiological fingerprint of this process is significant. Ground-measurements, acquired using dosimeters, measured more than 2 mSv h^{-1} directly in the vicinity of the "hoppers." Attempts at recording gamma spectrometry measurements were futile due to saturation issues.

The measurements collected by the aircraft at this point in space are significantly lower than the values measured by the ground team (3.3 μSv h^{-1} vs. 2 mSv h^{-1}). There may be a number of reasons for the discrepancy between these values. Firstly, the analysis performed on the results collected by the aircraft focuses solely on the ^{137}Cs signal, ignoring contributions from any other radionuclides (these are outside the scope of this study and will be investigated in future studies). The myriad of radioactive material released from the accident is highly complex and the measured contribution of ^{137}Cs is but a component of the total output (Smith and Beresford, 2005; Burtniak et al., 2018). Given that the "hopper" hot spot is so intensely radioactive, the on-ground measurements could be recording inputs from other radionuclides in addition to the measured ^{137}Cs signal. This could potentially include gamma-ray signals from ^{241}Am, which emits a low energy gamma-ray (0.06 MeV) that is more easily attenuated by the medium between the source and the detector (see **Figure 4**). These

kinds of signals are difficult to detect with any confidence at the altitudes used within this survey, especially because the incomplete transfer of energy between incoming photons and the detection crystal (very common for small-volume, room temperature detection systems Gilmore, 2008) creates a high background signal at the low energy range of the spectrum. Radionuclides other than ^{241}Am and ^{137}Cs are also expected to be present within the signal emanating from this region, including contributions from fission products from spent nuclear fuel.

Another reason for the observed disparity between the aerial measurements and the on-ground measurements is the short sampling time and FOV averaging at each point on the earth. As the system operates at a minimum velocity of 14 ms^{-1}, the measured signal consists of a sampling area of 1,120 m^2 for a single, 1-s measurement before the altitude correction is applied. Highly localized variations of a few meters in area will therefore not be fully resolved, instead being averaged with the surrounding area that constitutes a single measurement. This provides one possible explanation into the differences in expected dose-rate and measured dose-rate within this study. Once more, it must be noted that the dose-rate measured here considers only the dose-rate from ^{137}Cs, whereas raw ground measurements recorded by personal dosimeters are collecting information from all sources (including natural radionuclides and other sources released during the accident). The time between the occurrence of the accident and this study is also an important consideration. In the three decades since the accident, radioactive material has had time to penetrate the ground surface and new sediment has had time to be deposited on top of the original radionuclide deposition. The burial of material means that there is not only more attenuating material between the detector and its target source, but the material is also much denser than air, resulting in fewer interactions with unscattered ^{137}Cs photons with the active detection volume.

The overall mission objective to deploy the fixed-wing system within a real-world post-disaster environment has been achieved. As far as the authors are aware, this is the first time that this variation of a radiation mapping UAS has been deployed in a non-controlled situation to map terrestrial gamma radiation. Previous iterations of this detection unit for multi-rotor systems have been successfully utilized in similar environments within the Fukushima fallout zone in Japan, albeit these zones have been less intense than the levels of radioactivity experienced in the proximity of the "Red Forest." One of the advantages of using lightweight UASs is that payloads can be altered, or completely removed, with ease. Given the observed limitations of the detectors used within the system at relatively high gamma fluxes, these would likely be changed in future iterations. Cerium Bromide ($CeBr_3$) and Lanthanum Bromide ($LaBr_3$) detection systems are being considered for future systems as these provide excellent energy resolutions and optical yields, even at small-volumes (Lowdon et al., 2019).

In terms of short-term improvements, one of the two Sigma-50 units will be swapped for a smaller-volume GR1 unit (a CZT semi-conductor detector from Kromek Group PLC, County Durham, UK). This detector has a better tolerance for high gamma fluxes and an improved energy resolution when compared to the Sigma-50 unit used within the current system. It is however, more vibrationally sensitive and will require some efforts to dampen these effects within the UAS payload. A further survey specifically aiming to map the saturated zone of the "Red Forest" (see **Figure 7**) using the updated system is planned for October 2019 to improve upon the results collected herein.

4.2. System Evaluation and Wider Applications

The results from the radiological investigations of the CEZ suggest that the fixed-wing system presented within this study is effective at mapping ^{137}Cs distribution, although the significant radioactivity of the "Red Forest" proved to be too much for the detectors used within the payload. Whilst this conclusion satisfies the overall aim of this study, there are a few more considerations to be discussed before the system can be considered for use in more routine situations or be implemented into emergency procedures in the future.

UAS-based investigations are often at the mercy of the weather. Certain counter-measures can be implemented in some cases to overcome problems, for example, waterproofing the central electronic components can allow certain types of UASs to operate even in wet weather. However, the fixed-wing UAS used in this study is sensitive to variations in localized wind velocities. The surveys were conducted at a target velocity 14–18 ms^{-1}, but on some occasions ground-speeds of up to 25 ms^{-1} were recorded during survey lines orientated such that the aircraft experienced a tailwind. As previously mentioned within section 2.3, the differences in wind velocities experienced by the UAS during individual legs of the same survey create inconsistencies in recorded data. The system records raw data at 10 Hz before being resampled in the post-processing phase into 1 Hz intervals. Differences in the velocity of the aircraft mean that the effective sampling area of each measurement varies throughout the survey. As a result, the pixel size has been increased slightly to a lower resolution in the final map to encompass some of this variation.

Even though UASs with similar or greater ranges have been reported within the literature, the fixed-wing UAS used within this study is considerably lighter than these reported platforms. The Yamaha RMAX platform utilized within Sanada and Torii (2014) weighs 100.5 kg with the under-mounted radiation mapping payload attached, whilst the system presented herein weighs 8.5 kg by comparison. Whilst the radiation mapping payload used within Sanada and Torii (2014) was able to carry a larger payload (resulting in larger detection volumes), the extra weight is significant in terms of the operation of UAS in the real-world. Regulatory restrictions exist for the operation of heavier platforms around the world as they present a greater hazard to the environment through the increased energy involved in an impact (Connor et al., 2016). As a result, it is easier to deploy lighter platforms within surveys globally.

FIGURE 8 | Aerial photographs of the "hopper" hot spot, presented in a plan view **(Left)** and a areal view in context of the local area **(Right)**.

The concept of utilizing fixed-wing vehicle for this task aims to bridge the gap between manned aircraft and multi-rotor capabilities. Currently, multi-rotor UASs represent the high-resolution end-member of the airborne radiation mapping spectrum, often achieving sub-10 m pixel sizes (Martin et al., 2015, 2016; Burtniak et al., 2018; Connor et al., 2018a,b). Manned aircraft systems (MASs) represent the opposite end of the scale, operating at between 90 and 200 m agl and achieving spatial resolutions of 200–500 m (Pitkin and Duval, 1980; Sanderson and Cresswell, 2008). Within this survey, operating altitudes were maintained at 40–60 m agl, with a spatial resolution of 20 m after the post-processing procedure. This successfully provides a middle ground between the two end-members, both in terms of the resolution and total coverage capabilities of the system. It is worth noting that the flight line separation can be increased if the survey values absolute coverage over spatial resolution. As the detector FOV increases linearly with altitude, the flight line spacing could be increased up to two times the altitude of flight without incurring a loss of net spatial coverage.

As well as bridging the gap between current methods, there is also the potential that manned aircraft could be superseded by using fixed-wing UASs in the future, especially with sufficient improvements in battery technology. This is especially true when considering financial factors. The total cost of building and deploying the UAS used herein was $24,000 (including all parts, labor costs for build and deployment and insurance costs), whereas a manned survey would be considerably more expensive. The cost for repeat surveys following the initial investment totals at $9,000 for the equivalent survey conducted within this study. This is based on salary estimates and operational costs for a three man crew over 6 days of active operation. If the equipment is used multiple times, the cost-benefit of the system is significantly improved over utilizing manned aircraft.

Without much prior familiarity of operating within the CEZ, the fixed-wing system was successfully deployed at as low an altitude as reasonably possible using information obtained from on the fly pre-flight surveys. With a good knowledge of the survey area, it would be possible to achieve much more. Overall, the authors believe that there is extreme promise in widely utilizing these systems for a number of survey applications in the future after the implementation of the improvements suggested herein.

5. CONCLUSIONS AND FUTURE WORK

This study presents the most comprehensive radiation map of the CEZ ever produced from a UAS. Over the 6 days of active fieldwork with the fixed-wing system, 15 km^2 was investigated in a high spatial resolution (20 m $pixel^{-1}$). In total, more than 580 km were flown across the region in a total flight time of 09h:17m:37s. The system demonstrated that radiation mapping investigations using UASs can be launched from safe-zones outside contaminated regions and operated continuously for more than an hour before returning to the safe-zone to land. Some previous systems presented within the literature have been required to launch within the contaminated zones at the risk of the operators.

Due to issues with detector saturation, the area in the most radiologically intense portions of the "Red Forest" were not presented with the mapped CED as this information could not be reliably extracted from the spectra recorded over this area. This problem is hoped to be solved in the future by operating using a different dual detector set-up (Sigma-50 and GR1 combination). In this configuration, the Sigma-50 would be used to map the areas displaying lower contamination concentrations (as presented by more than 85% of the area mapped in this study) and the GR1 would be used to map the areas wherein the Sigma-50 was saturated.

One of the most interesting findings was the presence of the previously unreported, anthropogenically-enhanced hot spot located in the south-eastern corner of the surveyed area. With the knowledge that the 2 mSv h^{-1} hot spot exists, a coordinated ground sampling investigation will be conducted to determine the nature of the radionuclide content and correlated against the measurements collected by the aerial platform.

The work conducted within the CEZ was part of a multi-faceted field investigation using numerous types of radiation monitoring methods. These included both fixed-wing and low-altitude multi-rotor UAS surveys, as well as ground-based monitoring methods using both tracked robots and humans. The data presented herein will be combined with the measurements recorded using the other methods in future works to complete a comprehensive radiological survey of the CEZ using mobile radiation monitoring methods. The demonstration of this system

in this environment has further-reaching consequences than just the monitoring of post-disaster environments. With alterations to the included detection systems, using recommendations from Lowdon et al. (2019), this system could become a low-cost solution to monitoring large areas of land for mineral resources. This could be of particular interest to developing countries who currently struggle to conduct mineral reserve estimates due to the high expenditure involved in chartering manned-aircraft surveys. Further work within the CEZ is planned for October 2019 and April 2020.

DATA AVAILABILITY STATEMENT

The datasets generated for this study are available on request to the corresponding author.

AUTHOR CONTRIBUTIONS

Data collection was conducted by DC, KW, PM, YV, SG, IC, and TS. Continual aircraft development was conducted by KW and TR, with the piloting of the aircraft performed by KW. Data processing was conducted by DC with assistance from KW and DM-S, with necessary local information provided by IC and SK. Original manuscript preparation was conducted by DC. Further manuscript developments and improvements were performed by DC, KW, PM, TS, DM-S, SG, IC, TR, SK, and NS.

FUNDING

The authors would like to thank UK Research and Innovation (UKRI) for their support and funding which has come via the National Centre for Nuclear Robotics (NCRN); grant EP/R02572X/1). DC acknowledges funding from the EPSRC (EP/N509619/1) and the NNL (NNL/UA/037). NS is supported by a Royal Society Industry Fellowship and the NNL IR&D Programme.

ACKNOWLEDGMENTS

Special thanks are extended to all the members of the SSE Ecocentre who supervised and assisted the authors day to day within the CEZ. The authors would also like to extend sincere gratitude to the two reviewers of the manuscript whose insightful comments have greatly aided the development of the manuscript.

REFERENCES

Boudergui, K., Carrel, F., Domenech, T., Guenard, N., Poli, J.-P., Ravet, A., et al. (2011). "Development of a drone equipped with optimized sensors for nuclear and radiological risk characterization," in *2011 2nd International Conference on Advancements in Nuclear Instrumentation, Measurement Methods and their Applications* (Ghent), 1–9.

Burtniak, V., Zabulonov, Y., Stokolos, M., Bulavin, L., and Krasnoholovets, V. (2018). Application of a territorial remote radiation monitoring system at the Chornobyl nuclear accident site. *J. Appl. Remote Sens.* 12:046007. doi: 10.1117/1.JRS.12.046007

Connor, D., Martin, P. G., and Scott, T. B. (2016). Airborne radiation mapping: overview and application of current and future aerial systems. *Int. J. Remote Sens.* 37, 5953–5987. doi: 10.1080/01431161.2016.1252474

Connor, D. T., Martin, P., Smith, N., Payne, L., Hutton, C., Payton, O., et al. (2018a). Application of airborne photogrammetry for the visualisation and assessment of contamination migration arising from a Fukushima waste storage facility. *Environ. Pollut.* 234, 610–619. doi: 10.1016/j.envpol.2017.10.098

Connor, D. T., Martin, P. G., Pullin, H., Hallam, K. R., Payton, O. D., Yamashiki, Y., et al. (2018b). Radiological comparison of a FDNPP waste storage site during and after construction. *Environ. Pollut.* 243, 582–590. doi: 10.1016/j.envpol.2018.08.099

Cort, M. D., Dubois, G., Fridman, S. D., Germenchuk, M. G., Izrael, Y. A., and Janssens, A. (1998). *Atlas of Caesium Deposition on Europe after the Chernobyl Accident.* Luxembourg: Office of official publication of the european communities.

Erdi-Krausz, G., Matolin, M., Minty, B., Nicolet, J., Reford, W., and Schetselaar, E. (2003). *Guidelines for Radioelement Mapping Using Gamma Ray Spectrometry Data*. Vienna: IAEA (International Atomic Energy Agency).

Furutani, T., Uehara, K., Tanji, K., Usami, M., and Asano, T. (2013). "A study on micro-scale airborne radiation monitoring by unmanned aerial vehicle for rural area reform contaminated by radiation," in *Proceedings of the Disaster Management* (A Coruña), 1–9.

Gilmore, G. (2008). *Practical Gamma-Ray Spectrometry*. Chichester, UK: John Wiley & Sons.

Gudiksen, P., Harvey, T., and Lange, R. (1988). Chernobyl source term, atmospheric dispersion, and dose estimation. *Health Phys.* (1989) 57, 697–706.

International Atomic Energy Agency (1991). *Airborne Gamma Ray Spectrometer Surveying: Technical Reports Series No. 323*. Technical report.

Kashparov, V., Levchuk, S., Zhurba, M., Protsak, V., Khomutinin, Y., Beresford, N. A., et al. (2018). Spatial datasets of radionuclide contamination in the Ukrainian Chernobyl Exclusion Zone. *Earth Syst. Sci. Data* 10, 339–353.

Knoll, G. F. (2010). *Radiation Detection and Measurement, 3rd Edn.* New York, NY: John Wiley & Sons.

Kurvinen, K., Smolander, P., Pöllänen, R., Kuukankorpi, S., Kettunen, M., and Lyytinen, J. (2005). Design of a radiation surveillance unit for an unmanned aerial vehicle. *J. Environ. Radioactivity* 81, 1–10. doi: 10.1016/j.jenvrad.2004.10.009

Lowdon, M., Martin, P., Hubbard, M., Taggart, M., Connor, D., Verbelen, Y., et al. (2019). Evaluation of scintillator detection materials for application within airborne environmental radiation monitoring. *Sensors* 19:3828. doi: 10.3390/s19183828

MacFarlane, J. W., Payton, O. D., Keatley, A. C., Scott, G. P. T., Pullin, H., Crane, R. A., et al. (2014). Lightweight aerial vehicles for monitoring, assessment and mapping of radiation anomalies. *J. Environmental Radioactivity* 136, 127–130.

Martin, P. (2019). *The 2011 Fukushima Daiichi Nuclear Power Plant Accident: An Analysis From The Metre to The Nano-Metre Scale.* Springer Ltd.

Martin, P., Connor, D., Payton, O., Leal-Olloqui, M., Keatley, A., and Scott, T. (2018). Development and validation of a high-resolution mapping platform to aid in the public awareness of radiological hazards. *J. Radiol. Protect.* 38, 329–342. doi: 10.1088/1361-6498/aaa914

Martin, P., Kwong, S., Smith, N., Yamashiki, Y., Payton, O., Russell-Pavier, F., et al. (2016). 3D unmanned aerial vehicle radiation mapping for assessing contaminant distribution and mobility. *Int. J. Appl. Earth Observ. Geoinform.* 52, 12–19. doi: 10.1016/j.jag.2016.05.007

Martin, P., Payton, O., Fardoulis, J., Richards, D., and Scott, T. (2015). The use of unmanned aerial systems for the mapping of legacy uranium mines. *J. Environ. Radioactivity* 143, 135–140. doi: 10.1016/j.jenvrad.2015.02.004

Minty, B. (1997). Fundamentals of airborne gamma-ray spectrometry. *AGSO J. Austral. Geolo. Geophys.* 17, 19–50.

NASA JPL (2013). *NASA Shuttle Radar Topography Mission Global 1 Arc Second Number.* NASA EOSDIS Land Processes DAAC.

Pitkin, J. A. and Duval, J. S. (1980). Design parameters for aerial gamma-ray surveys. *Geophysics* 45, 1351–1441.

Pöllänen, R., Toivonen, H., Peräjärvi, K., Karhunen, T., Ilander, T., Lehtinen, J., et al. (2009). Radiation surveillance using an unmanned aerial vehicle. *Appl. Radiat. Isotopes* 67, 340–344. doi: 10.1016/j.apradiso.2008.10.008

Šálek, O., Matolín, M., and Gryc, L. (2018). Mapping of radiation anomalies using UAV mini-airborne gamma-ray spectrometry. *J. Environ. Radioactivity* 182, 101–107.

Sanada, Y., and Torii, T. (2014). Aerial radiation monitoring around the Fukushima Dai-ichi nuclear power plant using an unmanned helicopter. *J. Environ. Radioactivity* 139, 294–299.

Sanderson, D., and Cresswell, A. (2008). The effect of flight line spacing on radioactivity inventory and spatial feature characteristics of airborne gamma-ray spectrometry data. *Int. J. Remote Sens.* 29, 31–46. doi: 10.1080/01431160701268970

Smith, J., and Beresford, N. (2005). *Chernobyl: Catastrophe and Consequences*. Chichester, UK: Springer.

Towler, J., Krawiec, B., and Kochersberger, K. (2012). Terrain and radiation mapping in post-disaster environments using an autonomous helicopter. *Remote Sens.* 4, 1995–2015. doi: 10.3390/rs4071995

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

frontiers in Robotics and AI

CORRECTION
published: 28 February 2020
doi: 10.3389/frobt.2020.00030

Corrigendum: Radiological Mapping of Post-Disaster Nuclear Environments Using Fixed-Wing Unmanned Aerial Systems: A Study From Chornobyl

Dean T. Connor[1], Kieran Wood[2], Peter G. Martin[3], Sevda Goren[3], David Megson-Smith[1], Yannick Verbelen[3], Igor Chyzhevskyi[4], Serhii Kirieiev[4], Nick T. Smith[5], Tom Richardson[2] and Thomas B. Scott[3]*

[1] Interface Analysis Centre, University of Bristol, Bristol, United Kingdom, [2] Aerospace Engineering, University of Bristol, Bristol, United Kingdom, [3] South West Nuclear Hub, University of Bristol, Bristol, United Kingdom, [4] SSE Ecocentre, Ministry of Ecology, Chornobyl, Ukraine, [5] National Nuclear Laboratory, Workington, United Kingdom

Keywords: radiation, Chornobyl, UAS (unmanned aircraft system), fixed-wing aerial surveys, post-disaster, cesium, nuclear, drones (UAV)

OPEN ACCESS

Edited and reviewed by:
Chie Takahashi,
University of Cambridge,
United Kingdom

****Correspondence:***
Dean T. Connor
dean.connor@bristol.ac.uk

Specialty section:
This article was submitted to Robotic Control Systems, a section of the journal Frontiers in Robotics and AI

Received: *11 February 2020*
Accepted: *20 February 2020*
Published: *28 February 2020*

Citation:
Connor DT, Wood K, Martin PG, Goren S, Megson-Smith D, Verbelen Y, Chyzhevskyi I, Kirieiev S, Smith NT, Richardson T and Scott TB (2020) Corrigendum: Radiological Mapping of Post-Disaster Nuclear Environments Using Fixed-Wing Unmanned Aerial Systems: A Study From Chornobyl. Front. Robot. AI 7:30. doi: 10.3389/frobt.2020.00030

A Corrigendum on

Radiological Mapping of Post-Disaster Nuclear Environments Using Fixed-Wing Unmanned Aerial Systems : A Study From Chornobyl

by Connor, D. T., Wood, K., Martin, P. G., Goren, S., Megson-Smith, D., Verbelen, Y., et al. (2020). Front. Robot. AI 6:149. doi: 10.3389/frobt.2019.00149

In the original article, there was a mistake in **Figure 6** and **Figure 7** as published. The published dose-rates are incorrect due to a typing error within the code used to process the radiation data. Instead of correcting the measurements for the live time of the detector (TL), the typing error caused the intensity measurements to be adjusted according to the uncorrected sample time (TR). This affects the dose-rate magnitude of all the measurements presented but does not affect the trend or the reliability of the dataset. The corrected **Figure 6** and **Figure 7** appears below.

A correction has been made to Section 3.2 (Radiological Monitoring), Section 3.2.1 (Buryakivka), Paragraph 1:

"The results of the derived CED for the Buirakivka survey area are presented within **Figure 6**. The map within this figure is produced from three flights conducted at 40–45 m altitude agl, flights conducted at more elevated altitudes during the testing process have not been included within the map as many of these cover the same areas. An inverse distance weighting (IDW) interpolation algorithm has been applied to the data to produce the color-scaled CED overlay, which is presented at a pixel size of 20 × 20 m. This resolution was chosen as it is slightly coarser than the inline point spacing of the data set. The overall trend of the map follows the expected pattern from previous soil sampling investigations as presented within Kashparov et al. (2018), exhibiting a contaminant plume trending east to west, which drops off in intensity to the immediate north and south of the central line. The maximum CED measured within this area is 4.65 μSv h^{-1}, measured at 51.363198 N, 30.107020 E, which is more than 23 times greater than the average total background dose-rate of the UK (0.2 μSv h^{-1})."

A correction has also been made to Section 3.2 Radiological Monitoring, Section 3.2.2 Red Forest and Kopachi, Paragraphs 1 and 2:

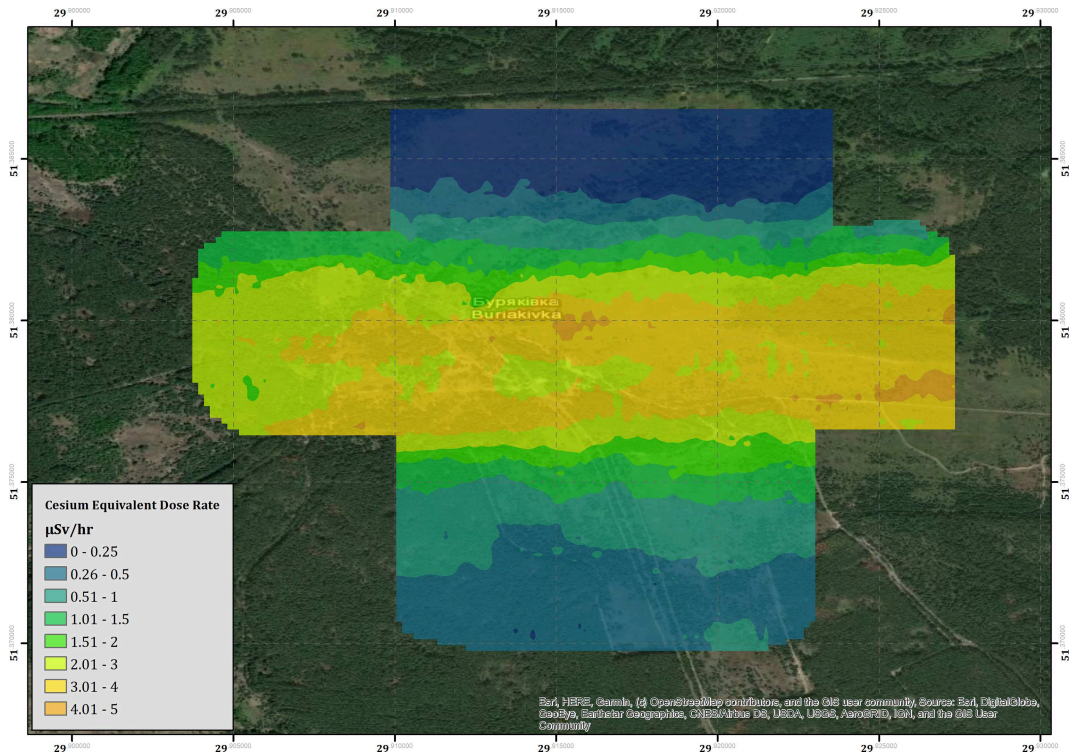

FIGURE 6 | The cesium equivalent dose-rate (CED) of the Buryakivka area.

"The measured CED for the region surrounding the ChNPP is presented within **Figure 7**. The combined survey amalgamates the data from seven flights conducted over 4 days of deployment. Contrary to the data collected within the Buryakivka region (section 3.2.1), all the surveys conducted within this area are included within the presented data set (see Table 2 for full flight details). The color-scaled overlay is once more presented at a pixel size of 20 × 20 m. As expected, the overall CED measured in the area surrounding the ChNPP is significantly larger than that measured in Buryakivka. The maximum CED successfully recorded by the fixed-wing system was 12.8 μSv h^{-1}, which is 2.8 times greater than the maximum CED recorded within the Buryakivka region. The map shows two main areas displaying elevated dose-rates. The first is a sharply delineated hot spot that extends immediately to the west of the ChNPP itself and covers the "Red Forest" area [51.379 N, 30.071 E]. The second is a much broader zone of elevated intensity, extending southwards from the plant toward the village of Kopachi [51.366 N, 30.100 E]. This overall trend is also depicted within the soil sampling investigations previously conducted by Kashparov et al. (2018), showing a general agreement between this dataset and previously published works from other institutions."

"Located at the south-eastern corner of the area is a region of elevated dose-rate (3.3 μSv h^{-1}) that lies within an area of relatively low dose-rate (0.51–1.0 μSv h^{-1}). The hot spot [51.343843 N, 30.110399 E] manifests in an almost idealized point-source geometry when compared to the broad spreading of radioactivity evident within the measurements collected in the area surrounding it. The shape and location of the hot spot suggest that its presence is the result of anthropogenic concentration of radioactivity rather than the natural deposition following the accident. Dose-rate information could not be extracted from the cross-hatched area within **Figure 7** due to detector saturation issues."

A correction has also been made to Section 4.1 Radiological Monitoring, Paragraphs 2 and 5:

"Previous surveys have measured dose-rates within a small portion of the "Red Forest" area to be up to around 170 μSv h^{-1} (Burtniak et al., 2018). These surveys were conducted within the portion of the "Red Forest" that could not be mapped by our system at much lower altitudes (5 m) and much slower velocities that are typical of multi-rotor surveys. Despite being inherently unreliable, the total-count data recorded by the fixed-wing system (Figure 3) reported a maximum count-rate of 12,436 cps at 45 m altitude. Even though the measurements were saturated, using this count-rate as a minimum value for the radiological intensity within this area would produce an expected dose rate of at least 95 μSv h^{-1} (based upon the approximate ratio of the altitude corrected total intensity to cesium dose-rate). As the detector is facing an overload during these measurements, the real total counts value would most likely be greater, producing a larger CED."

"The measurements collected by the aircraft at this point in space are significantly lower than the values measured by the ground team (3.3 μSv h^{-1} vs. 2 mSv h^{-1}). There may be a number of reasons for the discrepancy between these values.

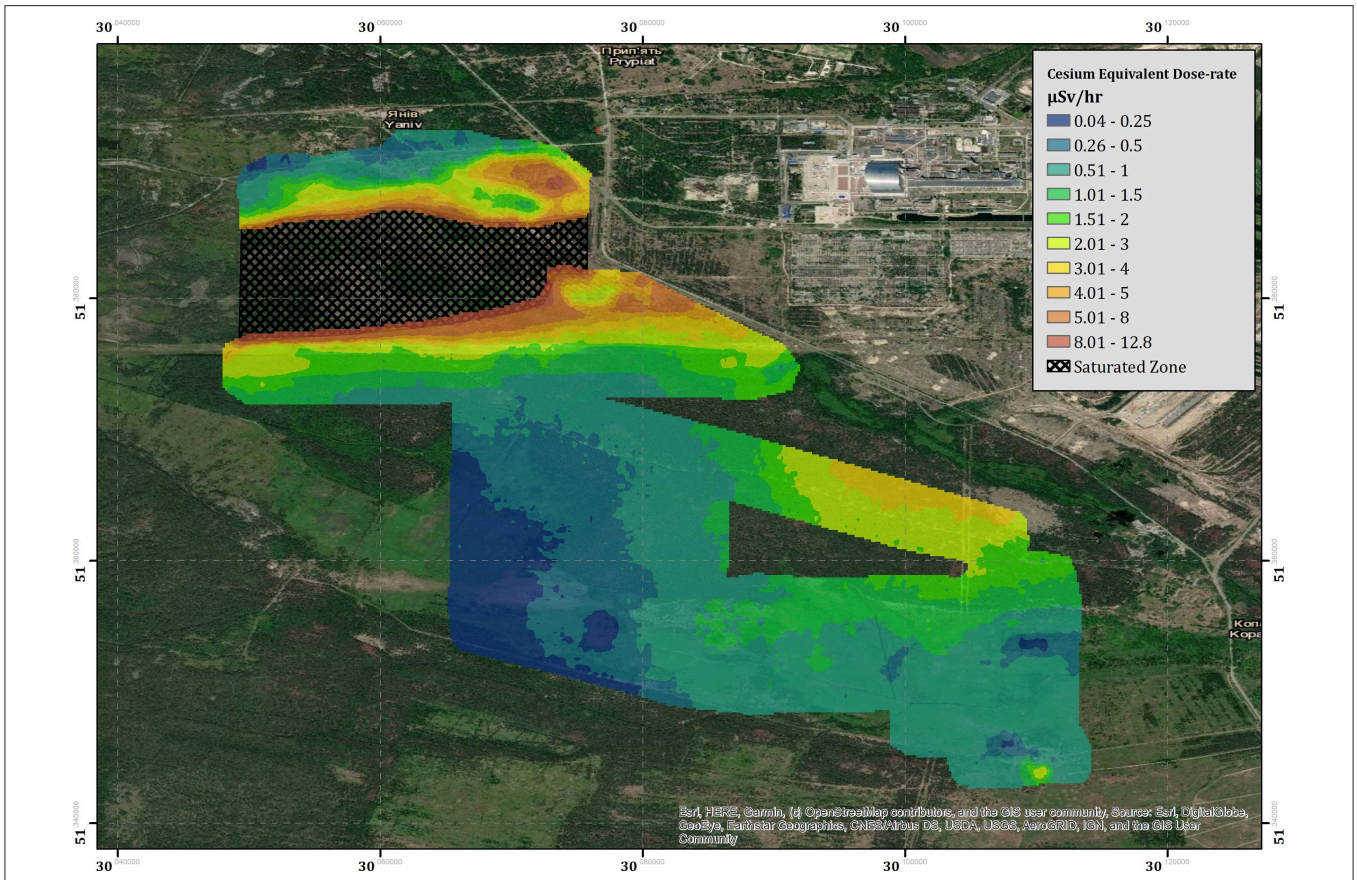

FIGURE 7 | The cesium equivalent dose-rate (CED) of the "Red Forest" area surrounding the ChNPP.

Firstly, the analysis performed on the results collected by the aircraft focuses solely on the ^{137}Cs signal, ignoring contributions from any other radionuclides (these are outside the scope of this study and will be investigated in future studies). The myriad of radioactive material released from the accident is highly complex and the measured contribution of ^{137}Cs is but a component of the total output (Smith and Beresford, 2005; Burtniak et al., 2018). Given that the "hopper" hot spot is so intensely radioactive, the on-ground measurements could be recording inputs from other radionuclides in addition to the measured ^{137}Cs signal. This could potentially include gamma-ray signals from ^{241}Am, which emits a low energy gamma-ray (0.06 MeV) that is more easily attenuated by the medium between the source and the detector (see Figure 4). These kinds of signals are difficult to detect with any confidence at the altitudes used within this survey, especially because the incomplete transfer of energy between incoming photons and the detection crystal (very common for small-volume, room temperature detection systems Gilmore, 2008) creates a high background signal at the low energy range of the spectrum. Radionuclides other than ^{241}Am and ^{137}Cs are also expected to be present within the signal emanating from this region, including contributions from fission products from spent nuclear fuel."

The authors apologize for these errors and state that they do not change the scientific conclusions of the article in any way. The original article has been updated.

REFERENCES

Burtniak, V., Zabulonov, Y., Stokolos, M., Bulavin, L., and Krasnoholovets, V. (2018). Application of a territorial remote radiation monitoring system at the Chornobyl nuclear accident site. *J. Appl. Remote Sens.* 12:046007. doi: 10.1117/1.JRS.12.046007

Gilmore, G. (2008). *Practical Gamma-Ray Spectrometry*. Chichester, UK: John Wiley & Sons.

Kashparov, V., Levchuk, S., Zhurba, M., Protsak, V., Khomutinin, Y., Beresford, N. A., et al. (2018). Spatial datasets of radionuclide contamination in the Ukrainian Chernobyl Exclusion Zone. *Earth Syst. Sci. Data* 10, 339–353.

Smith, J., and Beresford, N. (2005). *Chernobyl: Catastrophe and Consequences*. Chichester, UK: Springer.

frontiers in Robotics and AI

REVIEW
published: 06 February 2020
doi: 10.3389/frobt.2020.00008

Let's Push Things Forward: A Survey on Robot Pushing

Jochen Stüber[1], *Claudio Zito*[2]* *and Rustam Stolkin*[2]

[1] IRLab, School of Computer Science, University of Birmingham, Birmingham, United Kingdom, [2] Extreme Robotics Lab (ERL), University of Birmingham, Birmingham, United Kingdom

As robots make their way out of factories into human environments, outer space, and beyond, they require the skill to manipulate their environment in multifarious, unforeseeable circumstances. With this regard, pushing is an essential motion primitive that dramatically extends a robot's manipulation repertoire. In this work, we review the robotic pushing literature. While focusing on work concerned with predicting the motion of pushed objects, we also cover relevant applications of pushing for planning and control. Beginning with analytical approaches, under which we also subsume physics engines, we then proceed to discuss work on learning models from data. In doing so, we dedicate a separate section to deep learning approaches which have seen a recent upsurge in the literature. Concluding remarks and further research perspectives are given at the end of the paper.

Keywords: robotics, pushing, manipulation, forward models, motion prediction

OPEN ACCESS

Edited by:
Holger Voos,
University of Luxembourg,
Luxembourg

Reviewed by:
Mathias Hauan Arbo,
Norwegian University of Science and Technology, Norway
Senka Krivic,
King's College London,
United Kingdom

***Correspondence:**
Claudio Zito
c.zito@bham.ac.uk

Specialty section:
This article was submitted to Robotic Control Systems, a section of the journal Frontiers in Robotics and AI

Received: *08 March 2019*
Accepted: *20 January 2020*
Published: *06 February 2020*

Citation:
Stüber J, Zito C and Stolkin R (2020) Let's Push Things Forward: A Survey on Robot Pushing. Front. Robot. AI 7:8. doi: 10.3389/frobt.2020.00008

1. INTRODUCTION

We argue that pushing is an essential motion primitive in a robot's manipulative repertoire. Consider, for instance, a household robot reaching for a bottle of milk located in the back of the fridge. Instead of picking up every yogurt, egg carton, or jam jar obstructing the path to create space, the robot can use gentle pushes to create a corridor to its lactic target. Moving larger obstacles out of the way is even more important to mobile robots in environments as extreme as abandoned mines (Ferguson et al., 2004), the moon (King, 2016), or for rescue missions, such as for the Fukushima Daiichi Nuclear Power Plant. In order to save cost, space, or reduce payload, mobile robots are often not equipped with grippers, meaning that prehensile manipulation is not an option. Even in the presence of grippers, objects may be too large or too heavy to grasp.

In addition to the considered scenarios, pushing has numerous beneficial applications that come to mind less easily. For instance, pushing is effective at manipulating objects under uncertainty (Brost, 1988; Dogar and Srinivasa, 2010), and for pre-grasp manipulation, allowing robots to bring objects into configurations where they can be easily grasped (King et al., 2013). Dexterous pushing skills are also widely applied and applauded in robot soccer (Emery and Balch, 2001).

Humans perform skilful manipulation tasks from an early age, and are able to transfer behaviors learned on one object to objects of novel sizes, shapes, and physical properties. For robots, achieving those goals is challenging. This complexity arises from the fact that frictional forces are usually unknown but play a significant role for pushing (Zhou et al., 2016). Furthermore, the dynamics of pushing are highly non-linear, with literal tipping points, and sensitive to initial conditions (Yu et al., 2016). The large body of work on robotic pushing has produced many accurate models for predicting the outcome of a push, some analytical, and some data-driven. However, models that generalize to novel objects are scarce (Kopicki et al., 2017; Stüber et al., 2018), highlighting the demanding nature of the problem.

In this paper, we review the robotic pushing literature. We focus on work concerned with making predictions of the motion of pushed objects, but we also cover relevant applications of pushing for planning and control. This work is primarily targeted at newcomers to robotic pushing, such as Ph.D. students, interested in understanding the evolution of the field. While the main body of this paper focuses on a qualitative analysis of the presented methods, the mathematical treatment is delivered as a set of mini-lectures in the figures. We use the figures to "draw" on the blackboard, thus providing a geometrical intuition for important formalizations used across the literature. Each figure is accompanied by a caption explaining the mathematical content in an accessible yet careful way.

Related to our work is the survey conducted by Ruggiero et al. (2018) which covers the literature on planning and control for non-prehensile dynamic manipulation. Pushing is one of the motion primitives which they consider, among throwing, catching, and others.

In the next section, we provide the problem statement of this survey (section 2). Subsequently, we present the existing literature, beginning with analytical approaches (section 3), under which we also subsume physics engines. We then proceed to discuss data-driven approaches (section 4), including deep learning methods which have recently become very popular in the literature. Finally, we conclude by summarizing the presented approaches and by discussing open problems and promising directions for future research (section 5).

2. PROBLEM STATEMENT

Even in ideal conditions, such as *structured environments* where an agent has a complete model of the environment and perfect sensing abilities, the problems of robotic grasping and manipulation are not trivial. By a complete model of the environment we mean that physical and geometric properties of the world are exactly known, e.g., pose, shape, and friction parameters, as well as the mass of the object we wish to manipulate. In fact, the object to be manipulated is indirectly controlled by contacts with a robot manipulator (e.g., pushing by a contacting finger part). For planning and control, robots need either an *inverse model* (IM) or a *forward model* (FM). IMs compute the action that transforms the current state into the target state (see **Figure 1**). In contrast, FMs predict the next state resulting from applying an action in the current state (see **Figure 2**). Depending on the type of model used, a variety of planning and control strategies exist. For instance, an agent may use an FM to imagine the likely outcomes from all possible actions and then choose the action which achieves the most desirable end state (e.g., Zito et al., 2012). An example of an IM-based controller is the work of Igarashi et al. (2010) where a dipole-like vector field is used to compute the direction of motion of a robot pusher such that the object is pushed along a specified path. As manipulation and grasping problems are defined in continuous state and action spaces, finding an optimal continuous control input to achieve the desired state is often computationally intractable.

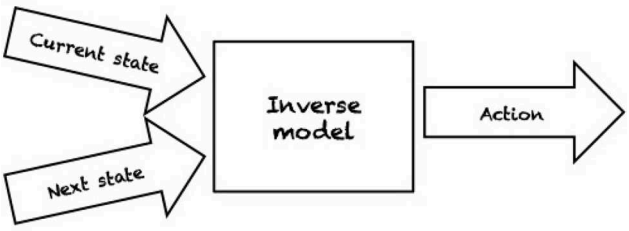

FIGURE 1 | An inverse model computes an action which will affect the environment such that the next desired state (or configuration) is achieved from the current state.

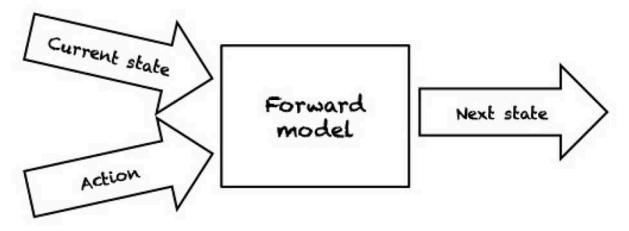

FIGURE 2 | A forward model makes a prediction on how an action will affect the current state of the environment by returning the configuration after the action is taken.

Even more challenging is the problem of grasping and manipulation in *unstructured environments*, where the ideal conditions of structured environments do not exist. There are several reasons why an agent may fail to build a complete description of the state of the environment: sensors are noisy, robots are difficult to calibrate, and actions' outcomes are unreliable due to unmodeled variables (e.g., friction, mass distribution). Uncertainty can be modeled in several ways, but in the case of manipulation there are typically two types of uncertainty:

- *Uncertainty in physical effects*: occurs when the robot acts on external bodies via physical actions (e.g., contact operations). This interaction transforms the current state of the world according to physical laws which are not fully predictable. For example, a pushed object may slide, rotate or topple with complex motions which are extremely difficult to predict, and involve physical parameters which may not be known. We can think of this as uncertainty on future states.
- *Uncertainty in sensory information*: occurs when some of the quantities that define the current state of the world are not directly accessible to the robot. Thus the necessity to develop strategies to allow the robot to complete tasks in partial ignorance by recovering knowledge of its environment. When executing robotic actions in such cases, sensory uncertainty may propagate to the result of the action.

This paper is concerned with the evolution of models to predict object motions and their application in robotics. **Table 1**

TABLE 1 | Summary of the literature a glance.

		Assumptions			Motion			Aim				Model		
		Quasi-static assumption	2D object	Known object	1D	planar	3D	Motion prediction	Parameter estimation	Path planning	Grasping	Analytical	Data-driven	Physics simulator
Purely analytical	Mason (1982)		✓	✓	✓		✓				✓			
	Mason (1986b)	✓	✓	✓		✓		✓				✓		
	Peshkin and Sanderson (1988a,b)	✓	✓	✓			✓		✓				✓	
	Goyal et al. (1991)	✓	✓	✓		✓			✓					
	Alexander and Maddocks (1993)	✓		✓		✓		✓				✓		
	Lee and Cutkosky (1991)	✓	✓	✓		✓		✓				✓		
	Lynch et al. (1992)	✓		✓		✓			✓			✓		
	Howe and Cutkosky (1996)			✓		✓		✓				✓		
	Mason (1990)			✓			✓			✓		✓		
	Mayeda and Wakatsuki (1991)			✓		✓				✓		✓		
	Akella and Mason (1992, 1998)			✓		✓				✓		✓		
	Narasimhan (1994)	✓	✓	✓		✓				✓		✓		
	Lynch and Mason (1996)	✓	✓	✓		✓				✓		✓		
	Agarwal et al. (1997)		✓	✓		✓				✓		✓		
	Nieuwenhuisen et al. (2005)	✓	✓	✓		✓				✓		✓		
	de Berg and Gerrits (2010)		✓	✓		✓				✓		✓		
	Miyazawa et al. (2005)		✓		✓					✓		✓		
	Cappelleri et al. (2006)	✓		✓		✓				✓		✓		
	Dogar and Srinivasa (2011)	✓		✓			✓				✓	✓		
	Cosgun et al. (2011)		✓	✓		✓				✓				
	Lee et al. (2015)		✓	✓		✓				✓				
	King (2016)			✓		✓				✓				✓
Hybrid	Lynch (1993)		✓			✓			✓			✓		
	Yoshikawa and Kurisu (1991)			✓		✓			✓			✓		
	Ruiz-Ugalde et al. (2010, 2011)			✓		✓			✓			✓		
	Zhu et al. (2017)			✓			✓			✓				
	Bauza and Rodriguez (2017)	✓		✓		✓				✓		✓		
Dynamic analysis	Brost (1992)					✓		✓				✓		
	Jia and Erdmann (1999)			✓		✓		✓				✓		
	Behrens (2013)			✓		✓		✓				✓		
	Chavan-Dafle and Rodriguez (2015)			✓			✓	✓			✓		✓	
	Woodruff and Lynch (2017)		✓	✓		✓				✓		✓		
Physic engine	Zito et al. (2012)	✓		✓			✓			✓			✓	
	Scholz et al. (2014)			✓		✓			✓				✓	✓
	Zhu et al. (2017)			✓			✓		✓					✓
Data driven	Moldovan et al. (2012)			✓		✓		✓					✓	
	Ridge et al. (2015)						✓	✓					✓	
	Zrimec and Mowforth (1991)			✓			✓	✓					✓	
	Salganicoff et al. (1993)			✓		✓		✓					✓	
	Walker and Salisbury (2008)					✓		✓					✓	
	Lau et al. (2011)					✓		✓					✓	
	Krivic and Piater (2019)	✓				✓				✓			✓	
	Kopicki et al. (2011, 2017)	✓		✓			✓	✓					✓	
	Stüber et al. (2018)	✓					✓	✓	✓				✓	
	Meriçli et al. (2015)	✓				✓		✓					✓	

(Continued)

TABLE 1 | Continued

		Assumptions			Motion			Aim				Model		
		Quasi-static assumption	2D object	Known object	1D	planar	3D	Motion prediction	Parameter estimation	Path planning	Grasping	Analytical	Data-driven	Physics simulator
Deep learning	Denil et al. (2016)						✓		✓				✓	
	Chang et al. (2016)					✓			✓				✓	
	Li et al. (2018)	✓				✓		✓	✓				✓	
	Watters et al. (2017)			✓		✓		✓					✓	
	Fragkiadaki et al. (2015)			✓		✓		✓					✓	
	Ehrhardt et al. (2017)			✓		✓		✓					✓	
	Byravan and Fox (2017)			✓			✓	✓					✓	
	Finn et al. (2016)						✓	✓					✓	

summarizes the literature at a glance. The papers are classified according to the type of approach implemented. We identify the following six classes.

1. **Purely analytical**. This is mostly seminal work drawn from classical mechanics that uses the quasi-static assumption. To be precise, some of these approaches also venture into dynamic analysis, but with many simplifying assumptions (section 3.1.1).
2. **Hybrid**. Works in this class extend analytical approaches with data-driven methods. Whilst the interactions between objects are still represented analytically, some quantities of interest are estimated based on observations, e.g., the coefficients of friction (section 3.1.2).
3. **Dynamic analysis**. These approaches integrate dynamics in the model (section 3.2.1).
4. **Physics engines**. Here we consider work that employs a physics engine as a "black box" to make predictions about the interactions (section 3.2.2).
5. **Data-driven**. Such models learn how to predict physical interaction from examples (sections 4.1 and 4.2).
6. **Deep learning**. As the data-driven approaches, such models learn how to construct an FM from examples. The key insight is that the deep learning approaches are based on feature extraction (section 4.3).

The features highlighted for each approach are as follows.

- The assumptions made by the authors on their approach. We highlight i) the quasi-static assumption in the model, ii) if it is a seminal work on 2D shapes, and iii) if the method required a known model of the object to be manipulated.
- The type of motion analyzed in the paper, such as 1D, planar (2D translation and 1D rotation around the $x-$axis), or full 3D (3D translation and 3D rotation).
- The aim of the paper. We distinguish between predicting the motion of the object, estimating physical parameters, planning pushes, and analysing a push to reach a stable grasp.
- The model. We distinguish between analytical, constructed from data, and by using a physics simulator.

3. ANALYTICAL APPROACHES

3.1. Quasi-Static Planar Pushing

Early work on robotic pushing focused on the problem of quasi-static planar pushing of sliding objects. In a first phase, several researchers, following pioneering work by Matthew T. Mason, approached the problem analytically, explicitly modeling the objects involved and their physical interactions whilst drawing on theories from classical mechanics. More recently, this tradition has moved to extend analytical models with more data-driven methods.

3.1.1. Purely Analytical Approaches

To briefly introduce the problem, *planar pushing* (Mason, 1982), refers to an agent pushing an object such that pushing forces lie in the horizontal support plane while gravity acts along the vertical. Both pusher and pushed object move only in the horizontal plane, effectively reducing the world to 2D. Meanwhile, the *quasi-static assumption* (Mason, 1986b) in this context means that the involved objects' velocities are small enough that inertial forces are negligible. In other words, objects only move when pushed by the robot. Instantaneous motion is then the consequence of the balance between contact forces, frictional forces, and gravity. The quasi-static assumption makes the problem more tractable, yielding simpler models. A key challenge in predicting the motion of a pushed object under manipulation is that the distribution of pressure at the contact between object and supporting surface is generally unknown. Hence, the system of frictional forces that arise at that contact is also indeterminate (Mason, 1982).

Mason (1982, 1986a) started the line of work on pushing, proposing the voting theorem as a fundamental result. It allows one to find the sense of rotation of a pushed object given the pushing direction and the center of friction without requiring knowledge of the pressure distribution. Drawing on this seminal work, Peshkin and Sanderson (1988a,b) found bounds on the rotation rate of the pushed object given a single-point push. Following that, Goyal et al. (1991) introduced the *limit surface* which describes the relationship between the motion of a sliding object and the associated support friction given that the support distribution is completely specified. Under

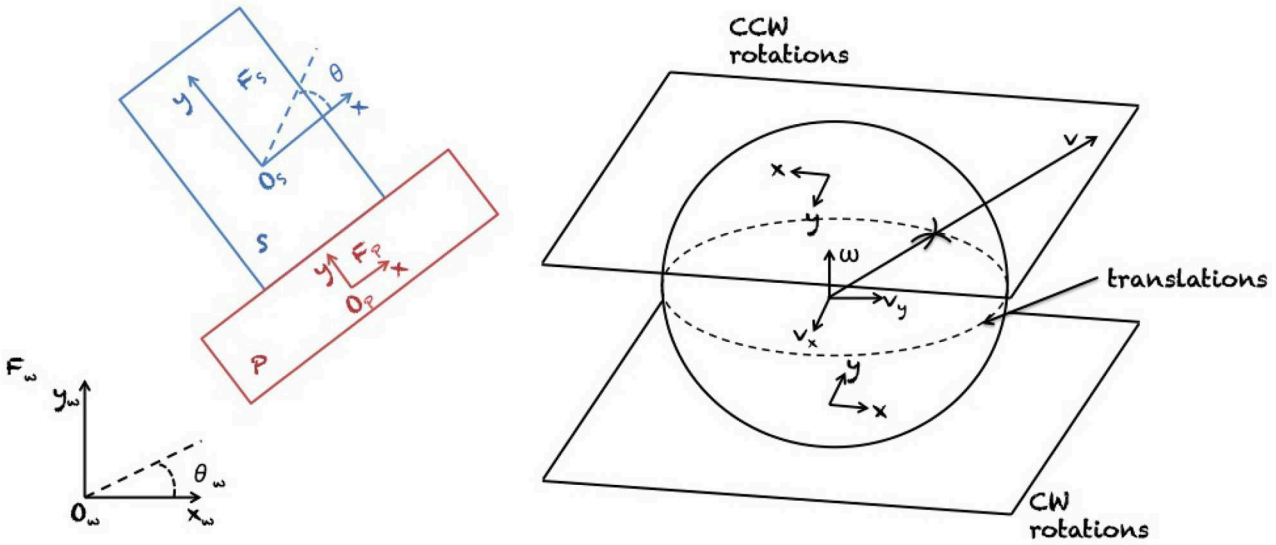

FIGURE 3 | The slider S (blue) is a rigid object in the plane $\mathbb{R}^2$, and its configuration space is $\mathbb{R}^2 \times \Theta$, i.e., 2D translation and one rotation over the $x-$axis. The slider is pushed by a rigid pusher P (red) at a point or set of points of contact. A world frame F_w with origin O_w is fixed in the plane, and a slider frame F_w with origin F_s is attached to the center of friction of the slider S. The configuration $(x_w, y_w, \theta_w)^\intercal$ describes the position and orientation of the slider frame F_s relative to the world frame F_w. Similarly, a pusher frame F_p with origin O_p and its configuration is computed. On the right side of the figure, the relation between the unit motion vector $\mathbf{v} = (v_x, v_y, \omega)^\intercal$ and the center of rotation of frame F_s is described by the projection shown from the unit motion sphere to the tangent planes (one for each rotation sense). The line at the equator of the sphere represents translations. Reproduced from Lynch and Mason (1996).

the quasi-static assumption, the limit surface allows one to convert the generalized force applied by a pusher at a contact to the instantaneous generalized velocity of the pushed object. Alexander and Maddocks (1993) considered the case when only the geometric extent of the support area is known, and described techniques to bound the possible motions of the pushed object. While the limit surface provides a powerful tool for determining the motion of a pushed object, there exists no convenient explicit form to construct it. In response to this challenge, Lee and Cutkosky (1991) proposed to approximate the limit surface as an ellipsoid to improve computational time. However, their approximation requires knowledge of the pressure distribution. Marking a milestone of planar pushing research, Lynch et al. (1992) applied the ellipsoidal approximation to derive a closed-form analytical solution for the kinematics of quasi-static single-point pushing, including both sticking and sliding behaviors. Subsequently, Howe and Cutkosky (1996) explored further methods for approximating limit surfaces, including guidance for selecting the appropriate approach based on the pressure distribution, computational cost, and accuracy.

Results on the mechanics of planar pushing have been used for *planning and control* of manipulator pushing operations. To begin with, Mason (1990) showed how to synthesize robot pushing motions to slide a block along a wall, a problem later also studied by Mayeda and Wakatsuki (1991). Akella and Mason (1992, 1998) analyzed the series of pushes needed to bring a convex polygon to a desired configuration. Narasimhan (1994) and Kurisu and Yoshikawa (1995) studied the problem of moving an object among obstacles by pushing with point contact. Lynch and Mason (1996) comprehensively studied stable pushing of a planar object with a fence-shaped finger, considering mechanics, control, and planning. First, they derived conditions for stable edge pushing, considering the case where the object will remain attached to the pusher without slipping or breaking contact. Based on this result, they then used best-first search to find a path to a specified goal location. **Figure 3** shows the proposed representation of motions by Lynch and Mason (1996). Agarwal et al. (1997) proposed an algorithm for computing a contact-preserving push plan for a point-sized pusher and a disk-shaped object. Nieuwenhuisen et al. (2005) utilized compliance of manipulated disk-shaped objects against walls to guide their motion. They presented an exact planning algorithm for 2D environments consisting of non-intersecting line segments. Subsequently, de Berg and Gerrits (2010) improved this approach from a computational perspective and presented push-planning methods both for the contact-preserving case and less restrictive scenarios. Miyazawa et al. (2005) used a rapidly-exploring random tree (RRT) (LaValle, 1998) for planning non-prehensile manipulation, including pushing, of a polyhedron with three degrees of freedom (DOF) by a robot with spherical fingers. They do not allow for sliding and rolling of robot fingers on the object surface. Cappelleri et al. (2006) have solved a millimeter scale 2D version of the peg in the hole problem, using Mason's models for quasi-static manipulation and an RRT-based approach for planning a sequence of pushes. **Figure 4** presents a graphical representation of planar motions and Coulomb's frictional

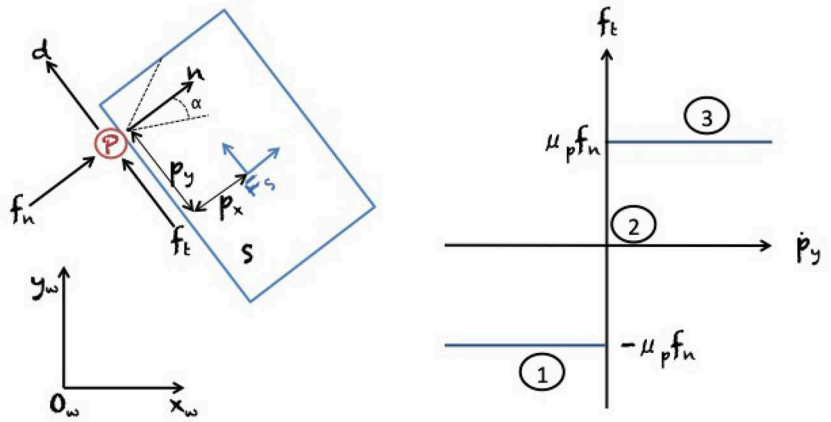

FIGURE 4 | (Left) Planar pushing system with world frame F_w (with origin O_w) and a slider S (blue) with frame F_w as described in **Figure 3**. The pusher P (red) is interacting with the slider on one point of contact. It impresses a normal force f_n, a tangential friction force f_t, and a torque τ about the center of mass. The normal force f_n is in the direction of the normal vector n of the contact point between pusher and slider, and $\alpha = \arctan \mu_p$ is the angle of the friction cone assuming μ_p as the coefficient of friction. The terms p_x and p_y describe, respectively, the normal and the tangential distance between the pusher P and the center of friction of the slider S. **(Right)** Coulomb's frictional law for the planar pushing system on the left-hand figure. Coulomb's law states that the normal and tangential forces are related by $f_t = \mu_p f_n$. Three contact modes are defined. (1) Sliding right in which friction acts as a force constraint; (2) Sticking in which friction acts as a kinematic constraint; and (3) Sliding left in which friction acts as a force constraint. Reproduced from Bauza et al. (2018).

law that governs such systems (see the figure caption for further details).

More recently, Dogar and Srinivasa (2011) employed the ellipsoidal approximation of the limit surface to plan robust push-grasp actions for dexterous hands and used them for rearrangement tasks in clutter. To use results for planar pushing, they assumed that objects do not topple easily. Furthermore, they assumed that the robot has access to 3D models of the objects involved. Cosgun et al. (2011) presented an algorithm for placing objects on cluttered table surfaces, thereby constructing a sequence of manipulation actions to create space for the object. However, focusing on planning, in their 2D manipulation they simply push objects at their center of mass in the desired direction. Lee et al. (2015) presented a three-stage hierarchical approach to planning sequences of non-prehensile and prehensile actions. First, they find a sequence of qualitative contact states of the moving object with other objects, then a feasible sequence of poses for the object, and lastly a sequence of contact points for the manipulators on the object.

In summary, although of fundamental importance for understanding the mechanisms of pushing, analytical approaches are limited by their own inherent complexity. The assumptions around which they are built do not hold in real applications, e.g., a robot link in contact with an object does not produce a single-point contact or the frictional forces are not constant over a supporting surface. Proofs of concept for demonstrating the validity and stability of such methods are generally confined to carefully chosen testing scenarios or special applications, e.g., the frictionless millimeter scale peg-in-the-hole scenario in Cappelleri et al. (2006). Extensions to non-convex or novel-shaped objects challenge analytical approaches. Yet, controllers and planners can easily be synthesized for specified objects and environments. Due to the deterministic nature of the models, they do not implicitly account for uncertainty in the state description or the predictions. Nonetheless, an analytical method can be employed as a black box to forward-simulate the effect of a given action within a planner. For instance, King (2016) developed a series of push planners for open-loop non-prehensile rearrangement tasks in cluttered environments. Before considering more complex scenarios, they used a simple analytical approach for forward-simulation of randomly sampled time-discrete controls within an RRT-based planner. They tested their planners on two real robotic platforms, the home care robot HERB with a seven DOF arm, and the NASA rover K-Rex.

3.1.2. Complementing Analytical Approaches With Data-Driven Methods (Hybrid)

Transitioning to the second phase of planar pushing research, multiple factors have contributed a shift toward more *data-driven approaches*. For one thing, much of the previous work makes minimal assumptions regarding the pressure distribution. While convenient, those methods lead to conservative strategies for planning and control, providing only worst case guarantees. Furthermore, while assumptions regarding the pressure distribution in previous work were often minimal, other strong assumptions were frequently made to derive results analytically. Hence, more recent work has set out to validate common assumptions such as the ubiquitous quasi-static assumption. Additionally, purely analytical models do not take into account the stochastic nature of pushing in the sense that pushes indistinguishable to sensor and actuator resolution have empirically been found to produce variable results (Yu et al., 2016). Instead of making minimal or strong assumptions about parameters, they can instead be *estimated* based on observations. Several researchers have explored this approach to deal with the inherent uncertainties of this problem (section 1). **Figure 5**

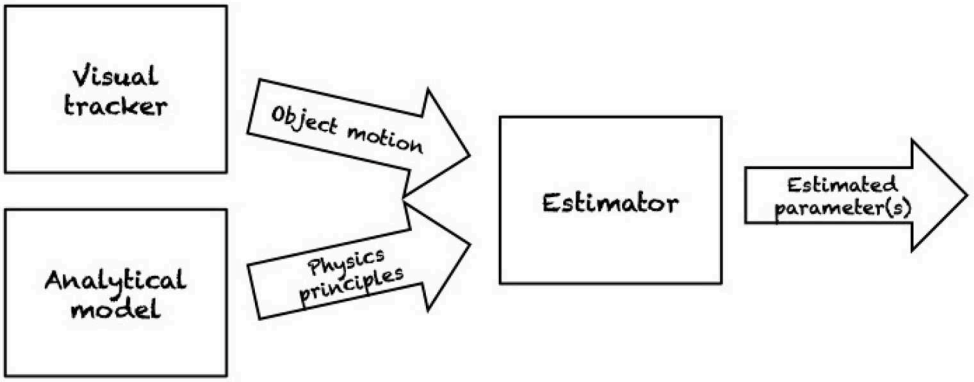

FIGURE 5 | A classical workflow for estimating relevant physical parameters of a pushed object. A robotic pusher performs a set of push operation on an object which is typically tracked using vision. Simpler approach employs markers on the object for more accurate estimations. An analytical model of the motion for the target object is also employed. Sensory data and physical principles are the inputs of the estimator. As output, the estimator provides with an estimate of the desired parameters, e.g., friction distribution or center of mass. In Lynch (1993) the estimated parameters are also used for recognizing objects based on their (estimated) physical properties.

summarizes a classical workflow for estimating relevant physical parameters of a pushed object.

Lynch (1993) presented methods both for estimating the relevant friction parameters by performing experimental pushes, and for recognizing objects based on their friction parameters. Similarly, Yoshikawa and Kurisu (1991) described how a mobile robot with a visual sensor can estimate the friction distribution of an object and the position of the center of friction by pushing and observing the result. Yet, both of these approaches discretise the contact patch into grids so that they are either imprecise if the approximation is too coarse or suffer from the curse of dimensionality when using a fine-grained approximation. Ruiz-Ugalde et al. (2010, 2011) formulated a compact mathematical model of planar pushing. Assuming that the object's base shape is given, their robot explored object-table and finger-object friction coefficient parameters. Zhou et al. (2016) developed a method for modeling planar friction, proposing a framework for representing planar sliding force-motion models using convex polynomials. Notably, they also showed that the ellipsoid approximation is a less accurate special case of this representation. Zhou et al. (2017) extended the convex polynomial model to associate a commanded position-controlled end effector motion to the instantaneous resultant object motion. They modeled the probabilistic nature of object-to-surface friction by sampling parameters from a set of distributions. They presented the motion equations for both single and multiple frictional contacts and validated their results with robotic pushing and grasping experiments on the dataset published by Yu et al. (2016). That dataset comprises planar pushing interactions with more than a million samples of positions of pusher and slider, as well as interaction forces. Push interaction is varied along six dimensions, namely surface material, shape of the pushed object, contact position, pushing direction, pushing speed, and pushing acceleration. Using their dataset, they characterized the variability of friction, and evaluated the most common assumptions and simplifications made by previous models of frictional pushing. They provide an insightful table that lists the assumptions and approximations made in much of the work cited in this section. More recently, Bauza et al. (2019) have published Omnipush, an extensive dataset of planar pushing behavior that extends their previous work. It comprises 250 pushes for each of 250 objects. The pushing velocity is constant and chosen so that the interaction is close to quasi-static. They improved on their previous dataset by providing RGB-D sensor data in addition to tracking data, increasing object diversity, adding controlled variation of the objects mass distribution, and creating benchmarks to evaluate models. Finally, Bauza and Rodriguez (2017) used a data-driven approach to model planar pushing interaction to predict both the most likely outcome of a push and, as a novelty, its expected variability. The learned models, also trained on the dataset by Yu et al. (2016), rely on a variation of Gaussian processes whilst avoiding and evaluating the quasi-static assumption by making the velocity of the pusher an input to the model. However, the learned models are specific to the particular object and material. Transfer learning is left for future work.

3.2. Physics Engines and Dynamic Analysis

While the quasi-static assumption may be reasonable in a variety of situations, other problems call for dynamic models of pushing. One popular approach to achieving this is using a physics engine. Before covering this field, we first consider work concerned with dynamic pushing that does not recur to physics engines.

3.2.1. Dynamic Analysis

Using dynamic analysis, Brost (1992) investigated the problem of catching an object by pushing it, i.e., determining the pushing motions that lead to a pusher-object equilibrium. This work was motivated by dealing with uncertainty in positioning, generating plans that work also in the worst case. Jia and Erdmann (1999) investigated dynamic pushing assuming frictionless interaction between pusher and object. Behrens (2013) instead studied

dynamic pushing but assumed infinite friction between pusher and object. Chavan-Dafle and Rodriguez (2015) considered planning non-prehensile in-hand manipulation with patch contacts. They described the quasi-dynamic motion of an object held by a set of frictional contacts when subject to forces exerted by the environment. Given a grasp configuration, gripping forces, and the location and motion of a pusher, they estimate both the instantaneous motion of the object and the minimum force required to push the object into the grasp. To this end, complex contact geometries are broken up into rigid networks of point contacts. More recently, Woodruff and Lynch (2017) extended earlier work by Lynch and Mason (1999) on dynamic manipulation primitives, including pushing, by proposing a method for motion planning and feedback control of hybrid, dynamic, and non-prehensile manipulation tasks.

3.2.2. Physics Engines

A large body of work related to pushing makes use of *physics engines*. Commonly used examples of such engines include Bullet Physics, the Dynamic Animation and Robotics Toolkit (DART), MuJoCo, the Open Dynamics Engine (ODE), NVIDIA PhysX, and Havok (Erez et al., 2015). Those engines allow for 3D simulation but 2D physics engines exist, as well, e.g., Box2D. While some physics engines have been designed for graphics and animation, others have been developed specifically for robotics. In the first category, visually-plausible simulations are key while physically-accurate simulations are essential for many robotics applications. Most physics engines today use impulse-based velocity-stepping methods to simulate contact dynamics. As this requires solving NP-hard problems at each simulation step, more tractable convex approximations have been developed, highlighting the trade-off between computational complexity and accuracy present in those engines (Erez et al., 2015). 3D physics engines use a Cartesian representation where each body has six DOF and joints are modeled as equality constraints in the joint configuration space of the bodies. In robotics, where joint constraints are ubiquitous, using generalized coordinates is computationally less expensive and prevents joint constraints from being violated.

For a *comparison* of physics engines, we refer the reader to two recent studies (Erez et al., 2015; Chung and Pollard, 2016). Erez et al. (2015) compared ODE, Bullet, PhysX, Havok, and MuJoCo. It should be noted that the study was written by the developers of MuJoCo. They introduced quantitative measures of simulation performance and focused their evaluation on challenges common in robotics. They concluded that each engine performs best on the type of system it was tailored to. For robotics, this is MuJoCo while gaming engines shine in gaming-related trials, whereby no engine emerges as a clear winner. Chung and Pollard (2016) compared Bullet, DART, MuJoCo, and ODE with regard to contact simulations whilst focusing on the predictability of behavior. Their main result is that the surveyed engines are sensitive to small changes in initial conditions, emphasizing that parameter tuning is important. Another evaluation of MuJoCo was carried out by Kolbert et al. (2017) who evaluated the contact model of MuJoCo with regard to predicting the motions and forces involved in three in-hand robotic manipulation primitives, among them pushing. In the course, they also evaluated the contact model proposed by Chavan-Dafle and Rodriguez (2015). They found that both models make useful yet not highly accurate predictions. Concerning MuJoCo, they state that its soft constraints increase efficiency but limit accuracy, especially in the cases of rigid contacts and transitions in sticking and slipping at contacts.

Researchers have applied physics engines in multifarious ways to study robotic pushing. To begin with, physics engines have been used in RRT-based planners to forward-simulate pushes. Zito et al. (2012) presented a two-level planner that combines a global RRT planner operating in the configuration space of the object, and a local planner that generates sequences of actions in the robot's joint space that will move the object between a pair of nodes in the RRT. In this work, the experimental set-up consists of a simulated model of a tabletop robot manipulator with a single rigid spherical fingertip which it uses to push a polyflap (Sloman, 2006) to a goal state. To achieve this, the randomized local planner utilizes a physics engine (PhysX) to predict the object's pose after a pushing action. Erroneous estimates and uncertainty in the motion is not directly taken into account by the planner. Hence, a re-planning stage is required when the actual motion differs from the prediction by more than a user-defined threshold. **Figure 6** shows a sequence of actions planned by the two-level planner for pushing a polyflap to a desired configuration (see caption for further details). Similarly, King (2016) incorporated a dynamic physics engine (Box2D) into an RRT-based planner to model dynamic motions such as a ball rolling. To reduce planning complexity, they considered only dynamic actions that lead to statically stable states, i.e., all considered objects need to come to rest before the next action. Another application of physics engines in robotic pushing was proposed by Scholz et al. (2014). In what they refer to as Physics-Based Reinforcement Learning, an agent uses a physics engine as a model representation. Hence, a physics engine can be seen as a hypothesis space for rigid-body dynamics. They introduced uncertainty using distributions over the engine's physical parameters and obtained transitions by taking the expectation of the simulator's output over those random variables. Finally, Zhu et al. (2017) utilized a physics engine for motion prediction, learning the physical parameters through black-box Bayesian optimization. First, a robot performs random pushing actions on an object in a tabletop set-up. Based on those observations, the Bayesian learning algorithm tries to identify the model parameters that maximize the similarity between the simulated and observed outcomes. To support working with different objects, a pre-trained object detector is used that maps observed objects to a library of 3D meshes and estimates the objects' poses on that basis. Once the physical parameters have been identified, they are used to simulate the results of new actions.

Finally, while physics engines and dynamic analysis offer great value for robotic applications, e.g., by taking into consideration dynamic interaction and 3D objects, they nevertheless require explicit object modeling and extensive parameter tuning. Another approach, which we consider next, is to learn how to predict object motions from data.

FIGURE 6 | Simulation of a Katana robot arm equipped with a spherical finger that plans a sequence of pushes to move an L-shaped object, called polyflap (Sloman, 2006) to a goal state. The plan is created by using a physics engine (PhysX) to predict the outcome of a push operation. Image 01 shows the initial pose. The wire-framed L-shaped polyflap is a "phantom" to indicate the desired goal state. The goal pose is translated from the initial pose by 28 cm and rotated by 90°. Image 02 shows the collision-free trajectory to bring the end effector to the start pose of the first push. Images 01–04 show the first push which makes the polyflap tip over. Images 05–09 show a series of pushes which culminate in the polyflap resting in an unstable equilibrium pose along its folded edge. Images 12 and 13 show a sideways push. Images 14 and 15 show the final frontal push which aligns the polyflap with the target configuration. Courtesy of Zito et al. (2012).

4. LEARNING TO PREDICT FROM EXAMPLES

This part of the literature is based on learning predictive models for robotic pushing from data. We first review work on qualitative models and then consider models that make metrically precise predictions. In both of those sections, we do not include work that uses deep learning techniques. We dedicate a separate section to such approaches, given the current research interest in that area and the large number of papers being published.

4.1. Qualitative Models

Much work on qualitative models revolves around the concept of *affordances*. The term affordance was invented by Gibson (1979) and generally refers to an action possibility that an object or environment provides to an organism. Although it has originated from psychology, the concept has been influential in various domains, among them robotics. Sahin et al. (2007) discussed affordances from a theoretical perspective while laying emphasis on their use in autonomous robotics. Min et al. (2016) provided a recent survey of affordance research in developmental robotics.

Although the concept of affordances is typically associated with learning "high-level" actions from contexts, e.g., pushing an object in a clutter scene when grasping is not available, in this paper we focus on investigations that extend affordances to the effects of an action too. Ugur et al. (2011) considered an anthropomorphic robot that learns object affordances as well as effect categories through self-interaction and self-observation. After learning an FM as a mapping between object affordances and effects, the proposed method can make plans to achieve desired goals, emulate end states of demonstrated actions, and cope with uncertainty in the physical effects by monitoring the plan execution and taking corrective actions using the perceptual structures employed or discovered during learning. While much previous work has focused on affordance models for individual objects, Moldovan et al. (2012) learned affordance models for configurations of multiple interacting objects with push, tap, and grasp actions for achieving desired displacements or rotations, as well as contacts between objects, for selecting the appropriate object and action for the subsequent step. Their model is capable of generalizing over objects and dealing with uncertainty in the physical effects. Ridge et al. (2015) developed a self-supervised online learning framework based on vector quantization for acquiring models of effect classes and their

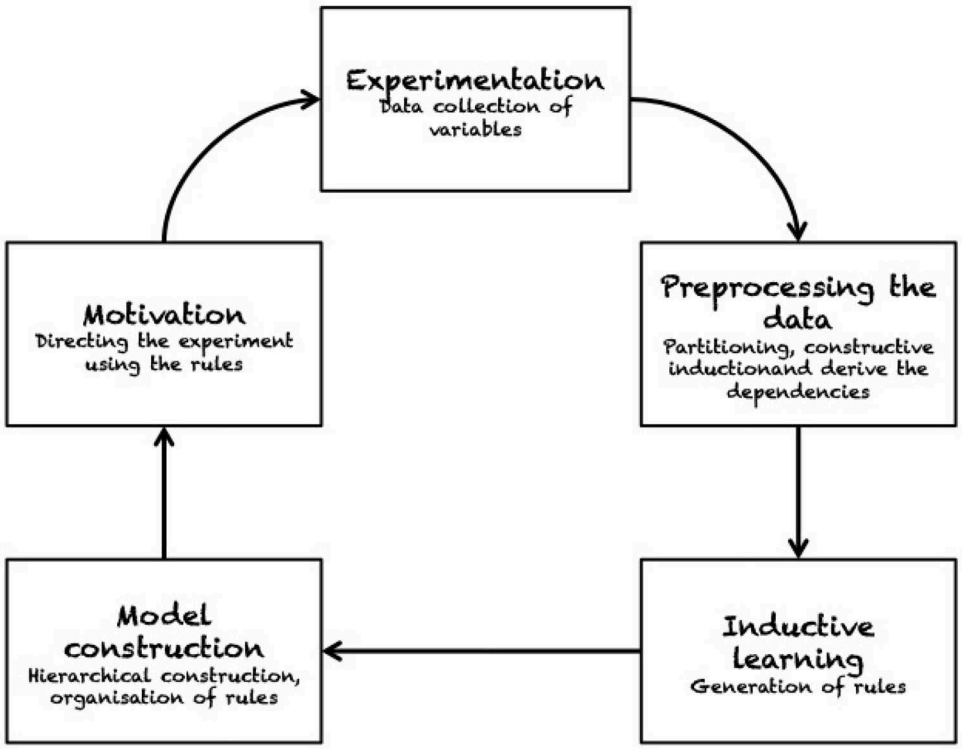

FIGURE 7 | The sequence of operations adopted by Zrimec and Mowforth (1991) to construct their causality learning model. The robot learns by interacting with the environment in an unsupervised fashion. The system can autonomously discover knowledge, as e.g., whether an action generates a push on an object. The "motivation" module guarantees that the system is driven toward acquiring more knowledge about the robot/environment interaction. Reproduced from Zrimec and Mowforth (1991).

associations with object features. Specifically, they considered robots pushing household objects and observing them with a camera. Limitations of such learning approaches are that they do not tend to generalize well to novel objects and actions. This is also due to a lack of interpretation and understanding of novel contexts. In fact, self-interaction and self-observation are mainly limited by the ability of the robot to discover novel scenarios and learning opportunities by itself (see also section 5.1).

Considering *other qualitative approaches* than those related to affordances, Zrimec and Mowforth (1991) developed an algorithm for knowledge extraction and representation to predict the effects of pushing. In their experiment, a robot performs random pushes and uses unsupervised learning on those observations. Their method involves partitioning, constructive induction and determination of dependencies (see **Figure 7**). Hermans et al. (2013) developed a method for predicting contact locations for pushing based on the global and local object shape. In exploratory trials, a robot pushes different objects, recording the objects' local and global shape features at the pushing contacts. For each observed trajectory, the robot computes a push-stability or rotate-push score and maps shape features to those scores by means of regression. Based on that mapping, the robot can search objects of novel shape for features associated with effective pushes. Experimental results are reported for a mobile manipulator robot pushing household objects in a tabletop set-up.

While learned affordances, and other qualitative models, can be useful in various scenarios, other applications require the ability to predict the effects of pushing more precisely, e.g., by explicitly predicting six DOF rigid body motions. We consider efforts made to achieve precise predictions in the next section.

4.2. Metrically Precise Models

Early seminal work by Salganicoff et al. (1993) presented a vision-based unsupervised learning method for robot manipulation. A robot pushes an object at a rotational point contact and learns an FM of the action effects in image space. Subsequently, they used the FM for stochastic action selection in manipulation planning and control. The scenarios considered in this work are relatively simple in that the pusher remains within the friction cone of the object and the contact only has one rotational DOF. Yet, this work takes an approach that is markedly different from analytical models discussed before. Instead of estimating parameters such as frictional coefficients explicitly, the authors *encode* that information *implicitly* in the mapping between actions and their effects in image space. Similarly, Walker and Salisbury (2008) learned a mapping between pushes and object motion as an alternative to explicitly modeling support friction. Set in a 2D tabletop environment, a robot with a single finger pushes objects

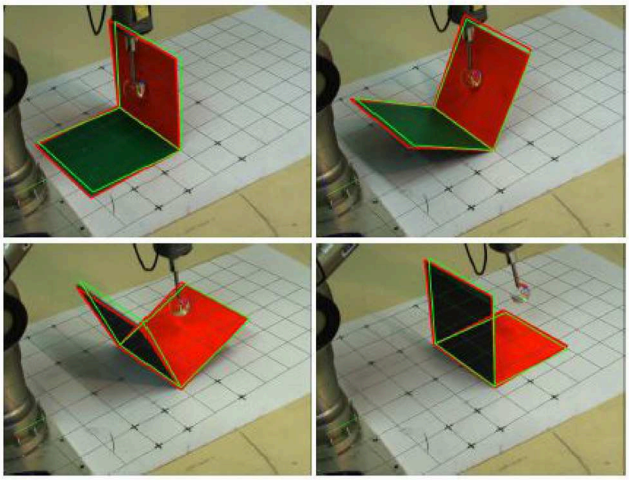

FIGURE 8 | The example shows the interaction between a 5-axis Katana robotic manipulator and an L-shape object, called polyflap (Sloman, 2006). A set of contact experts are learned as probability densities for encoding geometric relations between parts of objects under a push operation. This approach allows these experts to learn from demonstration physical properties, such as non-penetration between an object and a table top, without explicitly representing physics knowledge in the model. The green wire frame denotes the prediction whilst the red wire frame denotes the visual tracking. Courtesy of Kopicki et al. (2011).

and uses an online, memory-based local regression model to learn manipulation maps. To achieve this, they explicitly detect the object's shape using a proximity sensor and fit a shape to the thus obtained point cloud. A method for handling objects with more complex shapes was proposed by Lau et al. (2011). In their work, a robot, while being of simple circular shape itself, aims to deliver irregularly shaped flat objects to a goal position by pushing them. The objects that they consider are chosen to exhibit quasi-static properties. Collecting several hundred random example pushes as training data, an FM is learned using non-parametric regression, similar to the approach taken by Walker and Salisbury (2008). Also tackling the problem of object delivery, Krivic and Piater (2019) proposed a modular method for pushing objects in cluttered and dynamic environments that can work with unknown objects without prior experience. Drawing on their previous work (Krivic et al., 2016; Krivic and Piater, 2018), the authors' approach comprises a space-reasoning module, a strategy module, and an adaptive pushing controller which learns local IMs of robot-object interaction online. While their approach shows a high success rate of object delivery for objects with quasi-static properties in simulated and real-world environments, it does not take into account object orientation and depends on vision-based measurements of the object motion.

Kopicki et al. (2011) presented two data-driven probabilistic methods for predicting 3D motion of rigid bodies interacting under the quasi-static assumption. First, they formulated the problem as regression and subsequently as density estimation. **Figure 8** shows an example of the interaction between a 5-axis Katana robotic manipulator and a polyflap (Sloman, 2006). In Kopicki et al. (2017) they extended this work further. Their architecture is modular in that multiple object- and context-specific FMs are learned which represent different constraints on the object's motion. A product of experts is used which, contrary to mixture models, does not add but multiply different densities. Hence, all constraints, e.g., those imposed by the robot-object contact and multiple object-environment contacts, need to be satisfied so that a resulting object motion is considered probable. This formulation facilitates the transfer of learned motion models to objects of novel shape and to novel actions. In experiments with a robot arm, the method is compared with and found to outperform the physics engine PhysX tuned on the same data. For learning and prediction, their algorithms require access to a point cloud of the object. A further extension of this approach is presented in Stüber et al. (2018). In this work, the authors aim to contribute to endowing robots with versatile non-prehensile manipulation skills. To that end, an efficient data-driven approach to transfer learning for robotic push manipulation is proposed. This approach combines and extends two separate strings of research, one directly concerning pushing manipulation (Kopicki et al., 2017), and one originating from grasping research (Kopicki et al., 2016). The key idea is to learn motion models for robotic pushing that encode knowledge specific to a given type of contact, see the work by Kopicki et al. (2016) for further details. **Figure 9** presents a graphical representation of the feature-based predictors as well as resulting predictions across object shapes. In an previously unseen situation, when the robot needs to push a novel object, the system first establishes how to create a contact with the object's surface. Such a contact is selected among the learned models, e.g., a flat contact with a cube side or a contact with a cylindrical surface. At the generated contact, the system then applies the appropriate motion model for prediction, similarly to that of Kopicki et al. (2017). The underlying rationale for this approach to prediction is that predicting on familiar ground reduces the motion models' sample complexity while using local contact information for prediction increases their transferability (Krivic et al., 2016).

Meriçli et al. (2015) similarly presented a case-based approach to push-manipulation prediction and planning. Based on experience from self-exploration or demonstration, a robot learns multiple discrete probabilistic FMs for pushing complex 3D objects on caster wheels with a mobile base in cluttered environments. Subsequently, the case models are used for synthesizing a controller and planning pushes to navigate an object to a goal state whilst potentially pushing movable obstacles out of the way. In the process, the robot continues to observe the results of its actions and feeds that data back into the case models, allowing them to improve and adapt.

Metrically precise models have become a stable trend of research in the field of robotic push manipulation. Their probabilistic nature elegantly deals with state and motion uncertainties. Nonetheless, real applications of robot pushing may require higher levels of reasoning to be useful assistants. For example, a warehouse robot may need to fill shelves with many boxes via push operations. This would also require planning multiple sequences of actions where some earlier placements may lead to a sub-optimal final configuration.

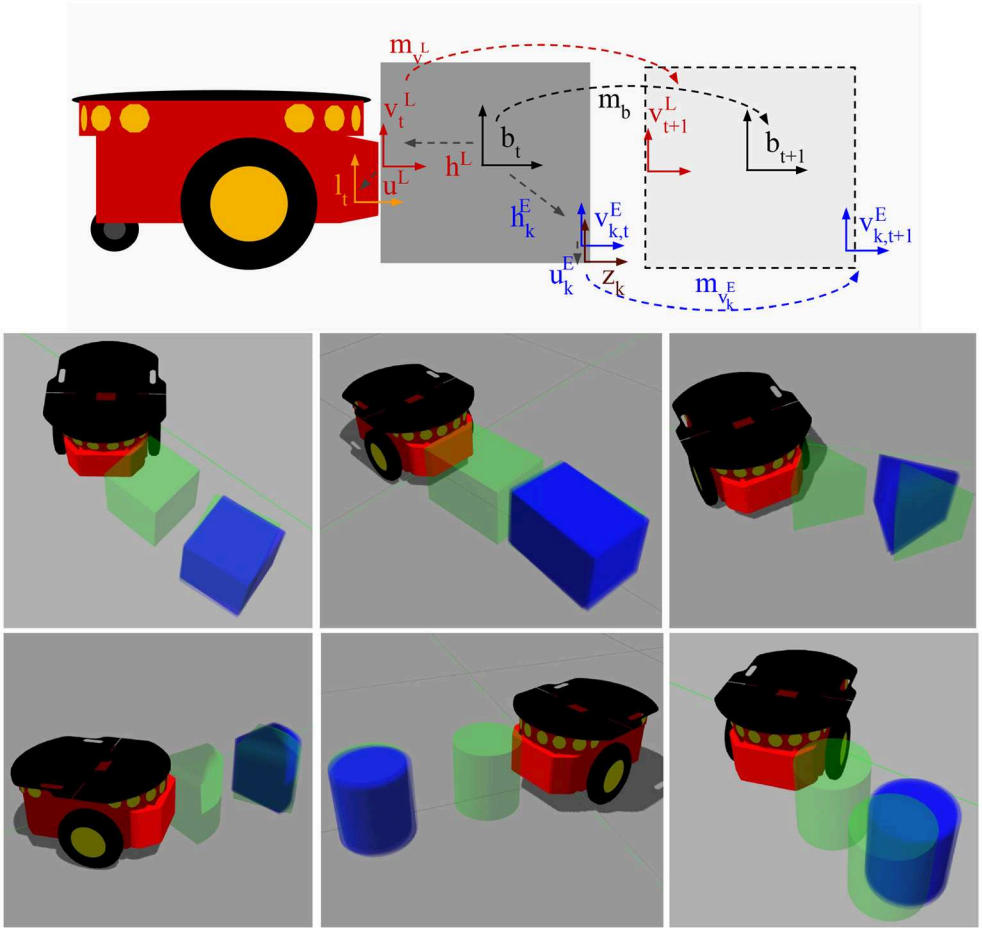

FIGURE 9 | (Top) graphical representation of the feature-based predictors for push operations. The global motion of the object after a push is described by the rigid body transformation m_b. This transformation is unknown to the robot. However, the robot can estimate it by learning a set of local predictors for the motions m_{v^L} and $m_{v_k^E}$, for $k = 1, \ldots, N_E$. The rigid body transformations h^L and h_k^E describe the estimated contacts on the object's surface w.r.t. the estimated global frame of the object, b. Since the object is assumed to be rigid, this relation does not change over time, thus once the local motions m_{v^L} and $m_{v_k^E}$ are estimated, b_{t+1} can be estimated by using the relations h^L and h_k^E. **(Bottom)** resulting predictions. initial object pose (green, in contact with robot), true final object pose (green, displaced), and predictions (blue). Courtesy of Stüber et al. (2018).

Efficient planning solutions would need to reason about gathering critical information concerning the task space or propagating the uncertainty in the action's effects to future states. We present some suggestions on how to deal with these types of problems in the final remarks (see sections 5.2 and 5.3).

4.3. Deep Learning Approaches

Deep learning commonly refers to methods that employ artificial neural networks to learn models from data. It has been used in robotic pushing to estimate physical parameters, predict the outcome of pushing actions, and for planning and control. Previously, we have seen work concerned with estimating physical parameters of the environment from data. Deep learning has been used to address the same problem. Denil et al. (2016) studied the learning of physical properties such as mass and cohesion of objects in a simulated environment. Using deep reinforcement learning, their robots learn different strategies that balance the cost of gathering information against the cost of inaccurate estimation.

Instead of explicitly estimating physical parameters, another approach is learning a *dynamics model*. Several studies have investigated learning general physical models or "physical intuition" directly from image data. Chang et al. (2016) presented the Neural Physics Engine, a deep learning framework for learning simple physics simulators. They factorize the environment into object-based representations and decompose dynamics into pairwise interaction between objects. However, their evaluation is limited to simple rigid body dynamics in 2D. Li et al. (2018) proposed Push-net, a deep recurrent neural network to tackle the problem of quasi-static planar pushing to re-orient and re-position objects. Their approach requires only visual camera images as input and remembers pushing interactions using a long short-term memory (LSTM) module. An auxiliary objective function estimates the COM of the object,

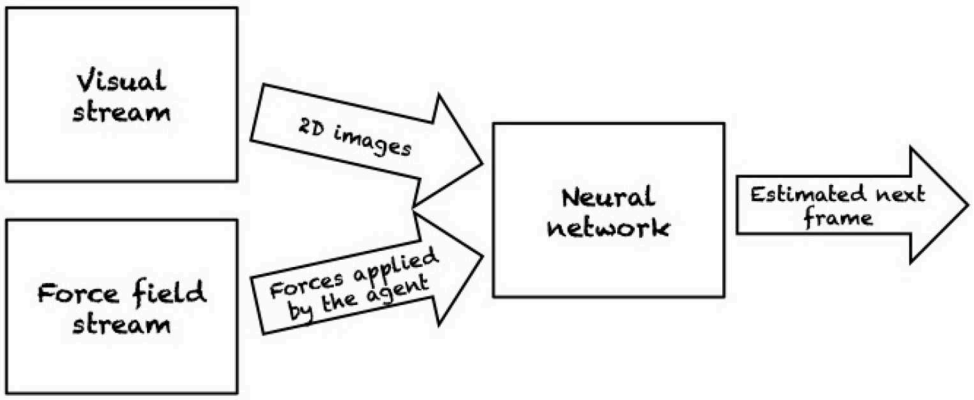

FIGURE 10 | Frame-centric model for motion prediction of billiard balls. The model takes as input the 2D image of the billiard and the forces applied by the agent to make predictions about the future configurations of the balls. Reproduced from Fragkiadaki et al. (2015).

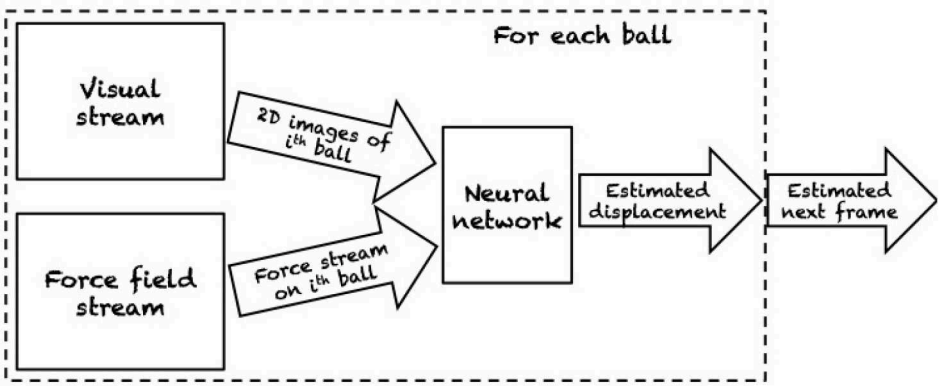

FIGURE 11 | Object-centric model for motion prediction of billiard balls. The system predicts the future configurations of the balls by individually modeling the temporal evolution of each ball. In this scenario, predicting the velocities of each ball is sufficient for computing the next configuration of the billiard. Reproduced from Fragkiadaki et al. (2015).

thus encoding physics into their model. They trained their model in simulation and tested it on various objects in simulation and on two real robots, with results indicating that Push-Net is capable of generalizing to novel objects.

Watters et al. (2017) introduced the Visual Interaction Network, a model for learning the dynamics of a physical system from raw visual observations. First, a convolutional neural network (CNN) generates a factored object representation from visual input. Then, a dynamics predictor based on interaction networks computes predicted trajectories of arbitrary length. They report accurate predictions of trajectories for several hundred time steps using only six input video frames. Yet, their experiments are also limited to rather simple environments, namely 2D simulations of colored objects on natural-image backgrounds. Similarly, Fragkiadaki et al. (2015) also used an object-centric formulation based on raw visual input for dynamics prediction. Based on object-centric visual glimpses (snippets of an image), the system predicts future states by individually modeling the behavior of each object. A graphical representation of this model is presented in **Figure 10**. After training in different environments by means of random interaction, they also use their model for planning actions in novel environments, in this case moving balls on a 2D table (**Figure 11**). Ehrhardt et al. (2017) constructed a neural network for end-to-end prediction of mechanical phenomena. Their architecture consists of three components: a CNN extracts features from images which are updated by a propagation module, and decoded by an estimation module. What their network outputs is a distribution over outcomes, thus explicitly modeling the inherent uncertainty in manipulation prediction. In terms of experiments, they study the relatively simple problem of a small object sliding down an inclined plane.

Moving toward more complex scenarios, Byravan and Fox (2017) introduced SE3-NETS, a deep neural network architecture for predicting 3D rigid body motions. Instead of RGB images, their network takes depth images as input, together with continuous action vectors, and associations between points in subsequent images. SE3-NETS segment point clouds into

object parts and predict their motion in the form of $SE(3)$ transformations. They report that their method outperforms flow-based networks on simulated depth data of a tabletop manipulation scenario. Furthermore, they demonstrate that it performs well on real depth images of a Baxter robot pushing objects. However, their approach requires that associations between depth points are provided. They aim to learn those automatically in future work and to apply SE3-NETS to non-rigid body motion, recurrent prediction, and control tasks. A different approach to learning dynamics from images was taken by Agrawal et al. (2016). They jointly learn FMs and IMs of dynamics of robotic arm operation that can be used for poking objects. In doing so, they extract features from raw images and make predictions in that feature space. In real-world experiments with Baxter, their model is used to move objects to target locations by poking. In order to cope with the real world, their model requires training on large amounts of data. By poking different objects for over 400 h, their robot observed more than 100, 000 actions.

Most of the studies presented this section make use of *object-centric representations* to model dynamics. Other approaches predict motion without such representations. For instance, Finn et al. (2016) developed an action-conditioned video prediction system which predicts a distribution over pixel motions only based on previous frames. No information concerning object appearance is provided to the model. It borrows that information from previous frames and merges it with model predictions. It is this mechanism that allows the model to generalize to previously unseen objects. By conditioning predictions on an action, the model can effectively imagine the action's consequences. As with previously presented deep learning models, this approach also requires large amounts of data to perform well in real-world situations. Hence, the authors have collected a dataset of 59, 000 robot pushing motions (frames associated with the action being applied) on different objects. While their results demonstrate that no object-centric representation is required for prediction, the authors argue that such representations are a promising direction for research as they provide concise state representations for use in reinforcement learning.

We have seen how artificial neural networks can be used to model the dynamics of physical systems. In addition to that, deep reinforcement learning has been used to learn control policies in the field of robotic pushing. Many of those approaches make use of dynamics models so that they can be seen as complementary to the work presented before. We do not provide a detailed review of this very active field here and refer the reader to Levine et al. (2015), Levine et al. (2016), Finn and Levine (2017), and Ghadirzadeh et al. (2017) for overviews of such work.

5. FINAL REMARKS

In this paper we have provided an overview of the problem of robot pushing and summarized the development of the state-of-the-art, focusing on the problem of motion prediction of the object to be pushed. We have also covered some aspects of relevant applications of pushing for planning and control.

Typical approaches have been classified as (i) purely analytical, (ii) hybrid, (iii) dynamic analysis, (iv) physics engines based, (v) data-driven, and (vi) deep learning. Representative work for each of these categories has been listed for readers to have a general overview of the field and its state-of-the-art from the earlier work in the 1980s to the most recent approaches.

A set of assumptions in the proposed methods have been highlighted. Earlier work has mostly investigated motion prediction with the quasi-static assumption to get rid of complex dynamics and provided the groundwork to understand the mechanics for pushing 2D shapes. This seminal work has been extended to more realistic scenarios involving 3D objects to be pushed. Nonetheless, as we have seen there are two types of uncertainty that affect manipulation problems: (i) prediction uncertainty and (ii) state uncertainty. Unfortunately, purely analytical approaches are computationally tractable only under the assumption that the geometrical properties of the object to be pushed are known *a priori* and the dynamics are negligible (e.g., Mason, 1990; Mayeda and Wakatsuki, 1991). Key physical properties that would affect the prediction, e.g., mass distribution or friction coefficients, were typically assumed to be known or possible to estimate on the fly, as in Yoshikawa and Kurisu (1991), by combining data-driven methods with the analytical mechanics of pushing.

More recently, a few efforts were made toward robot pushers that can also deal with state uncertainty. By relaxing the assumption that the model of the object to be pushed is known, the robot typically perceives the object as a point cloud or RGB image to estimate the geometric properties, such as pose and shape, before even attempting a motion prediction, see Fragkiadaki et al. (2015) and Stüber et al. (2018).

Two strands of approaches can be identified: data-driven and deep learning techniques. They are similar in that they both define (or extract) some informative features as a basis for learning and model predictions in a probabilistic framework to estimate an action's most likely outcome given the information available, e.g., an image of the scene or contact models.

Qualitative models have made use of the concept of affordances for learning a mapping between object features and candidate actions, which they then employ for planning. For manipulation tasks, however, the planner also needs to learn the relationship between actions and their effects by creating a mapping from actions to observable end states and their variability. End states can be represented as displacements after a push operation or a set of contacts for a prehensile operation. This enables us to synthesize controllers with multiple classes of actions and their expected effects which are then employed for planning. Affordances can be learned for 3D motions of a single or multiple interacting objects, but they do not generalize well to novel objects or actions. Erroneous predictions can be dealt with at execution time by triggering a re-planning procedure. In contrast to qualitative models, metrically precise models are concerned with directly learning a mapping between observable features (e.g., contacts or geometrical features) and their effects in the context of manipulation. Factorizing robot-object and object-environments contacts as (probabilistic) experts enables a robot pusher to generalize predictions across object categories, e.g., demonstrated by Kopicki et al. (2017). Physical properties such as impenetrability can also be learned implicitly by the experts and transferred to novel contexts. The quasi-static

assumption limits the model to predictions of object motion in the next time-step, but roll-out predictions can be executed to approximate continuous operations. A second, more recent strand is the application of deep learning techniques to learning a physical intuition of the mechanics of pushing from visual data, see Fragkiadaki et al. (2015). Automatic feature extraction can be used to estimate uncertain state information such as the COM from raw data (e.g., RGB images), as well as to make predictions on interaction dynamics. Controllers for specific robots can also be learned directly by the models, but there is a lack of evidence on whether these controllers could be transferred across robot platforms, or how their performance would degrade when dealing with novel objects or actions. Nevertheless, the main disadvantage of these approaches is the amount of data required. While Stüber et al. (2018) and Meriçli et al. (2015) can learn from as little as a few hundred pushes, Push-net and the other deep learning approaches require hundreds of times more data.

While some typical problems still require a better solution, new challenges and requirements are emerging in the field. To make pushing an essential motor primitive in practical robotics, the challenges are either currently under investigation in research group worldwide or need to be investigated in the future. Following we list some suggested trends of open problems that we have identified.

5.1. Understanding and Semantic Representation

The scene is typically perceived as an RGB image or a point cloud. However, for robot pushing, we need to be able to identify pushable objects from static ones. Labeling can be done but it is very expensive in terms of human labor. Converting from source image data to geometrical shapes, and from geometrical shape to semantic representation will be beneficial for the robot. Once the robot can identify probable dynamic objects within a semantic map of the environment it would be able to interact with the environment prioritizing those objects and improving its understanding.

5.2. Sensory Fusion and Feedback

Multiple sensor inputs are nowadays available for robotic system. Instead of solely relying on vision, other sources of information should be used to close the loop of the manipulation. Tactile, proprioception, and visual feedback should be fused together to enable the robot to perform complex manipulation and recover from failures. Generative models, such recurrent neural networks, can learn manipulative operations from multiple sensory sources. This enables the robot to compensate for missing or corrupted input data, as well as to predict the next sensory state and the associated expected error with respect to the next observed sensory state, which can then be used for implementing adaptive behaviors. For readers interested in sensory fusion for manipulative tasks, we refer to the work of Yang et al. (2016) as an interesting starting point.

5.3. Explicitly Modeling Uncertainty in the Model

Due to a lack of perfect perception abilities, it is not unusual that robots have to operate with an incomplete description of their environment. In robot pushing, but more generally in the problem of manipulation, the robot needs to generate a set of contacts to interact with other objects. When the pose of the object to be manipulated is unknown, what is the best way to create a robust set of contacts? In the case of planning for dexterous manipulation, our previous work in Zito et al. (2013) has demonstrated that approaching directions that maximize the likelihood of gathering (tactile) information are more likely to achieve a successful set of contacts for a grasp. This was tested in the case when, due to imperfect perception abilities, the pose of the object to be grasped remained uncertain. This empirically suggests that reasoning about the uncertainty leads to more robust reach-to-grasp trajectories with respect to object-pose uncertainty. Similarly, selecting an action for physical effects (e.g., pushing, push, and grasp) should benefit from incorporating state uncertainty with respect to the initial pose estimate of the object. Finally, we highlight the complexity of incorporating uncertainty in models for pushing which results from the multi-modal stochasticity inherent to the task demonstrated by Yu et al. (2016).

5.4. Cooperative Robots and Multiple Contacts Pushing

Moving large-scale objects is a common problem in warehouses that can be achieved with cooperative robots. Besides the problem of sharing information between them and coordinating the efforts, a new challenge arises from the manipulation point of view. Multiple contacts pushing is hard to predict, especially when the actions are carried by multiple agents. Scholars interested in the problem of multiple contacts pushing are referred to the works of Lynch (1992) and Erdmann (1998) as interesting starting points.

5.5. Real-World Applications

Although the theory behind motion prediction is well-established and applications to simple, structured scenarios have been made, the combination of the existing methods with any industrial applications has not yet been achieved. Robots in warehouses can navigate freely and deliver goods, however, no robotic system is capable of exploiting pushing operations for novel items in novel situations, such as inserting a box of varied produce onto an over-the-head store shelf. Theoretical solutions are rarely reliable in practical engineering applications, hence many sophisticated practical approaches will be needed in the future. Very recently, the robotics community has realized that one of the main issues that prevent a methodological and stable advancement in the field is the lack of standardized benchmarks and metrics for objective evaluations of different approaches. Following the example of fields such as computer vision and natural language process, large and diverse datasets are required to provide the equivalent richness for physical understanding and its application to robot manipulation tasks. A recent attempt in this direction is presented in Bauza et al. (2019).

AUTHOR CONTRIBUTIONS

JS was the main author of this paper and collected the literature. CZ was the leading supervisor of this work and he has co-written the paper. RS has co-supervised and funded this project.

FUNDING

This work was supported by UK Engineering and Physical Sciences Research Council (EPSRC No. EP/R02572X/1) for the National Centre for Nuclear Robotics (NCNR).

REFERENCES

Agarwal, P. K., Latombe, J. C., Motwani, R., and Raghavan, P. (1997). "Nonholonomic path planning for pushing a disk among obstacles," in *Proceedings of International Conference on Robotics and Automation*, Vol. 4 (Albuquerque, NM), 3124–3129.

Agrawal, P., Nair, A. V., Abbeel, P., Malik, J., and Levine, S. (2016). "Learning to poke by poking: experiential learning of intuitive physics," in *Advances in Neural Information Processing Systems 29*, eds D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett (Barcelona: Curran Associates, Inc.), 5074–5082.

Akella, S., and Mason, M. T. (1992). "Posing polygonal objects in the plane by pushing," in *Proceedings 1992 IEEE International Conference on Robotics and Automation*, Vol. 3, 2255–2262.

Akella, S., and Mason, M. T. (1998). Posing polygonal objects in the plane by pushing. *Int. J. Robot. Res.* 17, 70–88. doi: 10.1177/027836499801700107

Alexander, J., and Maddocks, J. (1993). Bounds on the friction-dominated motion of a pushed object. *Int. J. Robot. Res.* 12, 231–248. doi: 10.1177/027836499301200303

Bauza, M., Alet, F., Lin, Y.-C., Lozano-Perez, T., Kaelbling, L. P., Isola, P., et al. (2019). "Omnipush: accurate, diverse, real-world dataset of pushing dynamics with rgb-d video," in *International Conference on Intelligent Robots and Systems (IROS)* (Macau).

Bauza, M., Hogan, F. R., and Rodriguez, A. (2018). "A data-efficient approach to precise and controlled pushing," in *Proceedings of The 2nd Conference on Robot Learning*, volume 87 of *Proceedings of Machine Learning Research*, eds A. Billard, A. Dragan, J. Peters, and J. Morimoto (Zürich: PMLR), 336–345.

Bauza, M., and Rodriguez, A. (2017). "A probabilistic data-driven model for planar pushing," in *2017 IEEE International Conference on Robotics and Automation (ICRA)* (Marina Bay Sands), 3008–3015.

Behrens, M. J. (2013). *Robotic manipulation by pushing at a single point with constant velocity: Modeling and techniques* (Ph.D. thesis). University of Technology, Sydney, NSW, Australia.

Brost, R. C. (1988). Automatic grasp planning in the presence of uncertainty. *Int. J. Robot. Res.* 7, 3–17. doi: 10.1177/027836498800700101

Brost, R. C. (1992). "Dynamic analysis of planar manipulation tasks," in *Proceedings 1992 IEEE International Conference on Robotics and Automation*, Vol. 3 (Nice), 2247–2254.

Byravan, A., and Fox, D. (2017). "Se3-nets: learning rigid body motion using deep neural networks," in *2017 IEEE International Conference on Robotics and Automation (ICRA)* (Marina Bay Sands), 173–180.

Cappelleri, D. J., Fink, J., Mukundakrishnan, B., Kumar, V., and Trinkle, J. C. (2006). "Designing open-loop plans for planar micro-manipulation," in *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006, ICRA 2006* (Orlando, FL), 637–642.

Chang, M. B., Ullman, T., Torralba, A., and Tenenbaum, J. B. (2016). A compositional object-based approach to learning physical dynamics. *CoRR* abs/1612.00341.

Chavan-Dafle, N., and Rodriguez, A. (2015). "Prehensile pushing: in-hand manipulation with push-primitives," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Hamburg), 6215–6222.

Chung, S.-J., and Pollard, N. (2016). Predictable behavior during contact simulation: a comparison of selected physics engines. *Comput. Animat. Virt. Worlds* 27, 262–270. doi: 10.1002/cav.1712

Cosgun, A., Hermans, T., Emeli, V., and Stilman, M. (2011). "Push planning for object placement on cluttered table surfaces," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (San Francisco, CA) 4627–4632.

de Berg, M., and Gerrits, D. H. P. (2010). "Computing push plans for disk-shaped robots," in *2010 IEEE International Conference on Robotics and Automation* (Anchorage, AK), 4487–4492.

Denil, M., Agrawal, P., Kulkarni, T. D., Erez, T., Battaglia, P., and de Freitas, N. (2016). Learning to perform physics experiments via deep reinforcement learning. *arXiv[Preprint].arXiv:1611.01843*. Available online at: https://openreview.net/group?id=ICLR.cc/2017/conference

Dogar, M., and Srinivasa, S. (2011). "A framework for push-grasping in clutter," in *Proceedings of Robotics: Science and Systems VII*, eds Hugh N. R. Durrant-Whyte and P. Abbeel (Los Angeles, CA: MIT Press).

Dogar, M. R., and Srinivasa, S. S. (2010). "Push-grasping with dexterous hands: mechanics and a method," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Taipei), 2123–2130.

Ehrhardt, S., Monszpart, A., Mitra, N. J., and Vedaldi, A. (2017). Learning a physical long-term predictor. *arXiv[Preprint]arXiv:1703.00247*.

Emery, R., and Balch, T. (2001). "Behavior-based control of a non-holonomic robot in pushing tasks," in *Proceedings 2001 ICRA. IEEE International Conference on Robotics and Automation (Cat. No.01CH37164)*, Vol. 3 (Seoul), 2381–2388.

Erdmann, M. (1998). An exploration of nonprehensile two-palm manipulation. *Int. J. Robot. Res.* 17, 485–503. doi: 10.1177/027836499801700502

Erez, T., Tassa, Y., and Todorov, E. (2015). "Simulation tools for model-based robotics: comparison of bullet, havok, mujoco, ode and physx," in *2015 IEEE International Conference on Robotics and Automation (ICRA)* (Seattle, WA), 4397–4404.

Ferguson, D., Morris, A., Haehnel, D., Baker, C., Omohundro, Z., Reverte, C., et al. (2004). "An autonomous robotic system for mapping abandoned mines," in *Advances in Neural Information Processing Systems* (Vancouver, BC), 587–594.

Finn, C., Goodfellow, I., and Levine, S. (2016). "Unsupervised learning for physical interaction through video prediction," in *Advances in Neural Information Processing Systems* (Barcelona), 64–72.

Finn, C., and Levine, S. (2017). "Deep visual foresight for planning robot motion," in *2017 IEEE International Conference on Robotics and Automation (ICRA)* (Marina Bay Sands), 2786–2793.

Fragkiadaki, K., Agrawal, P., Levine, S., and Malik, J. (2015). Learning visual predictive models of physics for playing billiards. *arXiv[Preprint].arXiv:1511.07404*. Available online at: https://dblp.org/db/conf/iclr/iclr2016.html

Ghadirzadeh, A., Maki, A., Kragic, D., and Björkman, M. (2017). "Deep predictive policy training using reinforcement learning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Vancouver, BC), 2351–2358.

Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Houghton Mifflin.

Goyal, S., Ruina, A., and Papadopoulos, J. (1991). Planar sliding with dry friction part 1. Limit surface and moment function. *Wear* 143, 307–330. doi: 10.1016/0043-1648(91)90104-3

Hermans, T., Li, F., Rehg, J. M., and Bobick, A. F. (2013). "Learning contact locations for pushing and orienting unknown objects," in *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)* (Atlanta, GA), 435–442. doi: 10.1109/HUMANOIDS.2013.7030011

Howe, R. D., and Cutkosky, M. R. (1996). Practical force-motion models for sliding manipulation. *Int. J. Robot. Res.* 15, 557–572. doi: 10.1177/027836499601500603

Igarashi, T., Kamiyama, Y., and Inami, M. (2010). "A dipole field for object delivery by pushing on a flat surface," in *2010 IEEE International Conference on Robotics and Automation* (Anchorage, AK), 5114–5119.

Jia, Y.-B., and Erdmann, M. (1999). Pose and motion from contact. *Int. J. Robot. Res.* 18, 466–487. doi: 10.1177/027836499901800504

King, J. E. (2016). *Robust rearrangement planning using nonprehensile interaction* (Ph.D. thesis). Carnegie Mellon University, Pittsburgh, PA, United States.

King, J. E., Klingensmith, M., Dellin, C. M., Dogar, M. R., Velagapudi, P., Pollard, N. S., et al. (2013). "Pregrasp manipulation as trajectory optimization," in *Robotics: Science and Systems* (Berlin).

Kolbert, R., Chavan-Dafle, N., and Rodriguez, A. (2017). *Experimental Validation of Contact Dynamics for In-Hand Manipulation*. Cham: Springer International Publishing.

Kopicki, M., Detry, R., Adjigble, M., Stolkin, R., Leonardis, A., and Wyatt, J. L. (2016). One-shot learning and generation of dexterous grasps for novel objects. *Int. J. Robot. Res.* 35, 959–976. doi: 10.1177/0278364915594244

Kopicki, M., Zurek, S., Stolkin, R., Moerwald, T., and Wyatt, J. L. (2017). Learning modular and transferable forward models of the motions of push manipulated objects. *Auton. Robot.* 41, 1061–1082. doi: 10.1007/s10514-016-9571-3

Kopicki, M., Zurek, S., Stolkin, R., Mörwald, T., and Wyatt, J. (2011). "Learning to predict how rigid objects behave under simple manipulation," in *2011 IEEE International Conference on Robotics and Automation* (Shanghai), 5722–5729.

Krivic, S., and Piater, J. (2018). "Online adaptation of robot pushing control to object properties," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Madrid), 4614–4621.

Krivic, S., and Piater, J. (2019). Pushing corridors for delivering unknown objects with a mobile robot. *Auton. Robots* 43, 1435–1452. doi: 10.1007/s10514-018-9804-8

Krivic, S., Ugur, E., and Piater, J. (2016). "A robust pushing skill for object delivery between obstacles," in *2016 IEEE International Conference on Automation Science and Engineering (CASE)* (Fort Worth, TX), 1184–1189.

Kurisu, M., and Yoshikawa, T. (1995). Trajectory planning for an object in pushing operation. *J. Robot. Soc. Jpn.* 13, 1115–1121. doi: 10.7210/jrsj.13.1115

Lau, M., Mitani, J., and Igarashi, T. (2011). "Automatic learning of pushing strategy for delivery of irregular-shaped objects," in *2011 IEEE International Conference on Robotics and Automation* (Shanghai), 3733–3738.

LaValle, S. (1998). *Rapidly-Exploring Random Trees: A New Tool for Path Planning*. Technical report, CS Department, Iowa State University.

Lee, G., Lozano-Pérez, T., and Kaelbling, L. P. (2015). "Hierarchical planning for multi-contact non-prehensile manipulation," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Hamburg), 264–271.

Lee, S. H., and Cutkosky, M. R. (1991). Fixture planning with friction. *J. Manufact. Sci. Eng.* 113, 320–327. doi: 10.1115/1.2899703

Levine, S., Finn, C., Darrell, T., and Abbeel, P. (2016). End-to-end training of deep visuomotor policies. *J. Mach. Learn. Res.* 17, 1334–1373.

Levine, S., Wagener, N., and Abbeel, P. (2015). "Learning contact-rich manipulation skills with guided policy search," in *2015 IEEE International Conference on Robotics and Automation (ICRA)* (Seattle, WA), 156–163. doi: 10.1109/ICRA.2015.7138994

Li, J. K., Lee, W. S., and Hsu, D. (2018). "Push-net: deep planar pushing for objects with unknown physical properties," in *Robotics: Science and Systems XIV* (Pittsburgh, PA).

Lynch, K. M. (1992). "The mechanics of fine manipulation by pushing," in *Proceedings 1992 IEEE International Conference on Robotics and Automation*, Vol. 3 (Nice), 2269–2276.

Lynch, K. M. (1993). "Estimating the friction parameters of pushed objects," in *Proceedings of 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '93)*, Vol. 1 (Yokohama), 186–193.

Lynch, K. M., Maekawa, H., and Tanie, K. (1992). "Manipulation and active sensing by pushing using tactile feedback," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Vol. 1 (Raleigh, NC), 416–421.

Lynch, K. M., and Mason, M. T. (1996). Stable pushing: mechanics, controllability, and planning. *Int. J. Robot. Res.* 15, 533–556. doi: 10.1177/027836499601500602

Lynch, K. M., and Mason, M. T. (1999). Dynamic nonprehensile manipulation: controllability, planning, and experiments. *Int. J. Robot. Res.* 18, 64–92. doi: 10.1177/027836499901800105

Mason, M. T. (1982). *Manipulator grasping and pushing operations* (Ph.D. thesis). Massachusetts Institute of Technology, Cambridge, MA, United States.

Mason, M. T. (1986a). Mechanics and planning of manipulator pushing operations. *Int. J. Robot. Res.* 5, 53–71. doi: 10.1177/027836498600500303

Mason, M. T. (1986b). "On the scope of quasi-static pushing," in *International Symposium on Robotics Research* (Gouvieux-Chantilly: MIT Press), 229–233.

Mason, M. T. (1990). *Compliant Sliding of a Block Along a Wall*. Berlin; Heidelberg: Springer.

Mayeda, H., and Wakatsuki, Y. (1991). "Strategies for pushing a 3d block along a wall," in *Proceedings IROS '91. IEEE/RSJ International Workshop on Intelligent Robots and Systems '91. Intelligence for Mechanical Systems*, Vol. 2 (Osaka), 461–466.

Meriçli, T., Veloso, M., and Akın, H. L. (2015). Push-manipulation of complex passive mobile objects using experimentally acquired motion models. *Auton. Robots* 38, 317–329. doi: 10.1007/s10514-014-9414-z

Min, H., Yi, C., Luo, R., Zhu, J., and Bi, S. (2016). Affordance research in developmental robotics: a survey. *IEEE Trans. Cogn. Dev. Syst.* 8, 237–255. doi: 10.1109/TCDS.2016.2614992

Miyazawa, K., Maeda, Y., and Arai, T. (2005). "Planning of graspless manipulation based on rapidly-exploring random trees," in *(ISATP 2005). The 6th IEEE International Symposium on Assembly and Task Planning: From Nano to Macro Assembly and Manufacturing, 2005* (Montreal, QC), 7–12.

Moldovan, B., Moreno, P., van Otterlo, M., Santos-Victor, J., and De Raedt, L. (2012). "Learning relational affordance models for robots in multi-object manipulation tasks," in *2012 IEEE International Conference on Robotics and Automation (ICRA)* (St. Paul, MN: IEEE), 4373–4378.

Narasimhan, S. (1994). *Task level strategies for robots* (Ph.D. thesis). Massachusetts Institute of Technology, Cambridge, MA, United States.

Nieuwenhuisen, D., van der Stappen, A. F., and Overmars, M. H. (2005). "Path planning for pushing a disk using compliance," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Edmonton, AB), 714–720.

Peshkin, M. A., and Sanderson, A. C. (1988a). The motion of a pushed, sliding workpiece. *IEEE J. Robot. Automat.* 4, 569–598. doi: 10.1109/56.9297

Peshkin, M. A., and Sanderson, A. C. (1988b). Planning robotic manipulation strategies for workpieces that slide. *IEEE J. Robot. Automat.* 4, 524–531. doi: 10.1109/56.20437

Ridge, B., Leonardis, A., Ude, A., Deniša, M., and Skocaj, D. (2015). Self-supervised online learning of basic object push affordances. *Int. J. Adv. Robot. Syst.* 12:24. doi: 10.5772/59654

Ruggiero, F., Lippiello, V., and Siciliano, B. (2018). Nonprehensile dynamic manipulation: a survey. *IEEE Robot. Automat. Lett.* 3, 1711–1718. doi: 10.1109/LRA.2018.2801939

Ruiz-Ugalde, F., Cheng, G., and Beetz, M. (2010). "Prediction of action outcomes using an object model," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Taipei), 1708–1713. doi: 10.1109/IROS.2010.5649552

Ruiz-Ugalde, F., Cheng, G., and Beetz, M. (2011). "Fast adaptation for effect-aware pushing," in *2011 11th IEEE-RAS International Conference on Humanoid Robots* (Bled), 614–621. doi: 10.1109/Humanoids.2011.6100863

Şahin, E., Çakmak, M., Doğar, M. R., Uğur, E., and Üçoluk, G. (2007). To afford or not to afford: a new formalization of affordances toward affordance-based robot control. *Adapt. Behav.* 15, 447–472. doi: 10.1177/1059712307084689

Salganicoff, M., Metta, G., Oddera, A., Sandini, G., Salganico, M., Metta, G., et al. (1993). "A vision-based learning method for pushing manipulation," in *AAAI Fall Symposium Series on Machine Learning in Vision: What Why and How?* (Raleigh, NC).

Scholz, J., Levihn, M., Isbell, C., and Wingate, D. (2014). "A physics-based model prior for object-oriented MDPs," in *Proceedings of the 31st International Conference on Machine Learning (ICML-14)* (Beijing), 1089–1097.

Sloman, A. (2006). "Polyflaps as a domain for perceiving, acting and learning in a 3-d world," in *Position Papers for 2006 AAAI Fellows Symposium* (Boston, MA: AAAI).

Stüber, J., Kopicki, M., and Zito, C. (2018). "Feature-based transfer learning for robotic push manipulation," in *2018 IEEE International Conference on Robotics and Automation (ICRA)* (Brisbane, QLD), 5643–5650.

Ugur, E., Oztop, E., and Sahin, E. (2011). Goal emulation and planning in perceptual space using learned affordances. *Robot. Auton. Syst.* 59, 580–595. doi: 10.1016/j.robot.2011.04.005

Walker, S., and Salisbury, J. K. (2008). "Pushing using learned manipulation maps," in *2008 IEEE International Conference on Robotics and Automation* (Pasadena, CA), 3808–3813.

Watters, N., Zoran, D., Weber, T., Battaglia, P., Pascanu, R., and Tacchetti, A. (2017). "Visual interaction networks: learning a physics simulator from video," in *Advances in Neural Information Processing Systems 30*, eds I. Guyon, U. V.

Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Long Beach, CA: Curran Associates, Inc.), 4539–4547.

Woodruff, J. Z., and Lynch, K. M. (2017). "Planning and control for dynamic, nonprehensile, and hybrid manipulation tasks," in *2017 IEEE International Conference on Robotics and Automation (ICRA)* (Marina Bay Sands), 4066–4073.

Yang, P., Sasaki, K., Suzuki, K., Kase, K., Sugano, S., and Ogata, T. (2016). Repeatable folding task by humanoid robot worker using deep learning. *IEEE Robot. Automat. Lett.* 2, 397–403. doi: 10.1109/LRA.2016.2633383

Yoshikawa, T., and Kurisu, M. (1991). "Indentification of the center of friction from pushing an object by a mobile robot," in *Proceedings IROS '91. IEEE/RSJ International Workshop on Intelligent Robots and Systems '91. Intelligence for Mechanical Systems*, Vol. 2 (Osaka), 449–454.

Yu, K. T., Bauza, M., Fazeli, N., and Rodriguez, A. (2016). "More than a million ways to be pushed. A high-fidelity experimental dataset of planar pushing," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Daejeon), 30–37.

Zhou, J., Bagnell, J. A., and Mason, M. T. (2017). A fast stochastic contact model for planar pushing and grasping: Theory and experimental validation. *arXiv[Peprint]. arXiv: 1705.10664*. doi: 10.15607/RSS.2017.XIII.040

Zhou, J., Paolini, R., Bagnell, J. A., and Mason, M. T. (2016). "A convex polynomial force-motion model for planar sliding: identification and application," in *2016 IEEE International Conference on Robotics and Automation (ICRA)* (Stockholm), 372–377.

Zhu, S., Kimmel, A., and Boularias, A. (2017). Information-theoretic model identification and policy search using physics engines with application to robotic manipulation. *arXiv[Preprint].arXiv:1703.07822*.

Zito, C., Kopicki, M., Stolkin, R., Borst, C., Schmidt, F., Roa, M. A., et al. (2013). "Sequential trajectory re-planning with tactile information gain for dextrous grasping under object-pose uncertainty," in *Proceeding of IEEE International Conference on Intelligent Robots and Systems (IROS)* (Tokyo), 2013–2040.

Zito, C., Stolkin, R., Kopicki, M., and Wyatt, J. L. (2012). "Two-level rrt planning for robotic push manipulation," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Algarve), 678–685.

Zrimec, T., and Mowforth, P. (1991). Learning by an autonomous agent in the pushing domain. *Robot. Auton. Syst.* 8, 19–29. doi: 10.1016/0921-8890(91)90012-A

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

frontiers in Robotics and AI

ORIGINAL RESEARCH
published: 06 February 2020
doi: 10.3389/frobt.2020.00006

Radiation Tolerance Testing Methodology of Robotic Manipulator Prior to Nuclear Waste Handling

Kaiqiang Zhang[1], *Chris Hutson*[2]*, *James Knighton*[2], *Guido Herrmann*[3] *and Tom Scott*[2]

[1] Department of Mechanical Engineering, University of Bristol, University Walk, Bristol, United Kingdom, [2] H.H. Wills Physics Laboratory, University of Bristol, Bristol, United Kingdom, [3] Department of Electrical and Electronic Engineering, University of Manchester, Manchester, United Kingdom

Dramatic cost savings, safety improvements and accelerated nuclear decommissioning are all possible through the application of robotic solutions. Remotely-controlled systems with modern sensing capabilities, actuators and cutting tools have the potential for use in extremely hazardous environments, but operation in facilities used for handling radioactive material presents complex challenges for electronic components. We present a methodology and results obtained from testing in a radiation cell in which we demonstrate the operation of a robotic arm controlled using modern electronics exposed at 10 Gy/h to simulate radioactive conditions in the most hazardous nuclear waste handling facilities.

Keywords: robot, manipulator, radiation, tolerance, gamma, waste, radioactive, exposure

OPEN ACCESS

Edited by:
William R. Hamel,
The University of Tennessee, Knoxville, United States

Reviewed by:
Jonathan M. Aitken,
University of Sheffield, United Kingdom
Deepak Trivedi,
General Electric, United States

***Correspondence:**
Chris Hutson
chris.hutson@bristol.ac.uk

Specialty section:
This article was submitted to Robotic Control Systems, a section of the journal Frontiers in Robotics and AI

Received: *18 September 2019*
Accepted: *15 January 2020*
Published: *06 February 2020*

Citation:
Zhang K, Hutson C, Knighton J, Herrmann G and Scott T (2020) Radiation Tolerance Testing Methodology of Robotic Manipulator Prior to Nuclear Waste Handling. Front. Robot. AI 7:6. doi: 10.3389/frobt.2020.00006

1. INTRODUCTION

The nuclear industry in the UK and worldwide increasingly seeks cost-effective methods to implement remote technologies to enable decommissioning of legacy waste treatment and handling facilities to reduce residual hazard, a practice highly recommended by the International Atomic Energy Agency (Iqbal et al., 2012). Whilst modern robots are in ubiquitous use in other industries, such as manufacturing, significant uptake is yet to take place in the nuclear industry which would benefit greatly from increased use of robotics, if implemented to carry out work too hazardous or difficult for human workers. Remote operations are required in legacy nuclear facilities for the purposes of inspection, characterization, cutting, dismantling, sorting, and segregating hazardous waste prior to the demolition of buildings. Commonly-encountered hazards in legacy nuclear facilities include aggressive chemical species, radioactive materials emitting alpha, beta or gamma radiation, and asbestos. Often an understanding of the nature and distribution of these types of hazards is required before decommissioning can take place, and so providing this improved situational awareness, and the ability to carry out tasks remotely is where new robotic technology can offer benefits to the industry.

A conservative regulatory framework and high safety standards make nuclear operators conservative in their approach to adopting new technology, and so any new systems must have their capabilities demonstrated and rigorously tested in a simulated environment prior to use. Many facilities at the end of their useful life were built in the 1950s and 60s, based on simple but effective technology.

The use of simple remote systems is not new to the industry, but uptake has been slow and restricted: Houssay (2000) describes several systems already trialed, including for routine monitoring and surveillance at Savannah River in the USA, and for cleaning of steam generators in generating plant at Indian Point 2 as far back as 1989.

Modern electronics enables more intelligent use of robotic systems, but this comes with its own challenges. It has been perceived by many in the industry that radiation will render any modern electronic components immediately inoperable, and so electronics have generally been avoided. Famously several of the robots used to explore the Fukushima Daiichi stricken reactor cores were subject to extremely high doses resulting in their failure very quickly. Such extremes of radiation are not common in the field of nuclear decommissioning, and so this paper aims to challenge the received wisdom that radiation and electronics do not mix. In the current work we demonstrate that it is indeed possible to use electronic technology in radioactive environments with successful outcomes. The notion that "no electronic equipment can survive radiation" and so cannot be used for any decommissioning task is not correct, and our experiments using a high activity cobalt-60 radiation source have shown that useful work can be done using sophisticated robots.

The radiation tolerance of many electronic components is of course of the utmost importance for robotic functionality (Garg et al., 2006). Often, researchers or manufacturers have attempted to make their electronic devices more radiation tolerant by altering the designs of integrated circuits, increasing the signal amplitudes or by simple shielding (Ferlet-Cavrois, 2011). In some instances however, the most cost effective solution for the use of electronics in radioactive environments is to use standard components, and plan for their replacement. Clearly, there are instances in which replacement would not be possible (e.g., space exploration or long term insertion in a reprocessing tank), yet by understanding the exposure environment posed by each application and the required component longevity, it is possible to develop appropriate solutions. Central to the assessment of each application is knowledge of the radiation tolerance of each component and also the system as a whole.

Many irradiation tests of electronic components examine individual integrated circuits or chips (for example, the work of Katz and Some, 2003; Nagatani et al., 2011; Ducros et al., 2017), whereas our trial sought to test the whole robotic arm, to better simulate what might happen in a real industrial environment containing highly radioactive material. A methodology is suggested for designing experiments to assess the system performance of an industrial robot whilst the robot is carrying out a dynamic task. In this way, it is possible to observe the degradation of the robotic system during its operation, and so the test provides a more meaningful assessment of operational challenges, compared to a performance assessment carried out on individual components or when the robot is stationary.

This is the very first work that assesses the system performance of an off-the-shelf industrial robot arm. It was not clear to the manufacturers whether the KUKA iiwa LBR800 would be operable at all when exposed to the levels of radiation contained within facilities handling Intermediate Level Waste (ILW). In this paper we describe an experiment to test the radiation tolerance of a robotic arm as an exploratory test to determine the off-the-shelf radiation tolerance of such a system, and to understand what improvements might be carried out to increase its suitability for decommissioning applications.

2. METHODOLOGY FOR PERFORMANCE ASSESSMENT OF INDUSTRIAL ROBOTS

Robotic systems containing electronic components are likely to suffer some form of damage causing altered functionality when exposed to radiation, and the effect will be related to the exposure dose. For many decommissioning tasks, the materials present emitting radiation will not be well-defined, and so measurement of the environmental dose rate may well be one of the robotic inspection tasks.

The effect of radiation on components is material dependent, and has been well-studied. Metal-oxide semiconductors (MOS) have electronic properties altered (Ma and Dressendorfer, 1989), elastomer materials used in seals can become embrittled (Wündrich, 1984), and optical components have been known to change their transparency and refractive indices over time (Brichard et al., 2001). Changes in the magnitude of measured errors in motor control were observed by Howard et al. (2018). Given enough radiation, these changes in mechanical, optical, or electronic properties can ultimately cause failure in susceptible components, potentially causing the robot to break down.

The impact of radiation on a robotic system can depend on its operational state, and a "stationary" assessment method may not demonstrate performance changes as an entire system. For systems operated in dynamic motion it is critical to guarantee the system performance of the whole robot (Aitken et al., 2018; Tsitsimpelis et al., 2019) to enable the robot to complete its tasks for a specific mission. Recently, different off-the-shelf industrial robots have been proposed, for example the iiwa 14 LBR820 suggested by Aitken et al. (2018) and the prototype systems funded by the UK Nuclear Decommissioning Authority (2019). The deployment of off-the-shelf robots avoids the need to design and manufacture special robotic arms for particular requirements, speeding up deployment on nuclear sites but most off-the-shelf industrial robots have not been tested in a radiation environment. For compliance with strict regulatory controls and to build safety case arguments, it is essential to qualify the system performance of an industrial robot arm during exposure from radioactive materials to reduce the risk of an accident.

This paper proposes a systematic methodology to assess the performance of an industrial robot consisting of the following steps:

1. Identification of the critical positions in the robotic arm containing (often electronic) components potentially susceptible to radiation exposure.
2. Planning the robot motion for specific applications considering any safety and physical constraints.
3. Measurement of exposure dose rates at each critical position.
4. Data collection and performance monitoring of available parameters and indicators during a repetitive motion of a planned trajectory whilst exposed to radiation.
5. Following an observed performance degradation, assessment of the point of failure using the "stationary" method.

Step 1. Due to high complexity of off-the-shelf industrial robots, it is often very hard to establish analytical models of

the systems which may be influenced by radiation (Howard et al., 2018). However, it is possible to identify the components which are least radiation tolerant and estimate their failure dose through individual component irradiation testing. Typically, the control processors and the sensor's integrated circuits (ICs) are considered as the least radiation-hardened components, according to the results of stationary assessments (Katz and Some, 2003; Nagatani et al., 2011; Ducros et al., 2017). It might also be necessary to consider the degradation of other materials such as elastomer polymers (for structural integrity; Wündrich, 1984) and optical fibers (for device-to-device communication; Brichard et al., 2001).

Step 2. An industrial robot needs to be assessed in a radiation environment at a required dose rate over a given exposure time (Katz and Some, 2003; Nagatani et al., 2011; Ducros et al., 2017; Tsitsimpelis et al., 2019). This is easy to achieve in typical stationary assessments as individual electronic components can be exposed to a constant dose rate. For a closer simulation of real operation we recommend the use of a dynamic test, in which the robot undertakes a repetitive motion, providing firm evidence about the failure mode to be encountered during real operation. A comparison of operating parameters collected during each cycle of repeating actions during periods of exposure and prior to exposure can be used to show or predict changes in robot performance. The trajectory of the robot must avoid obstacles in the event of a malfunction, particularly with radioactive sources.

Step 3. In practice, exposure dose rates vary depending on the proximity of each component of the robot to the radiation source, and the differences can be dramatic. It is important to accurately measure the dose rate at locations of interest, particularly the less radiation-tolerant components, such that performance can be correlated with exposure dose. The dose rate at each position should be measured over multiple dynamic task-cycles, requiring a real time radiometric instrument such as a diamond radiation detector as described by Hutson (2018).

Step 4. All possible measurements related to the performance of a robotic system should be recorded for later analysis of degradation from radiation exposure. Howard et al. (2018) recorded all the input and output signals of each sensing and controlling component at a high sampling rate of 1.25 kHz, during their test exposure using X-rays. However, such hardware level signals are often difficult to access in industrial robots, but it may be possible to observe at least the control errors from robotic controllers.

A reference data set collected whilst moving in the planned trajectory, but prior to any exposure is desirable to enable analysis of any immediate effects radiation may have on a robot's operation.

Step 5. The last step for assessing the industrial robot is to run a stationary performance of the robot to identify the total radiation tolerance of the system at a configured dose rate until a system failure occurs (as in Nagatani et al., 2011; Ducros et al., 2017). As a result, the assessment step identifies the extreme total dose-tolerance of the industrial robot and any specific hardware/software issues limiting the overall system.

The proposed methodology has been used in the current study for assessing the system performance of a KUKA iiwa 7 LBR800 robot arm positioned ∼1.60 m from a 20 TBq source of cobalt-60 (^{60}Co).

3. EXPERIMENTAL SET-UP

3.1. Robot Under Test

3.1.1. KUKA iiwa 7 LBR800 Robot

The KUKA iiwa 7 LBR800 robotic arm manufactured by KUKA Deutschland GmbH (2019) has been proposed for several uses in the nuclear industry, including decontamination of glove boxes. It has 7 rotational joints providing 7 degrees of freedom, and has a maximum payload of 7 kg with a 926 mm reach. The robot is highly flexible, allowing it to easily avoid obstacles. **Figure 1** shows the location of the robot's joints. Within each joint, the robot has three different types of sensor to enable measurements of temperature, angular position, and force-torque. At each joint, two encoders used for angular position measurements, to achieve good positioning performance of its end-effector at milli-meter level. The use of two encoders provides some redundancy which could be used to detect sensor failure, and comparison of their feedback is used for a robot control safety mechanism. Force-torque sensors are used to detect any external forces applied to the robot, making it "human-safe" for work in which humans collaborate with the robot.

3.1.2. Identification of the Least Radiation Tolerate Components

In each joint there are several electronic components such as the integrated circuits of motor drives, encoders, and torque sensors, which are all enclosed by an aluminum casing. These

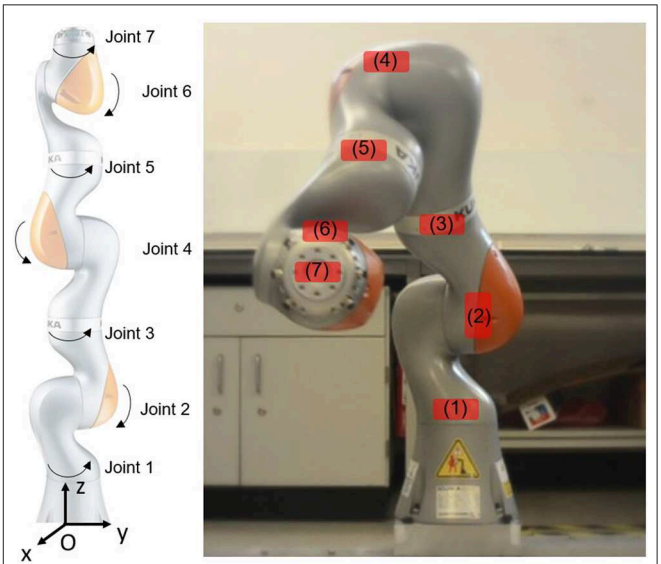

FIGURE 1 | The KUKA iiwa 7 LBR800 robot used in radiation tolerance testing. The left image shows the robot in its "zero" position status, with all the robot joints at their initial zero positions. The positive directions of each joint are indicated. The right-hand photo shows the robot during a mock-up test in which a periodic trajectory was programmed. The red colored areas show mounting locations used for the radiation detector.

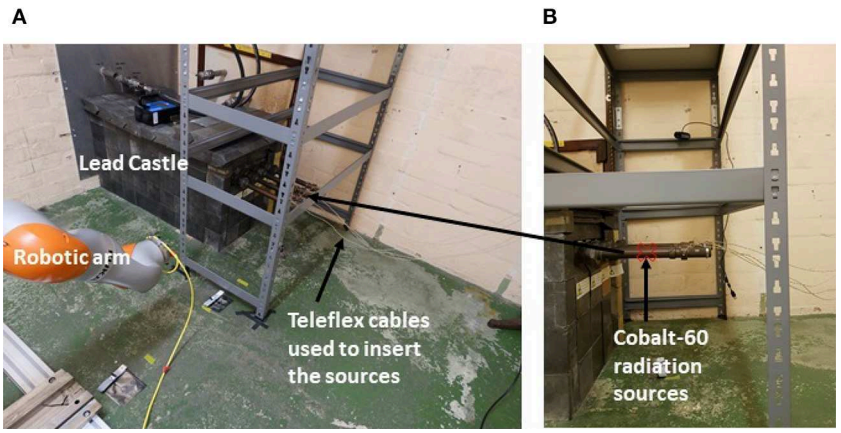

FIGURE 2 | (A) KUKA iiwa robotic arm in front of cobalt-60 source tubes **(B)**.

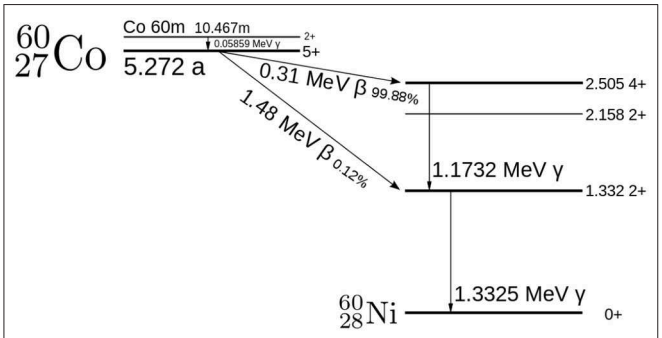

FIGURE 3 | Beta decay of $^{60}_{27}Co$ is followed by emission of gamma at 1.17 and 1.33 MeV from $^{60}_{28}Ni$ (Camp and Van Hise, 1976).

joint electronics constitute the locations of interest, and so dose rate measurements were made as shown in **Figure 1**. The most important of these was the "end-effector" responsible for carrying tools, grippers, and sensing packages. The end-effector is also likely to receive a higher dose than other components during operations dealing with radioactive material.

3.1.3. Irradiation Test Source

The radiation tests were carried out at the Medical Research Council (MRC) Harwell facility. As illustrated in **Figure 2**, the radiation source was made up of four ^{60}Co sources with a combined activity of around 20 TBq. The ^{60}Co sources produce intense gamma radiation via a decay to ^{60}Ni, as shown in the decay scheme in **Figure 3**. The source can roughly be considered a point source located at the position marked in **Figure 2B**.

During decay, beta particles produced were attenuated by the source tube casings, leaving just gamma photons with energies 1.17 and 1.33 MeV. ^{60}Co is used as a convenient radiation source because it is present in nuclear waste and emits photons close in energy to those emitted during the decay of ^{137}Cs, the main contributor to the gamma emitted by long-lived nuclear waste. To expose the cell, the four ^{60}Co sources were pushed with teleflex cables through containment tubes, passing into the cell via a lead castle into four protruding exposure tubes. The sources could be withdrawn at any time during the experiment.

3.2. Robot Trajectory Planning for Dynamic Performance Test

The robotic arm was anchored in place attached to a heavy base designed to stabilize the robot in the event of an unexpectedly high momentum should the robot control fail catastrophically. The robot was positioned so that all objects were out of reach as shown in **Figure 2**.

The robot's x-z plane was aligned with the sources and the source-tubes were parallel to the robot's y-axis (see the robot's original coordinates in **Figure 1**). The robot's original position was about 1.6 m from the center of the sources along its x-axis.

The robot is controlled to follow a defined trajectory, simulating the robot executing a repetitive task. The continuous motion of the robot arm is designed in a way to ensure each of the motors is active at all times. Therefore, this evaluates whether each joint could remain physically capable of moving during exposure (see results in section 4). The robot's end-effector is considered the most important component of the system, and has a target exposure of ∼10 Gy/h. Ideally, the robot's end-effector would be controlled to follow an arc-trajectory retaining a constant distance toward to the assumed point source. Practically, this is not possible since very small changes in the distance from end-effector to source result in noticeable changes in dose rate. These changes were easily overcome using dosimetry operated at a frequency of 20 Hz.

From the center of the point source, the flux of gamma radiation decays as distance from the source increases following the inverse-square law. The robot's end-effector trajectory was planned to follow an arc-trajectory of radius 1 m from the assumed point source. The end-effector was moved along the arc-trajectory at a speed of 20 mm/s in a time period of 1 min, sweeping around the radiation sources in a repeating pattern. The trajectories were computed via a Jacobian

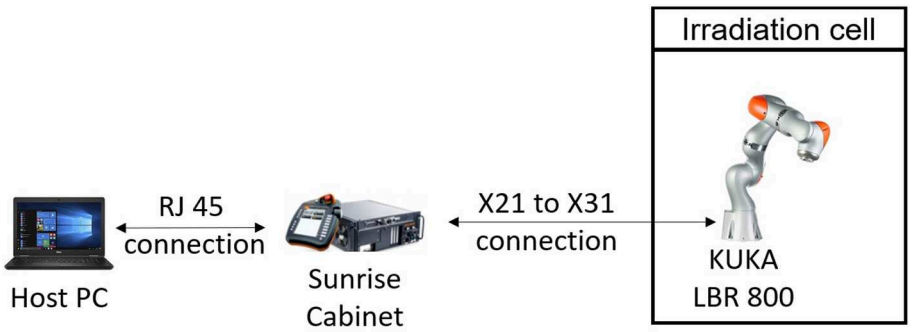

FIGURE 4 | Set-up enabling the control and data recording of the LBR800 robot.

matrix-based inverse kinematic algorithm as described by Meredith and Maddock (2004).

The robotic arm was connected by a robot control cable (X21-to-X31 cable) to a KUKA Sunrise cabinet control box, as shown in **Figure 4**. The control box was situated outside of the irradiation room (behind concrete walls), which meant it experienced little radiation dose. A host PC was connected via an RJ45 connection to the control box. In command of the 1-min cycle, the host PC instructed new position values for each joint to the control box, which passed these to the relevant joints. Specifically, a C++ control program was developed to enable the real-time communication between the standard host PC and the Sunrise Cabinet. The control program employed an application programming interface, called KUKA Fast Robot Interface (KUKA Deutschland GmbH, 2019). This design minimized the communication latency and, therefore, allowed for 100 Hz control/data-acquisition rate.

3.3. Dose Rate Measurement at Critical Positions

The primary purpose of the test was to understand radiation tolerance, and this required measurement of the robot's exposure dose. A real time measurement was required to accurately quantify the dose accrued prior to failure by each component of interest. A passive dosimetry measurement would not have been suitable for this test since the time to failure could not be predicted, and the failure could occur when experimenters were not present to withdraw the radiation sources and end the exposure.

The measurement required is particularly onerous for radiation detectors, since the exposure was sufficient to affect electronic performance. The most suitable available detector was a diamond radiation detector calibrated for air kerma dose rate measurement, a system already in use at Sellafield and described by Hutson (2018). This detector was chosen over other semiconductor detectors and scintillation detectors for its excellent radiation tolerance, and was already suitably calibrated for the target exposure dose rate. The detector used a single crystal diamond grown by chemical vapor deposition measuring 4.5 × 4.5 × 0.5 mm. The diamond detector system applies dose rates calibrations obtained using both caesium-137 and cobalt-60, and produces values for air kerma dose rates every 50 ms. Operating in current mode, the detector has no known upper dose rate limit and has proven extremely tolerant to gamma radiation.

The detector was placed at a series of positions along the robotic arm in turn, to quantify the exposure of each bundle of electronic circuitry and sensors, as shown in **Figure 5**. For each position, real time dose rates were measured over 10 min. Once the exposure to each joint had been measured, the diamond detector was attached to the robot's end effector for the remainder of the experiment to continue real time radiometric measurement. The end effector component was subject to the largest radiation dose and it was of most interest for the test, because in a real nuclear decommissioning operation would hold any cutting tool or sensing package.

3.4. Identifying Robot Performance Degradation

There were several indicators to help identify performance degradation and failure of the robotic arm:

1. Demanded and measured joint torque and position values were recorded in real time into a data spreadsheet at a sampling frequency of 100 Hz. Prior to the irradiation, the values were logged for the robot's arc trajectory for a reference set of values that could be compared to the values recorded during the irradiation. Any differences in the pre-irradiation and irradiation values would illustrate degradation in robot performance.
2. Two webcams recorded continuous video footage throughout the irradiation experiment. Example images are given in **Figure 6**, showing the footage could be used to monitor the robot in real time and identify any unexpected movements from the robotic arm.
3. The KUKA system displayed error messages of the robotic arm. It was expected that the robot's software would act to shut down the system for safety reasons in the event of a component failure.

In summary, we followed the methodology described herein consisting of: (1) identifying vulnerable components; (2) programming a repeating trajectory; (3) measuring the exposure

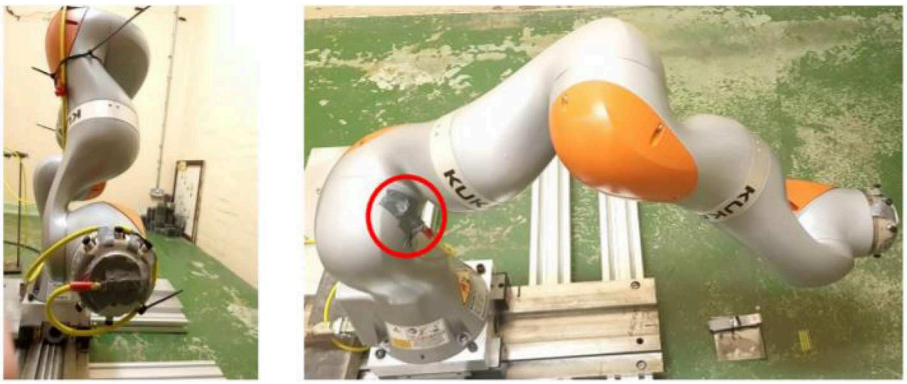

FIGURE 5 | Exposure dose rates varied across the length of the robotic arm, hence the diamond detector was mounted at different locations on the KUKA iiwa robotic arm during the radiation tolerance test, to enable the dose rate measurements shown in **Figure 7**.

FIGURE 6 | The robotic arm was observed during dynamic irradiation test using two webcams located inside the irradiation cell. Speckle may be observed in both images.

dose rates of vulnerable components; (4) measuring degradation; and (5) observing system failure.

4. RESULTS AND DISCUSSIONS

4.1. Exposure Dose Rate Measurement

Necessarily at different distances from the radiation sources, and moving around an arc, each set of sensors and actuators was exposed to different fluxes of radiation. Therefore, for each location of interest on the robot a separate dose rate profile was measured, to characterize the exposure conditions for each joint. Measured using a diamond radiation detector, these dose rate profiles are shown in **Figure 7**.

4.2. Robot Performance up to Failure

During the initial dynamic experiment, the robot was controlled to follow the planned trajectory repetitively for ∼6.3 h. This was a deliberately repetitive action, aimed to provide a detailed understanding of any changes in parameters occurring as a result of accumulated radiation damage. These chronic symptoms were important to understand in case they had any bearing on the overall control of the robot.

- No changes to the robot trajectory as a result of exposure were large enough to notice visually using the webcam images.
- No changes to the robot trajectory as a result of exposure were noticeable by changes in the dose rate profile by comparing a dosimetry profile at the beginning and end of the dynamic assessment.
- Subtle changes to the standard deviation of the Joint 2 control error began to occur after 5 h of this test (after a total of 8 h exposure) (shown in red in **Figure 8**).
- No other joints suffered the same increase in control error during irradiation, leading us to believe it may have been caused by a slight original defect in Joint 2, not necessarily caused solely by radiation damage.

4.3. Acute Failure

The robot failed during its stationary exposure, with the cause being a damaged optical encoder in the end effector (Joint 7). This component converts the angular position of the joint to a digital signal for feedback to the controller. The damaged component was diagnosed by the KUKA control system itself, with a message "encoder error torque sensor" and "safe axis position invalid," and as a result the controller prevented further operation of the

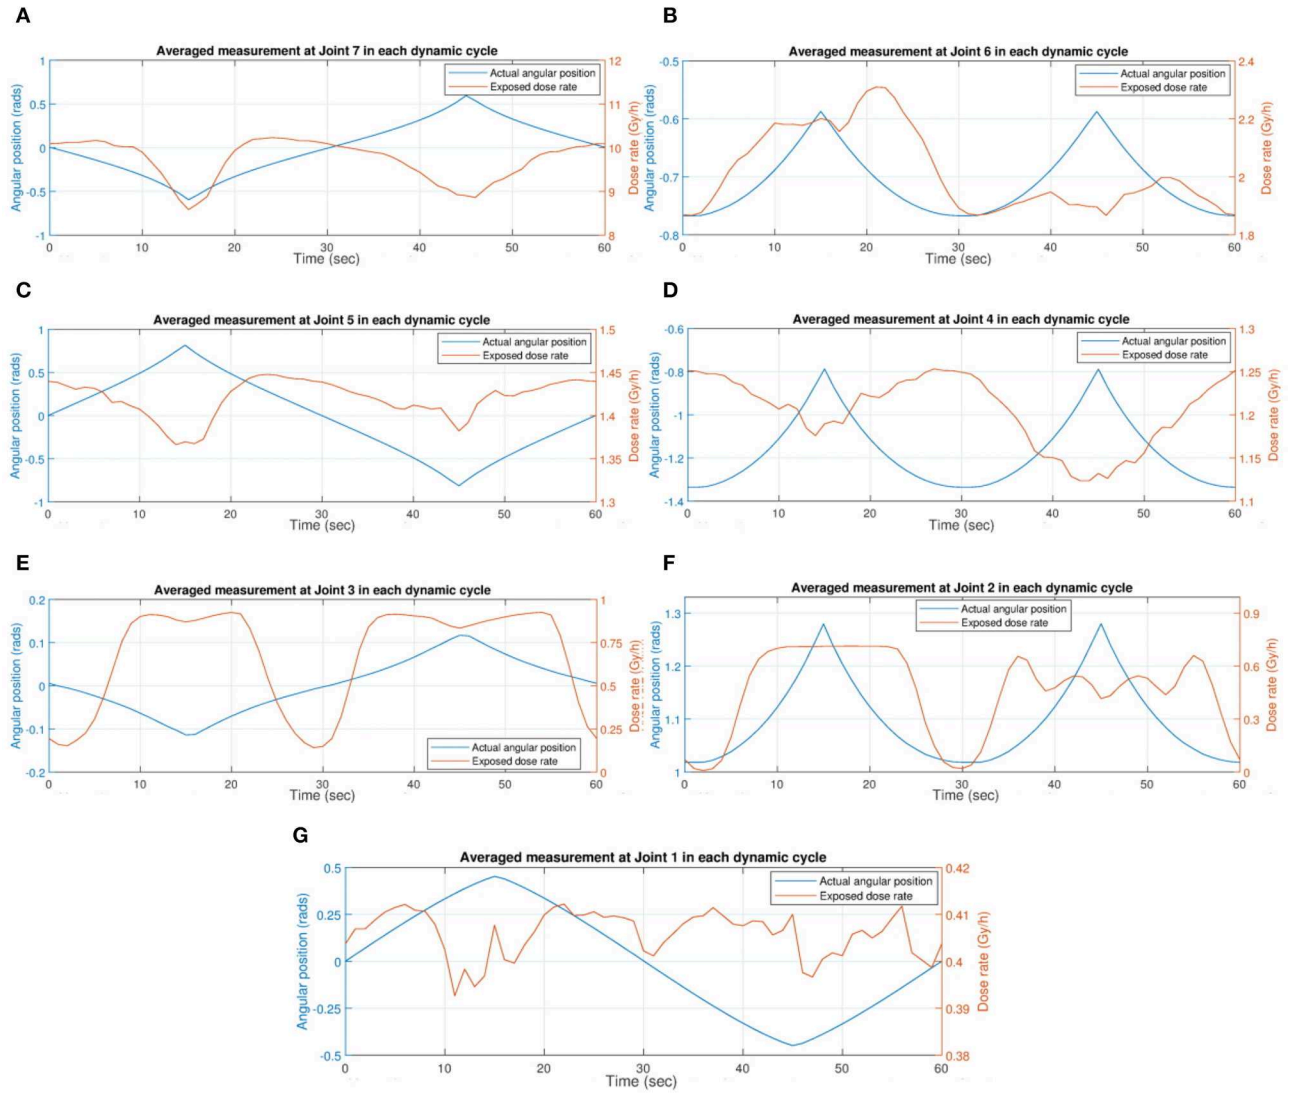

FIGURE 7 | The averaged dose rate and angular position measurements of each dynamic assessment cycle (for 60 s) over 10 cycles. The angular position and exposed dose rates are shown for each joint: **(A)** at Joint 7; **(B)** at Joint 6; **(C)** at Joint 5; **(D)** at Joint 4; **(E)** at Joint 3; **(F)** at Joint 2; and **(G)** at Joint 1.

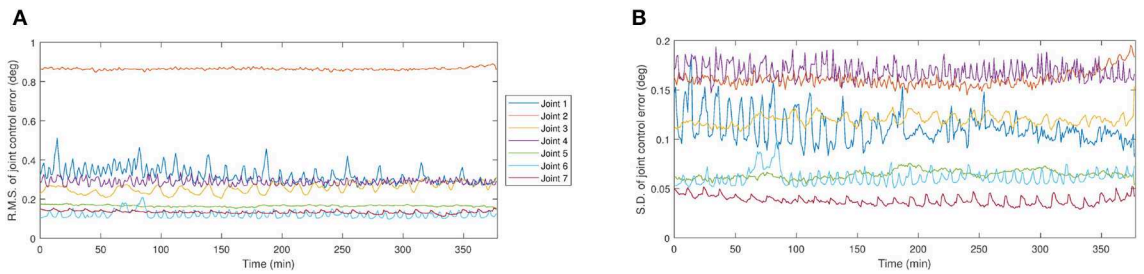

FIGURE 8 | Evaluation of the angular position control performance of the robot at each joint. **(A)** The root mean square (R.M.S.) of the control error at each joints. **(B)** The standard deviation (S.D.) of the control error at each joints.

robot. Attempts to reboot, remaster, and the passage of time have been unsuccessful in recovering robot movement: this failed component was permanently destroyed.

4.4. Exposure Dose Rate and Total Dose

The dose rates measured by the diamond detector dosimetry system are shown for each joint in **Table 1** for both dynamic and

TABLE 1 | The air kerma exposure dose-rate at each robot joint, measured using the diamond dose rate system.

	Exposed dose per dynamic cycle (Gy)	Dose rate during dynamic assessment (Gy/h)	Dose rate during stationary assessment (Gy/h)
Joint 7	0.1619	9.71	9.68
Joint 6	0.0336	2.02	1.93
Joint 5	0.0237	1.42	1.41
Joint 4	0.0201	1.21	1.16
Joint 3	0.0113	0.68	0.92
Joint 2	0.0077	0.46	0.53
Joint 1	0.0068	0.41	0.41

TABLE 2 | The air kerma exposure dose at each robot joint, measured using the diamond dose rate system.

	During dynamic assessment (Gy)	During stationary assessment (Gy)	Overall exposed dose before failure (Gy)
Joint 7	91.83	72.72	164.55
Joint 6	19.08	14.47	33.55
Joint 5	13.42	10.58	24.00
Joint 4	11.38	8.70	20.08
Joint 3	6.40	6.87	13.27
Joint 2	4.35	3.95	8.30
Joint 1	3.84	3.08	6.91

stationary assessments. These measurements have been used to calculate the overall exposure to each joint as detailed in **Table 2**.

The uncertainty of the dose measurement is ± 0.6 Gy in the dynamic assessment for about 9.3 h. Within the following static assessment, the exposure dose measurement has an uncertainty of ± 0.49 Gy over a period of about 7.5 h. The system exposure dose has approximate ± 1.09 Gy uncertainty.

5. DISCUSSION

The KUKA LBR800 robot stopped operating after a large radiation dose of 164.55($\pm$1.09) Gy to its end effector, and the component causing the failure was an optical encoder. The failure of this component was noted by the control software and smartPAD controller, which subsequently prevented the robot operating. The inbuilt smart software features were able to take excellent control of the situation, and we were able to demonstrate that once the encoder failed, Kuka's software locked down the robot in a safe state. This fail-safe mechanism in software would not have been possible in other more traditional types of robots using fewer electronic components, so this software approach should be considered as a significant safety benefit to any nuclear operator should such faults occur on a nuclear licensed site. Our test demonstrates that the standard safety features help ensure that a nuclear material handling accident due to a slowly failing system would not be possible.

It would be useful in a nuclear facility to maintain a cumulative dose reading on a joints (using diamond detectors or other similarly-sized miniature detectors) to ensure the system can be given preventative maintenance or component replacement, at say 75% of its dose-to-failure lifetime, rather than waiting for device failure at an inconvenient stage in a process.

The target dose rate of 10 Gy/h was chosen as a conservatively high value of exposure for ILW facilities, and would in reality more closely represent dose rates found in facilities dealing with high level waste. Contact dose rates encountered in ILW-handling facilities are generally 1 Gy/h and below, and so to simulate ILW facilities more closely it would be appropriate to use lower exposure doses for future testing programmes.

This paper presents a methodology to test off-the-shelf robots in radiation environments at a system level. Such system level tests are of significant benefits providing reference data to deploy the robot into practical operations, thereby building confidence that the system has potential to be used in radioactive environments. The system-level tests also allow for identifying the least radiation-tolerant component. Clearly, the system-level test is the necessary initial assessment in order to apply an off-the-shelf system in a radiation environment.

Note that investment required to buy complete off-the-shelf systems such as the robot tested in this work is significant. Therefore, although it is important to carry out system-level tests to build confidence in environments found in real applications, the cost of destroying large numbers of robots would be prohibitive. Hence, the following methodology is recommended to generate reasonable confidence of expected lifetime:

1. System level test to verify the off-the-shelf robot has the potential to satisfy initial design requirements in a radioactive environment.
2. Identification of particular components susceptible to radiation, through component-by-component analysis of irradiated robot.
3. Exposure of a statistically significant number of the identified susceptible components, allowing an estimation of system service life as a function of a variety of conditions, such as exposure at several dose rates, different temperatures and other robot load/operation modes. In a comparable study, two samples of each individual components were tested to develop a radiation hardened robot for the nuclear industry (Sharp and Decreton, 1996). In a good compromise between statistical rigor and cost, Oomichi et al. (2007) tested 7 to 30 samples of different components. For component-level testing, we recommend at least 20 samples be irradiated and analyzed to enable rigorous statistical analysis.

5.1. Recommendations for Further Work

Due to the conservative approach to utilizing new technology in the nuclear industry, it is of prime importance that the radiation tolerance of novel technology is understood fully prior to use. In this experiment the gamma radiation failure dose for the KUKA robot arm was 164 Gy. However, due to the probabilistic nature of photon interaction with matter, there is a probability (which is based on photon energy and electronic material) that when a photon is incident on an electronic device within the robotic arm, it will not be absorbed. Furthermore, the manufacturing

process of electronic components produces a distribution of characteristics, so if the experiment were repeated, the failure dose may be higher or lower than we measured. Therefore, further irradiation tests both of several complete robotic arm systems would be useful to confirm the failure dose result. Clearly, this would have significant cost implications, so we recommend that instead, irradiation testing of a large number of the least radiation tolerant electronic components (e.g. the optical encoder) could be performed. This would still provide an accurate figure on the radiation tolerance of the whole system.

A second recommendation is to fix/replace the broken encoder in the end effector of the robot and attempt to restore the robot's full functionality. After the robot is remastered, more irradiation tests can be performed. This would be useful to industry as it would demonstrate that the robot can be fixed and redeployed, and it will provide more irradiation data for the arm.

Consideration of shielding of particularly vulnerable components inside the robot would increase the system's lifetime in a radioactive environment. It would be beneficial to trial some micro-shields around components such as the optical encoder in each joint, constrained by the impact of the additional weight from shielding which will reduce the possible payload of any sensors or actuators (e.g., grippers or cutters). Selective shielding could, at limited cost, substantially increase the lifetime of the robot.

Component replacement with radiation tolerant alternatives should be considered if the off-the-shelf robot was unsuitable. For example, the function performed by the optical encoders could be carried out by rotary encoders, which are known to be less susceptible to radiation damage. A cost benefit analysis would consider the added cost to implement the new components and the money saved by having an increased lifetime of the robot.

Small alterations in the software safety features within the robotic arm could be altered to allow the robot to recover should a joint fail whilst in the middle of undertaking a task. A robot with as many as seven joints has significant kinematic redundancy in the positions within reach, so it is possible for the robot to still complete its task without a joint in operation. This would extend the lifetime of the robot further, alternatively it could enable the robot to complete its immediate tasks and afterwards return to a safe state ready for corrective maintenance.

Future studies could model/simulate the damage caused to the robot's electronics using Monte Carlo modeling software such as the Geant4 package developed by CERN. Previous research similar to this has been performed for the application of radiation damage of electronics used in space by Feng et al. (2007) and Xiao et al. (2018). Such modeling would require more detailed knowledge of the distribution of radiation sources than is generally available for decommissioning nuclear facilities.

6. CONCLUSIONS

The current work has investigated the controlled exposure of a KUKA iiwa LBR robot to gamma radiation to determine its tolerance and performance in highly radioactive environments analogous to nuclear waste processing and storage facilities.

The robot was exposed to gamma radiation from a 20 TBq cobalt-60 source and displayed significant radiation tolerance, with failure occurring in an optical encoder after a cumulative exposure of 164.55 Gy over a period of 16.8 h.

The results indicate that force-torque robots, which offer an enhanced level of finesse for manipulating objects, are potentially viable for nuclear waste processing applications. Used in appropriate applications, robotic technology using modern sensing and control software has the potential to make a large impact in the sector in terms of cost savings, safety and shorter decommissioning timescales. Future work should consider alternate radiation-tolerant replacements for optical encoders, and also examine methods of micro-shielding of vulnerable components to enhance the performance lifetime of the robotic systems.

DATA AVAILABILITY STATEMENT

All datasets generated for this study are included in the article/supplementary material.

AUTHOR CONTRIBUTIONS

TS, CH, and KZ contributed conception and design of the study. CH organized the irradiation exposure and measured the dose rates. JK and KZ programmed the control, data collection of the robotic arm, performed the data processing, and statistical analysis. CH, JK, and KZ each wrote major sections of the manuscript. Both KZ and CH supervised JK. TS and GH are the principle investigators associated with this project. All authors contributed to manuscript revision, read, and approved the submitted version.

FUNDING

We would like to thank UK Research and Innovation (UKRI) for their support and funding which has come via the National Centre for Nuclear Robotics (NCNR) grant EP/R02572X/1; from Robotics and Artificial Intelligence in Nuclear (RAIN) grant EP/R026084/1; and from Rad-Hard Diamond Detectors for Civil Nuclear Applications (grant ST/P001823). Funding is also acknowledged from Jacobs Engineering Group.

ACKNOWLEDGMENTS

The authors would like to thank the support from KUKA AG., providing the application interface software and suggesting potential points of failure. We would like to thank Bob Sokolowski of the Medical Research Council who supported experimental access to radiation facilities at Harwell. We thank Estelle Talfan Davies for proofreading this manuscript.

REFERENCES

Aitken, J. M., Veres, S. M., Shaukat, A., Gao, Y., Cucco, E., Dennis, L. A., et al. (2018). Autonomous nuclear waste management. *IEEE Intell. Syst.* 33, 47–55. doi: 10.1109/MIS.2018.111144814

Brichard, B., Borgermans, P., Fernandez, A. F., Lammens, K., and Decreton, A. (2001). Radiation effect in silica optical fiber exposed to intense mixed neutron-gamma radiation field. *IEEE Trans. Nucl. Sci.* 48, 2069–2073. doi: 10.1109/23.983174

Camp, D., and Van Hise, J. (1976). Weak gamma rays observed in the co 60 decay. *Phys. Rev. C* 14:261. doi: 10.1103/PhysRevC.14.261

Ducros, C., Hauser, G., Mahjoubi, N., Girones, P., Boisset, L., Sorin, A., et al. (2017). RICA: A tracked robot for sampling and radiological characterization in the nuclear field. *J. Field Robot.* 34, 583–599. doi: 10.1002/rob.21650

Feng, Y.-J., Hua, G.-X., and Liu, S.-F. (2007). Radiation hardness for space electronics. *J. Astronaut.* 28, 1071-1080

Ferlet-Cavrois, V. (2011). "Electronic radiation hardening—radiation hardness assurance and technology demonstration activities," in *JUICE Instrument Workshop* (Darmstadt).

Garg, R., Jayakumar, N., Khatri, S. P., and Choi, G. (2006). "A design approach for radiation-hard digital electronics," in *Proceedings of the 43rd Annual Design Automation Conference* (San Francisco, CA: ACM), 773–778.

Houssay, L. P. (2000). *Robotics and radiation hardening in the nuclear industry* (Ph.D. Thesis). State University of Florida, Gainesville, FL, United States.

Howard, J., Barth, E., Schrimpf, R., Reed, R., Adams, L., Vibbert, D., et al. (2018). Methodology for identifying radiation effects in robotic systems with mechanical and control performance variations. *IEEE Trans. Nucl. Sci.* 66, 184–189. doi: 10.1109/TNS.2018.2886242

Hutson, C. (2018). *Diamond-based dosimetry for measurements in highly radioactive nuclear environments* (Ph.D. Thesis). University of Bristol, Bristol, United Kingdom.

Iqbal, J., Tahir, A. M., ul Islam, R., et al. (2012). "Robotics for nuclear power plants—challenges and future perspectives," in *2012 2nd International Conference on Applied Robotics for the Power Industry (CARPI)* (Zurich: IEEE), 151–156.

Katz, D. S., and Some, R. R. (2003). NASA advances robotic space exploration. *Computer* 36, 52–61. doi: 10.1109/MC.2003.1160056

KUKA Deutschland GmbH (2019). *LBR iiwa 7 r800, LBR iiwa 14 R820 Specification*.

Ma, T.-P., and Dressendorfer, P. V. (1989). *Ionizing Radiation Effects in MOS Devices and Circuits*. John Wiley & Sons.

Meredith, M., and Maddock, S. (2004). *Real-Time Inverse Kinematics: The Return of the Jacobian*. Technical report, Technical Report No. CS-04-06, Department of Computer Science, University of Sheffield.

Nagatani, K., Kiribayashi, S., Okada, Y., Otake, K., Yoshida, K., Tadokoro, S., et al. (2011). "Gamma-ray irradiation test of electric components of rescue mobile robot quince," in *2011 IEEE International Symposium on Safety, Security, and Rescue Robotics* (Kyoto: IEEE), 56–60.

Nuclear Decommissioning Authority (2019). *Cleaning Up Our Nuclear Past: Faster, Safer and Sooner*.

Oomichi, T., Isozaki, Y., and Kojima, M. (2007). Practical design of robots operating in radiation environments. *Adv. Robot.* 21, 515–532. doi: 10.1163/156855307780108286

Sharp, R., and Decreton, M. (1996). Radiation tolerance of components and materials in nuclear robot applications. *Reliabil. Eng. Syst. Saf.* 53, 291–299. doi: 10.1016/S0951-8320(96)00054-3

Tsitsimpelis, I., Taylor, C. J., Lennox, B., and Joyce, M. J. (2019). A review of ground-based robotic systems for the characterization of nuclear environments. *Prog. Nucl. Energy* 111, 109–124. doi: 10.1016/j.pnucene.2018.10.023

Wündrich, K. (1984). A review of radiation resistance for plastic and elastomeric materials. *Rad. Phys. Chem.* 24, 503–510. doi: 10.1016/0146-5724(84)90185-7

Xiao, H., Hajdas, W., Beauvivre, S., Kraehenbuehl, D., Ziethe, R., and Banerji, N. (2018). "Simulations of radiation damage in spacecraft camera for ESA JUICE mission," in *2018 International Conference on Radiation Effects of Electronic Devices (ICREED)* (Heilongjiang: IEEE), 1–4.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

frontiers in Robotics and AI

ORIGINAL RESEARCH
published: 22 May 2020
doi: 10.3389/frobt.2020.00052

Simultaneous Material Segmentation and 3D Reconstruction in Industrial Scenarios

Cheng Zhao[1]*, *Li Sun*[2] *and Rustam Stolkin*[1]

[1] *Extreme Robotics Lab, University of Birmingham, Birmingham, United Kingdom,* [2] *Visual Computing Group, University of Sheffield, Sheffield, United Kingdom*

Recognizing material categories is one of the core challenges in robotic nuclear waste decommissioning. All nuclear waste should be sorted and segregated according to its materials, and then different disposal post-process can be applied. In this paper, we propose a novel transfer learning approach to learn boundary-aware material segmentation from a meta-dataset and weakly annotated data. The proposed method is data-efficient, leveraging a publically available dataset for general computer vision tasks and coarsely labeled material recognition data, with only a limited number of fine pixel-wise annotations required. Importantly, our approach is integrated with a Simultaneous Localization and Mapping (SLAM) system to fuse the per-frame understanding delicately into a 3D global semantic map to facilitate robot manipulation in self-occluded object heaps or robot navigation in disaster zones. We evaluate the proposed method on the *Materials in Context* dataset over 23 categories and that our integrated system delivers quasi-real-time 3D semantic mapping with high-resolution images. The trained model is also verified in an industrial environment as part of the EU RoMaNs project, and promising qualitative results are presented. A video demo and the newly generated data can be found at the project website[1] (**Supplementary Material**).

Keywords: material segmentation, 3D material reconstruction, transfer learning, deep neural network, nuclear applications

OPEN ACCESS

Edited by:
Claudio Semini,
Italian Institute of Technology (IIT), Italy

Reviewed by:
Jonathan M. Aitken,
University of Sheffield,
United Kingdom
Sudarshan Ramenahalli Govindaraju,
Johns Hopkins University,
United States

***Correspondence:**
Cheng Zhao
irobotcheng@gmail.com

Specialty section:
This article was submitted to Sensor Fusion and Machine Perception, a section of the journal Frontiers in Robotics and AI

Received: *19 November 2019*
Accepted: *23 March 2020*
Published: *22 May 2020*

Citation:
Zhao C, Sun L and Stolkin R (2020) Simultaneous Material Segmentation and 3D Reconstruction in Industrial Scenarios. Front. Robot. AI 7:52. doi: 10.3389/frobt.2020.00052

1. INTRODUCTION

Materials recognition is in high-demand in many industries, such as nuclear waste decommissioning and recycling in a circular economy. Take robotic nuclear waste decommissioning as an example. The legacy of nuclear waste clean-up is one of the largest environmental remediation problems in the UK as well as in Europe. An estimated over 100 billion pounds will be spent on waste clean-up over a few decades (of Commons Committee of Public Accounts, 2013). Humans can handle radioactive waste but only for limited periods and by wearing special air-fed protection suits, which then become contaminated. In other words, conventional nuclear waste decommissioning turns becomes an open-ended problem as more nuclear waste is generated. For these reasons, autonomous robotic nuclear waste sorting and segregation will be the only solution for reducing secondary waste.

[1]https://sites.google.com/view/dense-semantic-mapping/home

Recognizing the material of which waste objects are composed is important in nuclear waste decommissioning, as different post-process and levels of segregation will be applied according to the material. For example, combustible materials (e.g., wood and clothing) can be burned, and deformable materials (e.g., rubber and plastic) can be melted and compressed. Our team is part of NCNR (the National Center for Nuclear Research) and works closely with the National Nuclear Lab on advanced robot perception and manipulation for waste decommissioning. This paper uniquely tackles the material recognition problem for the nuclear industry, and we propose a visual-based semantic segmentation approach to identify waste material categories in cluttered scenes.

Deep learning-based semantic understanding is the state-of-the-art in fundamental computer vision challenges, and large-scale annotation is required to learn a robust model to deal with the variability of the real world. However, in novel robotic applications, e.g., nuclear waste material recognition, very limited data can be provided because confidential nuclear data is not publicly available. Hence, leveraging public datasets and transferring the knowledge from other vision tasks to this novel application is highly desirable. Moreover, the capability to perform boundary-aware annotation and 3D semantic reconstruction can provide high-level semantic knowledge to robots, which allows the manipulator to dexterously fetch or remove waste objects from highly self-occluded heap or bin.

Facing these challenges, in this paper, we mainly focus on the following two problems. (i) Recognizing the material categories pixel-wise and simultaneously fusing per-frame recognition into a dense 3D map for robotic applications in the nuclear industry. (ii) Transferring knowledge from meta computer vision data to the material recognition problem and transferring knowledge from a relatively simple task (i.e., material categorization) to a more challenging task (i.e., boundary-aware material segmentation).

Specifically, we present a material semantic reconstruction system that can perform real-time 3D reconstruction while simultaneously recognizing and labeling each voxel according to its material in the generated dense 3D map. We evaluate the proposed approach using both a public material dataset and real-world industrial data from qualitative, quantitative, and running-time perspectives to verify the feasibility of the proposed system. The main contributions of this paper can be summarized as follows:

1. To the best of our knowledge, this is the first system to achieve simultaneous material recognition and dense scene reconstruction. It can integrate high-level semantic knowledge into conventional 3D geometry reconstruction.
2. The pixel-wise material segmentation is achieved via transfer learning from general object recognition to specific material recognition and from an image-wise classification task to a pixel-wise segmentation task. The proposed approach is end-to-end learned, without requirements for hand-crafted features or post-processing optimization.
3. The running-time performance of the well-implemented system can be boosted to around 10 Hz using a standard GPU, which is enough to deploy quasi-real-time material semantic reconstruction in industrial scenarios.
4. Because the large-scale material dataset, i.e., Materials in Context (MINC) (Bell et al., 2015), only provided very coarse annotated data for the material classification and segmentation, we generated high-quality new data including RGB image patches (821,092 patches for training, and 96,747 patches for testing) and fully pixel-wise annotated RGB images (1,498 images for training, and 300 images for testing). Those data are released as an important supplement of the MINC dataset for benchmarking material classification and segmentation research.

2. RELATED WORK

Vision-based material understanding, including classification, segmentation, and reconstruction, has as yet been little investigated, even though it is highly desirable for industrial robotics applications, e.g., nuclear robotics. As it must deal with the variation in brightness and illumination in the real world as well as learning a generalizable model from observations, material recognition in unconstrained environments is known to be a challenging research task.

CURet (Dana et al., 1999) was the first material dataset to be established. This consists of 61 material categories, and each category is captured with images taken under 205 different illumination and pose conditions. Eric et al. proposed the KTH-TIPS (Hayman et al., 2004) and KTH-TIPS (Caputo et al., 2005) datasets as supplementary to CURet, providing variations in scale in addition to in pose and illumination. The Flickr Material Database (Sharan et al., 2009) provides 10 different material categories, with 100 different samples for each category. The GeoMat (DeGol et al., 2016) dataset is the first dataset to provide material images with 3D geometry data. However, all of the above datasets are built for material classification rather than the pixel-wise material segmentation. The Materials in Context (MINC) (Bell et al., 2015) dataset is the first large-scale material dataset that is of good diversity and is well-sampled across 23 category materials. It provides two kinds of annotated data: RGB image patches and pixel-wise labeled RGB images. Moreover, a 4D light field material dataset proposed by Wang et al. (2016) captures the material images from multiple viewpoints through a light-field camera.

For material classification, most previous research has employed hand-crafted visual features, e.g., reflectance-based edge features (Liu et al., 2010), pairwise local binary patterns (Qi et al., 2014), local binary patterns (Li and Fritz, 2012), and variances in oriented gradients (Hu et al., 2011) for classifiers such as Fisher Vector (FV) (DeGol et al., 2016) and Support Vector Machines (SVMs) (Hayman et al., 2004; Caputo et al., 2005). Since deep learning dominates the computer vision community, deep-learned features (Schwartz and Nishino, 2013; Cimpoi et al., 2014) are also adopted to achieve state-of-the-art accuracy of material classification. Moreover, DeGol et al. (2016) not only employ 2D visual features, e.g., texture and color, but also 3D geometrical features, e.g., surface normals,

to improve the material classification. However, this research can only perform material classification with RGB patches, and pixel-wise material recognition, i.e., semantic segmentation of materials, is not achieved.

In order to achieve pixel-wise material recognition, Bell et al. (2015) converted CNN classifiers trained on image patches into an efficient fully convolutional framework with a fully connected conditional random field (CRF) for the material segmentation. Schwartz and Nishino (2016) took advantage of the abilities of both CNN and RNN to perform superior segmentation using local appearance and separately recognized global contextual cues, e.g., objects and places. Cimpoi et al. (2015) proposed a new texture descriptor, FV-CNN, obtained by Fisher Vector pooling of a CNN filter bank and achieved state-of-the-art performance on the Flickr Material Database (Sharan et al., 2009). Wang et al. (2016) utilized proposed new 4D light-field images to train an FCN with post-processing to achieve material segmentation. Further research (Giben et al., 2015; Purri et al., 2019) achieved interesting material segmentation applications on the Satellite and Railway Track images, respectively.

In contrast to our proposed approach, all of the above studies focus on material classification or segmentation without reconstruction so that the semantic information of material cannot be integrated into the 3D geometry map. The proposed approach in this paper can perform material segmentation and reconstruction simultaneously to generate a 3D semantic map. With the assistance of this high-level semantic (material) knowledge, a robot can perform robot-environment interactive tasks or motion planning in industrial scenarios.

3. METHODOLOGY

3.1. System Overview

This paper proposes a fully integrated system for material segmentation and reconstruction. It can perform real-time 3D dense mapping while simultaneously recognizing and labeling each point cloud in the map according to its material category. As **Figure 1** shows, the system consists of two main parts: single-frame material segmentation and 3D semantic reconstruction (mapping). To be specific, the RGB image captured by the RGB-D camera is fed into a Deep Neural Network (DNN) to achieve pixel-level material segmentation. The semantic point cloud is then generated using the data pair of the semantically labeled RGB image and the corresponding depth image via back-projection. A sequence of semantic point clouds are incrementally combined via visual odometry, and meanwhile, the label probability of each point is refined by Bayesian updating. Finally, a dense 3D semantic map indicating voxel-wise material categories is generated. Please note that the color palette used in all of the Figures in this paper can be found in **Figure 2**.

3.2. Dataset and Data Preprocessing

The Materials in Context (MINC) dataset (Bell et al., 2015) is used to train and evaluate the proposed neural network. MINC is diverse and well-sampled across 23 categories, including ceramic, fabric, leather, stone, wood, etc.

Nuclear waste can be categorized into fuel waste and technical waste, and both are radioactive. The technical waste makes up more than 97% of the total waste and includes all types of waste produced during power generation, for example, liquid containers (such as bottles, cans), disposable protective items (e.g., suits, masks, gloves), and even construction materials (e.g., bricks and wood) used in the nuclear power station. Because nuclear waste containers are very expensive and space in a container is limited, the waste will be processed according to its materials to best utilize the space in containers. For example, wood and clothes can be burned, and the ashes can be stored, while objects like plastic bottles and metal cans can be compressed into blocks with small volumes. Therefore, material recognition is a critical task for nuclear waste sorting and segregation.

This paper focuses on recognizing the material categories of nuclear technical waste and the challenges of dealing with the variation of materials (i.e., inter-class dissimilarities) and the variability of the real world (e.g., brightness and illumination). In addition, we cannot obtain real nuclear technical waste at the current stage for network training. For these reasons, we choose to use a large-scale material dataset, MINC, which includes most kinds of materials found in technical waste, to train and evaluate our model.

MINC provides two different types of annotations for training: a set of RGB patches with class labels, as shown in **Figure 3A**, and another set of partially pixel-wise labeled RGB images, as shown in **Figure 3B**. However, neither of these can be used directly for training end-to-end semantic segmentation network. There are many NaN values (shown as gray parts in **Figure 3**) in the original RGB patches, which give rise to a need for strong error supervision to prevent the classification network from converging. On the other hand, in the partially pixel-wise labeled RGB image, only the foreground object is labeled, whereas the background objects are masked. Thus, these images cannot provide contextual information during the training of the segmentation network.

Therefore, data preprocessing is applied to the MINC dataset. We first resize the original RGB image (500×500) and then extract RGB patches of different sizes (56×56, 156×156, 256×256, and 356×356) from it. This ensures that there are no NaN values in extracted patches and that only one type of material is at the center of each patch. Finally, 821,092 patches with corresponding class labels are generated for training, and 96,747 patches with class labels are generated for testing.

Next, we combine all of the partially pixel-wise labeled images that belong to the same original image to generate a single fully pixel-wise labeled image, as shown in **Figure 3D**. Since not all pixels are labeled in original images, the newly generated ground truth images are not 100% labeled. We further label all the unlabeled pixels, and repeated labeled pixels are treated as a category to be ignored during the training process. Finally, we generate 1,498 pixel-wise labeled training images and 300 pixel-wise labeled testing images. The size of the pixel-wise labeled ground truth images is also set to 500×500.

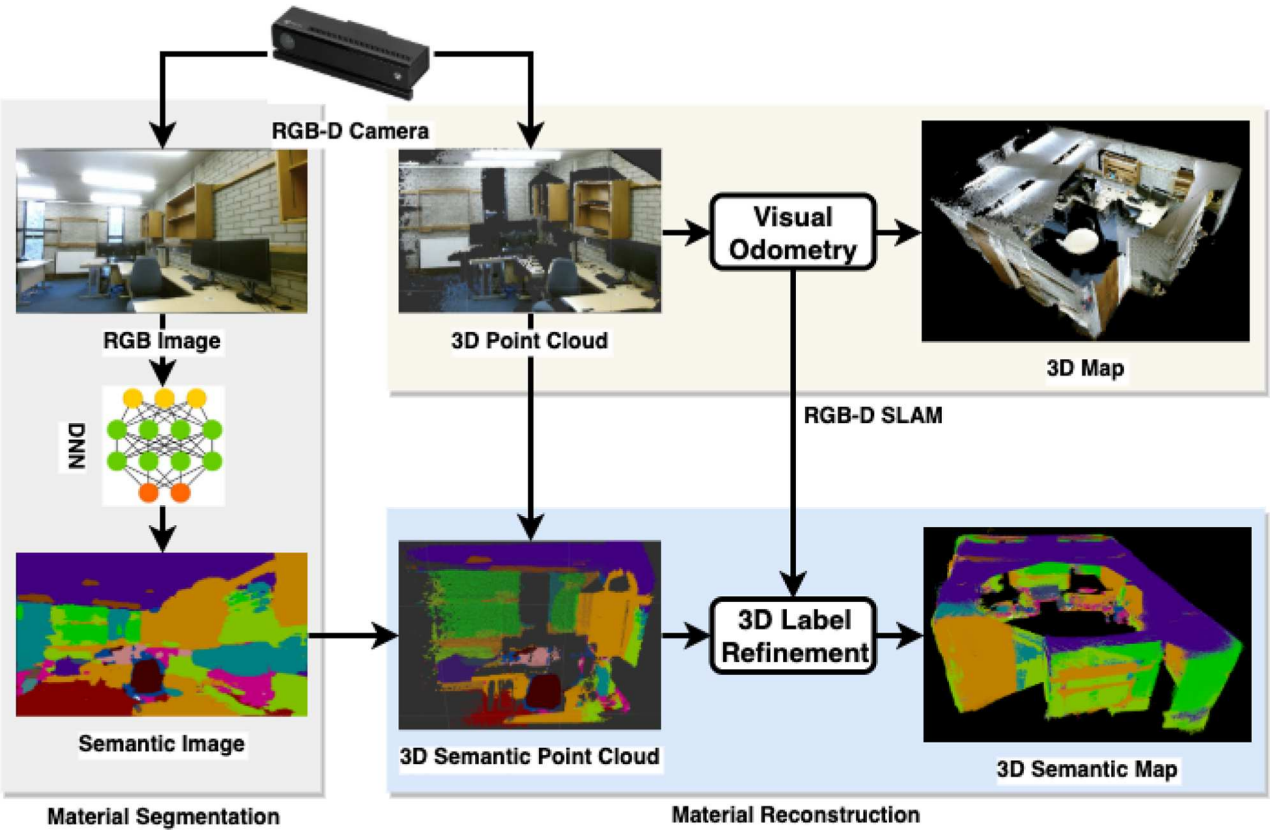

FIGURE 1 | The pipeline of the proposed system of simultaneous material segmentation and reconstruction.

FIGURE 2 | The color palette used in this paper.

3.3. Material Classification

We first train a deep classification network using the generated RGB patches with the corresponding class labels. The VGG-16 (Simonyan and Zisserman, 2014) network, consisting of five convolution stacks and three dense connect layers, is employed for the classification task. However, the VGG-16 network is designed for the ImageNet challenge[2] and thus can classify images into 1,000 object categories. We therefore modify the number of output nodes (i.e., the last dense connected layer of the VGG-16 network) to 23 instead of 1,000 for MINC material classification. Moreover, in order to accelerate the training, we transplant the weights of VGG-16 from the off-the-shelf pre-trained model[3] on ImageNet to our neural network.

We provide performance evaluation of the classification using different sizes of patches in the experiment section. For feature representation learning, small patches can provide more texture information, while fully annotated images can provide

[2]https://www.image-net.org/

[3]https://www.robots.ox.ac.uk/~vgg/research/very_deep/

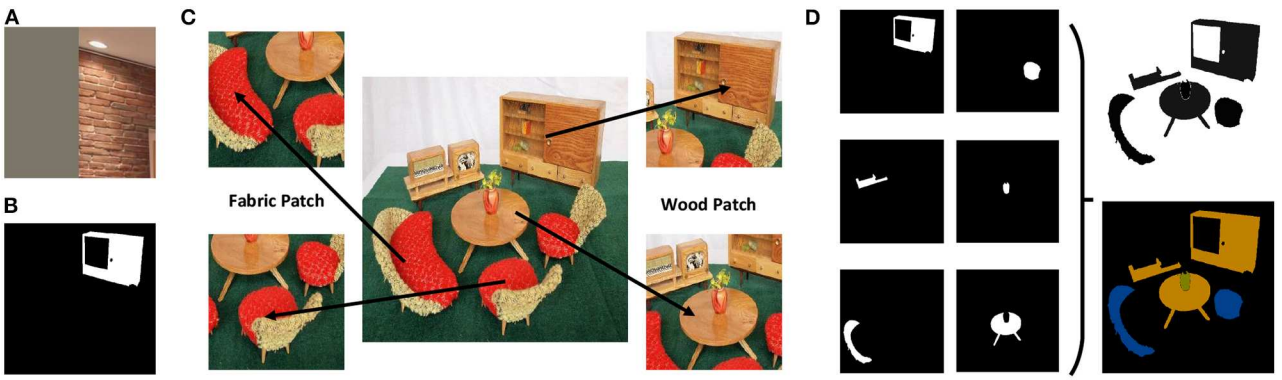

FIGURE 3 | Data preprocessing. **(A)** A patch with NaN values in MINC. **(B)** A partially pixel-wise labeled image in MINC. **(C)** Extracting new patches from original images in MINC. **(D)** Combining the partially pixel-wise labeled images to generate a fully pixel-wise labeled image.

more contextual information. Thus, the choice of patch size for the classification task is a trade-off between textural and contextual information.

3.4. Material Segmentation

Next, we train a segmentation network using the generated pixel-wise labeled RGB images. As **Figure 4** shows, the segmentation network consists of two sequential sub-networks: a Fully Convolutional Network (FCN) (Long et al., 2015) and Conditional Random Fields as Recurrent Neural Networks (CRF-RNN) (Zheng et al., 2015). FCN can provide a coarse semantic segmentation with prediction potentials to CRF-RNN, while CRF-RNN can smooth the label assignments between neighboring pixels to refine the coarse segmentation generated by FCN. Unlike the conventional approaches, which employ CRF as post-processing, we plugged in CRF-RNN after FCN as a unified framework that can be trained in an end-to-end way.

3.4.1. FCN

FCN is a widely used end-to-end and pixel-to-pixel semantic segmentation network that consists of a convolution stack, a deconvolution stack, and a skip architecture. The convolution stack has the same architecture as the VGG-16 network truncated after pooling five layers. It can learn high-level semantic features with context cues by enlarging the receptive fields. However, it cannot retain significant boundary information on objects and structures due to the application of a series of pooling layers. The deconvolution stack can transform a feature map of the same size as the input RGB image. The skip architecture combines high-level and coarse semantic features from deep layers with low-level and fine features from shallow layers. Therefore, FCN can improve the performance of semantic segmentation by fusing the feature maps from both deep and shallow layers. However, FCN does not incorporate smoothness constraints between neighboring pixels so that it can only give a coarse pixel-wise prediction with some blob-like shapes.

3.4.2. CRF-RNN

CRF-RNN means Conditional Random Fields as Recurrent Neural Networks, which is a hybrid model combining the learning property of CNN with the probabilistic graphical property of CRF. It can be inserted after a deep neural network to refine the coarse segmentation results generated.

Fully connected CRF (Krähenbühl and Koltun, 2011) takes account of contextual cues by minimizing the energy $E(x)$ function to generate the most likely label assignment x. There are unary energy components and pairwise energy components in the Energy function $E(x)$:

$$E(x) = \sum_i \psi_u(x_i) + \sum_{i<j} \psi_p(x_i, x_j). \quad (1)$$

The unary term $\psi_u(x_i)$ obtained from the FCN measures the inverse likelihood of each pixel i assigning the corresponding label. However, the predicted pixel labels do not consider the smoothness or consistency of label assignments between neighboring pixels j. In contrast, the pairwise term $\psi_p(x_i, x_j)$ can penalize similar pixels that have different labels and encourage similar labeling of pixels with similar properties.

Pairwise potentials can be modeled as a linear combination of M Gaussian edge potential kernels $k_G^{(m)}$ using different weights $\omega^{(m)}$:

$$\psi_p(x_i, x_j) = \mu(x_i, x_j) \sum_{m=1}^{M} \omega^{(m)} k_G^{(m)}(f_i, f_j). \quad (2)$$

The Gaussian kernel $k_G^{(m)}$ is applied to feature vectors f_i of pixel i, e.g., spatial or color information. The label compatibility function is described as a Potts model $\mu(x_i, x_j) = [x_i \neq x_j]$. The Gaussian kernel $k_G^{(m)}$ in the pairwise potentials consists of a bilateral appearance potential and a spatial smoothing potential ($M = 2$):

$$k(f_i, f_j)_G = \omega^{(1)} exp(-\frac{|p_i - p_j|^2}{2\theta_\alpha^2} - \frac{|I_i - I_j|^2}{2\theta_\beta^2}) + \omega^{(2)} exp(-\frac{|p_i - p_j|^2}{2\theta_\gamma^2}), \quad (3)$$

where p_i and p_j refer to the spatial feature x, y, z and I_i and I_j refer to the color feature R, G, B. The parameters of Gaussian kernels are described using $\theta_\alpha, \theta_\beta$, and θ_γ.

Due to the consideration of pairwise potentials over all pixel-pairs in the whole image, minimizing the energy function in the fully connected CRF exactly is intractable. Therefore, the mean-field approximation is adopted to approximate the maximum posterior marginal inference. In CRF-RNN, one mean-field iteration can be formulated as a stack of common neural layers. The Initialization, Message Passing, Weighting Filter Outputs, Compatibility Transform, Adding Unary Potentials, and Normalizing operations in the mean-field iteration are implemented through Softmax, Convolutional, Convolutional, Convolutional, Concatenated, and Softmax layers, respectively. The repeated multiple mean-field iterations can be further formulated as a Recurrent Neural Network via repeating the above stack of layers.

In this work, the CRF-RNN is plugged in after the FCN to form a unified framework, and it is trained in an end-to-end manner. During the training process, the error differentials of CRF-RNN are passed to FCN via backward propagation through time, so that the FCN is able to generate better unary potentials for CRF-RNN optimization via forward propagation. More importantly, the parameters in CRF, e.g., the weights of Gaussian kernels and the label compatibility function, are automatically optimized during the full network end-to-end training.

3.5. Transfer Learning

The public VGG-16 model is well-trained using the large-scale ImageNet dataset and can classify objects from daily life belonging to 1,000 different categories. The learned knowledge from object classification should be helpful for the material classification. On the other hand, there are a huge number of sparsely labeled RGB patches (821,092) but a limited number of pixel-wise labeled RGB images (1,498) generated from the MINC dataset. Hence, we transfer the learned knowledge of the classification network to enhance the performance of the segmentation network via transfer learning.

As shown in **Figure 4**, there are two steps of knowledge transfer during the overall training process. The first step transfers the learned weights of the VGG-16 network pre-trained on ImageNet to the material classification network. The second step transfers the learned weights of the classification network, i.e., the VGG-16 network truncated after pooling five layers, to the segmentation network, i.e., the convolution stack of FCN. Both of them are implemented by learned network weights initialization followed by network fine-tuning. The first transfer learning focuses on the same network architecture but transfers the learned knowledge from object classification to material classification, while the second transfer learning focuses on two different network architectures but transfers the learned knowledge from a classification network to a segmentation network.

3.6. Material Reconstruction

A graph-based SLAM, i.e., RGB-D SLAM (Endres et al., 2014), is employed to achieve dense 3D material reconstruction. Given a semantic labeled image with the corresponding depth image, a 3D semantic point cloud (X, Y, Z) can be generated through back projection:

$$d_{u,v}\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \begin{bmatrix} f_x\ s\ c_x \\ 0\ f_y\ c_y \\ 0\ 0\ 1 \end{bmatrix}\begin{bmatrix} X \\ Y \\ Z \end{bmatrix}, \tag{4}$$

where (u, v) refer to the pixel position in the image plane and $d_{u,v}$ refer to the corresponding depth value. f_x, f_y refer to the focal length, and (c_x, c_y) refer to the principal point offset. s refers to the axis skew.

The visual odometry of RGB-D SLAM can estimate the ego-motion between two adjacent semantic point clouds and further enable an incremental semantic label fusion. Finally, using the global trajectory provided by the visual odometry, all of the semantic point clouds are combined together to generate a global semantic map. A Bayesian update is used for label hypothesis fusion using the multi-view semantic point clouds. Each voxel in a semantic point cloud stores the predicted label with the corresponding discrete probability. The voxel's label probability distribution is updated by means of a recursive Bayesian update:

$$P(x = l_i|I_{1,\ldots,k}) = \frac{1}{Z}P(x = l_i|I_{1,\ldots,k-1})P(x = l_i|I_k), \tag{5}$$

where l_i refers to the predicted label, I_k refer to the kth image, and Z refers to the constant for distribution normalization.

4. EXPERIMENTS

In this section, the details of the network training process are first introduced. We then present performance evaluations of three different experiments: material classification, material segmentation on the MINC dataset, and material semantic reconstruction in an industrial scenario.

4.1. Network Training

We first train the VGG-16-based classification network using the newly generated 821,092 RGB patches of four different sizes, 56 × 56, 156 × 156, 256 × 256, and 356 × 356. The network weights are initialized using the public VGG-16 model pre-trained on ImageNet. Secondly, we train the FCN-32s, FCN-16s, and FCN-8s segmentation networks step by step using the newly generated 1,498 pixel-wise labeled 500 × 500 RGB images. The weights of the convolution stack in FCN are inherited from the fine-tuned VGG-16 model truncated after pooling five layers.

Finally, we insert the CRF-RNN after FCN as the bottom part of the whole network. After inheriting the learned FCN weights, the FCN with CRF-RNN network is trained again using the pixel-wise labeled RGB images in an end-to-end manner. During the training process, we set the number of mean-field iterations T to 5 in the CRF-RNN. This can reduce the training time and mitigate the vanishing gradient problem. During the test process, we set the number of mean-field iterations T to 5 or increase it to 10 according to the run-time required.

The parameters of all the trained networks, i.e., the learning rate, momentum, batch size, weight decay, and the type of training data, can be found in **Table 1**.

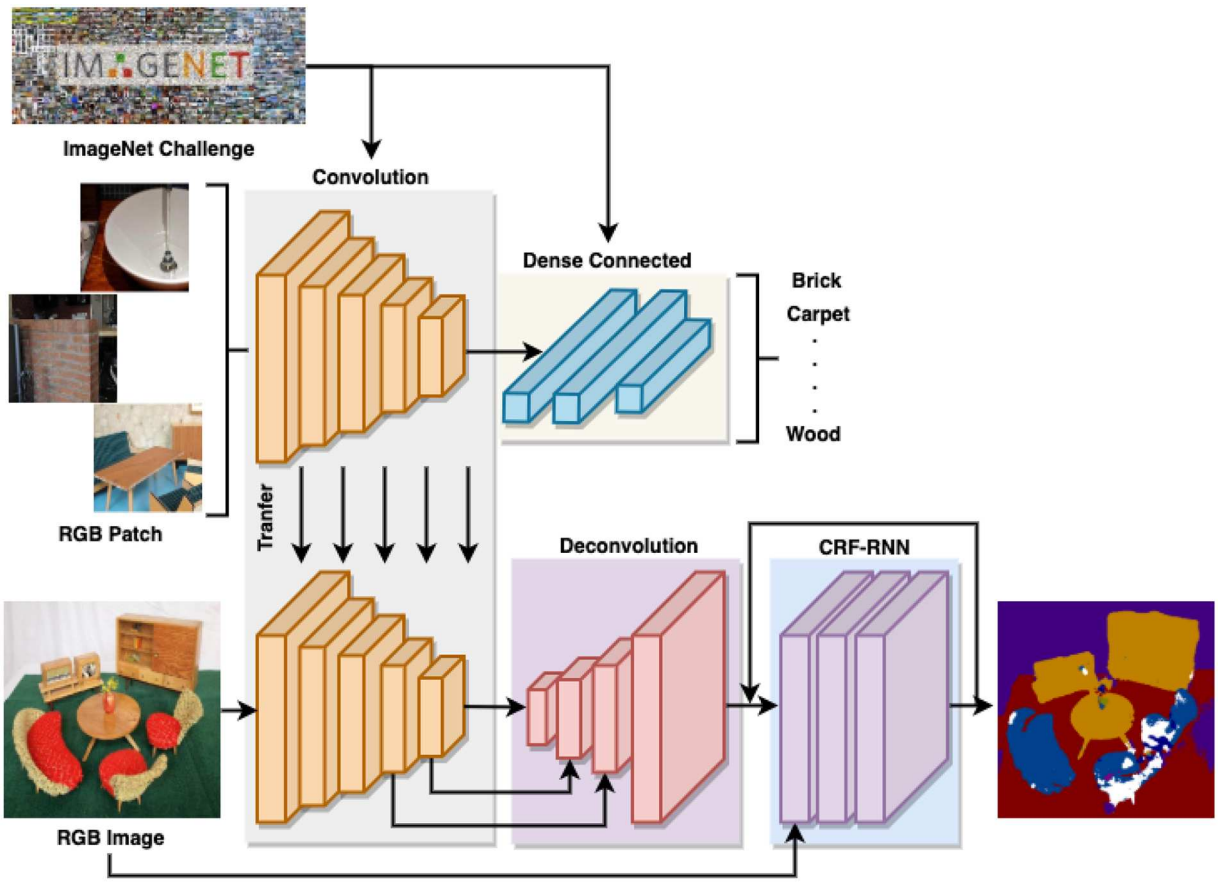

FIGURE 4 | The architecture of the proposed network and transfer learning.

TABLE 1 | The parameters of network training.

	Learning rate	Momentum	Batch size	Weight decay	Training data
VGG-16	1e-4 reduction with 0.1	0.95	50	0.0005	256×256 RGB patch
FCN-32s	1e-10	0.99	1	0.0005	500×500 RGB image
FCN-16s	1e-12	0.99	1	0.0005	500×500 RGB image
FCN-8s	1e-14	0.99	1	0.0005	500×500 RGB image
FCN with CRF-RNN	1e-12	0.99	1	0.0005	500 × 500 RGB image

TABLE 2 | The accuracy of material classification vs. patch size.

Patch size	56×56	156×156	256×256	356×356
Accuracy	69.20%	**81.06%**	80.18%	73.40%

The best performance is in bold.

4.2. Material Classification

The newly generated 96,747 RGB patches are employed for the material classification evaluation. We present the experimental results for the VGG-16 network trained by four differently sized patches in **Table 2**. It can be seen that the accuracy of classification initially increases but then decreases with increasing patch size. The optimal accuracy is reached when the patch size accounts for about 30–50% of the original image. The reason for the accuracy increasing initially is that more contextual cues become available with growth in the patch size, while the reason for the accuracy then decreasing is that there is a loss in spatial resolution with the growth of the path size. The best trade-off patch size for balancing the spatial resolution and contextual information is between 156 × 156 and 256 × 256 for the 500 × 500 images.

4.3. Material Segmentation

The newly generated 300 pixel-wise labeled images are employed for the material segmentation evaluation. Following (Long et al., 2015), the standard parameters for semantic segmentation evaluation, namely pixel accuracy,

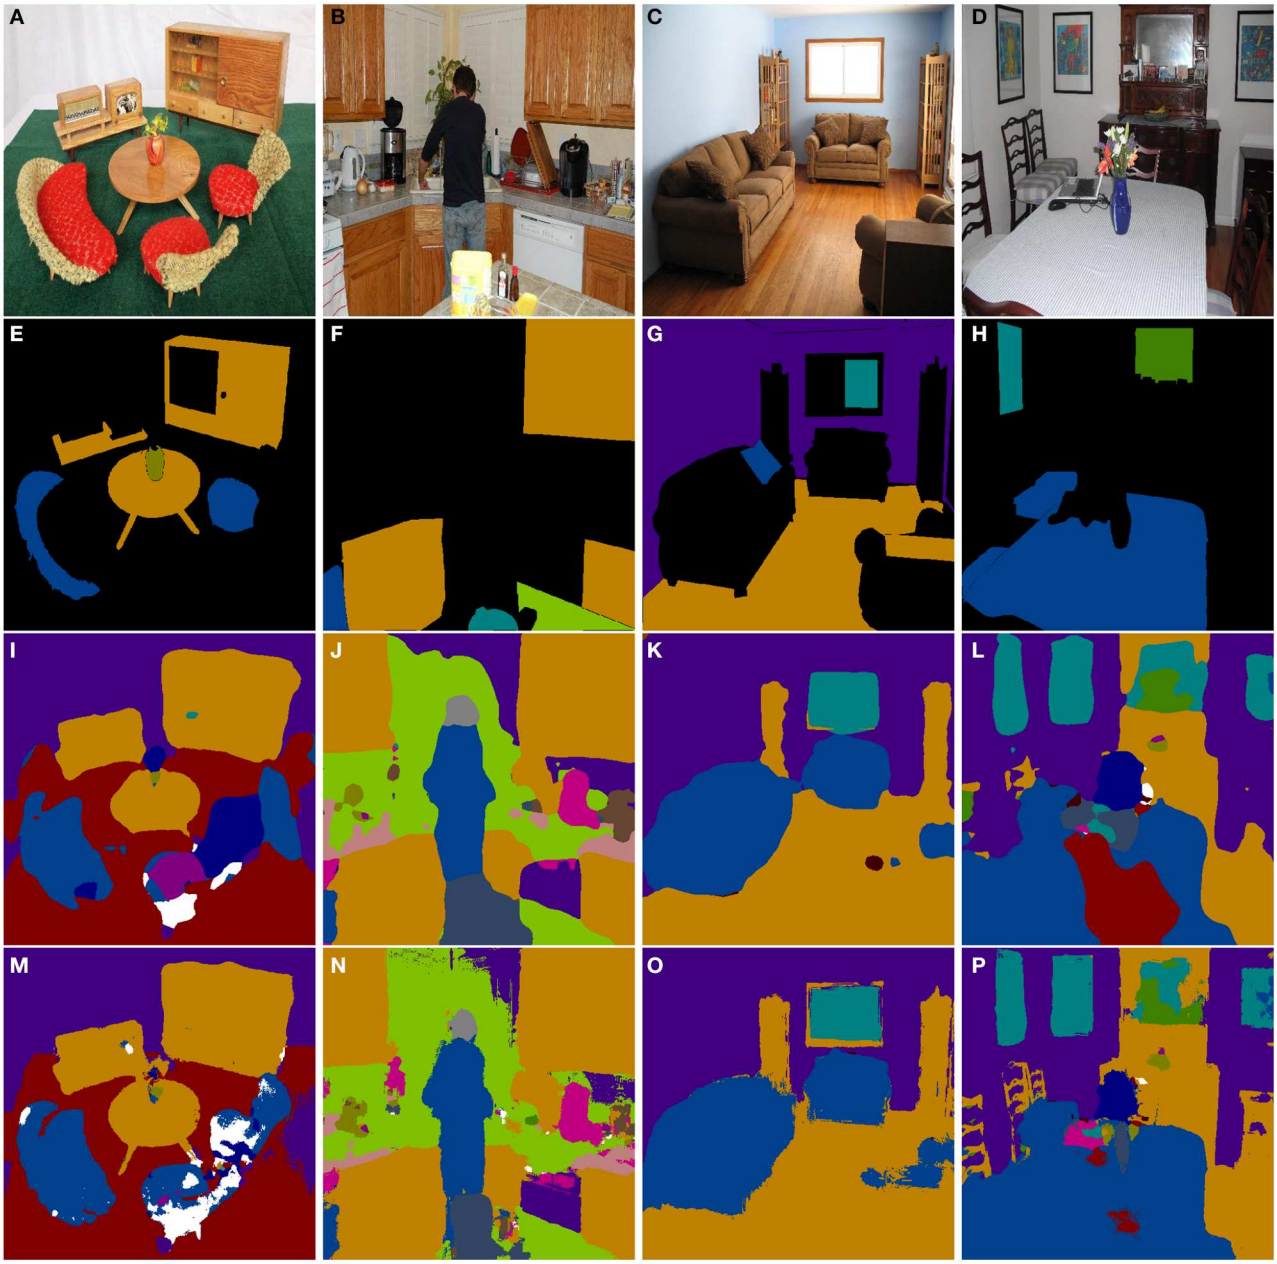

FIGURE 5 | Material segmentation in MINC dataset. From left to right, top to bottom, the IDs of the sub-figures are **(A–P)**. The first row, i.e., **(A–D)**, are original RGB images, the second row, i.e., **(E–H)**, are ground truth images, the third row, i.e., **(I–L)**, are semantic segmentation results of FCN, and the fourth row, i.e., **(M–P)**, are semantic segmentation results of FCN with CRF-RNN.

mean accuracy, mean intersection over union (IoU), and frequency weighed intersection over union (IoU), are adopted for performance analysis. These metrics are defined as:

- Pixel accuracy: $\sum_i n_{ii} / \sum_i t_i$,
- Mean accuracy: $(1/n_{cl}) \sum_i n_{ii}/t_i$,
- Mean IoU: $(1/n_{cl}) \sum_i n_{ii}/(t_i + \sum_j n_{ji} - n_{ii})$,
- Frequency weighted IoU: $(\sum_k t_k)^{-1} \sum_i t_i n_{ii}/(t_i + \sum_j n_{ji} - n_{ii})$.

Here, n_{cl} refers to the number of classes, n_{ij} refers to the number of pixels of class i classified as class j, and $t_i = \sum_j n_{ij}$ refers to the total number of pixels belonging to class i.

4.3.1. Qualitative Analysis

The qualitative results of material segmentation on the MINC dataset are given in **Figure 5**. Due to the lack of neighborhood consistency constraints, there are a lot of non-sharp boundaries in the segmentation results of FCN. After plugging in CRF-RNN

after FCN for the label assignment smoothing, the boundaries of the segmentation results are much clear compared with when using only FCN.

The first and second rows in **Figure 5** show the original and ground-truth images on MINC. The third and fourth rows in **Figure 5** show the 2D semantic segmentation results of FCN and of FCN with CRF-RNN, respectively. It can be seen that FCN with CRF-RNN generates semantic results with much clearer shapes than FCN alone, e.g., table legs in (**Figure 5M**), a person in (**Figure 5N**), a sofa in (**Figure 5O**), and a chair back and vase in (**Figure 5P**). In (**Figure 5L**), a large section of "fabric" is erroneously labeled as "carpet," while the size of this erroneous area greatly decreases in (P) because of the neighborhood consistency constraints of the CRF-RNN optimization.

4.3.2. Quantitative Analysis

Quantitative results for the overall performance and class-wise accuracy of material segmentation on the MINC dataset are given in **Tables 3**, **4**, respectively. As **Table 3** shows, FCN with CRF-RNN achieves 81.94, 74.19, 61.13, and 69.99% for the pixel accuracy, mean accuracy, mean IoU, and frequency weighed IoU, respectively, on the MINC dataset. Compared to FCN without CRF-RNN, FCN with CRF-RNN exhibits an improvement of 3.53, 2.28, 4.62, and 3.92%, respectively, for the pixel accuracy, mean accuracy, mean IoU, and frequency weighed IoU. As **Table 4** shows, the class-wise accuracy for most classes is satisfactory, e.g., Hair (92.11%), Sky (96.71%), and Water (99.07%), but the performances for several classes are still inferior, especially Plastic (35.94%), due to the limited amount of training data. After introducing CRF-RNN following FCN-8s, the class-wise accuracy of each class increases by around 3–6%.

To the best of our knowledge, material segmentation is currently a less-studied research topic, and no good benchmark ranking has yet been deployed on the large-scale material datasets. The MINC dataset is the most popular material dataset for deep learning research, but it is a very coarse dataset, so a lot of data preprocessing and generation are required. The newly generated data in this paper are released as an important supplement to the MINC dataset, and the results provided can be employed as a baseline for future research. We hope that these can improve the benchmarking of research with respect to material classification and segmentation.

4.3.3. Running-Time Analysis

We also provide the running-time performance of the proposed network in **Table 5**. The network is deployed using the 500 × 500 RGB images from the MINC dataset on a computer with an i7-6800k (3.4 Hz) 8-core CPU and NVIDIA TITAN X GPU (12 G).

TABLE 3 | The overall performance of material segmentation on the MINC dataset.

	Pixel acc. (%)	Mean acc. (%)	Mean IU (%)	f.w. IU (%)
FCN	78.41	71.91	56.51	66.07
FCN with CRF-RNN	**81.94**	**74.19**	**61.13**	**69.99**

The best performance is in bold.

TABLE 5 | The running-time performance on the MINC dataset.

	Running-time (s)	CRF iterations	Image size
FCN	0.13–0.15	–	500 × 500
FCN with CRF-RNN	0.4–0.6	10	500 × 500
FCN with CRF-RNN	0.2–0.3	5	500 × 500

TABLE 4 | Comparison of the class-wise accuracy on the MINC dataset.

Category	Brick	Carpet	Ceramic	Fabric	Foliage	Food
FCN	61.02%	84.87%	72.95%	80.88%	**78.62%**	**65.04%**
FCN with CRF-RNN	**63.82%**	**86.18%**	**80.84%**	**84.26%**	77.38%	63.86%
Category	**Glass**	**Hair**	**Leather**	**Metal**	**Mirror**	**Other**
FCN	**67.28%**	92.08%	**72.05%**	72.35%	63.45%	39.44%
FCN with CRF-RNN	62.66%	**92.11%**	71.91%	**76.25%**	**69.81%**	**65.19%**
Category	**Painted**	**Paper**	**Plastic**	**P-Stone**	**Skin**	**Sky**
FCN	**90.62%**	56.83%	**43.94%**	51.75%	81.72%	95.96%
FCN with CRF-RNN	89.35%	**62.82%**	35.94%	**65.12%**	**83.37%**	**96.71%**
Category	**Stone**	**Tile**	**Wallpaper**	**Water**	**Wood**	**Mean**
FCN	62.68%	**66.16%**	**77.11%**	97.82%	79.28%	71.91%
FCN with CRF-RNN	**63.73%**	63.98%	69.17%	**99.07%**	**82.89%**	**74.19%**

The best performance among the compared methods is in bold. P-Stone, Polished Stone.

The running-time of FCN-based segmentation costs 0.13–0.15 s, and that of FCN with CRF-RNN costs 0.4s–0.6 s with 10 iterations or 0.2–0.3 s with five iterations. The running-time performance can be improved greatly if a smaller RGB image is used, which can enable real-time or near-real-time application of material segmentation.

4.4. Material Reconstruction

As well as the evaluation on the MINC dataset, we further evaluate the proposed system in an industrial scenario, i.e., a real industrial room containing many different materials such as wood, brick, paper, metal, carpet, painted surfaces, and others. The system deploys a real-time 3D mapping of the room while simultaneously recognizing and labeling each point according to its material in the built 3D map. The network used in the system is only trained using the MINC dataset without fine-tuning on the real industrial data.

4.4.1. Qualitative Analysis

We give the qualitative results of each step generated by the proposed system, i.e., original RGB image, material segmentation image, 3D point clouds, and 3D semantic point clouds in **Figure 6**. We also provide the local/global 3D map and local/global 3D semantic map of the industrial room in **Figures 7**, **8**, respectively.

We can see that most of the materials are correctly classified and segmented in the dense 3D semantic map. However, some small objects are not labeled correctly due to there not being enough pixels provided in the original RGB image. The pixels at the border between two different materials are more easily assigned to the wrong labels. The domain variances, e.g., varying field of view, varying illumination, different imaging devices between the training and test data, also result in some wrong label predictions.

4.4.2. Quantitative Analysis

We provide the quantitative results evaluated via pixel accuracy, mean accuracy, mean IoU, and frequency weighed IoU in **Table 6**. First, 40 key frames of 3D reconstruction in the industrial room were captured from RGB-D SLAM. Next, all the key frames were densely annotated according to the kind of material via JS Segment Annotator[4]. Finally, pixel-wise false or true numbers were counted between the corresponding pixels from predicted and ground-truth images.

As **Table 6** shows, we achieve 80.10, 58.75, 39.45, and 68.76% for the pixel accuracy, mean accuracy, mean IoU, and frequency weighed IoU, respectively, tested in the industrial room. The pixel accuracy (80.10%) achieves a satisfying level, but the mean accuracy (58.75%) is much lower than the reported result for MINC evaluation (76.87%). Because we only tested 40 samples, there is a large variance in material detection rates. The pixel-wise recognition and segmentation accuracy of some materials, e.g., Paper (6.78%) and Mirror (0%) is very low. However, a mirror appears in only one instance, so failure to recognize just one instance of Mirror generates an accuracy of 0% for that category, which misleadingly skews the overall mean accuracy score toward a low value. In addition, the domain variances, e.g., varying field of view, varying illumination, and different imaging devices between the training and test data, also decrease the performance tested in the industrial room because the network is only trained using the MINC dataset without fine-tuning using the industrial data.

TABLE 6 | The overall performance of material semantic reconstruction in an industrial scenario.

	Pixel acc. (%)	Mean acc. (%)	Mean IU (%)	f.w. IU (%)
3D semantic reconstruction	80.10	58.75	39.45	68.76

TABLE 7 | The running-time performance of the proposed system.

	Running-time (Hz)	CRF iterations	Image size
FCN with CRF-RNN	∼ 2	10	500 × 500
FCN with CRF-RNN	∼ 4	5	500 × 500
FCN with CRF-RNN	∼ 10	5	224 × 224

4.4.3. Running-Time Analysis

The whole system is deployed on a computer with an i7-6800k(3.4 Hz) 8-core CPU and NVIDIA TITAN X GPU (12 G). The IAI Kinect2 package[5] is adopted to interface with ROS and calibrate the RGB and depth cameras of Kinect2. The network is implemented based on the Caffe[6] toolbox and accelerated by CUDA and CUDNN. The overall system is implemented through C++ and GPU programming within the ROS[7] framework.

We provide the running-time performance of the whole system in **Table 7**. The system running-time performance is about 2 Hz (10 iterations) or 4 Hz (5 iterations) using the QHD RGB and depth images from Kinect2. The 540 × 960 RGB images are first reduced to 500 × 500 RGB images for material segmentation and are then recovered to 540 × 960 RGB images for semantic reconstruction. The running-time performance can be boosted to around 10 Hz when the QHD RGB images are decreased to 224 × 224 RGB images, using five CRF-RNN iterations for material segmentation.

As mentioned (Hermans et al., 2014), there is no necessity to segment all the frames for the RGB-D SLAM because most of the frames are abandoned and only a few key frames (about 20%) are used for dense 3D mapping. In this case, a 5–10 Hz running-time

[4]http://kyamagu.github.io/js-segment-annotator/

[5]https://github.com/code-iai/iai_kinect2/

[6]http://caffe.berkeleyvision.org/

[7]https://www.ros.org

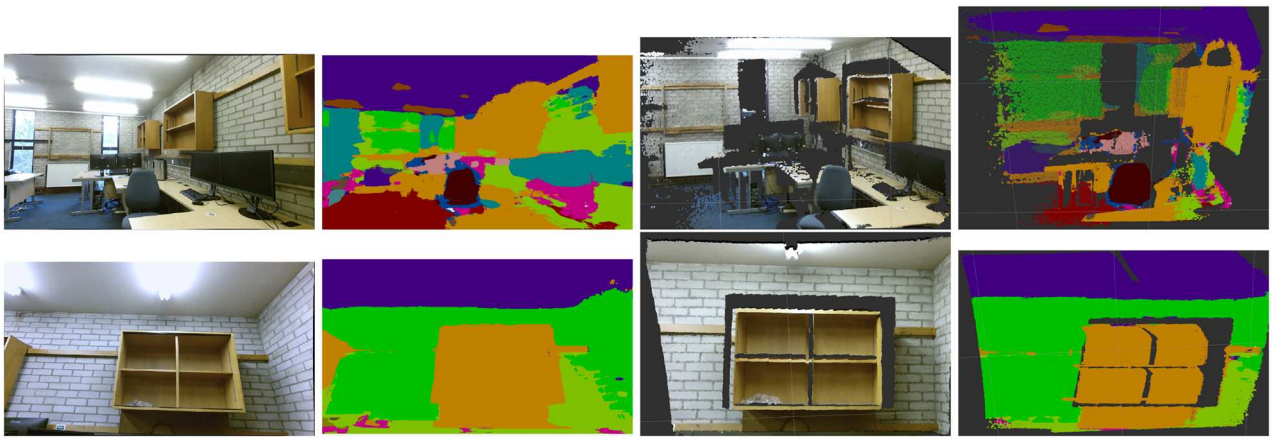

FIGURE 6 | The qualitative results of each step generated by the material segmentation and reconstruction system. The first column are RGB images from Kinect2, the second column are material segmentation images, the third column are 3D point clouds, and the fourth column are 3D semantic point clouds.

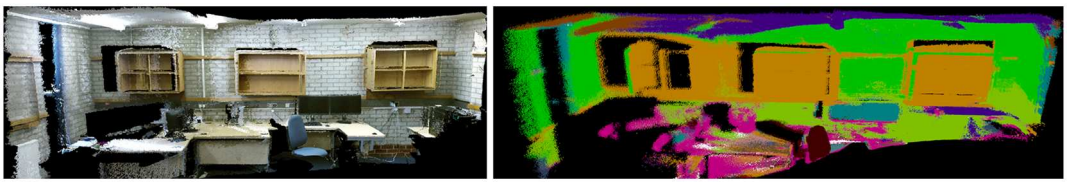

FIGURE 7 | Material segmentation and reconstruction in an industrial scenario. **(Left)** Local 3D map. **(Right)** Local 3D semantic map.

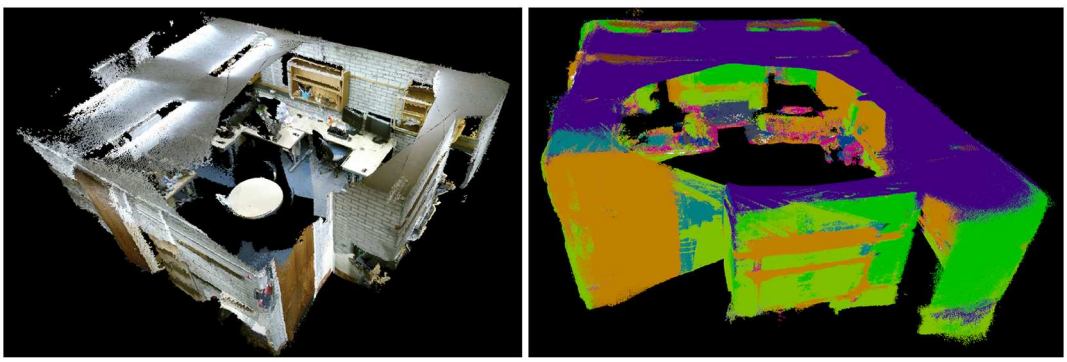

FIGURE 8 | Material segmentation and reconstruction in an industrial scenario. **(Left)** Global 3D map. **(Right)** Global 3D semantic map.

performance can basically meet the requirement of a real-time material semantic reconstruction.

5. CONCLUSIONS

In this paper, we propose a novel transfer learning method to determine material categories from RGB images. Our approach is data-efficient, with maximization of the utility of a fundamental computer vision dataset and coarse annotated data. Consequently, our approach shows strong effectiveness in solving real-world problems, where large-scale training datasets are not available.

Moreover, the material understanding proposed by the neural network is integrated with 3D dense reconstruction, and incremental dense material labeling of a 3D scene is performed. The running-time performance of the whole system can be boosted to approximately 10 Hz to satisfy the requirement of real-time applications. We report qualitative, quantitative, and running-time evaluation analysis of the proposed approach using both the public material dataset and real-world industrial data

to verify the resultant segmentation accuracy and running-time performance.

The newly generated high-quality dataset, including RGB image patches and fully pixel-wise annotated RGB images, is released as an important supplement for the MINC dataset. Our approach has a good alignment with industrial applications, especially nuclear robotics. As an essential part of the EU H2020 RoMaNs project, the proposed system has the potential to demonstrate its capability of guiding robots to navigate in industrial scenes and manipulate objects in a self-occluded heap.

DATA AVAILABILITY STATEMENT

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation, to any qualified researcher.

AUTHOR CONTRIBUTIONS

CZ proposed the main method and designed the experiments under the supervision of LS and RS.

FUNDING

The authors were sponsored by the DISTINCTIVE Scholarship, EPSRC AdMaLL project (FAIR-SPACE Innovation Award R/162435-11-1), and EU H2020 Project RoMaNS (645582).

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: https://www.frontiersin.org/articles/10.3389/frobt.2020.00052/full#supplementary-material

REFERENCES

Bell, S., Upchurch, P., Snavely, N., and Bala, K. (2015). "Material recognition in the wild with the materials in context database," in *IEEE Conference on Computer Vision and Pattern Recognition*, 3479–3487.

Caputo, B., Hayman, E., and Mallikarjuna, P. (2005). "Class-specific material categorization," in *IEEE International Conference on Computer Vision, Vol. 2*, 1597–1604.

Cimpoi, M., Maji, S., Kokkinos, I., Mohamed, S., and Vedaldi, A. (2014). "Describing textures in the wild," in *IEEE Conference on Computer Vision and Pattern Recognition*, 3606–3613.

Cimpoi, M., Maji, S., and Vedaldi, A. (2015). "Deep filter banks for texture recognition and segmentation," in *IEEE Conference on Computer Vision and Pattern Recognition*, 3828–3836.

Dana, K. J., Van Ginneken, B., Nayar, S. K., and Koenderink, J. J. (1999). Reflectance and texture of real-world surfaces. *ACM Trans. Graph.* 18, 1–34.

DeGol, J., Golparvar-Fard, M., and Hoiem, D. (2016). "Geometry-informed material recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 1554–1562.

Endres, F., Hess, J., Sturm, J., Cremers, D., and Burgard, W. (2014). 3-D mapping with an RGB-D camera. *IEEE Trans. Robot.* 30, 177–187. doi: 10.1109/TRO.2013.2279412

Giben, X., Patel, V. M., and Chellappa, R. (2015). "Material classification and semantic segmentation of railway track images with deep convolutional neural networks," in *IEEE International Conference on Image Processing*, 621–625.

Hayman, E., Caputo, B., Fritz, M., and Eklundh, J.-O. (2004). "On the significance of real-world conditions for material classification," in *European Conference on Computer Vision* (Springer), 253–266.

Hermans, A., Floros, G., and Leibe, B. (2014). "Dense 3D semantic mapping of indoor scenes from RGB-D images," in *IEEE International Conference on Robotics and Automation*, 2631–2638.

House of Commons Committee of Public Accounts (2013). *Nuclear Decommissioning Authority: Managing Risk at Sellafield,* 784–784.

Hu, D., Bo, L., and Ren, X. (2011). "Toward robust material recognition for everyday objects," in *The British Machine Vision Association, Vol. 2*.

Krähenbühl, P., and Koltun, V. (2011). "Efficient inference in fully connected CRFs with Gaussian edge potentials," in *Advances in Neural Information Processing Systems*, 109–117.

Li, W., and Fritz, M. (2012). "Recognizing materials from virtual examples," in *European Conference on Computer Vision* (Springer), 345–358.

Liu, C., Sharan, L., Adelson, E. H., and Rosenholtz, R. (2010). "Exploring features in a Bayesian framework for material recognition," in *IEEE Conference on Computer Vision and Pattern Recognition*, 239–246.

Long, J., Shelhamer, E., and Darrell, T. (2015). "Fully convolutional networks for semantic segmentation," in *IEEE Conference on Computer Vision and Pattern Recognition*, 3431–3440.

Purri, M., Xue, J., Dana, K., Leotta, M., Lipsa, D., Li, Z., et al. (2019). Material segmentation of multi-view satellite imagery. *arXiv preprint arXiv:1904.08537*.

Qi, X., Xiao, R., Li, C.-G., Qiao, Y., Guo, J., and Tang, X. (2014). Pairwise rotation invariant co-occurrence local binary pattern. *IEEE Trans. Pattern Anal. Mach. Intell.* 36, 2199–2213. doi: 10.1109/TPAMI.2014.2316826

Schwartz, G., and Nishino, K. (2013). "Visual material traits: recognizing per-pixel material context," in *IEEE International Conference on Computer Vision Workshops*, 883–890.

Schwartz, G., and Nishino, K. (2016). Material recognition from local appearance in global context. *arXiv preprint arXiv:1611.09394*.

Sharan, L., Rosenholtz, R., and Adelson, E. (2009). Material perception: what can you see in a brief glance? *J. Vis.* 9:784. doi: 10.1167/9.8.784

Simonyan, K., and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.

Wang, T.-C., Zhu, J.-Y., Hiroaki, E., Chandraker, M., Efros, A. A., and Ramamoorthi, R. (2016). "A 4D light-field dataset and CNN architectures for material recognition," in *European Conference on Computer Vision* (Springer), 121–138.

Zheng, S., Jayasumana, S., Romera-Paredes, B., Vineet, V., Su, Z., Du, D., et al. (2015). "Conditional random fields as recurrent neural networks," in *IEEE International Conference on Computer Vision*, 1529–1537.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

frontiers in Robotics and AI

ORIGINAL RESEARCH
published: 27 October 2020
doi: 10.3389/frobt.2020.549716

BVLOS UAS Operations in Highly-Turbulent Volcanic Plumes

Kieran Wood[1], Emma J. Liu[2], Tom Richardson[1], Robert Clarke[1], Jim Freer[3,4], Alessandro Aiuppa[5], Gaetano Giudice[6], Marcello Bitetto[5], Kila Mulina[7] and Ima Itikarai[7]*

[1] Department of Aerospace Engineering, University of Bristol, Bristol, United Kingdom, [2] Department of Earth Sciences, University College London, London, United Kingdom, [3] School of Geographical Sciences, University of Bristol, Bristol, United Kingdom, [4] University of Saskatchewan Centre for Hydrology, Canmore, AB, Canada, [5] Dipartimento di Scienze della Terra e del Mare, University of Palermo, Palermo, Italy, [6] Istituto Nazionale di Geofisica e Vulcanologia, Osservatorio Etneo, Sezione di Catania, Catania, Italy, [7] Rabaul Volcanological Observatory, Rabaul, Papua New Guinea

Long-range, high-altitude Unoccupied Aerial System (UAS) operations now enable *in-situ* measurements of volcanic gas chemistry at globally-significant active volcanoes. However, the extreme environments encountered within volcanic plumes present significant challenges for both air frame development and in-flight control. As part of a multi-disciplinary field deployment in May 2019, we flew fixed wing UAS Beyond Visual Line of Sight (BVLOS) over Manam volcano, Papua New Guinea, to measure real-time gas concentrations within the volcanic plume. By integrating aerial gas measurements with ground- and satellite-based sensors, our aim was to collect data that would constrain the emission rate of environmentally-important volcanic gases, such as carbon dioxide, whilst providing critical insight into the state of the subsurface volcanic system. Here, we present a detailed analysis of three BVLOS flights into the plume of Manam volcano and discuss the challenges involved in operating in highly turbulent volcanic plumes. Specifically, we report a detailed description of the system, including ground and air components, and flight plans. We present logged flight data for two successful flights to evaluate the aircraft performance under the atmospheric conditions experienced during plume traverses. Further, by reconstructing the sequence of events that led to the failure of the third flight, we identify a number of lessons learned and propose appropriate recommendations to reduce risk in future flight operations.

Keywords: unmanned aircraft system (UAS), UAV, aerial robotic, volcano, plume, BVLOS, Manam, gas sensing

OPEN ACCESS

Edited by:
Chie Takahashi, University of Cambridge, United Kingdom

Reviewed by:
Jonathan M. Aitken, The University of Sheffield, United Kingdom
Michael Schmuker, University of Hertfordshire, United Kingdom

****Correspondence:***
Kieran Wood
kieran.wood@bristol.ac.uk

Specialty section:
This article was submitted to Robotic Control Systems, a section of the journal Frontiers in Robotics and AI

Received: *07 April 2020*
Accepted: *24 August 2020*
Published: *27 October 2020*

Citation:
Wood K, Liu EJ, Richardson T, Clarke R, Freer J, Aiuppa A, Giudice G, Bitetto M, Mulina K and Itikarai I (2020) BVLOS UAS Operations in Highly-Turbulent Volcanic Plumes. Front. Robot. AI 7:549716. doi: 10.3389/frobt.2020.549716

1. INTRODUCTION

The application of instrumented small UAS (Unoccupied Aerial Systems), or alternatively "drones," has had a transformational influence on volcanological research over the past decade, particularly in recent years where the miniaturization of scientific instrumentation has begun to approach the rapid progression of UAS technology (Jordan, 2019; James et al., 2020). Driven largely by the consumer market, UAS control systems and hardware have now advanced to the point where relatively little training is required to operate multi-rotor platforms equipped with complex sensors. Aerial robotic systems are being developed and deployed increasingly for a range of environmental applications (Fladeland et al., 2011; Vivoni et al., 2014; Detweiler et al., 2015; Klemas, 2015; Pajares, 2015; Bhardwaj et al., 2016). In particular, significant traction is being realized in the areas of remote

sensing (Immerzeel et al., 2014; Tamminga et al., 2015), mapping 2D/3D structures (Nagai et al., 2009; Stöcker et al., 2015; Zweig et al., 2015) and atmospheric sampling (Cassano, 2013; Villa et al., 2016; Greatwood et al., 2017) using a range of emerging sensor technologies (Wildmann et al., 2013; Detert and Weitbrecht, 2015; Hill and Clemens, 2015). Atmospheric sampling has been performed either by multi-rotor UAS at lower altitudes in the 500–1,000 m range (Cassano, 2013; Peng et al., 2015) or by fixed wing platforms capable of long-range flight but that require considerable resources to deploy (Ramana et al., 2007; Corrigan et al., 2008; de Boer et al., 2016). The greatest limitation to multi-rotor UAS is often the battery technology, which determines the flight time and therefore distance and altitude (flight envelope). The use of fixed wing UAS can increase the flight time for a given payload and Maximum Take-Off Weight (MTOW), but with additional challenges in terms of launch and recovery, particularly in remote locations and vegetated/mountainous terrain typical of volcanic environments.

In volcanology, remote measurements using UAS now enable the collection of scientific data in previously inaccessible volcanic plumes (McGonigle et al., 2008; Shinohara, 2013; Di Stefano et al., 2018; Liu et al., 2019), or where large areal coverage is required (Darmawan et al., 2018; Favalli et al., 2018), whilst prioritizing the safety of the operator. To this end, aerial observations are now becoming integrated routinely within volcanic crisis response procedures (Turner et al., 2017; Nadeau et al., 2018; de Moor et al., 2019; Syahbana et al., 2019). Notably, most volcanological operations are typically conducted within Visual Line Of Sight (VLOS) and at relatively low altitudes. Critically, however, there remain significant gaps in our knowledge of some of the most active, yet inaccessible, volcanoes where Beyond Visual Line Of Sight (BVLOS) operations are the only way to obtain the data required (Schellenberg et al., 2019; Syahbana et al., 2019; Liu et al., 2020). Here, we focus on BVLOS operations at Manam volcano, Papua New Guinea, in the context of an international scientific effort to characterize the chemistry of the volcanic gases being released from this globally significant volcanic emission source.

Specifically, we present a detailed account and analysis of the platform development ("Titan" SUAS) and the operational procedures needed to realize safe and repeatable operations to an altitude of 2,300 m Above Mean Sea Level (AMSL) and a horizontal distance of 5 km from the take-off point. We analyse logged flight data for three flights to explore parameters related to aircraft performance, turbulence within the volcanic plume, and energy budgets. Although scientific data were collected from all three flights, we critically evaluate the event sequence that resulted in loss of the airframe during the third flight. Through the lessons learned and insights into plume conditions presented, our results will contribute to the continued development and operation of robust fixed wing sensor platforms for the volcanological community, and in extreme environments more generally.

1.1. Motivation

Measurements of volcanic gases are critical for the assessment of volcanic hazard (Aiuppa et al., 2007; de Moor et al., 2016, 2019) and for constraining global emissions of environmentally-important gases, such as carbon dioxide (Aiuppa et al., 2019; Fischer et al., 2019; Werner et al., 2019). Volcanic environments present challenging environments in which to make scientific measurements, particularly at high altitude, densely vegetated, or highly active volcanoes. These sampling limitations have led to significant bias in estimates of global volcanic gas emissions toward a relatively small number of accessible, passively degassing volcanoes (Fischer and Aiuppa, 2020). By enabling proximal sampling of remote or hazardously accessible volcanic plumes, instrumented UAS are now targeting gaps in our knowledge of gas emissions at some of the major remaining "known unknown" volcanic emitters.

Manam (**Figure 1**) is one of the most active volcanoes in Papua New Guinea (Palfreyman and Cooke, 1976), and has experienced five major eruptions in the past year alone (GVP, 2019). A series of climactic eruptions in 2004 devastated large sectors of the island and displaced the local population to the mainland. Mild to moderate explosive activity has continued sporadically at Manam since the 2004–2006 eruptions, causing continued social and environmental disruption (Mercer and Kelman, 2010). In a broader context, Manam is a globally-significant source of sulfur dioxide to the atmosphere (Carn et al., 2017) as measured by satellites, and yet its carbon dioxide emissions are previously uncharacterized. Aerial-based Observations of Volcanic Emissions (ABOVE), of which this study is a part, is an internationally-collaborative and cross-disciplinary endeavor to integrate novel UAS technology with state-of-the-art gas sensing instrumentation to improve our ability to measure the gas chemistry and emission rate at remote and inaccessible volcanoes, such as Manam. In this contribution, we focus on the engineering and control required to achieve long-range, high altitude fixed wing flights through the volcanic plume. The resulting scientific data are presented in a companion publication (Liu et al., 2020).

The requirements of volcanic plume intersection, at long-range or otherwise, present considerable challenges in terms of both hardware engineering and flight control. Volcanic plumes are energetic and thermally-buoyant mixtures of gas and (sometimes) ash particles, which are often emitted in a pulsatory manner (Woods, 2010). In the case of Manam, the plume can rise to altitudes of several kilometers above the vent under its own thermal energy before dispersing laterally with the wind (Liu et al., 2020). Yet, our knowledge of the conditions encountered within a volcanic plume is incomplete, leading to considerable uncertainty when designing an appropriate air frame and optimal flight path. To achieve plume intersection, we developed an instrumented fixed wing (7.5 kg UAS, hereafter referred to as the "Titan" aircraft). The Titan system is capable of carrying a payload of 1 kg up to an altitude of 2300 m ASL, for a distance of more than 10 km. The aircraft is also capable of a hand launch in zero wind, which combined with a parachute recovery makes it ideal for operation in rough terrain, or areas with very little flat ground for a conventional landing. Quad-plane type airframes can also operate under these conditions, but typically lack the required performance for long range missions with large ascent requirements.

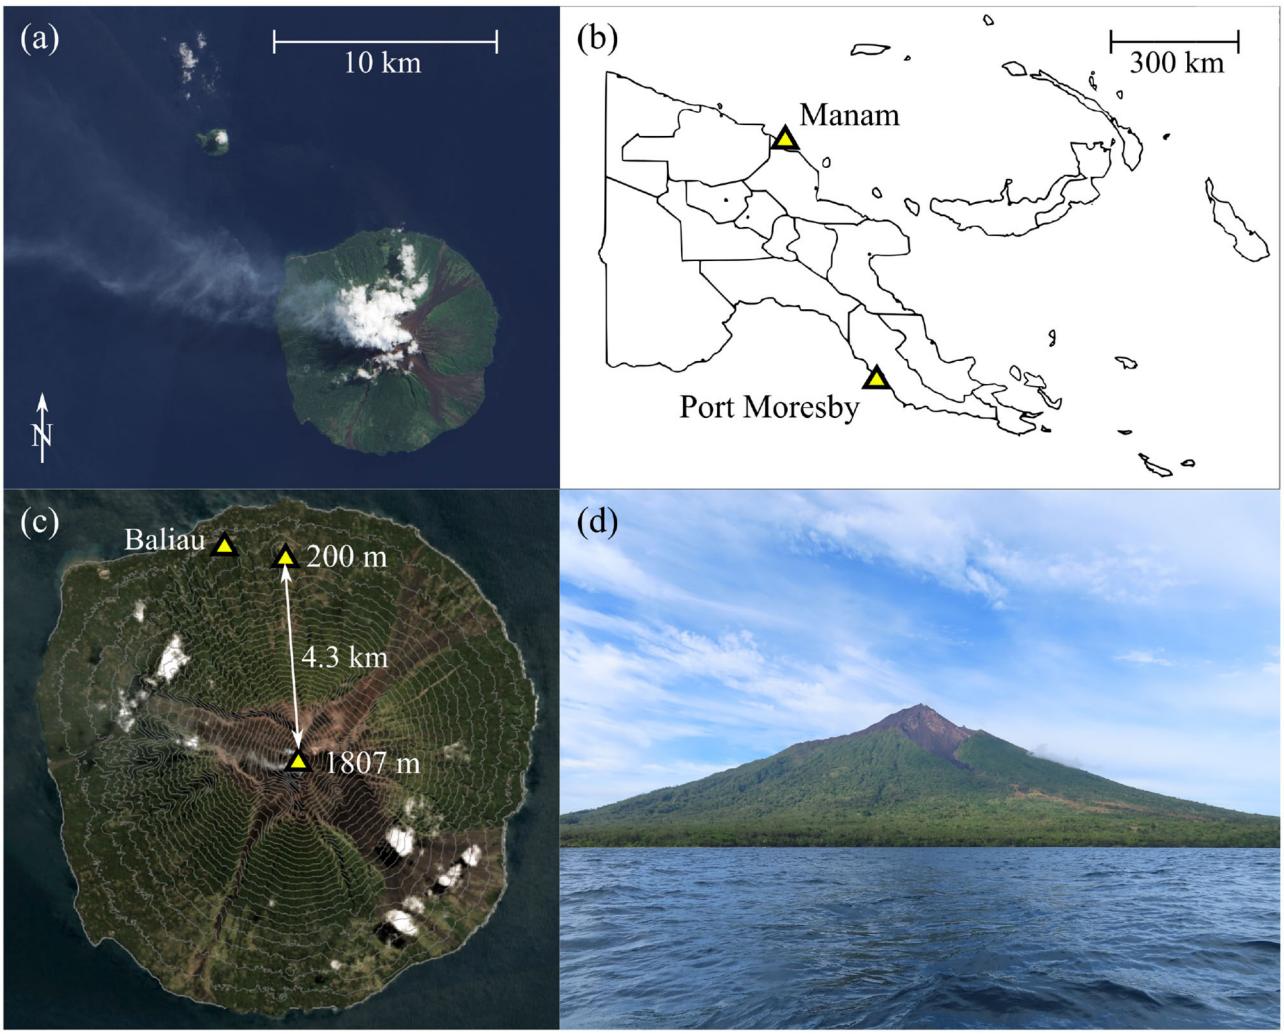

FIGURE 1 | (a) Satellite view of Manam with a visible plume drifting North-West. **(b)** Manam volcano is located on the Northern coast of mainland Papua new Guinea. **(c)** Each mission required flying to the summit from a small satellite cone located 4.3 km from the summit crater, near the village of Baliau. Note the satellite image has been overlayed with contour lines indicating the steep terrain. **(d)** A clear view of Manam volcano from the approach by boat.

By analyzing flight data from two successful flights at Manam, we present novel insight into the atmospheric conditions and the resulting airframe stresses encountered within the extreme environment of the plume of an active volcano. Further, by reconstructing the sequence of events that led to the failure of the third flight, we identify a number of lessons learned and propose appropriate recommendations for future flight operations and aircraft design requirements.

1.2. Manam Volcano (Papua New Guinea)

Manam volcano is located 13 km off the northern coast of mainland Papua New Guinea (**Figure 1b**). Most of the volcano is submerged but the exposed sub-aerial part of the volcanic edifice forms an island ∼10 km in diameter. Current volcanic activity involves persistent passive gas release, punctuated by occasional large explosive eruptions (GVP, 2019). A distinctive gas plume is often visible from both the ground and space (**Figure 1a**). With an almost equatorial latitude, the climate is tropical with temperatures of ∼30 °C and frequent rainfall. The flanks of the volcano are often obscured by cloud, especially from late morning through to mid-afternoon, although the summit can be clear above the cloud level.

The topography of the island is generally mountainous with small patches of level ground near the coast. The flanks of the volcano are densely vegetated and incised by four radial avalanche valleys (**Figure 1c**) that channel debris flows during large eruptions. The summit altitude is 1,807 m AMSL (GVP, 2019). A small volcanic cone (known locally by the name "Godagi") is located on the northern coast of the island, and has a summit altitude of 200 m AMSL. We selected Godagi cone as the base for fixed-wing operations due to its prominent topographic position, clear lines of sight in all direction unobscured by vegetation, and the altitude advantage. We identified an area of

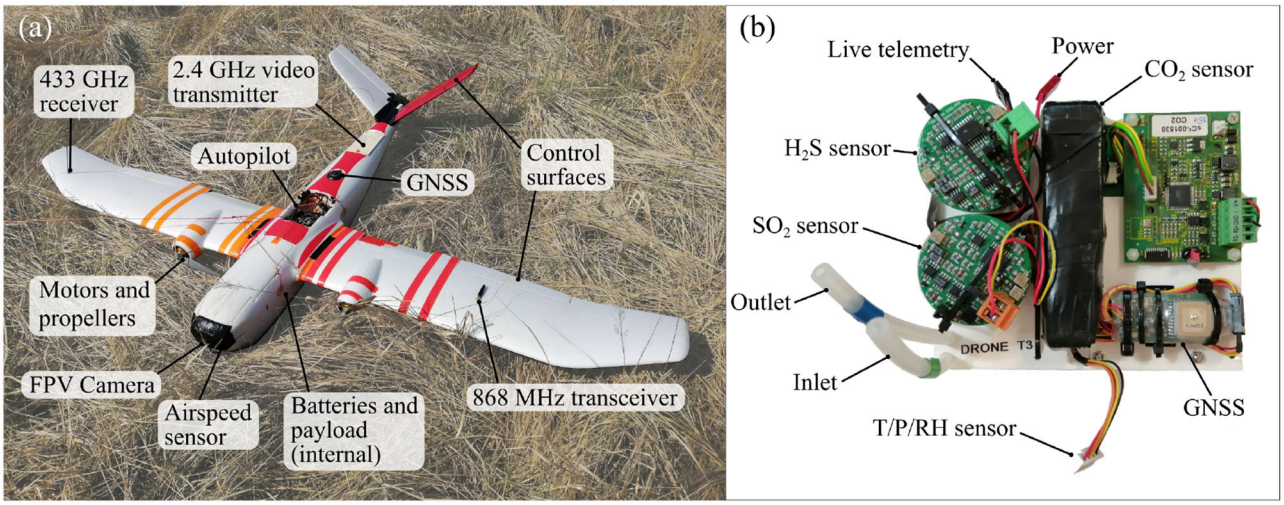

FIGURE 2 | (a) The "Titan" fixed-wing UAS. The radio frequency transmitters have been positioned as separated as possible to ensure clear transmissions without blocking from the fuselage or interference. **(b)** The multi-species gas sensor unit. The sensor was additionally shielded with metal foil (not shown) when installed in the fuselage to prevent interference from the vehicle RF transmitters.

tall grass ~10 × 10 m square as an appropriate landing zone for parachute recovery.

2. INSTRUMENTS AND METHODS

2.1. Titan SUAS

The fixed wing platform chosen for this project was the so called "Titan" —a twin-propeller, v-tail vehicle based upon the airframe kit of the same name (**Figure 2a**). A full list of avionics and specifications is given in **Table 1**. The aircraft has a wingspan of 2.1 m and a take-off weight of 8.5 kg (including 1 kg payload). This particular system was advantageous because it could be hand-launched and recovered by parachute into confined areas where a "skid" landing would have been impossible. The twin propeller design allowed for the installation of oversized motors which are essential to achieve acceptable climb rates. Power was provided by a 12.75 A h, 6S 22.2 V lithium polymer (LiPo) battery set (comprising three 4.25 A h to allow for international travel), giving a flight duration of 25–35 min depending on each mission's altitude-gain and airspeed requirements—nominally 2,100 m above takeoff, and 18 m s^{-1} equivalent air speed (EAS). The maximum thrust was measured in the laboratory to be 7 kg, hence the vehicle in this configuration had a thrust to weight of 82%. This was essential for the hand launch and to ensure the motors were operating at a sustainable power of ~40% during the long climbs.

The Titan featured a full autopilot flight computer with supporting sensors (GNSS, barometric altitude, airspeed indicator, and IMU). Running the open source ArduPlane software, the autopilot was capable of navigating the aircraft along pre-planned waypoint missions. Three wireless links were used to interact with the vehicle during flight. The pilot safety link, operating on the 433 MHz frequency, was used for initializing the automatic flight and for manual control during the plume intersections and parachute landing. The second link was a bi-directional telemetry modem operating on the 868 MHz frequency and was used for monitoring flight statistics, to issue updated commands to the autopilot, and also relay live gas concentration measurements to the ground station. The third link was a live first-person-video (FPV) stream from a camera in the nose of the aircraft operating on the 2.4 GHz frequency. The interconnection of the avionics systems is shown in **Figure 3**.

There are three internal cabins within the body of the "Titan." The first is the fuselage, which housed the control systems and batteries. The second is the payload bay, which contained the gas sensor and a downwards-orientated camera. The final cabin, located toward the tail of the aircraft, contained the parachute landing system and video transmitter.

The autopilot logged flight data at frequencies between 10 and 50 Hz, including signals, such as altitude, airspeed, orientation, servo commands, and GPS location. A reduced rate version of these signals is telemetered to the ground control station where they are also stored. Both the high-fidelity onboard and low-fidelity ground station logs can be analyzed post-flight alongside the recorded videos. These log files are often overlooked at the end of a successful mission, however they provide a rich source of additional information when analyzed more thoroughly. In section 4, we extract selected signals for detailed analysis to inform future UAS design requirements.

The ground station comprised receivers for the telemetry and video links, a laptop, a display screen, and pilot controller. The telemetry data are received, decoded, and displayed, with the live gas concentration data shown in a custom application. The live video is displayed and recorded on a handheld monitor screen. All items are battery-powered and portable. During flight, live data streams of parameters, such as battery voltage, airspeed, GPS-location, and gas concentration were monitored

by the operation crew, which included a pilot, co-pilot, and payload specialist. The pilot held the safety link controller, which was used to trigger mode changes, deploy the parachute, and maneuver the aircraft manually whilst in Fly By Wire (FBW) mode. When maneuvering manually at ranges beyond visual line of sight, the pilot used a video headset to view the First Person View (FPV) stream and direct the aircraft. The co-pilot monitored the vehicle telemetry data, verbally relaying essential flight data to the pilot for situational awareness, and, whenever necessary, adjusted mission parameters under instruction from the pilot. The payload specialist monitored the telemetered gas concentration data and FPV video, providing guidance on the quality of the data collected and suggesting modifications to the flight path based on the incoming data. Decisions to deviate from the pre-planned mission were agreed by all crew before they were executed.

TABLE 1 | Parts list for the "Titan" aircraft.

Spec./Part	Details
Maximum flight time	1 h (100 m ascent), 30 min (2,000 m ascent)
Take-off mass	8.5 kg
Wingspan	2,160 mm
Airframe	Skywalker Titan (China)
Battery (×3)	Overlander Supersport Pro 22.2V 4250 mAh 35 C (UK)
Main Motors (×2)	AXi 2826/13 v2 (Czech Republic)
Speed Controllers (×2) (ESC)	Jeti Spin Pro 66 OPTO (Czech Republic)
AutoPilot	UnmannedTech Pixhawk v1 (UK)
Autopilot Software	ArduPlane V3.9.7
Propellers (×2)	APC-E 12 × 6 Thin Electric (USA)
All Servos (×5)	Hitec HS-5065MG Digital (Japan)
Safety (Pilot) control link	DragonLink V3 Advanced (433 MHz) (USA)
Ground telemetry link	RFD 868x (868 MHz) (Australia)
FPV link	ImmersionRC 700 mW (2.4 GHz) (Hong Kong)
FPV Camera	RunCam Eagle 2 Pro (Hong Kong)
Parachute	Skywalker Landing Umbrella 8 kg (China)
Companion computer	PJRC Teensy 3.6 micro-controller (USA)

The payload comprised two miniature high definition cameras and a multi-component gas analyser system (multi-GAS). Two 4K video cameras (120 g each) were installed in the vehicle: one in the nose with a forward view and one in the payload bay with a nadir view. The multi-GAS is a miniaturized version of the established ground-based volcanic monitoring system developed by the University of Palermo-INGV (Aiuppa et al., 2007); (**Figure 2b**) and described in (Liu et al., 2019). The multi-GAS unit has dimensions of 150 × 130 × 90 mm and a weight of 550 g. Air is sampled from outside the fuselage and passed through a filter, two electro-chemical sensors (SO_2, H_2S), and a NDIR sensor (CO_2) before being expelled back into the freestream airflow. A separate pressure, temperature, and relative humidity sensor was also mounted on the exterior of the airframe to measure the air conditions. All payload data is stored on micro-SD cards, with the gas measurement data additionally telemetered to the ground station using an onboard companion computer to interpret and forward the essential values.

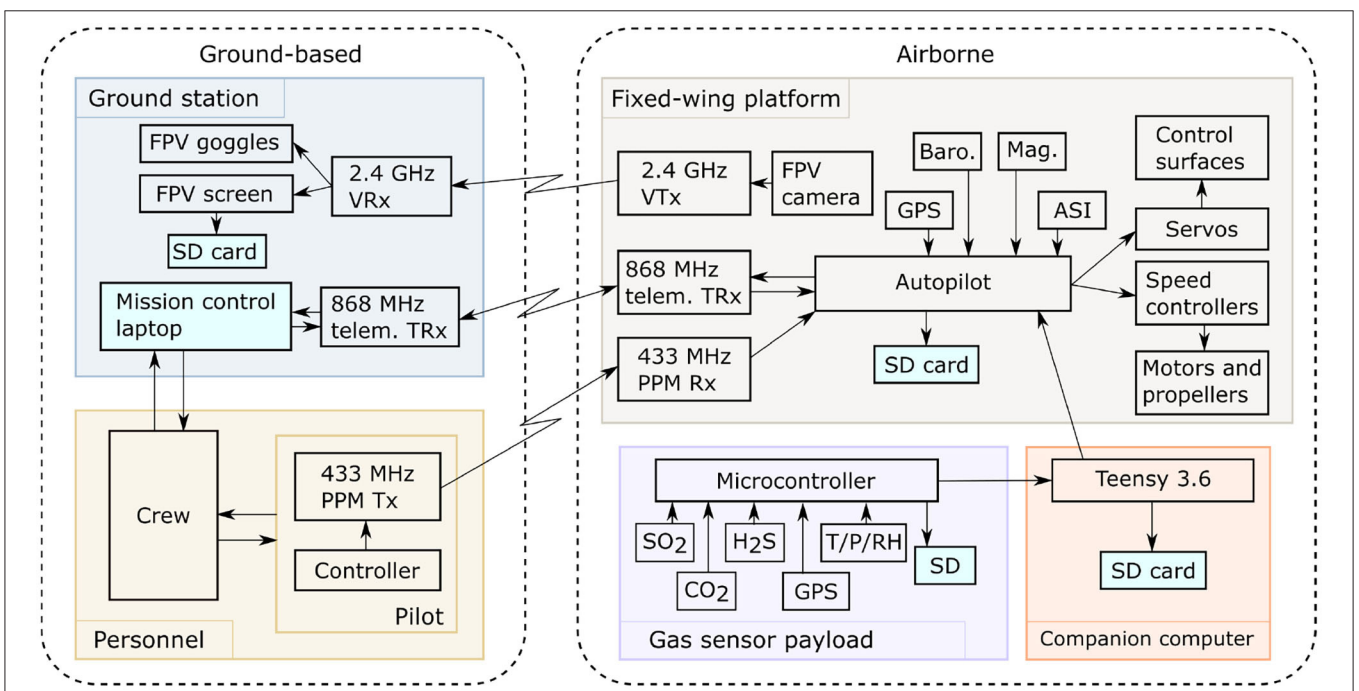

FIGURE 3 | Interconnection of the UAS sub-systems. Live sensor data is relayed from the sensors, via the companion computer, autopilot, and ground station before being displayed live to allow guidance during FBW flight segments.

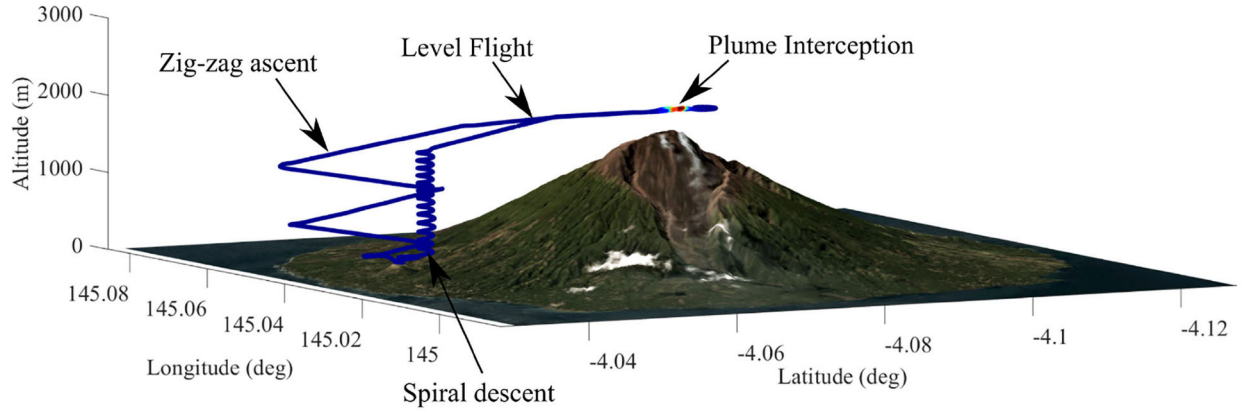

FIGURE 4 | The mission profile comprised a zig-zag ascent, level flight during the plume interceptions, and a spiral descent.

2.2. Flight Planning and Deployment

Flight operations included both automatic and manual flight segments. Initially, each mission was pre-programmed as a series of 3D waypoints based upon visual observations and coordinates taken from a high-accuracy digital elevation model (WorldDEM provided by Airbus Defense and Space GmbH). For reference, the coordinates of the take-off location and summit were [−4.0407N, 145.0356E] and [−4.0776N, 145.0384E], respectively, which are separated by ∼4.3 km horizontal distance and 1600 m vertical ascent. The Titan has a proven performance history for long-range missions having previously been deployed for low altitude survey missions where a flight duration of 1 h was achieved with a similar payload mass (Connor et al., 2020). Here, the expected flight duration was reduced to ∼30 min to accommodate the increased power consumption during the initial climb flight segment. A typical volcanic gas sensing mission is shown in **Figure 4**.

The flight segments were:

1. **Automatic take-off:** The vehicle is hand-launched, during which the auto-pilot keeps the wings level and the vehicle at full thrust until an ascent threshold of 15 m is achieved. At this point the aircraft has achieved cruising air-speed and begins the waypoint mission.
2. **Main ascent:** The auto-pilot guides the aircraft along a series of large "zig-zag" ramps. The ramp angle is set in the range 10–13° which is an acceptable indefinite motor load (<50% throttle, 40 A) for the hardware chosen. The final waypoint of the ascent is set at the desired plume intercept altitude, and has a horizontal offset of at least 1 km from the summit to ensure the aircraft enters a steady straight and level fight clear of any terrain. The summit overflight altitude is planned very conservatively to minimize risks due to inaccuracies in terrain data or poor weather conditions.
3. **Plume Intercept:** there are two options:
   (a) **Automatic:** If the pre-programmed mission successfully intercepts the plume, the auto-pilot is left in automatic mode (hereinafter referred to as AUTO) to perform a series of back-and-forth transects. Successful plume interception is identified in real-time by a rise in SO_2 concentrations, which are monitored at the ground station.
   (b) **Manual:** If the vehicle fails to intercept the plume, the pilot can choose to take manual control of the aircraft with fly-by-wire mode (hereinafter referred to as FBW), using the live video stream from the forward facing camera to visually direct the vehicle toward the plume. Essential flight statistics (airspeed and altitude) are relayed to the pilot by the ground station operators.
4. **Descent:** After the plume transects are completed or a low-battery threshold is reached, whichever occurs first, automatic flight is resumed for the descent. The descent profile is a large spiral path (**Figure 4**) to an altitude of 60 m above the landing point. The vehicle then circles the landing point indefinitely until the pilot resumes control. In this environment, the decent can be as steep as required since the power requirements are minimal, however, must still be within the stable flight envelope.
5. **Landing:** The vehicle is recovered using a deployable parachute. Due to inaccuracies in GNSS positioning and drift of the barometric altimeter, this process is flown manually by the pilot who aligns the vehicle over the landing zone (flying upwind), cuts the throttle (to avoid the lines being caught in the propellers), and triggers the parachute release. The descent rate is $\sim 5\,m\,s^{-1}$ until touchdown and the airframe is sturdy enough to width-stand impact on hard ground.

Full permissions were issued by the Civil Aviation Safety Authority of Papua New Guinea (CASA PNG) with exemptions issued for Beyond Visual Line of Sight (BVLOS) operations at altitudes above the summit height. A Notice to Airmen (NOTAM) was also in place during the entire expedition period to ensure other airspace users were aware UAS were operating.

3. RESULTS

Three BVLOS flights over the summit were conducted: one (flight A) on 22 May 2019 and two (flights B and C) on 23

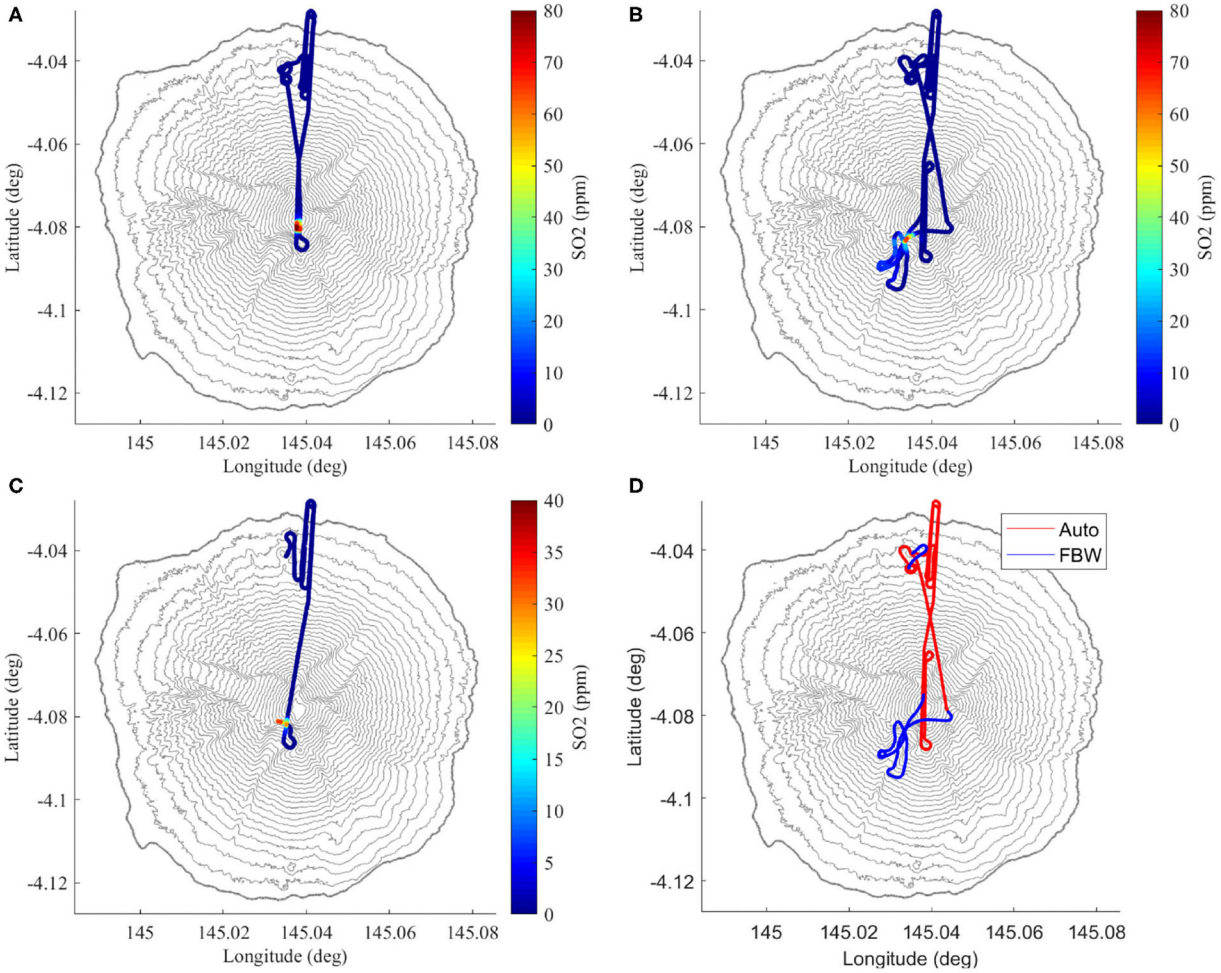

FIGURE 5 | (A–C) The three flight paths colored by SO_2 concentrations. Note how the highest SO_2 concentrations change location near the summit due to varying wind conditions. **(D)** The changing flight modes for flight B with AUTO (takeoff, ascent) and FBW (plume intercept, landing).

May 2019 (see **Table 2**). The timings of the flights were in part dictated by when the summit was clear of meteorological cloud—generally either morning, or late afternoon. All flights had pre-planned waypoint missions with a maximum altitude of 2,100 m above the take-off location and a path directly over the summit. This altitude was chosen to place the vehicle ∼600 m above the summit, as the buoyant plume typically ascended vertically for several kilometers before dispersing laterally with the wind.

The complete flight paths and measured SO_2 values are shown in **Figures 5A–C**. Non-zero SO_2 values are an indicator of plume interception, due the negligible concentrations in the background atmosphere. During flight A, the vehicle remained in AUTO during the plume transects. The plume was near-vertical at this time, so the flight intercepted the central region of the plume twice. For flight B, the pilot partially used FBW to more accurately penetrate the densest part of the plume for several transects before returning to AUTO for the descent (**Figure 5D**). The plume was slightly inclined with the wind at this time, meaning the automatic flight path passed tangentially to the plume position. Similarly, FBW was used again during flight C. In this case, however, a failure occurred during the transect and the vehicle was not recovered (section 4.3).

TABLE 2 | Details of the three flights.

Flight #	A	B	C
Date	22-05-2019	23-05-2019	23-05-2019
Local time in PNG (hh:mm)	16:49	09:17	11:48
Duration (mm:ss)	19:58	28:29	12:38
Battery consumption (mAh)	6,510	8,940	5,190
Notes	Auto plume intercept	Manual plume intercept	Vehicle lost

In section 4, we interrogate the autopilot log files to (a) quantify the plume aerodynamic conditions, (b) calculate the energy of the plume up-draft, and (c) decipher the events that

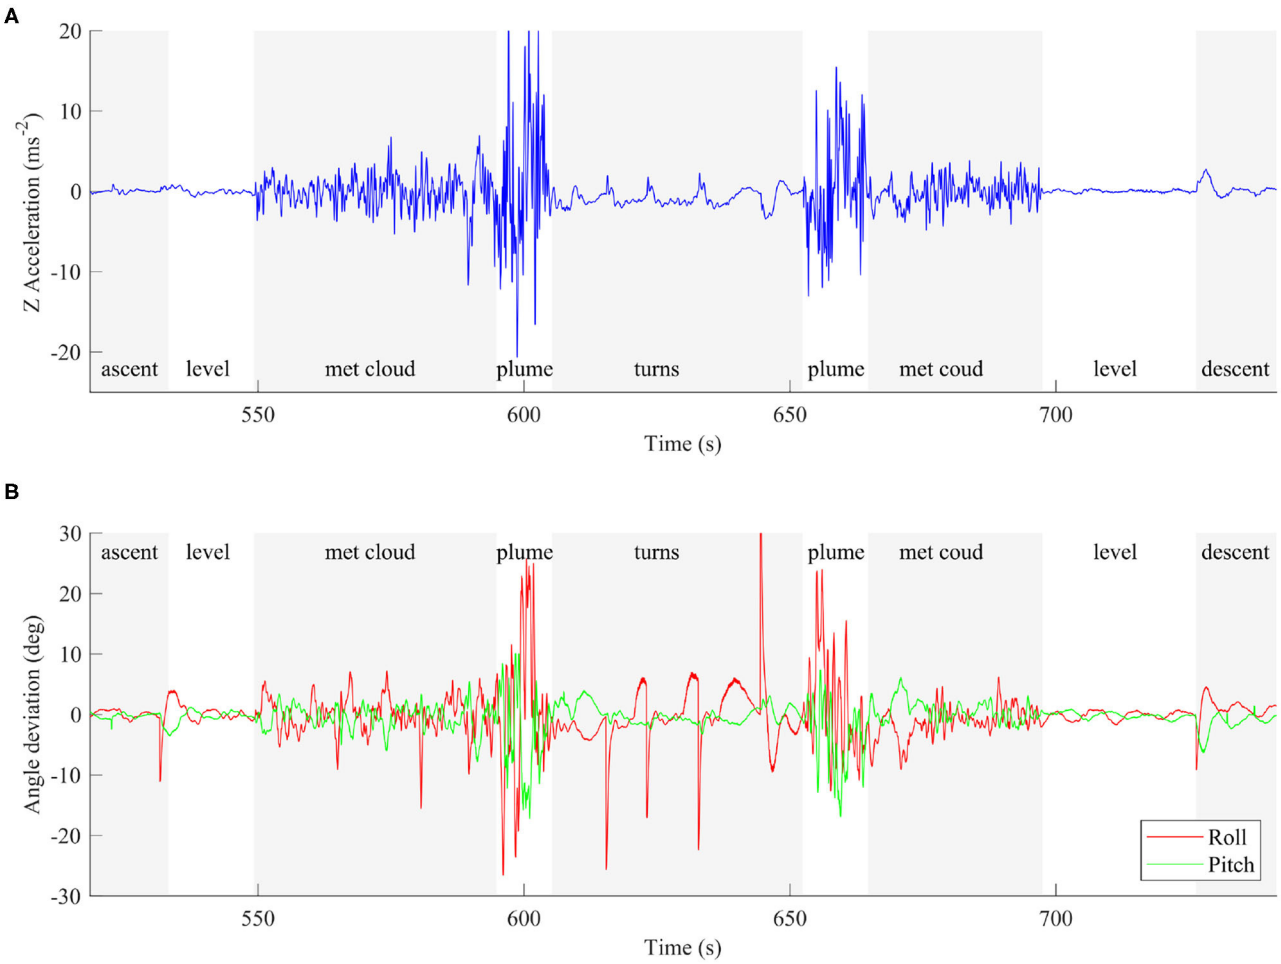

FIGURE 6 | (A) The body Z axis accelerations during Flight A. Note the difference in turbulence experienced by the aircraft as it encounters the various air conditions. **(B)** The roll and pitch deviations calculated as the difference between the orientation and its target. Larger values indicate the vehicle has been perturbed further from its trimmed condition.

led to the loss of the vehicle. For flights A and B, the full-rate log files were downloaded shortly after landing. For flight C, however, only the lower fidelity ground station telemetry log is available.

4. ANALYSIS AND DISCUSSION

4.1. Plume Conditions

Visual observations and theory predict that the conditions within a volcanic plume will be more turbulent than free air, however few data relating to quantification of these conditions exist. Here, we analyse the on-board autopilot sensors to interrogate the plume conditions encountered during flight A in detail.

The body frame accelerations provide a good indication of the conditions the vehicle was experiencing. During the summit overpass in Flight A the vehicle experienced several different air conditions, each characterized by distinctive Z (vertical) axis accelerations. The various segments have been determined by judgement, however the sharp rise in SO_2 also gives an indication of the main volcanic plume boundaries. Alternatively, methods to automatically determine the plume interception could be applied (Schellenberg et al., 2019). **Figure 6A** shows time series data for body accelerations during the summit overpass and has been labeled with the various air conditions encountered. A 5 Hz low-pass filter has been applied to remove high frequency noise. Specifically, we filter the data to ensure the accelerations are representative of the whole body accelerations rather than the small motions of the autopilot module on its flexible vibration isolation mount. The vehicle was in a state of straight and level flight ("level"; **Figure 6A**) for ~44 s, and during this time encountered a maximum acceleration of 0.82 m s^{-2} with a standard deviation of 0.23 m s^{-2}. Prior to and following plume interception, the aircraft passed through meteorological cloud surrounding the summit ("met cloud"; **Figure 6A**). In total, the vehicle was in the met cloud state for ~78 s, encountered a maximum acceleration of 11.7 m s^{-2} and a greater standard deviation of 1.88 m s^{-2} compared to level flight. The two plume traverses ("plume"; **Figure 6A**) are delimited by a step change in the magnitude

of the body accelerations encountered. The vehicle was inside the volcanic plume for ~22 s, during which the maximum vertical acceleration was ~25.1 m s^{-2}. This acceleration translates to effectively a 2.5G (where hereinafter G refers to g-force) additional loading once the offset of gravity (local gravity assumed to be 9.77 m s^{-2}) has been accounted for, and the standard deviation increases to 6.89 m s^{-2}. Following the first plume traverse, the vehicle entered a turning phase, which involves a wide 180° turn to reverse the flight path. The turn segment is not analyzed in detail since the aircraft was maneuvering actively, hence larger accelerations than level flight are expected.

We also consider the deviation of the body angles away from straight and level flight as a further indicator of plume conditions. **Figure 6B** shows the time series of the errors, where greater angle deviations represents larger differences in either pitch or roll from the target orientation. For roll this will be wings level, and for pitch it will be the cruise trim attitude. Again, we observe changes in the characteristic of both signals as the aircraft enters the different air masses described above. The pitch is controlled by the autopilot in order to return to the correct altitude, hence is expected to vary when passing though the plume, but the only cause for the roll deviations is turbulence. The roll deviations are 1–2° in clean air, 5–10° in meteorological cloud, and up to 25° in the plume. If vehicle was already turning at its maximum 35° FBW bank angle when a gust hit, it could force the aircraft into a 60° roll angle, which is outside of the tested flight envelope.

Based on the accelerations and attitude deviations observed, the vehicle is using a significant portion of its flight envelope to remain on course. Any maneuvering during a plume transect could add additional loads that move the air-frame and control algorithm outside of the tested flight envelope where failure might occur more easily. It is recommended that all plume transects are in straight lines with turns outside the plume boundaries to ensure the maximum stability and strength margins are available for the most turbulent plumes.

4.2. Energy From Plume

The thermally buoyant plume can be considered a source of additional energy to maintain flight. During flight A and B, the aircraft gained altitude when passing through the ascending air mass above the summit vent. To quantify the energy "gained" during the plume traverses we evaluate the total energy deviation of the aircraft, including the sum of the potential and kinetic energy deviations from the expected cruise conditions, and the energy not used by the propulsion system.

The potential energy deviation, E_p, was calculated by from the difference between the current height h above the target altitude h_C, hence:

$$E_p = mg\left(h - h_C\right) \tag{1}$$

The kinetic energy deviation, E_k, is calculated as the difference between the energy at cruise speed (V_C) and the current airspeed (V_T). Note the altitude-adjusted true air-speed (TAS) is used since this is the speed of the vehicle relative to the air-mass (Jimenez et al., 2017).

$$E_k = \frac{1}{2}m\left(V_T^2 - V_C^2\right) \tag{2}$$

We also account for the energy consumed by the motors, since any increase in altitude or speed may be due to increased thrust and not plume buoyancy. The power consumed (or not consumed) by the motors is calculated from the measured current (A_p) and voltage (V_p). By integrating the difference between the current power consumption and a cruise power condition (P_C), we can compare this parameter directly to the potential and kinetic energies derived above. The cruise power condition is found by averaging the power consumption during straight and level flight segments outside of the plume. The integral is taken over the time period (t_1) to (t_2) (annotated on **Figure 8**), corresponding to a subsection of the total flight from when the aircraft has finished the ascent to immediately before the descent commences.

$$E_m = \int_{t_1}^{t_2} \left[\left(V_p A_p\right) - P_C\right] dt \tag{3}$$

Auto mode was engaged during plume traverses, during which the flight computer attempts to maintain course, speed, and altitude. **Figure 7** shows time series data for altitude, throttle, airspeed, and pitch on a common time axis for the first transect. Airspeed is maintained at ~20 m s^{-1} TAS and the throttle is cut to zero (indicated by a PWM value of 1,100) whilst the autopilot demands a nose down pitch at the maximum FBW angle of −25°. Yet, despite the autopilot response, the vehicle still ascends by ~45 m above the target altitude, therefore gaining energy in a similar manner to a glider loitering in a thermal. Once the vehicle emerges from the plume and returns to less turbulent air, the autopilot energy control algorithm successfully returns the aircraft to cruise at the set speed and altitude. This is apparent from **Figure 8A**, where the potential and kinetic energy of the aircraft return to the outside of plume value. The reduction in throttle, however, was significant during the transect with the main motors stopped allowing the aircraft to briefly glide. Even with the varying throttle commands, when the total power consumption deviation is integrated over the time interval, there is an overall reduction of energy consumed compared to that consumed if the aircraft had flown in clean air. **Figure 8B** shows the power consumed and a cumulative integral of the difference between the current power and assumed cruise power of $P_C = 420$W. The final value of the integration is $E_m = -1,015$ mW h where the negative sign indicates an energy saving. This method is sensitive to the selected value of P_C, therefore the integral was also calculated for $P_C = 420 \pm 2\%$ resulting in $E_m = -1,015 \pm 394$ mW h. This equates to 44 mA h battery capacity with an assumed 22.5 V battery voltage.

Although the energy saved is only a small proportion of the total energy consumed, this evaluation demonstrates quantitatively the potential to harvest energy from a volcanic plume. Optimization of the autopilot response and mission

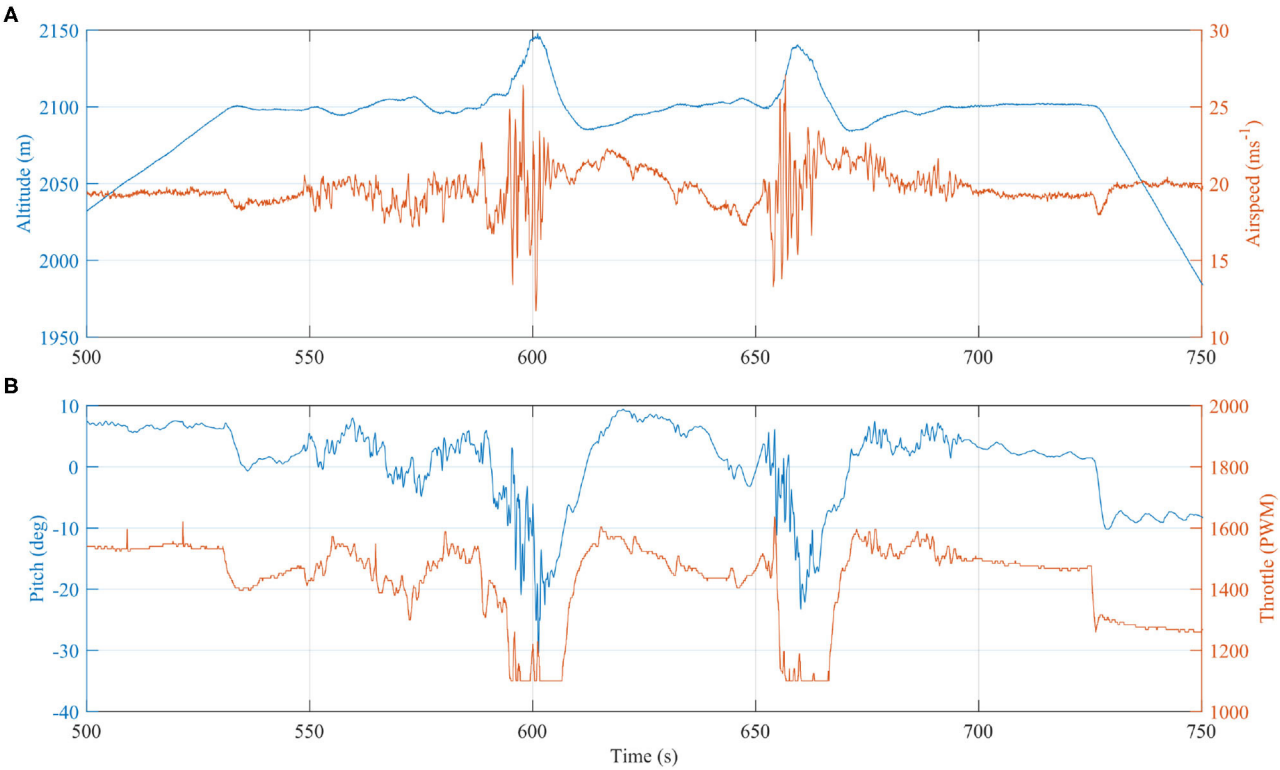

FIGURE 7 | Time series of **(A)** altitude and airspeed, and **(B)** throttle and pitch, with a common time axis as the vehicle encountered the rising plume. Due to the hot rising buoyant plume air mass, the vehicle gains altitude despite reducing throttle to zero and orientating nose down with a negative pitch angle.

plan could increase further the energy saved, and therefore the harvesting potential. For example, the steep pitch down attitude used to return the vehicle to the set point altitude will have placed the aircraft in an unfavorable aerodynamic state with increased drag losses wasting the potential gains. Also, the turning segment of the flight was planned conservatively, meaning the aircraft could have turned earlier and thereby reduced the time spend outside the buoyant plume. The aircraft is within the plume for 10 s on the first pass and 12 s on the second (**Figure 6**), which, with a true airspeed of 20.8 m s^{-1}, equates to a plume width of between 208 and 242 m. The Titan aircraft has a tightest turning circle of 120 m, therefore it may be possible to loiter within the bounds of the buoyant plume indefinitely. Extending the time spent within the plume is critical to the scientific application (i.e., volcanic gas measurements), where the associated uncertainties are for the most part related to the measurement duration. Differences in sensor response times between gas species introduce uncertainty for derived gas ratios for ground-based measurements (e.g., Roberts et al., 2017), and this effect is amplified for UAS-mounted instruments due to the comparatively short measurement periods (Liu et al., 2019). Response times, in the form of the T90 rise time (the time required for the sensor to equilibrate to 90%, when exposed to a step change in concentration), are generally on the order of tens of seconds for both the electrochemical and NDIR sensors used here. During plume traverses also on the order of tens of seconds, sensors may not have time to approach equilibrium, thus resulting in a signal that is truncated in amplitude relative to the true signal. Harvesting the thermal energy from the plume to extend flight endurance is a critical avenue for future research, and is especially relevant to long-range BVLOS operations.

Flight B has not been analyzed by the same method because the plume transects were not obvious. This was due to the auto mode missing the dense plume on the first pass and the subsequent FBW mode, which produced a more erratic flight path. The increased time spent intercepting the plume causes the power analysis integral analysis to becomes even more sensitive to the assumed value of P_C. This increases the errors to an unacceptable magnitude to be confident in drawing conclusions from the data.

4.3. Loss of Aircraft

An aircraft was lost during flight C, unfortunately. By analyzing the ground station telemetry logs, we infer the potential cause of failure, in particular, the conditions encountered in the volcanic plume and the order of events. The conclusions presented are somewhat speculative due to the limited data available, however the identified lessons learnt are still valuable for planning future operations and setting requirements for future airframe designs.

We initiated flight C as soon as possible following the successful landing of the previous flight to ensure comparable cloud conditions over the summit. Turn around time was 1 h, with tasks including downloading the sensor data files, swapping

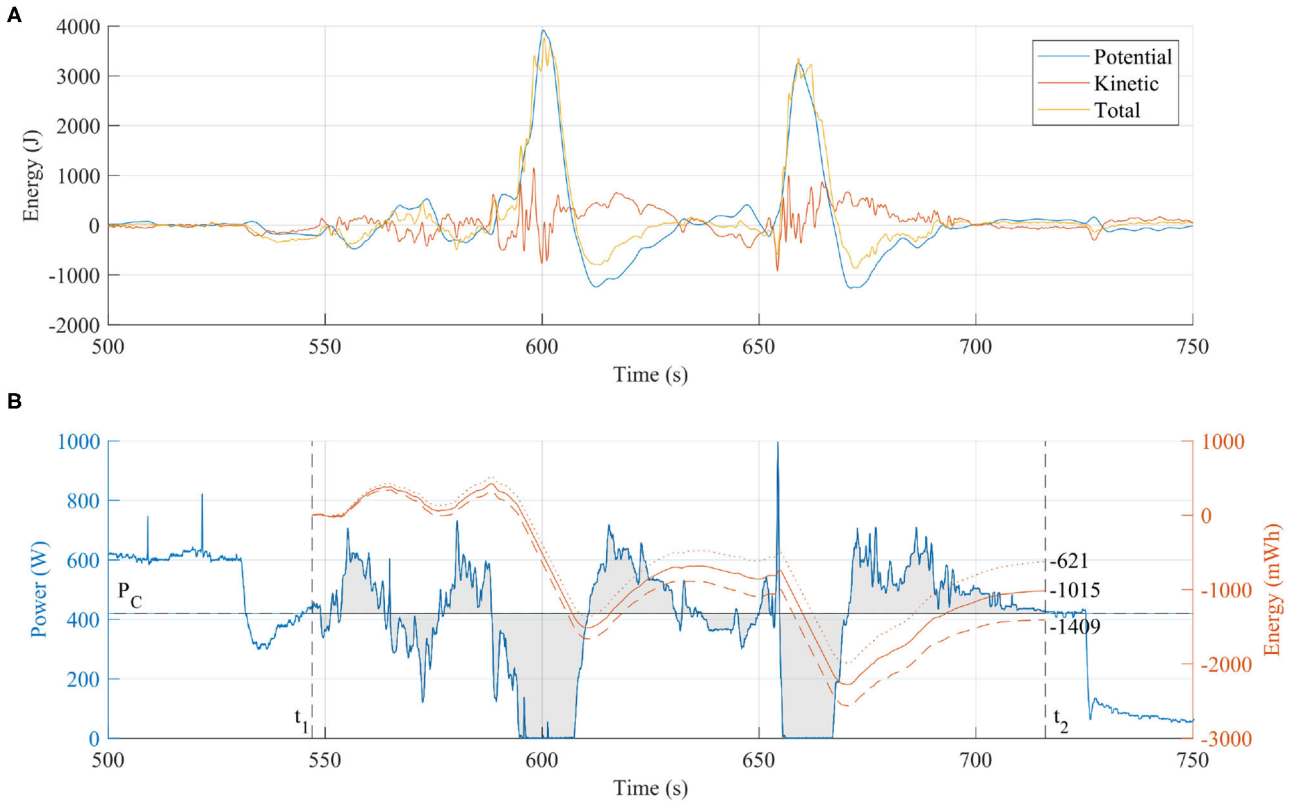

FIGURE 8 | (A) The kinetic and potential energy are almost equal before and after the plume transect because the autopilot corrected errors in speed and altitude. However, less electrical energy was consumed due to the plume updraft. **(B)** By integrating the power consumed during the plume transect period, we can calculate the battery energy "saved" in mWh. The effect of small variations (±2%) in the nominal cruise power, P_C, are shown by the dashed and dotted lines, respectively in **(B)**, which indicate the method is sensitive to this value.

battery packs, repacking the parachute, and completing all pre-flight checklists. Visually, plume conditions did not appear to change between the two flights.

The flight began with an ascent profile and plume approach identical to the previous flights (**Figures 5A,B**). After passing tangentially to the plume during the first traverse, as in flight B, FBW mode was engaged to guide the direction of the aircraft into the densest region of the plume. However, on re-entering the plume after the first manual turn, there was a catastrophic event that triggered the sequence that ultimately led to the loss of the vehicle. Referring to **Figure 9** and **Table 3**, we highlight key indicators that reveal the sequence of events.

Although the vehicle was not in a trimmed level flight condition, it was also not in a dangerous maneuver at the time of the event. We therefore suggest that the loss was caused by a sudden, high magnitude, change in external conditions. Real-time SO_2 concentrations exceeded 30 ppm immediately prior to the failure (**Figure 5C**), and therefore we confidently conclude that the aircraft had entered the main region of the plume column. In light of this, the most likely cause of the large vertical acceleration is an energetic up-draft of thermally-buoyant gas from the main volcanic vent. Nadir images from a summit overpass during flight A confirm the presence of shallow magma within the vent crater (Liu et al., 2020). Further, visual observations throughout the field deployment indicate that gas was emitted from the vent in a pulsatory manner, with sporadic pulses of fast-ascending gas superimposed on a background stable emission. At volcanoes where the magma viscosity is sufficiently low to allow decoupling between rising gas bubbles and the magma (as is the case at Manam), outgassing takes place either passively, where gas simply exsolves from the surface of the magma body, or actively, where large bubbles of overpressured gas burst more explosively at the surface (e.g., Edmonds, 2008; Tamburello et al., 2012). Pulsatory gas emissions are common to many volcanoes; the time scale of the periodicity reflects the fluid dynamics that govern gas bubble ascent (Pering et al., 2019). At Manam, the pulses of energetic, thermally-buoyant gas associated with active outgassing would have generated transient conditions of extreme turbulence within the plume column. We propose that entering one of these up-drafts over-stressed the airframe and ultimately resulted in the loss of either a wing-tip or one of the V-tail stabilizers. The loss of either component would generate significantly unbalanced aerodynamic forces, consistent with the rapid rotations indicated by roll and yaw rates. The video link continued transmission for a significant time after the loss

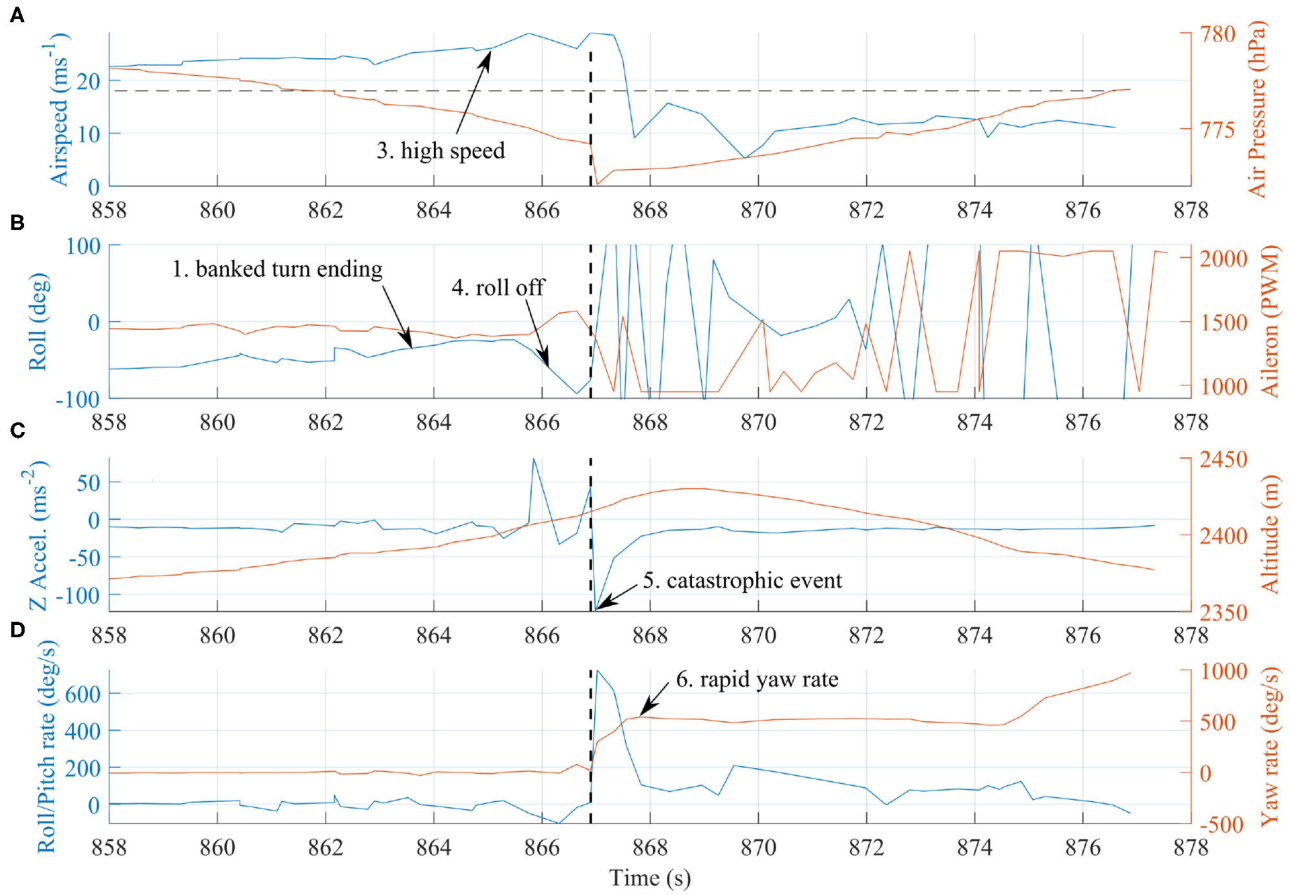

FIGURE 9 | Indicative signals from the telemetry log file of flight C for the final 20 s of transmission. Signals **(A)** airspeed and air pressure, **(B)** roll angle and aileron input command, **(C)** vertical acceleration and altitude, and **(D)** roll, pitch, and yaw angular rates, have been plotted against a common time-axis. A catastrophic event occurs at 867 s followed by a rapid (>1 Hz) spin. See **Table 3** for further details of the order of events.

TABLE 3 | Summary of the sequence of events that led to loss of the vehicle during flight C. Time is reported relative to takeoff time and is equivalent to **Figure 9**.

	Time (s)	Event
1.	858.0–865.0	Vehicle finished a relatively high speed (23 ms^{-1}) turn and was returning toward wings level.
2.	865.0	Due to the higher airspeed, the aircraft climbed ~40 m above the cruise set-point, but was near level in pitch orientation.
3.	865.0–867.0	In the 2 s period immediately before the event, the vehicle increased in speed to 26 ms^{-1} due to a high throttle command.
4.	865.8	At this point the vehicle experienced a sudden ~7G upwards acceleration (71ms^{-2}) and started to roll rapidly at a rate of 100° s^{-1}. The hypothesis is that at this point the aircraft experienced a significant up-draft from the main volcanic vent. The aileron moves to oppose the motion, but does not have a significant effect.
5.	867.0	There are several indicators of a catastrophic event. The body vertical acceleration suddenly changed sign and magnitude to ~11G (-113ms^{-2}). There is also a sudden change in static air pressure, and an abrupt reversal in roll rate.
6.	867–868	Very rapid yaw rate develops indicating a flat spin. Observations from the live video link also confirmed the vehicle was no longer maintaining orientation.
7.	867–877	The main telemetry link was lost 10 s after the event. The vehicle was still high above the summit at this time and falling at ~11 ms^{-1}. The aircraft's altitude was still high enough for a clear line of sight hence the loss of telemetry is thought to be due to the aircraft breaking up. In this case, we suggest the loss of a wing containing the telemetry module.

of telemetry, suggesting multiple stages of failure following the initial event.

From this investigation we present recommendations for the required aircraft strength based upon the loading encountered. The ~7G up-gust was the most likely cause of failure, therefore applying a reserve factor of 2 produces a requirement of 14G upwards load. A reserve factor of 2 is greater than the typical 1.5 used within aerospace design processes, however there are very limited data available quantifying volcanic plume conditions, hence a greater reserve margin is preferred. The

TABLE 4 | Recommended design requirements for fixed-wing UAS, applied to volcanic plume measurements.

Requirement	Description
Payload capacity	Minimum 1 kg to accommodate high sensitivity instrumentation and appropriate shielding. With careful design iterations a sensor mass might be minimized below this value, however this is a general starting recommendation for airframes of this size.
Ascent capability	Minimum 500 m above the summit height. This allows for conservative flight planning and also mitigates potential variations in the height requirements. The exact height of a volcanic summit might be unknown, or the activity may have modified the summit topography since the most recent survey. Variations in air density due to weather might also reduce the maximum ascent capability of a vehicle.
Airframe strength	14G. This value was determined from the conditions encountered in Flight C including a reserve factor of 2 applied. Structural strengthening should be applied to areas of load concentration, such as the main wing, and where control surfaces attach to the fuselage.
Airspeed capability	Minimum 20 ms^{-1}. A reduced airspeed could be advantageous to allow the aircraft to pass through the plume more slowly and therefore collect more data points. However, strong, topographically-enhanced winds can occur around prominent volcanic peaks, hence the top speed of the vehicle must be fast enough to overcome these. Note that the windspeed measured at the takeoff location can often be lower than that encountered at plume altitude.
Structure	At least partial sealing to minimize airflow over the flight avionics. Volcanic plumes contain acidic gases at high humidity, which cause corrosion and failure of electronics. Sealed enclosures may not be possible for components requiring airflow cooling (e.g., main motors), so these items should be inspected regularly.
Flight control modes	Automatic for large ascents to maintain optimal trim conditions, but with a visually-guided (First Person View; FPV) Fly By Wire option to make course modifications that ensure plume interception, as necessary. Once the aircraft is more than 500 m away from the pilot, it is unlikely they will be able to operate using eyesight alone. Augmenting the pilot situational awareness via FPV and flight telemetry data is important for accurate maneuvers and rapid response to unexpected situations.
Maximum roll angle	The maximum roll angle has a direct impact on the structural loading during turn maneuvers. By reducing the maximum roll angle, the loads created during maneuvers can be reduced, leaving a greater strength margin for unexpected turbulence-based loads. However, a reduction below $<20^{\circ}$ would compromise the handling qualities for manual flight control, hence a balance must be found. It is also recommended that plume transects only be attempted when in straight and level flight. For the Titan aircraft, the AUTO and FBW flight modes were set with a maximum roll angle of 45°.

downward load can be more conservative since the aircraft is not expected to be flown inverted or have significant down-gusts. The speed of the aircraft was within the range of values under which the vehicle had been tested, however the additional wing loading from such high speeds would have reduced the structural strength reserve. Consequently, a second operational recommendation is a more advanced FBW control system that maintains altitude and speed, but allows the pilot to "drive" the aircraft's direction.

Key recommended design criteria have been summarized in **Table 4**. These criteria are based upon the numerical values derived during the above analysis, the vehicle setup parameters, and from experimental field experience.

5. CONCLUSION

Volcanic environments present many challenges for aerial robotics, from the vehicle design through flight planning to the conditions encountered during the flight itself. Yet, despite these obstacles, instrumented UAS are stimulating transformative advances in volcanological research, motivating further engineering development to respond to these challenges. Here, we describe a series of fixed-wing flights BVLOS over the summit of Manam volcano, Papua New Guinea, to measure real-time gas concentrations within the volcanic plume. Our aim was to collect data that would constrain the emission rate of environmentally-important volcanic gases, such as carbon dioxide. However, the insights contributed by this study are also relevant more generally to other plume sampling applications. Specifically, we show that (a) the "Titan" aircraft is a versatile aircraft suited to BVLOS missions in difficult terrain; (b) an air frame can reasonably expect to be subjected to a 2.5G loading when traversing a thermally-buoyant volcanic plume, and that this may increase to 7G in more extreme, but transient, cases; and (c) energy harvesting from the volcanic plume presents a tractable means to enhance flight endurance, and thus extend the duration of scientific measurements. We describe the physical parameters and propulsion systems used in our aircraft design at a level of detail sufficient to guide future air craft design, and recommend that vehicles are strength tested up to 14G to ensure a factor of 2 reserve against the upper end-member condition. Further, for large plumes, such as that encountered at Manam, the flight path could be optimized to ensure maximum additional energy gain. From on-board flight parameters, we reconstruct the sequence of events that ultimately led to catastrophic vehicle failure and attribute the cause of failure to interaction with an energetic thermal updraft from the main volcanic vent, which is a prevalent characteristic of outgassing at similar volcanoes globally and therefore needs to be taken into account during full systems testing. If recent trends continue, scientific applications will increasing look to aerial robotics to enable sensor placement in hazardous environments. The large spatial scales and shifting targets involved (i.e., the plume is in constant motion, vertically and often horizontally) make FBW mode essential to ensure optimal data collection. Therefore, looking forward, the development of FBW modes with more autonomy for the speed and altitude loops, or the introduction of plume-hunting algorithms capable of processing sensor data in real-time for complete automation will be critical to continued advance in this field. Further, the use of quantitative ground-based measurements of plume parameters (for example,

plume rise speed, transport direction, and dimensions) to inform flight planning would contribute to both risk reduction and flight efficiency.

DATA AVAILABILITY STATEMENT

The raw data supporting the conclusions of this article will be made available by the authors, without reservation.

AUTHOR CONTRIBUTIONS

KW, EJL, and TR drafted the manuscript. KW, EJL, TR, RC, and JF built and operated the UAS during the field work. AA, GG, and MB developed the gas sensor, supported its integration into the UAS, and contributed to data analysis. KM and II supported the data collection and assisted with in-country logistics. All authors were involved in the data collection, read and revised the manuscript, and approved the submitted version.

FUNDING

This research was enabled through the Alfred P. Sloan Foundation's support of the Deep Carbon Observatory Deep Earth Carbon Degassing program (DECADE). Part funding also came from the EPSRC CASCADE programme grant (EP/R009953/1). EJL was supported by a Leverhulme Trust Early Career Fellowship. KW was supported by the National Center for Nuclear Robotics (NCNR) EPSRC grant (EP/R02572X/1).

ACKNOWLEDGMENTS

This work would not have been possible without the generous assistance provided by the staff of the Rabaul Volcanological Observatory (RVO), and the kind hospitality of the Manam community. We also thank the Civil Aviation Safety Autority of Papua New Guinea (CASA PNG), for help and guidance planning BVLOS UAS operations. The authors would like to thank all contributors to ArduPilot, which has enabled this work. We thank all members of the ABOVE field team.

REFERENCES

Aiuppa, A., Fischer, T. P., Plank, T., and Bani, P. (2019). CO_2 flux emissions from the Earth's most actively degassing volcanoes, 2005–2015. *Sci. Rep*. 9:5442. doi: 10.1038/s41598-019-41901-y

Aiuppa, A., Moretti, R., Federico, C., Giudice, G., Gurrieri, S., Liuzzo, M., et al. (2007). Forecasting Etna eruptions by real-time observation of volcanic gas composition. *Geology* 35:1115. doi: 10.1130/G24149A.1

Bhardwaj, A., Sam, L., Akanksha., Martín-Torres, F. J., and Kumar, R. (2016). UAVs as remote sensing platform in glaciology: present applications and future prospects. *Rem. Sens. Environ*. 175, 196–204. doi: 10.1016/j.rse.2015.12.029

Carn, S. A., Fioletov, V. E., McLinden, C. A., Li, C., and Krotkov, N. A. (2017). A decade of global volcanic SO_2 emissions measured from space. *Sci. Rep*. 7:44095. doi: 10.1038/srep44095

Cassano, J. J. (2013). Observations of atmospheric boundary layer temperature profiles with a small unmanned aerial vehicle. *Antarctic Sci*. 26, 205–213. doi: 10.1017/S0954102013000539

Connor, D. T., Wood, K., Martin, P. G., Goren, S., Megson-Smith, D., Verbelen, Y., et al. (2020). Radiological mapping of post-disaster nuclear environments using fixed-wing unmanned aerial systems: a study from chornobyl. *Front. Robot. AI* 6:149. doi: 10.3389/frobt.2019.00149

Corrigan, C. E., Roberts, G. C., Ramana, M. V., Kim, D., and Ramanathan, V. (2008). Capturing vertical profiles of aerosols and black carbon over the Indian Ocean using autonomous unmanned aerial vehicles. *Atmos. Chem. Phys*. 8, 737–747. doi: 10.5194/acp-8-737-2008

Darmawan, H., Walter, T. R., Brotopuspito, K. S., Subandriyo., and Nandaka, I. G. M. A. (2018). Morphological and structural changes at the Merapi lava dome monitored in 2012–15 using unmanned aerial vehicles (UAVs). *J. Volcanol. Geothermal Res*. 349, 256–267. doi: 10.1016/j.jvolgeores.2017.11.006

de Boer, G., Palo, S., Argrow, B., LoDolce, G., Mack, J., Gao, R.-S., et al. (2016). The Pilatus unmanned aircraft system for lower atmospheric research. *Atmos. Meas. Techn*. 9, 1845–1857. doi: 10.5194/amt-9-1845-2016

de Moor, J. M., Aiuppa, A., Avard, G., Wehrmann, H., Dunbar, N., Muller, C., et al. (2016). Turmoil at Turrialba Volcano (Costa Rica): degassing and eruptive processes inferred from high-frequency gas monitoring. *J. Geophys. Res. Solid Earth* 121, 5761–5775. doi: 10.1002/2016JB013150

de Moor, J. M., Stix, J., Avard, G., Muller, C., Corrales, E., Diaz, J. A., et al. (2019). Insights on hydrothermal-magmatic interactions and eruptive processes at Poás Volcano (Costa Rica) from high-frequency gas monitoring and drone measurements. *Geophys. Res. Lett*. 46, 1293–1302. doi: 10.1029/2018GL080301

Detert, M., and Weitbrecht, V. (2015). A low-cost airborne velocimetry system: proof of concept. *J. Hydraul. Res*. 53, 532–539. doi: 10.1080/00221686.2015.1054322

Detweiler, C., Ore, J.-P., Anthony, D., Elbaum, S., Burgin, A., and Lorenz, A. (2015). Environmental reviews and case studies: bringing unmanned aerial systems closer to the environment. *Environ. Pract*. 17, 188–200. doi: 10.1017/S1466046615000174

Di Stefano, G., Romeo, G., Mazzini, A., Iarocci, A., Hadi, S., and Pelphrey, S. (2018). The Lusi drone: a multidisciplinary tool to access extreme environments. *Mar. Petrol. Geol*. 90, 26–37. doi: 10.1016/j.marpetgeo.2017.07.006

Edmonds, M. (2008). New geochemical insights into volcanic degassing. *Philos. Trans. R. Soci. A Math. Phys. Eng. Sci*. 366, 4559–4579. doi: 10.1098/rsta.2008.0185

Favalli, M., Fornaciai, A., Nannipieri, L., Harris, A., Calvari, S., and Lormand, C. (2018). UAV-based remote sensing surveys of lava flow fields: a case study from Etna's 1974 channel-fed lava flows. *Bull. Volcanol*. 80:29. doi: 10.1007/s00445-018-1192-6

Fischer, T. P., and Aiuppa, A. (2020). AGU centennial grand challenge: volcanoes and deep carbon global CO_2 emissions from subaerial volcanism—recent progress and future challenges. *Geochem. Geophys. Geosyst*. 21:e2019GC008690. doi: 10.1029/2019GC008690

Fischer, T. P., Arellano, S., Carn, S., Aiuppa, A., Galle, B., Allard, P., et al. (2019). The emissions of CO_2 and other volatiles from the world's subaerial volcanoes. *Sci. Rep*. 9:18716. doi: 10.1038/s41598-019-54682-1

Fladeland, M., Sumich, M., Lobitz, B., Kolyer, R., Herlth, D., Berthold, R., et al. (2011). The NASA SIERRA science demonstration programme and the role of small–medium unmanned aircraft for earth science investigations. *Geocarto Int*. 26, 157–163. doi: 10.1080/10106049.2010.537375

Global Volcanism Program (2019). Report on Manam (Papua New Guinea). *Bull. Glob. Volcan. Netw*. 44:10. doi: 10.5479/si.GVP.BGVN201902-251020

Greatwood, C., Richardson, T., Freer, J., Thomas, R., MacKenzie, A., Brownlow, R., et al. (2017). Atmospheric sampling on Ascension Island using multirotor UAVs. *Sensors* 17:1189. doi: 10.3390/s17061189

Hill, S. L., and Clemens, P. (2015). "Miniaturization of high spectral spatial resolution hyperspectral imagers on unmanned aerial systems," in *Next-Generation Spectroscopic Technologies VIII*, Vol. 9482, eds M. A. Druy, R. A. Crocombe, and D. P. Bannon (SPIE: Baltimore, MD), 345–359. doi: 10.1117/12.2193706

Immerzeel, W. W., Kraaijenbrink, P. D. A., Shea, J. M., Shrestha, A. B., Pellicciotti, F., Bierkens, M. F. P., et al. (2014). High-resolution monitoring of Himalayan

glacier dynamics using unmanned aerial vehicles. *Rem. Sens. Environ.* 150, 93–103. doi: 10.1016/j.rse.2014.04.025

James, M. R., Carr, B. B., D'Arcy, F., Diefenbach, A. K., Dietterich, H. R., Fornaciai, A., et al. (2020). Volcanological applications of unoccupied aircraft systems(UAS): developments, strategies, and future challenges. *Volcanica* 3, 67–114. doi: 10.30909/vol.03.01.67114

Jimenez, P., Silva, J. P., and Hernandez, J. (2017). "Experimental validation of unmanned aerial vehicles to tune PID controllers in open source autopilots," in *Proceedings of the 7th European Conference for Aeronautics and Space Sciences* (Milano).

Jordan, B. R. (2019). Collecting field data in volcanic landscapes using small UAS (sUAS)/drones. *J. Volcanol. Geotherm. Res.* 385, 231–241. doi: 10.1016/j.jvolgeores.2019.07.006

Klemas, V. V. (2015). Coastal and environmental remote sensing from unmanned aerial vehicles: an overview. *J. Coast. Res.* 315, 1260–1267. doi: 10.2112/JCOASTRES-D-15-00005.1

Liu, E. J., Aiuppa, A., Alan, A., Arellano, S., Bitetto, M., Bobrowski, N., et al. (2020). Aerial strategies advance volcanic gas measurements at inaccessible, strongly degassing volcanoes. *Sci. Adv.* (in press). doi: 10.1126/sciadv.abb9103

Liu, E. J., Wood, K., Mason, E., Edmonds, M., Aiuppa, A., Giudice, G., et al. (2019). Dynamics of outgassing and plume transport revealed by proximal unmanned aerial system (UAS) measurements at Volcán Villarrica, Chile. *Geochem. Geophys. Geosyst.* 20, 730–750. doi: 10.1029/2018GC007692

McGonigle, A. J. S., Aiuppa, A., Giudice, G., Tamburello, G., Hodson, A. J., and Gurrieri, S. (2008). Unmanned aerial vehicle measurements of volcanic carbon dioxide fluxes. *Geophys. Res. Lett.* 35:L06303. doi: 10.1029/2007GL032508

Mercer, J., and Kelman, I. (2010). Living alongside a volcano in Baliau, Papua New Guinea. *Disaster Prev. Manag.* 19, 412–422. doi: 10.1108/09653561011070349

Nadeau, P. A., Elias, T., Kern, C., Lerner, A. H., Werner, C. A., Cappos, M., et al. (2018). "The 2018 eruption of Kīlauea volcano: tales from a gas perspective," in *AGU Fall Meeting Abstracts*, Vol. 2018, V21B-07 (Washington, DC).

Nagai, M., Chen, T., Shibasaki, R., Kumagai, H., and Ahmed, A. (2009). UAV-Borne 3-D mapping system by multisensor integration. *IEEE Trans. Geosci. Rem. Sens.* 47, 701–708. doi: 10.1109/TGRS.2008.2010314

Pajares, G. (2015). Overview and current status of remote sensing applications based on unmanned aerial vehicles (UAVs). *Photogramm. Eng. Rem. Sens.* 81, 281–330. doi: 10.14358/PERS.81.4.281

Palfreyman, W. D., and Cooke, R. J. S. (1976). "Eruptive history of Manam volcano, Papua New Guinea," in *Volcanism in Australasia*, ed R. W. Johnson (Amsterdam: Elsevier), 117–131.

Peng, Z.-R., Wang, D., Wang, Z., Gao, Y., and Lu, S. (2015). A study of vertical distribution patterns of PM2.5 concentrations based on ambient monitoring with unmanned aerial vehicles: a case in Hangzhou, China. *Atmos. Environ.* 123, 357–369. doi: 10.1016/j.atmosenv.2015.10.074

Pering, T. D., Ilanko, T., and Liu, E. J. (2019). Periodicity in volcanic gas plumes: a review and analysis. *Geosciences* 9:394. doi: 10.3390/geosciences9090394

Ramana, M. V., Ramanathan, V., Kim, D., Roberts, G. C., and Corrigan, C. E. (2007). Albedo, atmospheric solar absorption and heating rate measurements with stacked UAVs. *Q. J. R. Meteorol. Soc.* 133, 1913–1931. doi: 10.1002/qj.172

Roberts, T. J., Lurton, T., Giudice, G., Liuzzo, M., Aiuppa, A., Coltelli, M., et al. (2017). Validation of a novel multi-gas sensor for volcanic HCl alongside H_2S and SO_2 at Mt. Etna. *Bull. Volcanol.* 79:36. doi: 10.1007/s00445-017-1114-z

Schellenberg, B., Richardson, T., Watson, M., Greatwood, C., Clarke, R., Thomas, R., et al. (2019). Remote sensing and identification of volcanic plumes using fixed-wing UAVs over Volcán de Fuego, Guatemala. *J. Field Robot.* 36, 1192–1211. doi: 10.1002/rob.21896

Shinohara, H. (2013). Composition of volcanic gases emitted during repeating vulcanian eruption stage of Shinmoedake, Kirishima volcano, Japan. *Earth Planets Space* 65, 667–675. doi: 10.5047/eps.2012.11.001

Stöcker, C., Eltner, A., and Karrasch, P. (2015). Measuring gullies by synergetic application of UAV and close range photogrammetry A case study from Andalusia, Spain. *CATENA* 132, 1–11. doi: 10.1016/j.catena.2015.04.004

Syahbana, D. K., Kasbani, K., Suantika, G., Prambada, O., Andreas, A. S., Saing, U. B., et al. (2019). The 2017–19 activity at Mount Agung in Bali (Indonesia): Intense unrest, monitoring, crisis response, evacuation, and eruption. *Sci. Rep.* 9:8848. doi: 10.1038/s41598-019-45295-9

Tamburello, G., Aiuppa, A., Kantzas, E. P., McGonigle, A. J. S., and Ripepe, M. (2012). Passive vs. active degassing modes at an open-vent volcano (Stromboli, Italy). *Earth Planet. Sci. Lett.* 359–360, 106–116. doi: 10.1016/j.epsl.2012.09.050

Tamminga, A. D., Eaton, B. C., and Hugenholtz, C. H. (2015). UAS-based remote sensing of fluvial change following an extreme flood event. *Earth Surf. Process. Landf.* 40, 1464–1476. doi: 10.1002/esp.3728

Turner, N. R., Perroy, R. L., and Hon, K. (2017). Lava flow hazard prediction and monitoring with UAS: a case study from the 2014–2015 Pāhoa lava flow crisis, Hawai'i. *J. Appl. Volcanol.* 6:17. doi: 10.1186/s13617-017-0068-3

Villa, T., Gonzalez, F., Miljievic, B., Ristovski, Z., and Morawska, L. (2016). An overview of small unmanned aerial vehicles for air quality measurements: present applications and future prospectives. *Sensors* 16:1072. doi: 10.3390/s16071072

Vivoni, E. R., Rango, A., Anderson, C. A., Pierini, N. A., Schreiner-McGraw, A. P., Saripalli, S., et al. (2014). Ecohydrology with unmanned aerial vehicles. *Ecosphere* 5:art130. doi: 10.1890/ES14-00217.1

Werner, C., Fischer, T. P., Aiuppa, A., Edmonds, M., Cardellini, C., Carn, S., et al. (2019). "Carbon dioxide emissions from subaerial volcanic regions," in *Deep Carbon* (Cambridge: Cambridge University Press), 188–236. doi: 10.1017/9781108677950.008

Wildmann, N., Mauz, M., and Bange, J. (2013). Two fast temperature sensors for probing of the atmospheric boundary layer using small remotely piloted aircraft (RPA). *Atmos. Meas. Tech.* 6, 2101–2113. doi: 10.5194/amt-6-2101-2013

Woods, A. W. (2010). Turbulent plumes in nature. *Annu. Rev. Fluid Mech.* 42, 391–412. doi: 10.1146/annurev-fluid-121108-145430

Zweig, C. L., Burgess, M. A., Percival, H. F., and Kitchens, W. M. (2015). Use of unmanned aircraft systems to delineate fine-scale wetland vegetation communities. *Wetlands* 35, 303–309. doi: 10.1007/s13157-014-0612-4

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

frontiers in Robotics and AI

ORIGINAL RESEARCH
published: 26 November 2020
doi: 10.3389/frobt.2020.499056

Radiation Mapping and Laser Profiling Using a Robotic Manipulator

Samuel R. White[1], David A. Megson-Smith[1], Kaiqiang Zhang[1], Dean T. Connor[1], Peter G. Martin[1], Chris Hutson[1], Guido Herrmann[2,3], John Dilworth[4] and Thomas B. Scott[1,5]*

[1] Department of Physics, Interface Analysis Centre, University of Bristol, Bristol, United Kingdom, [2] Department of Mechanical Engineering, University of Bristol, Bristol, United Kingdom, [3] Department of Electrical and Electronic Engineering, The University of Manchester, Manchester, United Kingdom, [4] KUKA Systems UK Ltd, Halesowen, United Kingdom, [5] RACE UKAEA, Cullham Science Centre, Abingdon, United Kingdom

The use of a robotic arm manipulator as a platform for coincident radiation mapping and laser profiling of radioactive sources on a flat surface is investigated in this work. A combined scanning head, integrating a micro-gamma spectrometer and Time of Flight (ToF) sensor were moved in a raster scan pattern across the surface, autonomously undertaken by the robot arm over a 600 × 260 mm survey area. A series of radioactive sources of different emission intensities were scanned in different configurations to test the accuracy and sensitivity of the system. We demonstrate that in each test configuration the system was able to generate a centimeter accurate 3D model complete with an overlaid radiation map detailing the emitted radiation intensity and the corrected surface dose rate.

Keywords: radiation mapping, 3D modeling, spectrometry, gamma scanning, nuclear waste, robotic manipulator

OPEN ACCESS

Edited by:
Chie Takahashi, University of Cambridge, United Kingdom

Reviewed by:
C. James Taylor, Lancaster University, United Kingdom
Sean Andersson, Boston University, United States

****Correspondence:***
Samuel R. White
sam.white@bristol.ac.uk

Specialty section:
This article was submitted to Sensor Fusion and Machine Perception, a section of the journal Frontiers in Robotics and AI

Received: *18 September 2019*
Accepted: *04 September 2020*
Published: *26 November 2020*

Citation:
White SR, Megson-Smith DA, Zhang K, Connor DT, Martin PG, Hutson C, Herrmann G, Dilworth J and Scott TB (2020) Radiation Mapping and Laser Profiling Using a Robotic Manipulator. Front. Robot. AI 7:499056. doi: 10.3389/frobt.2020.499056

1. INTRODUCTION

The global nuclear industry is facing significant challenges in decommissioning and nuclear waste management owing to an ever-increasing amount of nuclear waste awaiting to be processed and prepared for long-term storage. In 2018, *Status and Trends in Spent Fuel and Radioactive Waste Management* from the IAEA Nuclear Energy Series, reported that globally there is some 6,317,000 m^3 of nuclear waste in storage awaiting a long term disposal solution (IAEA, 2018). Each waste category has a different associated disposal cost, for example Low Level Waste (LLW) in the UK has an attributable cost of £2.9k per m^3, whilst Intermediate Level Waste (ILW) has a cost of £46k per m^3 based on 2008 data (UK Government Department of Energy & Climate Change, 2011). It may be considered logical that higher activity wastes should cost more to manage because they represent a higher hazard to humans. ILW and HLW have activities which are sufficiently extreme that humans cannot come into close contact with them, mandating their remote handling and inspection. Accordingly, on a fiscal and safety basis it is important that waste materials are not inadvertently processed into the wrong waste category. Thresholds between these waste types are clearly defined (in terms of activity per unit mass) and there is a significant cost difference in managing each waste type. Nuclear decommissioning and waste management is therefore in urgent need for technologies that deliver high accuracy and automated radiation and 3D surveying to help waste sorting. There are numerous "sort and segregation" activities which seek to characterize mixed nuclear wastes into their correct streams, both repeatedly and with a high throughput; whilst avoiding the need for cost increases through excessive conservatism (Horizon2020, 2018). There is a current funding competition call from Sellafield to address the problem of sort and segregation tables, in which technology can identify radioisotopes and objects on tables is desired to solve this problem (Gov, 2020). In addition, there is a concurrent critical need to develop and implement

technologies with the ability to scan packaged wastes held in storage, to check for external radiation hot spots and/or signs of surface deformation or corrosion. In the UK, Sellafield Ltd accommodates one of the largest inventories of ILW. This must be routinely checked until such time as a Geological Disposal Facility (GDF) becomes available. An inability to routinely check waste in storage could present a multitude of issues.

An alternative approach, which is only now becoming possible, is in the use of robotic manipulators equipped with micro gamma-spectrometers to scan waste packages in a more dynamic way. Such solid-state detection units are usually very compact with detector crystals of 1–30 cm^3 and fast counting rates (typically 20,000 cps) able to discern different gamma-emitting radioisotopes based on their differing decay energies. Conceptually, they enable radiation scans to be performed robotically at much smaller stand-off distances (<10 cm) than segmented gamma scanning, yielding a much higher spatial resolution and sensitivity. However, in order to conduct such close-proximity scanning, a method of determining the sensor stand-off distance must also be integrated. To touch the waste material could potentially contaminate or damage the detector and hence this needs to be prevented at all costs.

There are therefore numerous robotic technologies and sensors which are capable of being combined to achieve this target of combined gamma scanning and 3D profiling of nuclear waste objects—yet to our knowledge this has not previously been reported. The use of robotics for gamma inspection in the nuclear industry is not new, but has typically been employed for plant inspection and not waste assay. Tsitsimpelis et al. (2019) discuss further developments in a recent review paper on ground-based robotic systems for the characterization of nuclear environments, highlighting numerous robotic systems which have been deployed for radiation monitoring within the last 50 years. In 1994, Redus et al. (1994) published a paper on the use of video and gamma ray imaging systems for inspection robots in nuclear environments. The group used a robot with a mounted gamma spectrometer and camera to record video footage with super-imposed gamma ray imaging enabling the identification of radioactive sources in a room. This is the first example of a robotic radiation mapping procedure. Since then Bird et al. (2019) have developed this concept by researching the use of mobile robotic platforms for the routine inspection of nuclear facilities. The Continuous Autonomous Radiation Monitoring Assistance (CARMA) robot uses LiDAR sensors for Simultaneous Localisation and Mapping (SLAM), generating a 2D plan map of the room. In addition to this, they obtain gamma radiation intensity mapping data from a Thermo Fisher Scientific RadEye. In a similar way, a considerable amount of radiation mapping research has been conducted with the use of Unmaned Aerial Vehicles (UAVs) to map radiation at nuclear sites including the Chernobyl and Fukushima fallout zones. Martin et al. (2017) presented "High-Resolution Aerial Radiation Mapping for Nuclear Decontamination and Decommissioning". The UAV flew autonomously along GPS defined flight paths above the Sellafield nuclear decommissioning facility, using a sensor package which simultaneously recorded GPS position, above ground height using a ranging LiDAR and gamma spectrometry data. This data was used to produce an accurate radiation map of each survey area studied, as well as using the spectral information from gamma measurements data to identify various radioactive isotopes in different facilities (Martin et al., 2016b). All of these techniques connect radiometric data to positional data to generate a radiation map.

The use of point cloud data in the formation of 3D models is another recent innovation that the field of robotics has been quick to adopt. Various techniques are routinely being applied to generate point cloud data, which in turn can be transformed into a 3D model. Within nuclear robotics, LiDAR scanning is already an established technique used in 3D environment reconstruction. Aerial radiation mapping routinely relates recorded radiometric data to a 3D model collected by either LiDAR or photogrammetry to produce a combined 3D representation (Connor et al., 2016; Martin et al., 2016a). There are numerous examples of high quality 3D models being generated for robotic systems (Marturi et al., 2018; Sarker et al., 2019; Barone et al., 2020). However they are all reliant on Charged Coupled Device (CCD) cameras. The main issue here is that within a highly radioactive environment CCD based devices fail, as the gamma radiation causes damage to their internal Metal Oxide Semi-conductor (MOS) capacitors. More novel are Time of Flight (ToF) scanners and cameras that are often able to generate data from complex objects for 3D reconstruction. The advantage here is ToF sensors often have small size and therefore small gamma interaction cross sections, making them more radiation hard. Hoegg et al. (2013) managed to reconstruct 3D models of a selection of cars using a ToF camera, whilst Gutierrez-Villalobos et al. (2017) created an accurate 3D model of a plastic cup using a cheap off the shelf VL53L0X ToF scanner. In this latter work, a VL53L0X was located in a fixed position and a plastic cup rotated and moved vertically in front of it. By recording the position of the cup relative to the detector and building cloud point data, the team successfully generated a 3D model of the cup.

The research in the current manuscript adopts a similar process to Martin et al. (2017) to perform close scanning of simulated nuclear waste but using a robot arm as the mobile platform. We use the co-ordinates generated by a robotic arm, combined with ToF ranging at the front of the sensor to determine its position relative to the target objects. The ToF sensor is for exact stand-off distance measurement to permit accurate radiation dose conversions. Other standard techniques could be used to make a higher resolution 3D model, for example photogrammetry or 3D lidar, from a greater distance where dose is lower. This novel scanning work builds toward the end goal of more accurately scanning mixed wastes produced during decommissioning as well as existing packages for routine assay. The use of industrial robots, such as those made by KUKA, is well-understood and many are already in use on nuclear sites around the world. The high spatial precision and repeatability of KUKA manipulators means that they can be accurate at the sub-millimeter scale, making them an ideal candidate platform for high precision detector research.

2. EXPERIMENTAL SETUP

Robotic manipulation and instrument scanning must be well-synchronized to provide effective integrated measurements. For radiation scanning a Kromek^{TM} Sigma was utilized. The Sigma incorporates a Thallium doped Caesium Iodide (CsI(TI)) scintillator crystal, 25 × 25 × 50 mm (Kromek, 2019), contained inside a 1 mm thick aluminum casing. The maximum count rate of the Sigma is 5,000 counts per second (CPS) (Connor et al., 2018) recording gamma photons over a 50 keV to 2 MeV energy range (Kromek, 2019). It operates effectively at room temperature and does not require active cooling, unlike other more classical semiconductor detectors such as HPGe or Si (Cherry et al., 2012); this makes it an ideal choice for this application. A lead (Pb) collimator was designed to surround the detector, to reduce as far as possible the higher angle extraneous gamma counts incident on the detector. A square opening on the front face of the collimator enabled the counts to be received from a limited solid angle, perpendicular to the scan surface. The radiation detection software was programmed on a Raspberry Pi by ImiTec Ltd as part of their Remote Isotopic Analysis System (RIAS) and recorded the total number of counts received within a given exposure time. This was sent via a server communication to LabVIEW across 4096 energy bins which could request detector data at various time intervals (typically 1 or 10 Hz). The system was programmed to receive total counts and spectral data every 100 ms. Using the Robot Sensor Interface (RSI) software provided by KUKA robotics, a 6 figure co-ordinate detailing the position of the robot flange (end piece of the robot) was attained with a time stamp. The orientation of the end-flange was fixed, such that it was constantly parallel with the work surface. The two readings were time stamped and synchronized within the LabVIEW software. The LabVIEW code generated a CSV file which included the x-y position of each measurement location on the table, the distance to the point from the robot arm (from the ToF sensor) and the number of gamma counts collected. This CSV file was then processed by a python script to interpolate the data into a 3D radiation map. The data was approximated to a series of points, with 2 mm lateral spacings using a linear interpolation. These 2D points were given a third dimension by using the ToF readings. The ToF points were similarly interpolated and the resulting data is displayed in **Figures 1–4**, **5–7**. The data collection process is shown in a flowchart in **Figure 8**. A VL53L0X time of flight (ToF) sensor was used to collect point cloud data to visualize the data in 3D. The ToF sensor transmits infrared light to measure distance. The solid angle produced by the stock ToF sensor is around 45 degrees, consequently it was found to regularly return anomalous data in testing, in particular where surface topography changes were more drastic. Hence a 3 mm collimating ball lens was fitted to the VL53L0X for this project, to collimate the beam and increase the spatial resolution of the scan. The radiation response of the detector was predicted to be good, as it has a small interaction cross section of around 0.25 cm^2. The time of flight scanner takes numerous measurements of the scene whilst moving on the arm. It was mounted on the scanning head, remaining perpendicular to the table at all times. The detector was longer in one axis, giving it greater sensitivity in one dimension. This dimension was kept perpendicular to the scan table at all times to ensure a consistent radiation map was generated at the highest resolution possible. For each 100 ms exposure time, the average reading collected from the sensor was returned via the LabVIEW code.

The data collected by the arm contained a full spectrum of 4096 energy bins ranging from 50 keV to 2 MeV. This allows for the distinct Cs-137 peak at 662 keV to be identified. A python script was written to calculate the counts identified within the peak range and subtract from that the baseline reading to correct for background. A multiplication factor was applied to convert

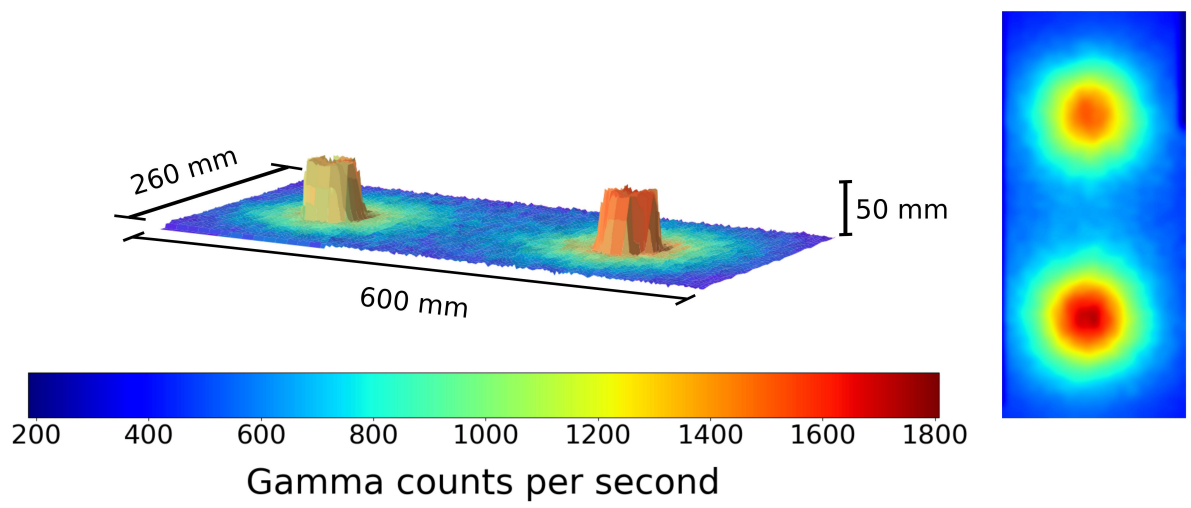

FIGURE 1 | A figure showing the 3D model and overlaid radiation map generated when the robot scan was completed at a 1 cm standoff above 2 Cs-137 sources 30 cm apart. Left to right the source activity is 7.5 and 10 μSvh^{-1}, respectively. The color represents the gamma radiation counts in counts per second for a given 2 × 2 mm pixel.

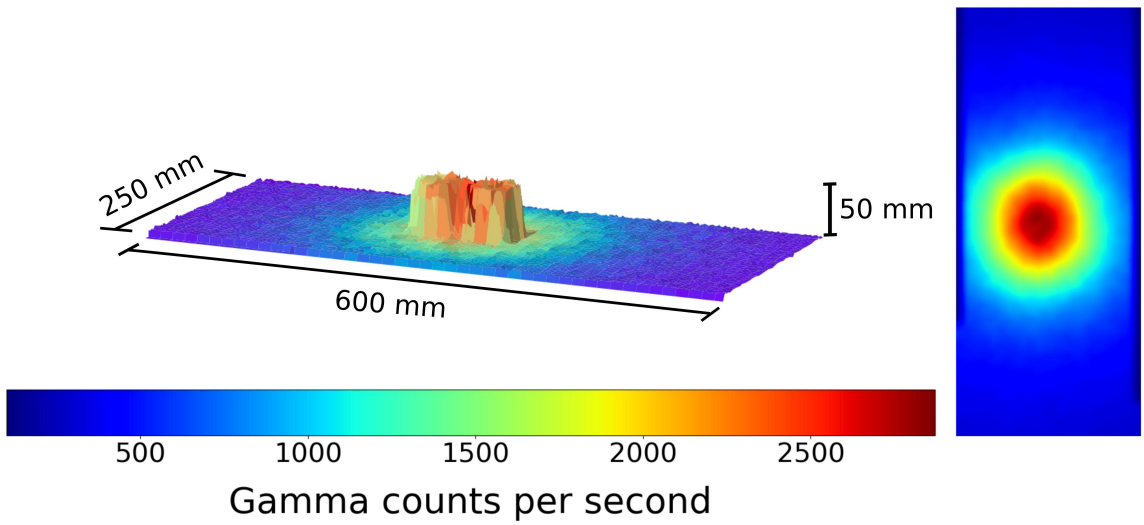

FIGURE 2 | A figure showing the 3D model and overlaid radiation map generated when the robot scan was completed at a 1 cm standoff above 2 Cs-137 sources directly adjacent. Left to right the source activity is 7.5 and 10 μSvh^{-1}, respectively. The color represents the gamma radiation counts in counts per second for a given 2 × 2 mm pixel.

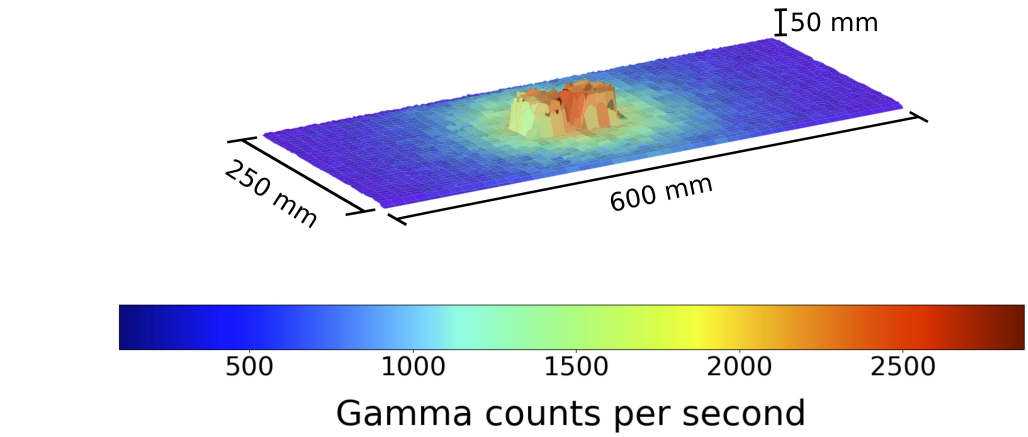

FIGURE 3 | A figure showing the 3D model and overlaid radiation map generated when the robot scan was completed at a 1 cm standoff above 2 Cs-137 sources directly adjacent, from a second angle to aid visual clarity of the distinguished pucks. Left to right the source activity is 7.5 and 10 μSvh^{-1} respectively. The color represents the gamma radiation counts in counts per second for a given 2 × 2 mm pixel.

the raw count value within the energy range to a dose rate in μSvh^{-1}, using the method described in Connor et al. (2020).

For scanning tests the robot was programmed to perform a basic raster pattern scan of the "scan surface", which was a 0.6 × 0.26 m area. The raster scan had a step length of 1 cm and the speed of the scan could be varied, depending on the activity of the test sources, less active sources require longer counting times (slower scans) to achieve adequate detection. Sealed radioactive sources containing caesium-137 (Cs-137), one of 7.5 μSvh^{-1} contact dose and the other at 10 μSvh^{-1} contact dose were used, alongside naturally occurring uranium (pitchblende) sealed sources of 4.3 and 4.5 μSvh^{-1} contact dose rate to test the radiation response of the system. The robot arm was set up to move at a consistent speed in a continuous linear motion of 10 mm per second, with a single scan taking approximately 30 min. Scan time could be reduced if higher activity sources were used, but as a proof of concept where timing is not restricted, greater scan time is able to yield a higher resolution radiation map, with a longer exposure for each collection interval. A photograph of the scanning system is shown in **Figure 9**.

3. RESULTS

To test the system, several different scanning scenarios were set up using the Cs-137 and Pitchblende sources available. The first used 2 Cs-137 sources separated apart by a distance of 30 cm,

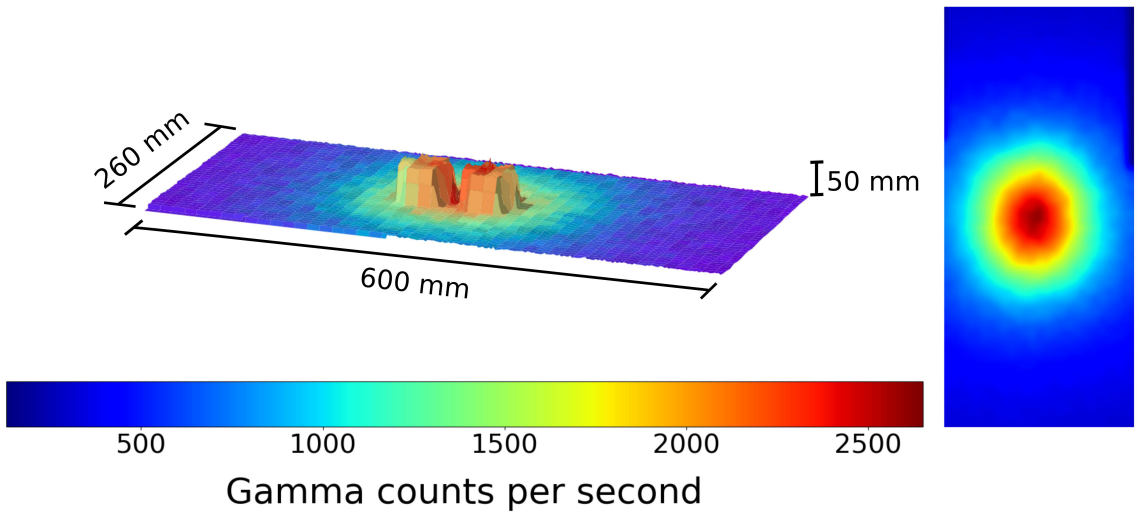

FIGURE 4 | A figure showing the 3D model and overlaid radiation map generated when the robot scan was completed at a 1 cm standoff above 2 Cs-137 sources 1 cm apart. Left to right the source activity is 7.5 and 10 μSvh^{-1}, respectively. The color represents the gamma radiation counts in counts per second for a given 2 × 2 mm pixel.

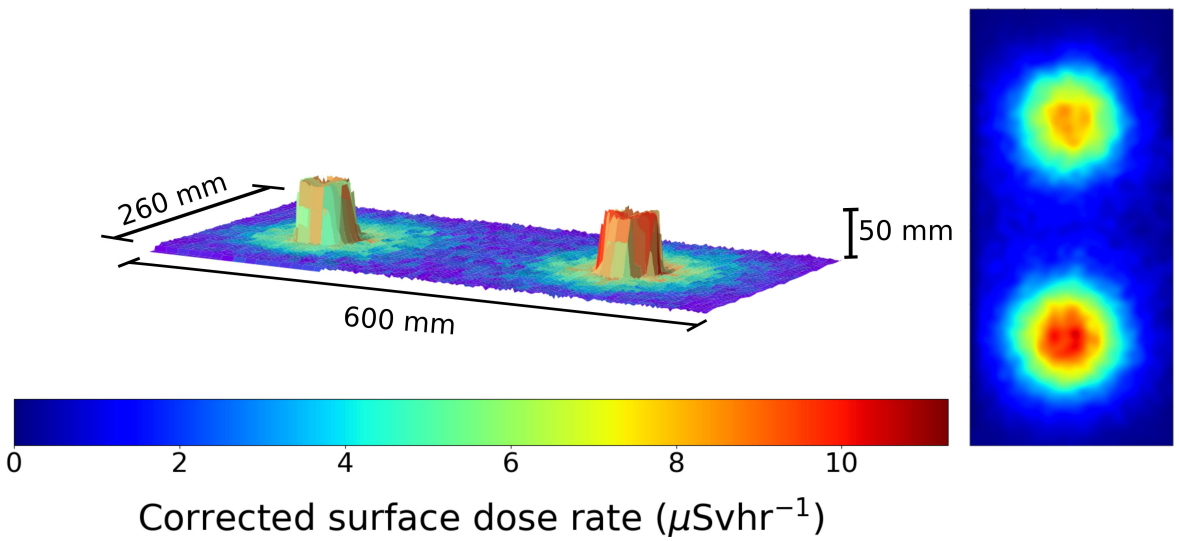

FIGURE 5 | A figure showing the 3D model and overlaid radiation map generated when the robot scan was completed at a 1 cm standoff above 2 Cs-137 sources 30 cm apart. Left to right the source activity is 7.5 and 10 μSvh^{-1}, respectively. The color represents the radiation dose rate in μSvh^{-1} for a given 2 × 2 mm pixel.

center to center. The scanning head was programmed to trace the raster scan path designed at a rate 10 mm per second. The resulting data can be seen in **Figure 1**.

This result demonstrates the capability of the 3D model generation using the ToF sensor, as it generates an identifiable geometric representation of the source pucks. The data is displayed with the counts recorded by the detector for a given 100 ms exposure. Following this experiment, the two sources were placed directly next to each other. This served as a test of the ToF mapping procedure. The resulting figure is shown in **Figure 2**.

From this we can clearly identify the radioactive hot spot. In addition the physical 3D separation of the sources can be comprehended. It is easier to distinguish on software which enables the rotation of the generated model. In order to aid the visual clarity in distinguishing the pucks, **Figure 3** presents the same data in **Figure 2**, but from a new observation angle.

To further monitor this 3D modeling routine, the sources were next placed 1 cm apart from each other. The result of this test is shown in **Figure 4**.

Here the centimeter gap is visible and both sources may be physically distinguished. This concept works and is sufficient

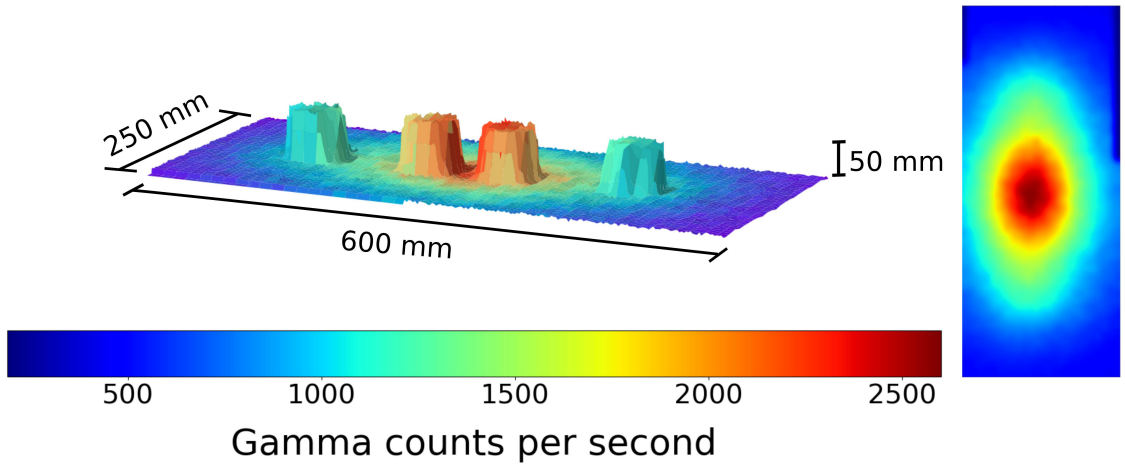

FIGURE 6 | A figure showing the 3D model and overlaid radiation map generated when the robot scan was completed at a 1 cm standoff above 2 Cs-137 sources 1 cm apart and 2 Pitchblende sources 30 cm apart. Left to right the source activity is 4.5 μSvh^{-1} Pitchblende, 7.5 μSvh^{-1} Cs-137, 10 μSvh^{-1} Cs-137 and 4.3 μSvh^{-1}, respectively. The color represents the gamma radiation counts in counts per second for a given 2 × 2 mm pixel.

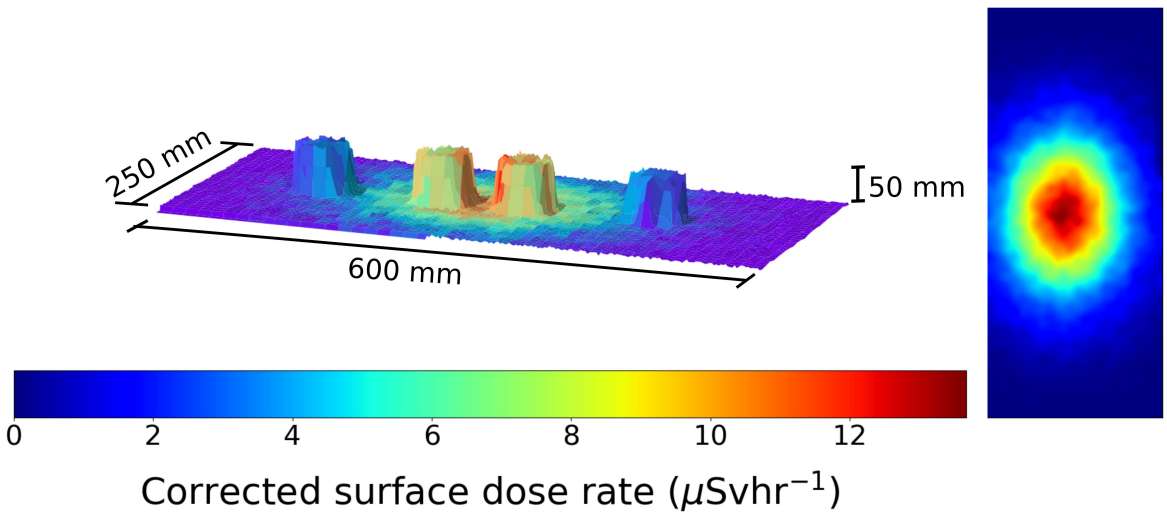

FIGURE 7 | A figure showing the 3D model and overlaid radiation map generated when the robot scan was completed at a 1 cm standoff above 2 Cs-137 sources 1 cm apart and 2 Pitchblende sources 30 cm apart. Left to right the source activity is 4.5 μSvh^{-1} Pitchblende, 7.5 μSvh^{-1} Cs-137, 10 μSvh^{-1} Cs-137, and 4.3 μSvh^{-1}, respectively. The color represents the radiation dose rate in μSvh^{-1} for a given 2 × 2 mm pixel.

for identifying radiation hotspots present in given scan and sort scenarios. However it does not provide any dose rate information. The data was subsequently processed to give an estimate of the dose rate of the given sources. The sources comprised of Cs-137 containing moss samples collected from Fukushima and were hence not perfect point-source emitters. Instead we approximated each source as a point emitter where the origin of the point was 3 cm beneath the puck surface, which corresponds to the thickness of the perspex and a small air gap, as shown in **Figure 10**.

For our calculations we assume that the perspex is gamma transparent due to its low density and Z number. Hence we may apply a correction factor to the dose rate method described in the above section, to display the data in surface dose rate format. The two Cs-137 sources positioned at 30 cm apart may be seen on a dose rate map in **Figure 5**.

The dose rates recorded at the 1 cm standoff are in good agreement with the actual recorded contact dose rates measured for each source: 7.5 and 10 μSvh^{-1}. This is very promising in demonstrating that the system sensitivity is sufficient to discriminate radioactive objects that classify at the Very Low Level Waste (VLLW) to Low Level Waste (LLW) threshold. In the UK VLLW is normally regarded as material with a specific activity up to 100 Bqg^{-1} (RWM, 2010).

An ability to measure the full gamma radiation spectrum enables different radioisotopes to be identified. **Figure 6** shows

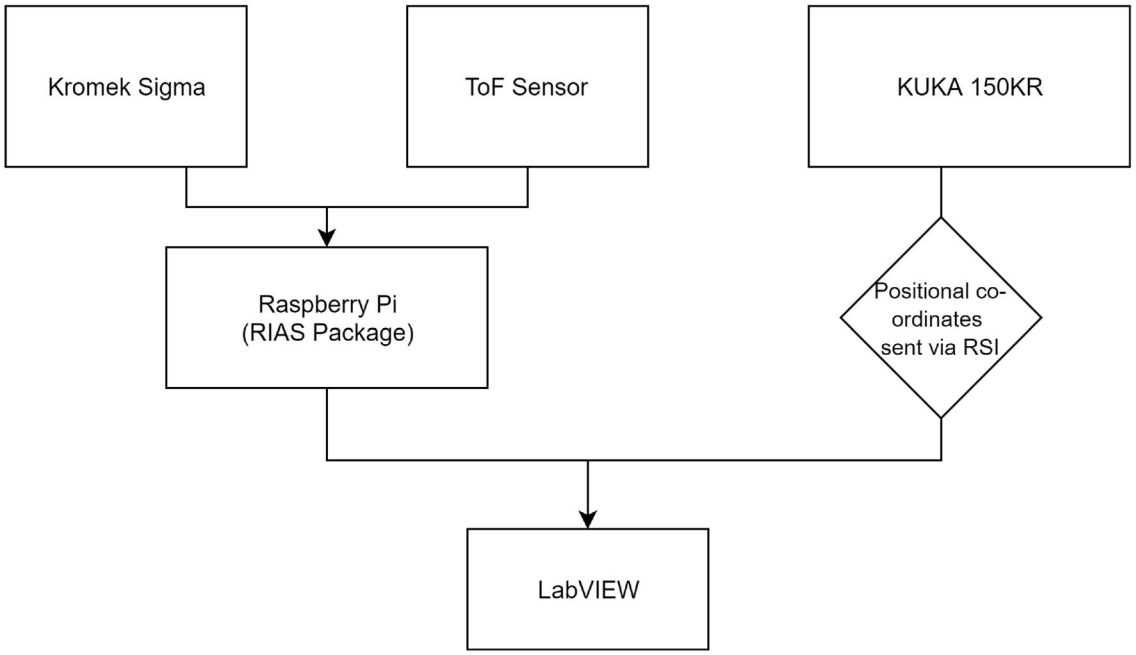

FIGURE 8 | A flowchart showing how the system was integrated.

the radiation intensity map of 2 NORM pitchblende sources and two Cs-137 sources.

One can determine they are all radioactive, with different intensities, but not tell which source is which. By restricting the spectral window to only the 662 keV gamma photons, one generates, as a Cs-137 specific plot, as in **Figure 7**.

Using spectral gating like this the Pitchblende sources become invisible to the radiation scanning system. This is key for the nuclear industry, as it would enable radionuclide characterization and separation to be carried out autonomously.

4. DISCUSSION

This paper introduces a novel integration of technologies that facilitates the scanning of radioactive materials and waste-forms, creating a 3D model of the object or environment and adding an overlaid radiation map. The setup is comprised of certified commercial off-the-shelf (COTS) components that in consequence require little control performance verification. Each component is integrated in a modular manner system, allowing for a highly flexible system design. The nature of the integrated system means that a multitude of sensor packages and grippers could be added alongside or in place of the combined detector unit used in this paper. Even the robotic arm used for this project could be replaced with an alternate choice as there are numerous different robotic arm systems commercially available of different sizes, lift capabilities, reaches and radiation tolerances that could enable a range of different scale applications, from sorting bulk rubble or pipework to sifting sediments for hot micro-particles.

Our experiments demonstrate that it is possible for such a scanning system to make very accurate, high sensitivity, high spatial resolution radiation maps for resolving nuclear waste materials from each other on the basis of emitted gamma intensity. The result was also successful from a 3D modeling perspective, as it clearly identifies the sources as separate objects despite their close proximity. The spherically symmetric radiative flux which is emitted from the radioactive sources means that there are limitations on identifying which physical shape corresponds to which emission on the radiation map. This is something that could be improved by an algebraic reconstruction technique and a comprehensive understanding of the detectors response. This forms an important part of the future work this project will require. The radiation sources used in this work were relatively weak compared to real ILW and LLW. For real waste scenarios it would be expected that a smaller micro gamma spectrometer with greater peak dose measurement capability e.g., CZT or GaAs, could be utilized. Scan times would also be dramatically reduced with increasing radioactivity levels.

The next proposed step in development is to use the laser profiling to generate a volume for the object and then based on an assumed density, e.g., for concrete, the system could be developed to automatically threshold objects and assign them as either VLLW, LLW, or ILW based on their radioactivity and calculated mass. This thresholding programme would need to utilize an algorithm based on the inverse square law to calculate the intensity of the emitted radiation at the surface of the objects being scanned, accounting for the efficiency and solid angle of the radiation detector being used. It is crucial to get an accurate distance from the detector to the target object

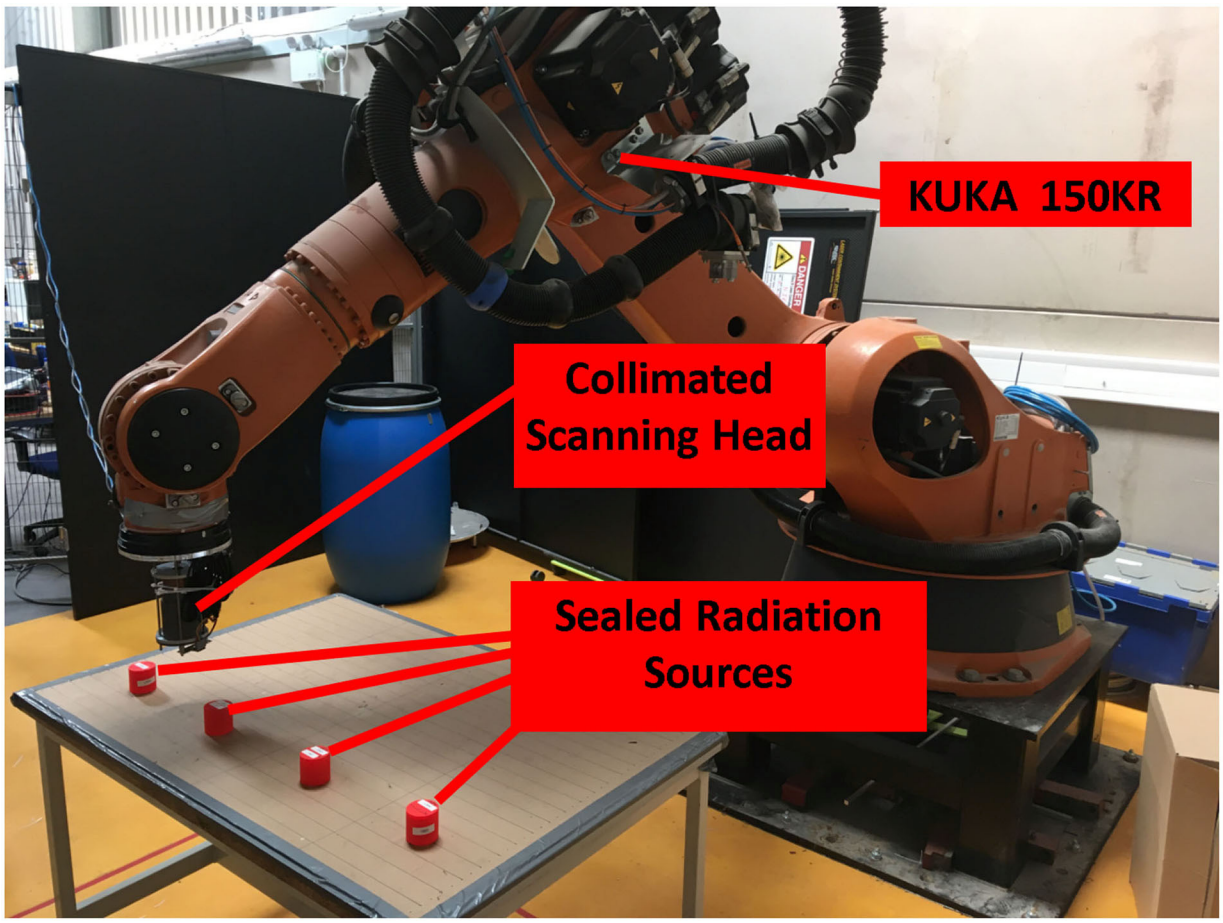

FIGURE 9 | A photograph of the scanning system in action.

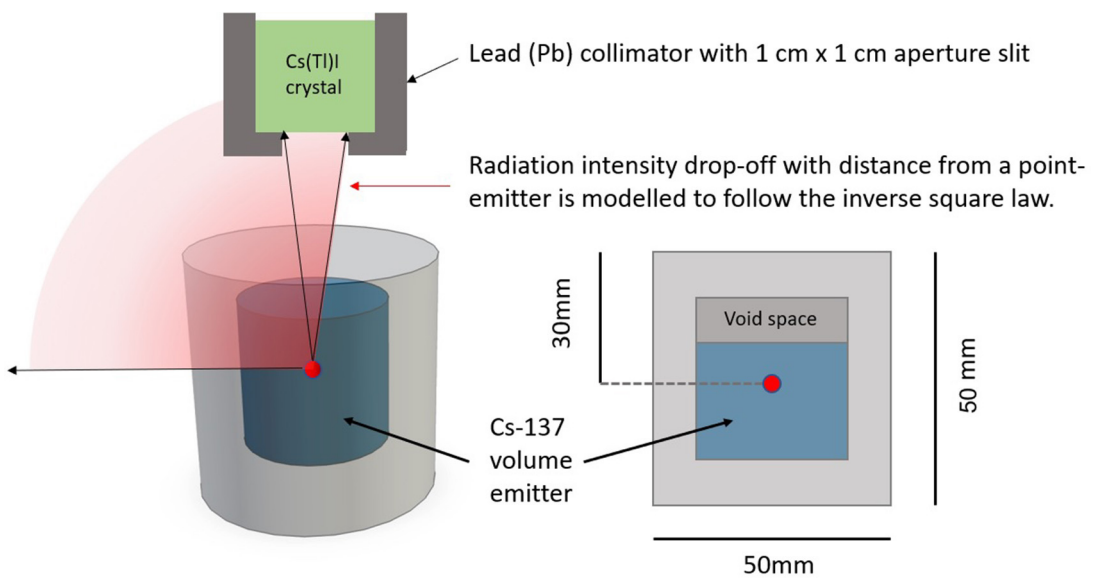

FIGURE 10 | A diagram explaining the point source modeling used to invoke the inverse square law.

because contact dose calculations are based on the inverse square law. This means any discrepancy in reading will significantly affect the corrected dose calculation. Ultimately, the full gamma spectrometry capability of the system could be used to distinguish different gamma emitters too. This would add a further level of finesse for separating mixed nuclear wastes that might be expected to arise during nuclear decommissioning activities. The accuracy of the 3D model would facilitate a robotic manipulator for grasping in addition, fulfilling the requirements of the sort and segregation table.

The system's scanning methodology could also be substantially refined vs. the current simple raster scanning we have demonstrated. For example, an initial survey scan could conduct rapid raster scan of a scene to determine the degree of variability in scene topology and from that, determine automatically a more detailed scanning path that would maintain a safe but close scanning proximity to the waste objects. The initial survey scan would also locate any strong radiation emitters, which the adaptive path plan could deferentially focus on to provide a more detailed scan of that specific area of the scene.

5. CONCLUSION

This work demonstrates the use of a combined laser profiling and gamma-scanning sensor unit, mounted on a robot arm, to form an accurate 3D profile of a series of test objects on a scanning table, with a coincident overlay of the mapped radiation intensity. Radiation maps are successfully created by the system, which is able to correctly identify radioactive sources of different intensities on a flat scan surface. The generated 3D surface model reveals an accurate visualization of the tested scene and is accurate to within a centimeter. In addition to this an estimate of the surface dose rate produced by the radioactive emitters is made based on the scan data received to a good level of accuracy, correctly identifying the dose rate of two radioactive Cs sources to within 1 μSvh^{-1}.

DATA AVAILABILITY STATEMENT

The raw (unprocessed) data that supports the findings from this study are available from Mendeley Data with the link: doi: 10.17632/gts8x9f84r.1.

AUTHOR CONTRIBUTIONS

All authors contributed technically, both in the data collection and writing of this paper.

FUNDING

We would like to thank UK Research and Innovation (UKRI) and Engineering and Physical Sciences Research Council (EPSRC) for their support and funding which has come via the National Centre for Nuclear Robotics (NCNR); grant EP/R02572X/1.

ACKNOWLEDGMENTS

The authors would like to thank John Jowsey and Alex Jenkins from Sellafield Ltd., for their valuable technical input. We would also like to thank KUKA AG., for providing technical support for the project and Jim Brooke, DRPS for the School of Physics for his assistance with the radiation sources used in this experiment.

REFERENCES

Barone, S., Neri, P., Paoli, A., and Razionale, A. (2020). 3D acquisition and stereo-camera calibration by active devices: A unique structured light encoding framework. *Opt. Lasers Eng.* 127:105989. doi: 10.1016/j.optlaseng.2019.105989

Bird, B., Griffiths, A., Martin, H., Codres, E., Jones, J., Stancu, A., et al. (2019). A robot to monitor nuclear facilities: using autonomous radiation-monitoring assistance to reduce risk and cost. *IEEE Robot. Automat. Mag.* 26, 35–43. doi: 10.1109/MRA.2018.2879755

Cherry, S. R., Sorenson, J. A., and Phelps, M. E. (2012). "Chapter 7: Radiation detectors," in *Physics in Nuclear Medicine, 4th Edn*, eds S. R. Cherry, J. A. Sorenson, and M. E. Phelps (Philadelphia, PA: W.B. Saunders), 87–106. doi: 10.1016/B978-1-4160-5198-5.00007-1

Connor, D., Martin, P., and Scott, T. (2016). Airborne radiation mapping: overview and application of current and future aerial systems. *Int. J. Remote Sens.* 37, 5953–5987. doi: 10.1080/01431161.2016.1252474

Connor, D., Martin, P., Smith, N., Payne, L., Hutson, C., Payton, O., et al. (2018). Application of airborne photogrammetry for the visualisation and assessment of contamination migration arising from a Fukushima waste storage facility. *Environ. Pollut.* 234, 610–619. doi: 10.1016/j.envpol.2017.10.098

Connor, D., Wood, K., Martin, P., Goren, S., Megson-Smith, D., Verbelen, Y., et al. (2020). Radiological mapping of post-disaster nuclear environments using fixed-wing unmanned aerial systems: a study from chernobyl. *Front. Robot. AI* 6:149. doi: 10.3389/frobt.2019.00149

Gov, U. (2020). *Sort and Segregate Nuclear Waste: Phase 1 (SBRI Competition)*. Available Online at: https://apply-for-innovation-funding.service.gov.uk/competition/685/overview#supporting-information (accessed August 14, 2020).

Gutierrez-Villalobos, J. M., Dimas, T., and Mora-Vazquez, J. C. (2017). "Simple and low cost scanner 3D system based on a time-of-flight ranging sensor," in *2017 XIII International Engineering Congress (CONIIN)* (Santiago de Queretaro), 1–5. doi: 10.1109/CONIIN.2017.7968193

Hoegg, T., Lefloch, D., and Kolb, A. (2013). Time-of-flight camera based 3d point cloud reconstruction of a car. *Comput. Indus.* 64, 1099–1114. doi: 10.1016/j.compind.2013.06.002

Horizon2020 (2018). *Robotic Manipulation for Nuclear Sort and Segregation*. Available Online at: https://cordis.europa.eu/project/id/645582 (accessed August 04, 2020).

IAEA (2018). *Status and Trends in Spent Fuel and Radioactive Waste Management*. Available online at: https://www.iaea.org/publications/11173/status-and-trends-in-spent-fuel-and-radioactive-waste-management (accessed May 24, 2019).

Kromek (2019). *SIGMA Compact CsI Scintillation Spectrometer Specification Sheet*. Available online at: https://www.kromek.com/product/sigma-scintillator-detectors/ (accessed September 06, 2018).

Martin, P., Kwong, S., Smith, N., Yamashiki, Y., Payton, O., Russell-Pavier, F., et al. (2016a). 3D unmanned aerial vehicle radiation mapping for assessing contaminant distribution and mobility. *Int. J. Appl. Earth Observ. Geoinform.* 52, 12–19. doi: 10.1016/j.jag.2016.05.007

Martin, P., Moore, J., Fardoulis, J., Payton, O., and Scott, T. (2016b). Radiological assessment on interest areas on the sellafield nuclear site via unmanned aerial vehicle. *Remote Sens.* 8:913. doi: 10.3390/rs81 10913

Martin, P., Scott, T., Payton, O., and Fardoulis, J. (2017). High-resolution aerial radiation mapping for nuclear decontamination and decommissioning. *WMS J.* 2.

Marturi, N., Kopicki, M., Rastegarpanah, A., Rajasekaran, V., Adjigble, M., Stolkin, R., et al. (2018). Dynamic grasp and trajectory planning for moving objects. *Auton. Robots* 43, 1241–1256. doi: 10.1007/s10514-018-9799-1

Redus, R., Squillante, M., Gordon, J., Knoll, G., and Wehe, D. (1994). A combined video and gamma ray imaging system for robots in nuclear environments. *Nuclear Instrum. Methods Phys. Res. A* 353, 324–327. doi: 10.1016/0168-9002(94)91667-5

RWM (2010). *Radioactive Wastes in the UK: A Summary of the 2016 Inventory*. Available online at: http://ukinventory.nda.gov.uk/wp-content/uploads/sites/18/2017/03/High-Level-Summary-UK-Radwaste-Inventory-2016.pdf (accessed July 06, 2019).

Sarker, M., da Costa, D. D., and Hadigheh, S. A. (2019). Multi-scale 3D roughness quantification of concrete interfaces and pavement surfaces with a single-camera set-up. *Construct. Build. Mater.* 222, 511–521. doi: 10.1016/j.conbuildmat.2019.06.157

Tsitsimpelis, I., Taylor, C. J., Lennox, B., and Joyce, M. J. (2019). A review of ground-based robotic systems for the characterization of nuclear environments. *Prog. Nuclear Energy* 111, 109–124. doi: 10.1016/j.pnucene.2018.10.023

UK Government Department of Energy & Climate Change (2011). *Waste Transfer Pricing Methodology for the Disposal of Higher Activity Waste From New Nuclear Power Stations*. Available online at: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/42629/3798-waste-transfer-pricing-methodology.pdf

Conflict of Interest: JD was employed by the company KUKA Systems UK Ltd.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

frontiers in Robotics and AI

ORIGINAL RESEARCH
published: 02 December 2020
doi: 10.3389/frobt.2020.513004

Automatic Fracture Characterization Using Tactile and Proximity Optical Sensing

Francesca Palermo[1]*, *Jelizaveta Konstantinova*[1,2], *Kaspar Althoefer*[1,3], *Stefan Poslad*[1] *and Ildar Farkhatdinov*[1,3,4]*

[1] School of Electronic Engineering and Computer Science, Queen Mary University of London, London, United Kingdom, [2] Robotics Research, Ocado Technology, London, United Kingdom, [3] The Alan Turing Institute, Programme – Artificial Intelligence, London, United Kingdom, [4] Department of Bioengineering, Imperial College of Science, Technology and Medicine, London, United Kingdom

This paper demonstrates how tactile and proximity sensing can be used to perform automatic mechanical fractures detection (surface cracks). For this purpose, a custom-designed integrated tactile and proximity sensor has been implemented. With the help of fiber optics, the sensor measures the deformation of its body, when interacting with the physical environment, and the distance to the environment's objects. This sensor slides across different surfaces and records data which are then analyzed to detect and classify fractures and other mechanical features. The proposed method implements machine learning techniques (handcrafted features, and state of the art classification algorithms). An average crack detection accuracy of ~94% and width classification accuracy of ~80% is achieved. Kruskal-Wallis results ($p < 0.001$) indicate statistically significant differences among results obtained when analysing only integrated deformation measurements, only proximity measurements and both deformation and proximity data. A real-time classification method has been implemented for online classification of explored surfaces. In contrast to previous techniques, which mainly rely on visual modality, the proposed approach based on optical fibers might be more suitable for operation in extreme environments (such as nuclear facilities) where radiation may damage electronic components of commonly employed sensing devices, such as standard force sensors based on strain gauges and video cameras.

Keywords: sensing, haptic exploration, crack recognition, extreme environment, optical sensing, fiber-optics

OPEN ACCESS

Edited by:
Kushal Mukherjee, United Technologies Research Center, Ireland

Reviewed by:
Zhanat Kappassov, Nazarbayev University, Kazakhstan
Mohsen Kaboli, Radboud University Nijmegen, Netherlands

***Correspondence:**
Francesca Palermo
f.palermo@qmul.ac.uk
Ildar Farkhatdinov
i.farkhatdinov@qmul.ac.uk

Specialty section:
This article was submitted to Sensor Fusion and Machine Perception, a section of the journal Frontiers in Robotics and AI

Received: *26 November 2019*
Accepted: *19 October 2020*
Published: *02 December 2020*

Citation:
Palermo F, Konstantinova J, Althoefer K, Poslad S and Farkhatdinov I (2020) Automatic Fracture Characterization Using Tactile and Proximity Optical Sensing. Front. Robot. AI 7:513004. doi: 10.3389/frobt.2020.513004

1. INTRODUCTION

An important task often performed in remote hazardous environments is the detection of mechanical fractures on the objects, such as containers, tanks, pipes, and other technical systems used for keeping chemical and radioactive waste. A crack may be caused by physical damage or material degradation over time or environment changes (e.g., temperature or pressure). The effects of non-detected fractures may lead to larger macro-scale catastrophic failures making the cracked surface mechanically weak to perform its function.

Conventional automatic crack detection methods applied in industry to inspect large mechanical structures rely on acoustic methods (Chakraborty et al., 2019), use X-ray scanning (Barhli et al., 2017; Naragani et al., 2017), apply eddy currents techniques (Yao et al., 2014), or explore changes in a system's motion dynamics (Lu and Chu, 2011; Nicoletti et al., 2018). Such techniques require

specialized and costly equipment and well-trained technical staff making their usage in extreme environments (i.e., decommissioning of radioactive waste) less beneficial or even impossible.

Rapid development of computer vision and machine learning led to the introduction of multiple vision-based tools for mechanical fracture detection that we briefly review below. Chen and Jahanshahi (2017) proposes a fusion between a convolutional neural network and a Naive Bayes to analyse video frames for crack detection in nuclear reactors. The framework achieves a 98.3% hit rate against 0.1 false positives per frame. Schmugge et al. (2016) suggested a crack detection method for nuclear power plant inspection videos by fine-tuning a deep neural network for detecting local patches containing cracks which are then grouped in spatial-temporal space for group-level classification which obtains an increase of 40% in the F1-Score with respect to the compared methods. Iliopoulos et al. (2015) analyzed the evolution of a cracked concrete structure obtained by applying Digital Image Correlation, Acoustic Emission, and Ultrasonic Pulse Velocity techniques. The results highlight the time of onset and location that the crack started to form as well as the width and depth of the cracks.

Vision based methods demonstrate high detection accuracy and they are easy to implement in telerobotics applications as cameras are essential parts of the remote inspection robots. However, vision-based methods can fail in remote environments with limited luminosity and video-cameras cannot operate in presence of strong radiation. Furthermore, vision-based methods are not capable of acquiring material properties, such as texture and hardness.

Our work proposes to use tactile and proximity sensing for mechanical cracks detection. In contrast to the visual modality, tactile, and proximity sensing can provide important information on material properties, such as shape, texture, and hardness (Huet et al., 2017; Yuan et al., 2017; Kaboli and Cheng, 2018). Tactile sensors were efficiently used to characterize different materials in robotic teleoperation. Liu et al. (2012, 2015) implemented a 6-axis force/torque finger-shaped sensor capable of estimating the instantaneous friction force and normal force to recognize physical properties of the surface of unknown objects. Average classification accuracy of 88.5% is obtained when implementing a naïve Bayes classifier on 12 different texture surfaces. Feng et al. (2018) proposed a new method, called Active Prior Tactile Knowledge Transfer (APTKT) to re-implement tactile knowledge of previously explored objects which improves the discrimination accuracy by over 20%. A multi-modal tactile sensor (BioTac, developed by SynTouch[1]) was used by Wong et al. (2014) to estimate the order of curvature and footprint dimensions explored with various movements (distal-proximal stroke, radial-ulnar stroke, etc.) of the robotic finger. Fishel and Loeb (2012) proposed a Bayesian exploration which selects the optimal movements based on previous experience to recognize 117 different textures. Kaboli et al. (2016) propose an online tactile transfer learning method to re-use previously learned tactile models to discriminate new textures with limited numbers of training samples. An expanded tactile sensors module was implemented for recognizing the alphanumeric characters inscribed on rubber stamps in Lee et al. (2006). The stiffness of objects was investigated by Konstantinova et al. (2017) implementing a hybrid force and proximity finger-shaped sensor achieving 87% classification accuracy on a set of household objects with different stiffness values. Drimus et al. (2014) proposed a method to classify objects into rigid and deformable using dynamic time warping to compare the distance between time series of signals. An optical sensor was implemented by Huang et al. (2018) to detect target objects in dynamic environments prior to contact allowing the teleoperator to feel the object without an actual contact improving the benefits of touch interaction to the operator, without negative consequences of the robot contacting unknown geometric structures. Tomo et al. (2017, 2018) introduced uSkin, a soft-skin based sensor, which measures the applied force based on changes in the magnetic field for object shape recognition. Not many approaches use tactile sensing for crack detection and characterization. For additional research on tactile sensing and texture recognition, please refer to Kappassov et al. (2015) and Luo et al. (2017).

In this work, we propose a novel tactile sensing-based technique for mechanical fractures detection with the potential application to nuclear-decommissioning tasks performed by remotely operated robots. The nuclear power industry has been among the slowest to adopt advanced technologies (Wood, 2004; Bogue, 2011). Any instrumentation to be used in the nuclear environment must show robustness under the influence of nuclear radiation, match safety requirements and satisfy the highest industrial standards. The effects of radiation greatly vary and depend on several parameters, including the type of radiation and the total dose (Bogue, 2013). Our approach relies on optical fibers for data transmission from the sensor's measurement elements to the remotely located electronic unit. Optical fibers are among the devices that are less influenced in a nuclear environment since gamma radiation does not interfere with their basic sensing mechanism (Berghmans et al., 1999; Inaudi et al., 2001; Phéron et al., 2012). Berghmans and Decreton (1998) compared the gamma radiation response of three types of optical fiber temperature sensors. For the three sensor types, the transducer mechanism does not seem to be affected by gamma radiation. Fiber optic cables are expected to see greater use in the nuclear power industry, replacing electrical cables (Berthold, 1994; Hashemian, 2009). Several applications implementing fiber optic cables are already been realized. Kim et al. (2017) developed a fiber-optic based monitoring system for water temperature, water level and radiation level of spent nuclear fuel pool (SNFP) at a nuclear power plant. The performance test results show that individual sensors can measure the changes in real-time. Ball et al. (2012) described several measurement technologies with potential application to gas reactors. Among these, an optical-based pressure sensor based on the trajectory of the light in

[1]https://www.syntouchinc.com/en/sensor-technology/

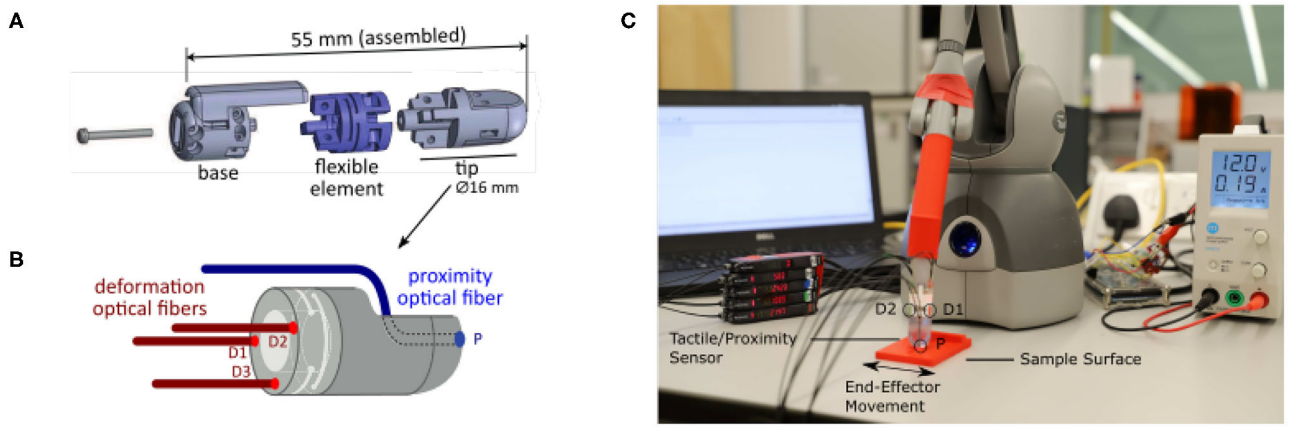

FIGURE 1 | Hybrid Fiber Optical Force/Proximity fingertip sensor: **(A)** Visualization of the three different components of the sensor; **(B)** Close up visualization of the fiber optics operating principles. D1, D2, D3 indicates the three deformation optical fibers. P the proximity optical fiber. **(C)** The complete setup for the data acquisition. From left to right: Laptop, Keyence sensors, Touch Haptic device with 3d printed end effector, and Hybrid Fiber Optical Force/Proximity Sensor, Arduino board, and Power Supply.

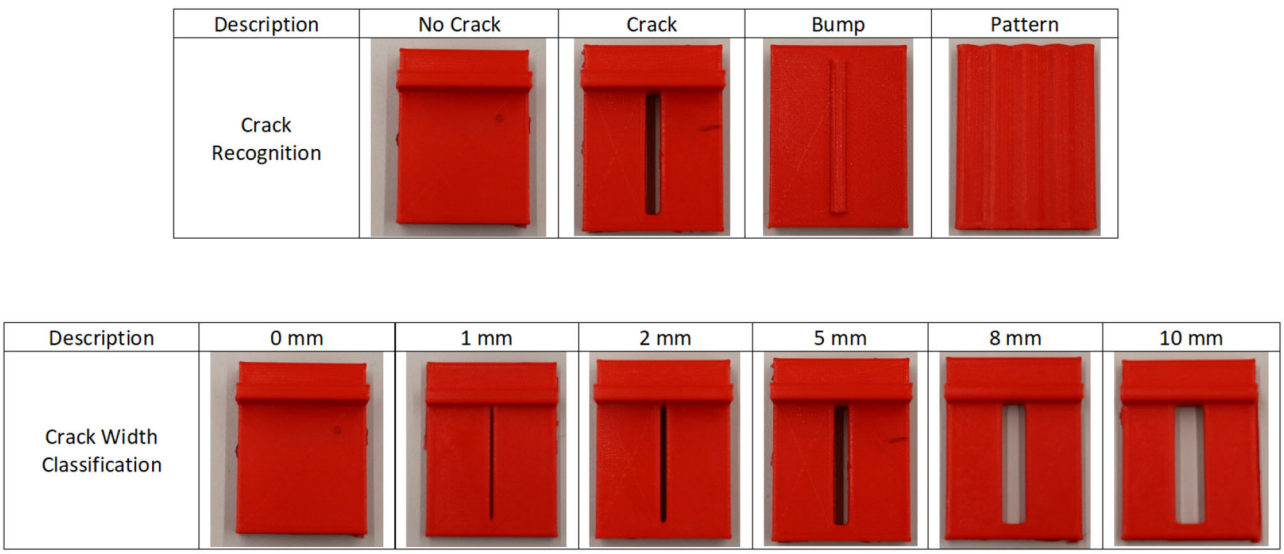

FIGURE 2 | Visualization of the set of objects explored during the experiments. The set for the Crack Recognition Analysis is formed by no crack, crack, bump, and wavy pattern surfaces. The series for the Crack Width classification experiment is made up of the same fractured surface with distinct widths of 0, 1, 2, 5, 8, and 10 mm.

glass is analyzed. The polarization of light crossing the glass is created through stress-induced in the glass as a result of pressure. Through the fiber optic sensor, the pressure measurement can be found through the polarized light intensity.

Present work demonstrates how tactile and proximity sensing can be efficiently used to perform automatic crack detection. The proposed method uses machine learning techniques to detect cracks and bumps based on the deformation and proximity signals which are recorded during physical interaction between a custom-designed robotic finger and the remote environment, Konstantinova et al. (2017). In case a crack is detected, the proposed automated technique classifies its width. Both offline and online classifications are performed. A fiber optic sensor has been implemented for data acquisition because of the reduced dimensions (~55 mm), weight (~200 g), low cost, the strong immunity to electromagnetic interference and the improved environmental resistance. This approach may be implemented also in extreme environments (e.g., in nuclear plants), since gamma radiation does not interfere with the basic sensing mechanism of fiber optic-based sensors (Berghmans et al., 1999). In addition, the nylon component of the implemented sensor can be used in irradiation conditions with limitations as Morita and Seguchi (1983) presented. To the best of the authors'

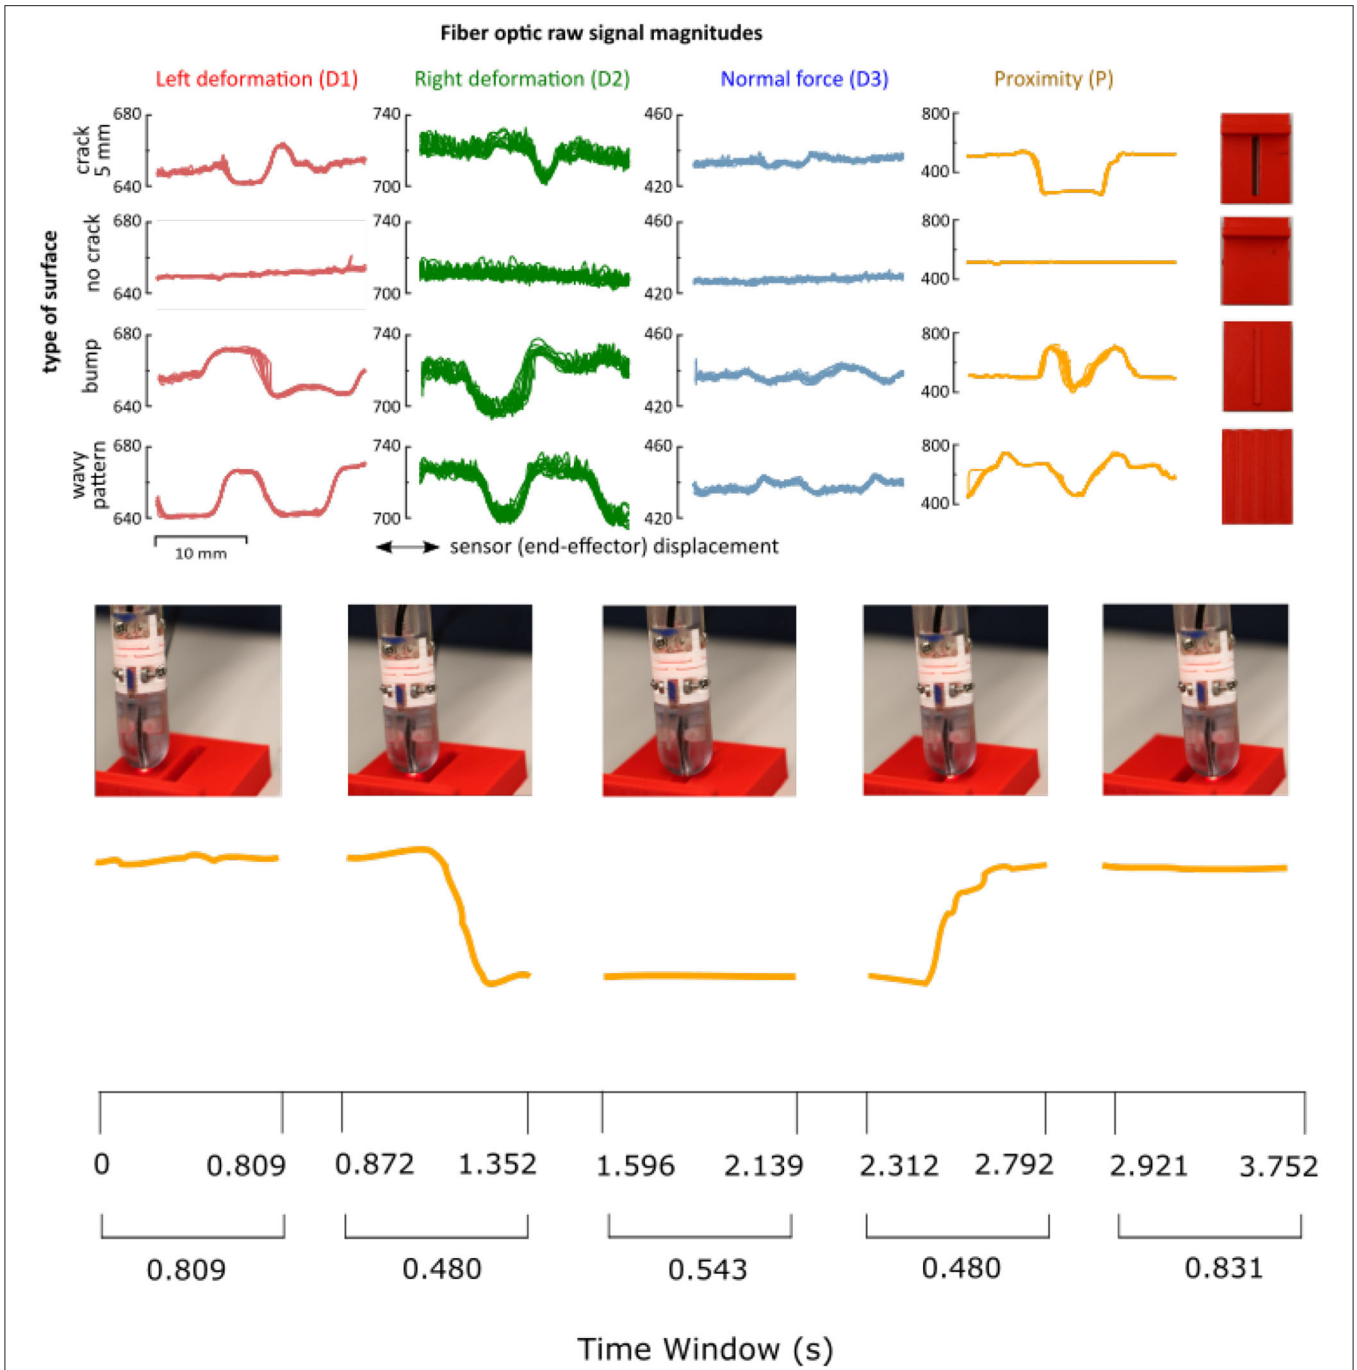

FIGURE 3 | (Top) Raw measurements from the four sensing elements of the sensor (deformations D1, D2, D3, and proximity P) for the set of surface patterns: "no crack," "crack," "bump," and "wavy pattern." Each column shows a different optical fiber signal. In red, the results of the left deformation (D1). In green, the data of the right deformation (D2). In blue, the normal deformation (D3). In yellow, the proximity data (P). **(Bottom)** The movement of the sensor is shown together with the corresponding proximity data.

knowledge, this is one of the first works on fracture recognition based on hybrid fiber optical force/proximity sensors. The Present work is based on our previous results (Palermo et al., 2020) demonstrating the feasibility of a tactile sensor for cracks detection. The novelty of this work is the implementation of more accurate mechanical fracture detection and classification methods, and a corresponding comparative study. Additionally, this paper provides a detailed description of the tactile data collection, processing and real-time classification implementation.

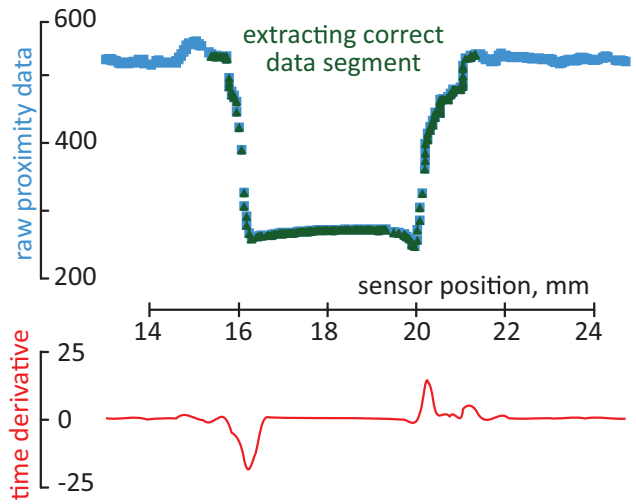

FIGURE 4 | Example of correct segment extraction. The complete measurement trial is displayed in blue. In green, the data points corresponding to the changes in the sensor-sample interaction mechanics. In red discrete-time derivative for proximity measurement.

2. EXPERIMENTAL METHODOLOGY

2.1. Tactile and Proximity Sensor

In this work, the integrated force and proximity finger-shaped sensor described by Konstantinova et al. (2017) is used. The sensor is made of 3D printed rigid (VeroClear Glossy) and soft (Nylon—PA2200) components allowing it to bend during interaction with the environment, as shown in **Figures 1A,B**. All the components are printed with an SLS printer EOS P100. The sensor employs three pairs of optical fiber cables (D1, D2, D3) to measure the deformation of the flexible middle part based on the changes of the reflected light intensity. The fourth pair of optical fiber cables (P) is used to measure proximity, i.e., the distance between the tip of the finger sensor and nearby objects. The implemented proximity permits the shape recreation in 2D of the explored surface. The sensor is capable of measuring bending torque and normal contact force during physical interaction with the environment. As described in Konstantinova et al. (2017), the implemented sensor is able to detect three-axis force/torque signals and measure the distance to the explored object. The sensor measures normal force up to 4.5 N. The lateral torque values (around the x- and y-axes) reach a maximum of ±18 N/mm. The usage of nylon to print the flexible structure led to low hysteresis and high robustness. The proximity sensor (P) can measure distances up to 30 mm. The calibration method has been described in Konstantinova et al. (2016). Each pair of the sensor's fiber optic cables is attached to a Keyence FS-N11MN light-to-voltage transducer. Thus, the change of light intensity modulation is measured and, using a calibration matrix, converted to force, torque, and distance measurements.

2.2. Experimental Setup

To collect data and test the proposed crack detection algorithm, the tactile and proximity sensor, described in section 2.1, has been attached to the end-effector of a Touch desktop haptic interface (formerly known as Phantom Omni Geomagic) as shown in **Figure 1C**. The Phantom Omni was programmed to slide the tactile sensor along a static sample surface following a programmed periodic movement. Data from tactile and proximity sensors were recorded through an Arduino Mega ADK micro-controller, connected via a USB port, at 400 Hz. These data were later synchronized with the absolute position of the tip of the tactile sensor calculated through the encoder readings of the Phantom Omni. Data acquisition and control were implemented through dedicated software libraries (OpenHaptics and Robotic Operating System) running on an Ubuntu desktop computer. The material samples, as well as the Phantom Omni interface, were fixed to a laboratory desk to minimize any vibration and unwanted displacements.

2.3. Data Acquisition Protocol

In this work, machine learning techniques are employed for crack detection and crack width classification. A set of 10 objects with different surfaces (no crack, cracks of different widths, a bump and a wavy pattern) were manufactured with PLA plastic using 3D printing technology (Ultimaker III, 0.2 mm layer height, 0.4 mm nozzle diameter). The wavy pattern consists of a repeated pattern of sine waves of 1mm amplitude and 5mm magnitude. The samples are shown in **Figure 2**. The types of these sample objects correspond to the classes implemented for training and testing the classifier. The Phantom Omni moved the tactile sensor across the sample objects: the periodic sliding has a magnitude of 1.6 cm and a frequency of 1,000 Hz. The average sliding velocity was 3.89 mm/s. The initial position of the tactile sensors was not controlled and varied from trial to trial and was set at ∼5–10 mm from the crack edge. No normal force was applied by the sensor to the sampled surfaces except the force caused by the sensor weight (∼200 g). Tactile and proximity signals were recorded for 12 repeated continuous sliding movements. This continuous recording was repeated five times. **Figure 3** shows an example of raw data acquired on "no crack," "crack," "bump," and "wavy pattern" for a continuous recording. For brevity, only the data acquired during sliding on different surfaces are shown.

2.4. Experimental Dataset

The data-set generated and used in this study is publicly available on the figshare repository "Automatic Fracture Database"[2]. The data is organized in a nine column format, corresponding to the following chronological measurements: sensor displacement, sensing elements signals (D1, D2, D3, P), the identification number of the current experiment, the number of the measurement trials (single sensors movement), the type of the surface explored (0 = no crack, 1 = crack, 2 = bump, and 3 = wavy pattern), and the direction of movement (0 = right, 1 = left).

[2]https://figshare.com/s/14deb00d874400e34d67

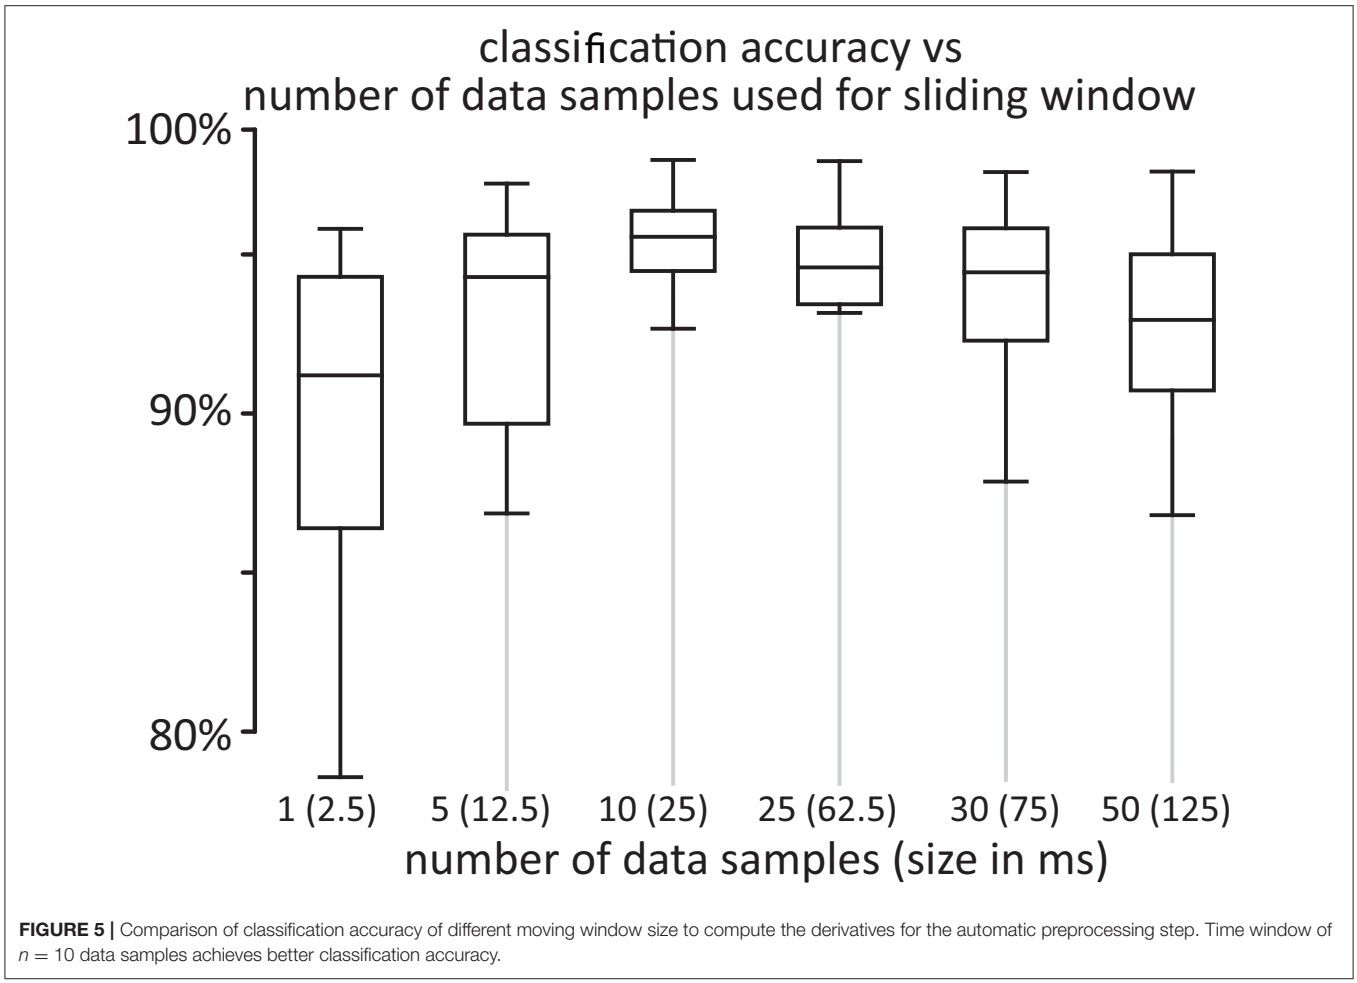

FIGURE 5 | Comparison of classification accuracy of different moving window size to compute the derivatives for the automatic preprocessing step. Time window of $n = 10$ data samples achieves better classification accuracy.

3. DATA ANALYSIS FOR CRACK DETECTION

The goal of the proposed algorithm is to detect and characterize mechanical fractures, such as cracks, based on the deformation and proximity data recorded from the sensors. The time history of the deformation and proximity data is recorded. Preprocessing step and features extraction are performed. The resultant output is used as an input for the classification algorithm.

3.1. Pre-processing

The goal of preprocessing was to prepare datasets containing information for the mechanical features, such as crack, bump, and wavy pattern and exclude not relevant datapoints (i.e., recording of the sensors sliding on a flat surface before and after interacting with the bump or crack). The preprocessed labelled datasets were then used for training process and cross-validation test. The preprocessing was fully automatic and was performed on the data collected from the haptic manipulator (sensor displacement) and optical sensing elements integrated in the sensor (deformations D1, D2, D3 and proximity P). The sensor's position data (obtained from the Geomagic haptic device) and the sensing elements data (D1–D3, P) were synchronized and sampled at 400 Hz. Prior the prepossessing stage the measurements of each trial (single sensor movement along the explored surface) were arranged in the following matrix:

$$\underset{m\times 5}{\mathbf{M}} = [\mathbf{x}^\top \ \mathbf{d_1}^\top \ \mathbf{d_2}^\top \ \mathbf{d_3}^\top \ \mathbf{p}^\top], \tag{1}$$

with $\mathbf{x}$, $\mathbf{d_1}$, $\mathbf{d_2}$, $\mathbf{d_3}$, and $\mathbf{p}$ vectors of size $m \times 1$ representing single trial recordings (time history) of the sensor's displacement, three deformation signals and proximity signal, correspondingly, and m the number of data points in a specific trial. The proximity data ($\mathbf{p}$) of each measurement trial was used to extract the data points corresponding directly to a specific mechanical feature (crack, bump, wavy pattern). This allowed to create a subset of data containing only the information specific to the mechanical feature, and to exclude the data points at the start and the end of the recording. This process was performed automatically, based on the analysis of the discrete-time derivative of proximity sensing for a given time window, and extracting the data for which the derivative exceeded a pre-defined threshold. The average discrete-time derivative for proximity measurement was

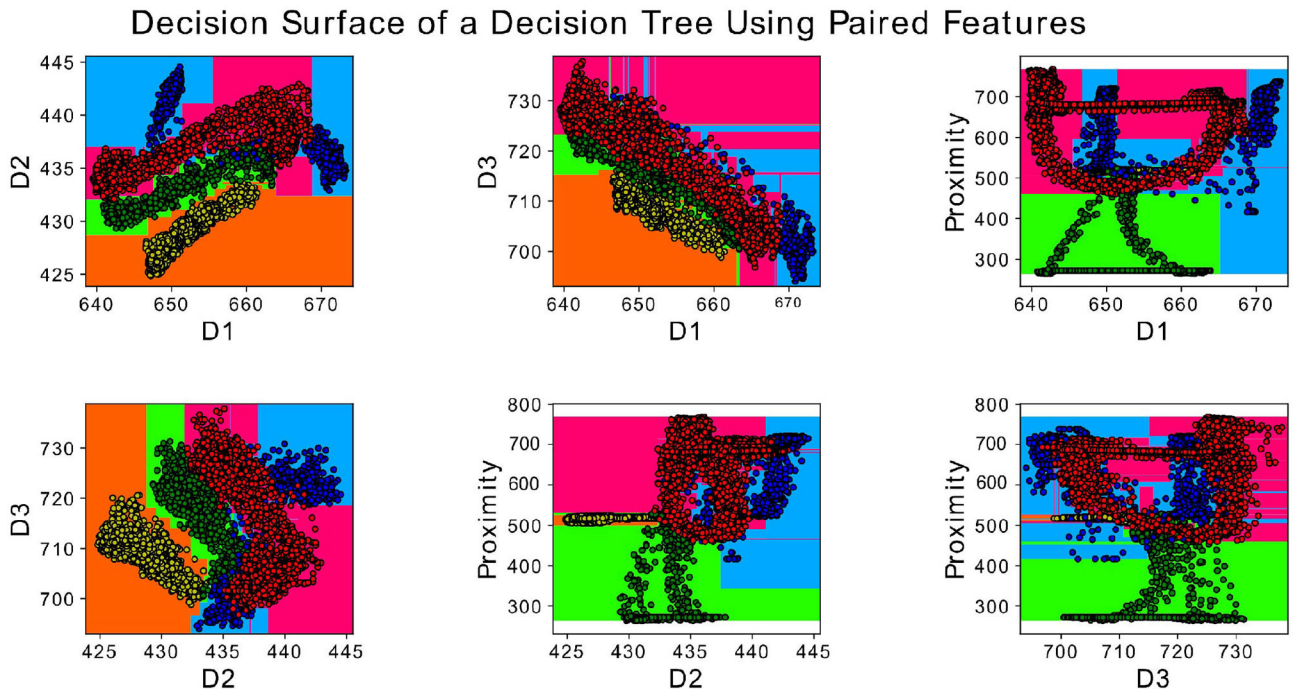

FIGURE 6 | Decision surface of random forest classifier with four classes (no crack, crack, bump, and wavy pattern) for paired features: D1, D2, D3, and Proximity.

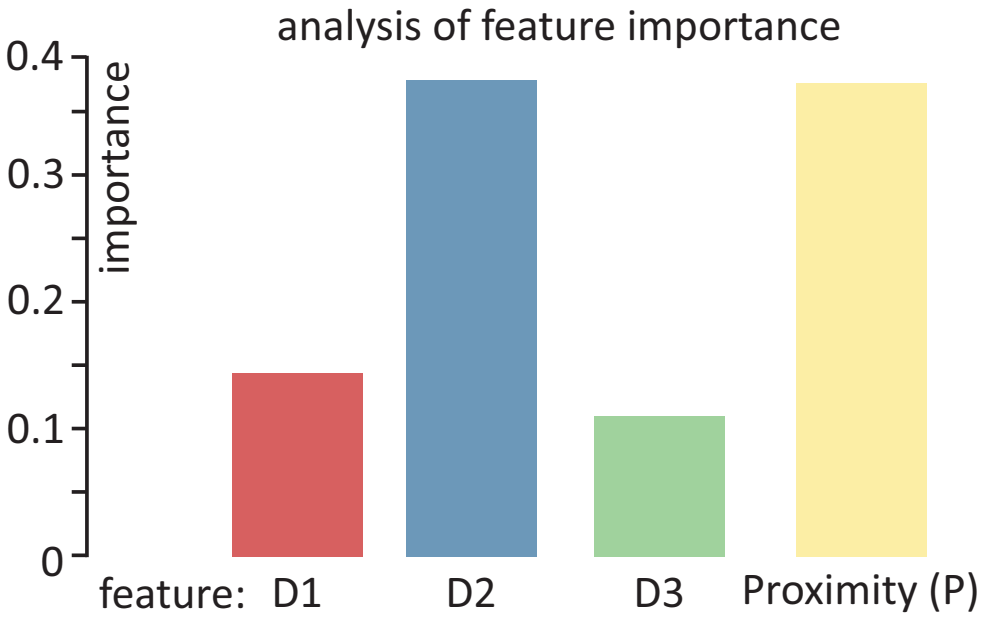

FIGURE 7 | Feature importance analysis for the random forest classifier.

computed as

$$\Delta p_i = \frac{\mathbf{p}_{i+n} - \mathbf{p}_i}{n}, \qquad \overline{\Delta \mathbf{p}} = \frac{1}{m-n} \sum_{i=1}^{m-n} \Delta p_i \tag{2}$$

with Δp_i a local discrete derivative of ith proximity signal at measurement based on n data points, $\mathbf{p_i}$ representing ith element of proximity measurement vector, and $\overline{\Delta \mathbf{p}}$ representing the averaged discrete-time derivative of the proximity measurements. Then, the data points of all measurement signals for which $\Delta p_i > |\overline{\Delta \mathbf{p}}|$ are extracted from each trial as

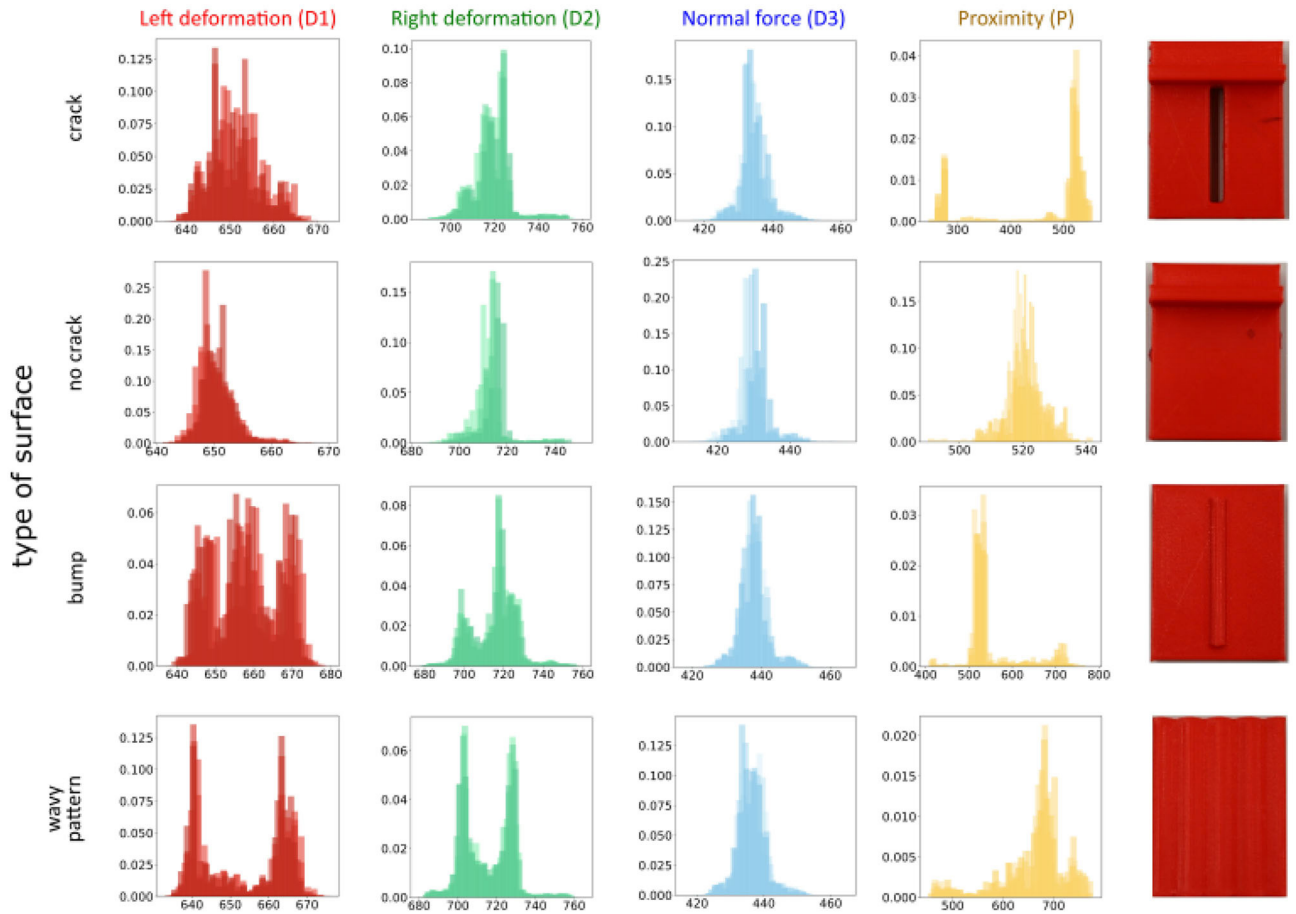

FIGURE 8 | Measurements distribution for sensing elements D1, D2, D3, and P for "crack," "no crack," "bump," and "wavy pattern" samples. Each plot shows distribution collected from six measurement trials.

they represent the changes in the sensor-sample interaction mechanics. Additional 10% of the original measurement data is added before and after the extracted points to ensure that the data is complete and represents the explored mechanical feature well. **Figure 4** shows an example of the data extraction for one sliding movement on a crack object. The complete measurement trial is displayed in blue. In green, the data points corresponding to the changes in the sensor-sample interaction mechanics. In red, these represent the discrete-time derivative for proximity measurement. To determine the appropriate moving window size for computing the derivatives we performed sample classification tests with different sliding window sizes. **Figure 5** shows the results of this test which demonstrated that a time window of 25 ms (containing $n = 10$ data samples), is sufficient to achieve better good classification accuracy.

3.2. Feature Extraction

Feature extraction was performed on each successive 25 ms time window with an increment of 5 ms. The size of the time window was selected based on the sampling frequency. Feature extraction is executed on windows of 10 data points with a window shift of 2 data points. The window length was empirically chosen through a grid search analysis. Time-domain features, including Mean Absolute Value (MAV) and Root Mean Square (RMS), were computed. The advantage of time-domain features is that they are fast to calculate since they do not require any mathematical transformation, e.g., into the frequency domain. On the other hand, they are sensitive to noise. These feature demonstrated high performance in previous surface Electromyography (sEMG) works of (Hakonen et al., 2015; Palermo et al., 2017).

3.3. Classification Algorithm

A set of classifiers was employed for the classification step: Random Forest with 100 trees, K-Nearest Neighbors (KNN) with five neighbors, and Quadratic Discriminant Analysis (QDA). Random Forest classifier (Breiman, 2001) can successfully handle high data dimensionality since it is both fast and insensitive to over-fitting. In addition, it was evaluated for remote sensing Belgiu and Drăguţ (2016). First, Random Forest with 1,000 trees was tested but it resulted in non-statistical relevant differences

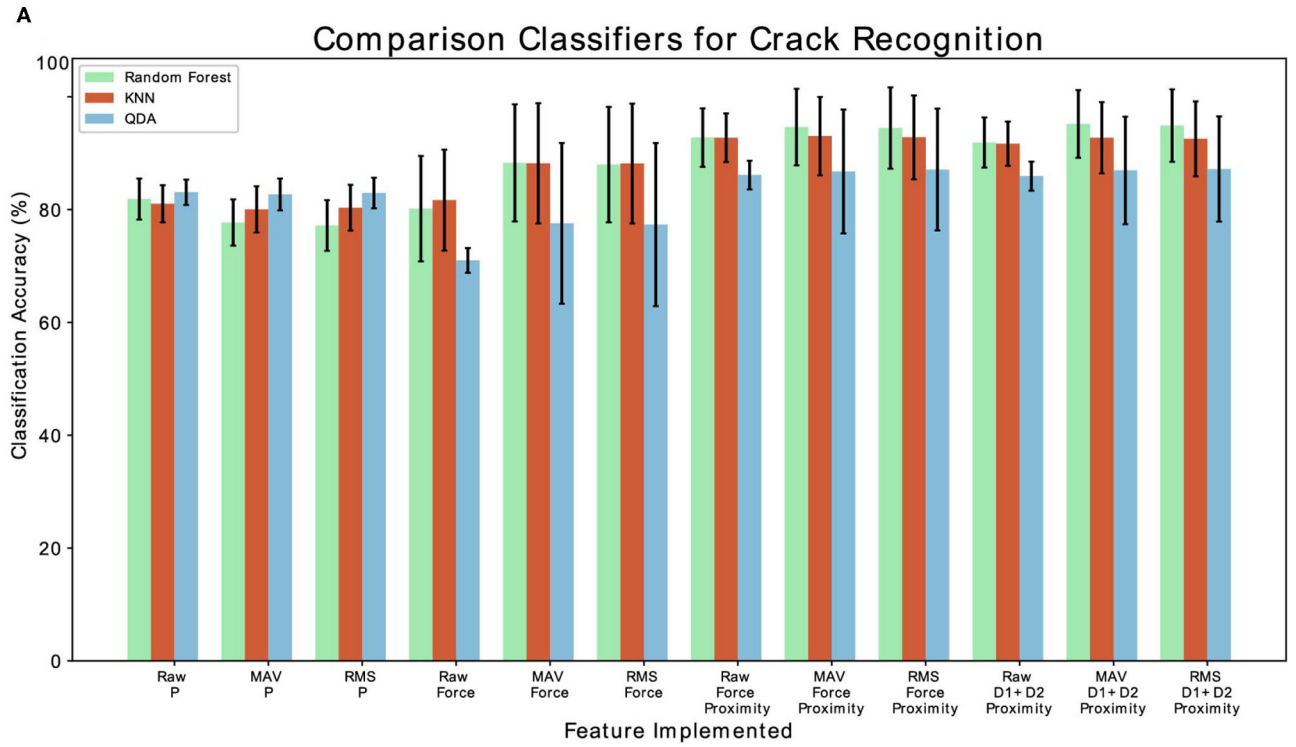

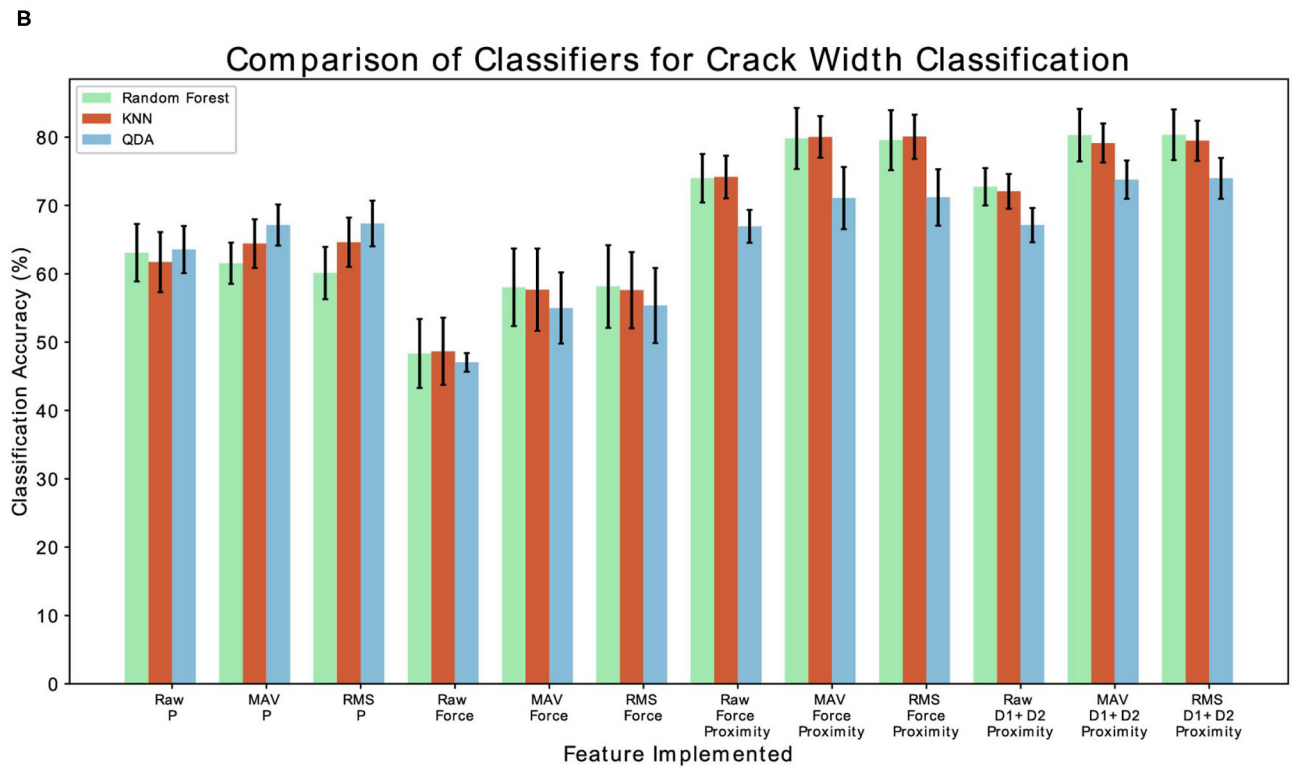

FIGURE 9 | (A) Comparison of the three implemented classifiers (Random Forest, QDA, and KNN) for fracture recognition classification. **(B)** Comparison of the three implemented classifiers (Random Forest, QDA, KNN) for crack width classification.

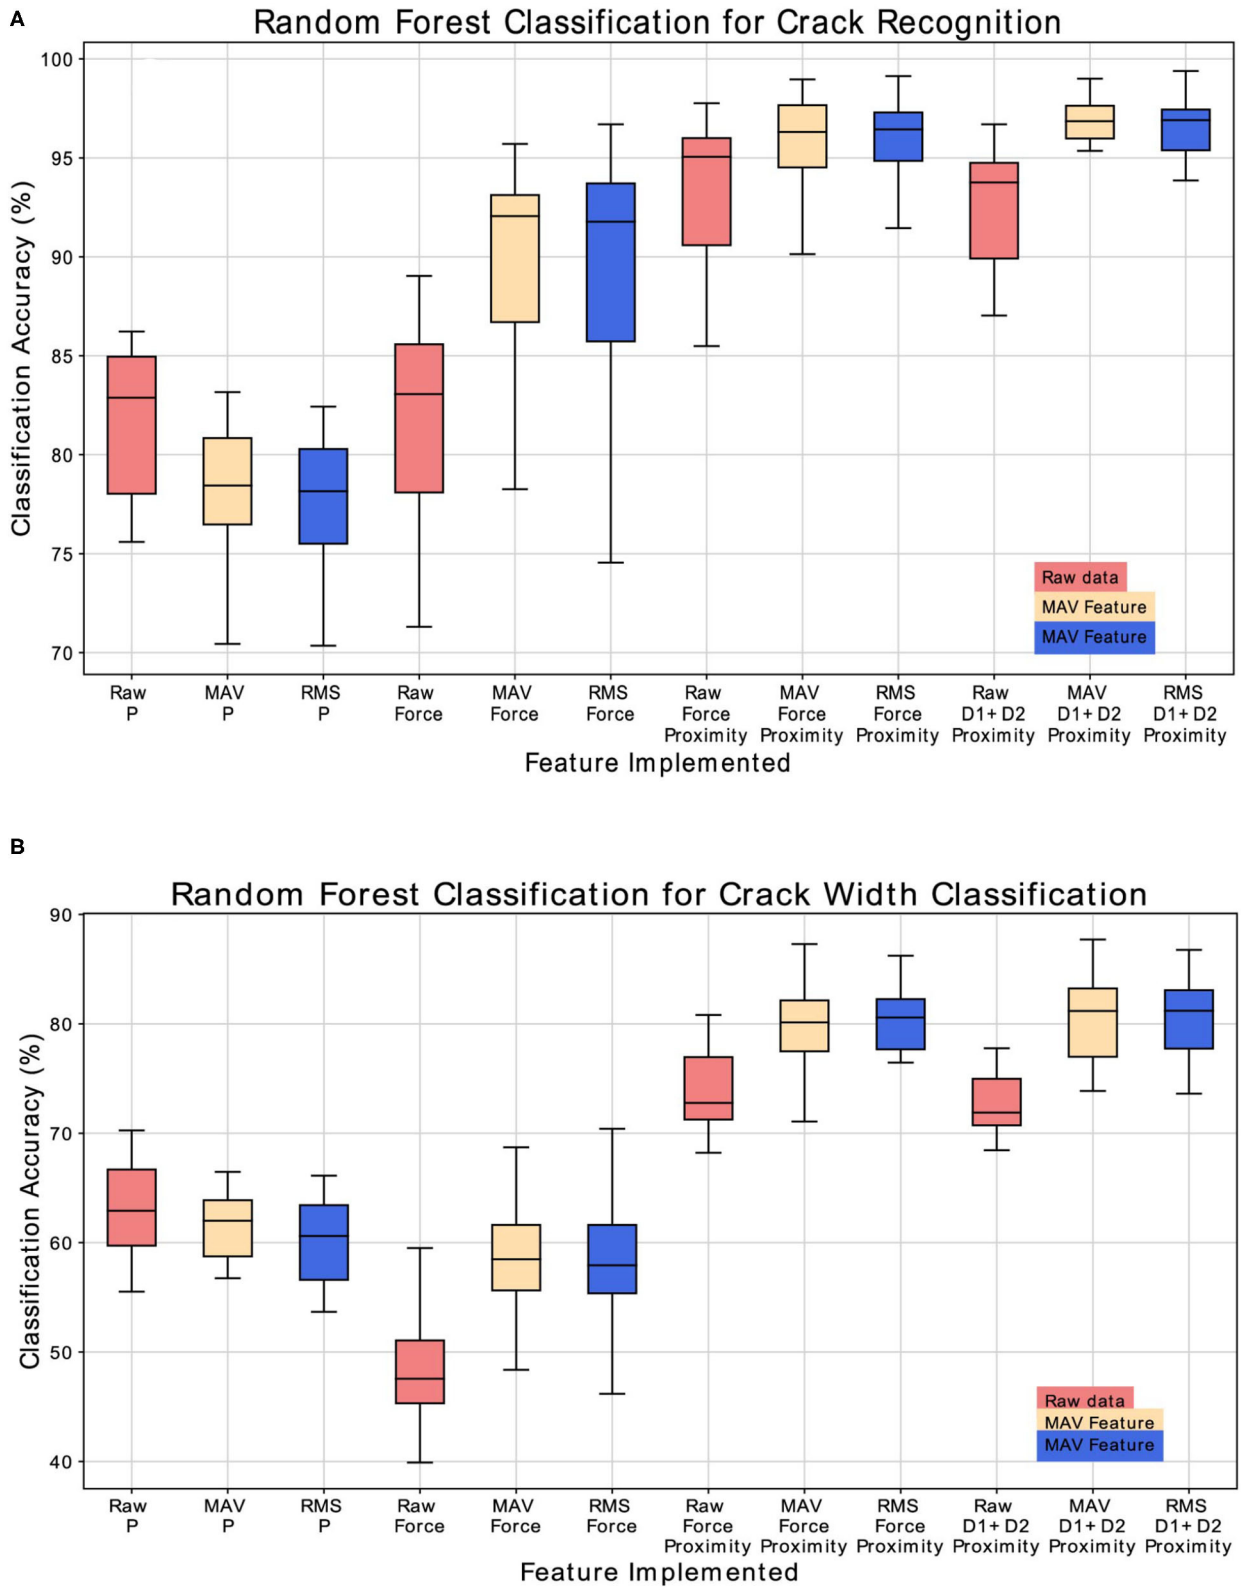

FIGURE 10 | (A) Classification accuracy results for the crack recognition. **(B)** Classification accuracy results for the crack width classification analysis.

TABLE 1 | Complete classification accuracy for crack recognition experiment with Random Forest classifier.

Implemented feature	Mean (%)	Standard deviation	Precision score (%)	Recall score (%)
RAW—Proximity	81.86	3.64	82.66	81.86
MAV—Proximity	77.70	4.10	79.25	77.70
RMS—Proximity	77.17	4.48	79.07	77.18
RAW—Deformation (D1, D2, D3)	80.16	9.34	81.97	80.16
MAV—Deformation (D1, D2, D3)	88.27	10.37	90.71	88.27
RMS—Deformation (D1, D2, D3)	87.96	10.23	90.43	87.96
RAW—Deformation + Proximity	92.75	5.17	93.83	92.75
MAV—Deformation + Proximity	94.64	6.79	95.93	94.64
RMS—Deformation + Proximity	94.48	7.21	95.88	94.48
RAW—Deformation (D1, D2) + Proximity	91.88	4.44	92.76	91.88
MAV—Deformation (D1, D2) + Proximity	95.17	5.99	96.10	95.17
RMS—Deformation (D1, D2) + Proximity	94.90	6.42	95.93	94.90

in respect to a Random Forest with 100 trees. It was then decided to discard it and use the Random Forest with 100 trees to increase the speed of the classification. The classification classes are equal to the type of surface explored (no crack, crack, bump and wavy pattern) for the surface crack recognition experiment and the width of the crack (0, 1, 2, 5, 8, and 10 mm) for the crack width classification. The complete data-set was then split 70% for training test and 30% for testing. **Figure 6** shows an example of the decision surface of one of a decision tree of the Random Forest for paired features of Proximity (P) and Deformation data (D1, D2 and D3) with MAV feature. First, raw, MAV and RMS data were classified using only the proximity data (P) or the deformation signals (D1, D2, D3). During the experiments, it was found that implementing the four dimensionality features together (P, D1, D2, D3) over-fitted the classifier. The features importance analysis was performed to avoid over-fitting. **Figure 7** shows the calculated feature importance. Among the four features. D3 is the least decisive one for the random forest classifier. Thus, the random forest was later trained and tested on proximity data (P) together with D1 and D2 deformation signals. Each observation was trained on itself and tested against the rest of the set one at a time (e.g., observation 2 was trained on itself and tested against observations 1, 3, 4, and 5) for intersession investigation. In total, 20 results for each analyzed feature were obtained. Kruskal-Wallis statistical analysis, which indicates if the data samples come from the same distribution, was performed on the whole set of results.

4. RESULTS

Figure 8 shows the distribution of data on the different surfaces of 6 of the 12 repetitions, for brevity. The common response among the repetitions permits to have no dependence on the sensor starting position and movement.

4.1. Crack Recognition

The goal of the Crack Recognition experiment is to recognize the presence of a crack in the object. **Figure 9A** shows the results of the classification with the implemented classifiers. Random Forest achieves the best classification accuracy using the implemented feature of MAV and RMS and considering the left and right displacement of the sensor (D1, D2) together with the proximity data (P). The second best classifier is KNN, which is expected, since the various class (nocrack, crack, bump, wavy pattern) data are distributed in close proximity to each other, as shown in **Figure 6**. For brevity, only the results of the Random Forest classifier are shown in the following tables and figures. **Figure 10A** shows the complete results for the classification analysis. **Table 1** shows that the lowest classification accuracy of 77% is obtained when classifying MAV or RMS data only considering the proximity data. Whereas, the best classification accuracy of 94% is achieved when implementing the MAV or RMS feature for the left and right displacement of the sensor and the proximity data. Using only deformation or only proximity data may be sufficient to train the classifier. However, better results are obtained when increasing the dimensionality of the classifier and considering proximity (P) together with the left and right displacement of the sensor (D1, D2). Thus, Implementing the whole deformation signals together with the proximity data brings little or nothing improvement to the classification accuracy in respect to using a feature with one less dimensionality. **Figure 11A** shows the results for the crack recognition. The most difficult surface to classify for the algorithm is the bump surface since it is comparable to the wavy pattern one. The Kruskal-Wallis test was performed on the results of the classification analysis of the different features and the value obtained ($p < 0.001$) indicates that the null hypothesis of having all data samples from the same distribution is rejected. Thus, there are significant differences between the implemented features.

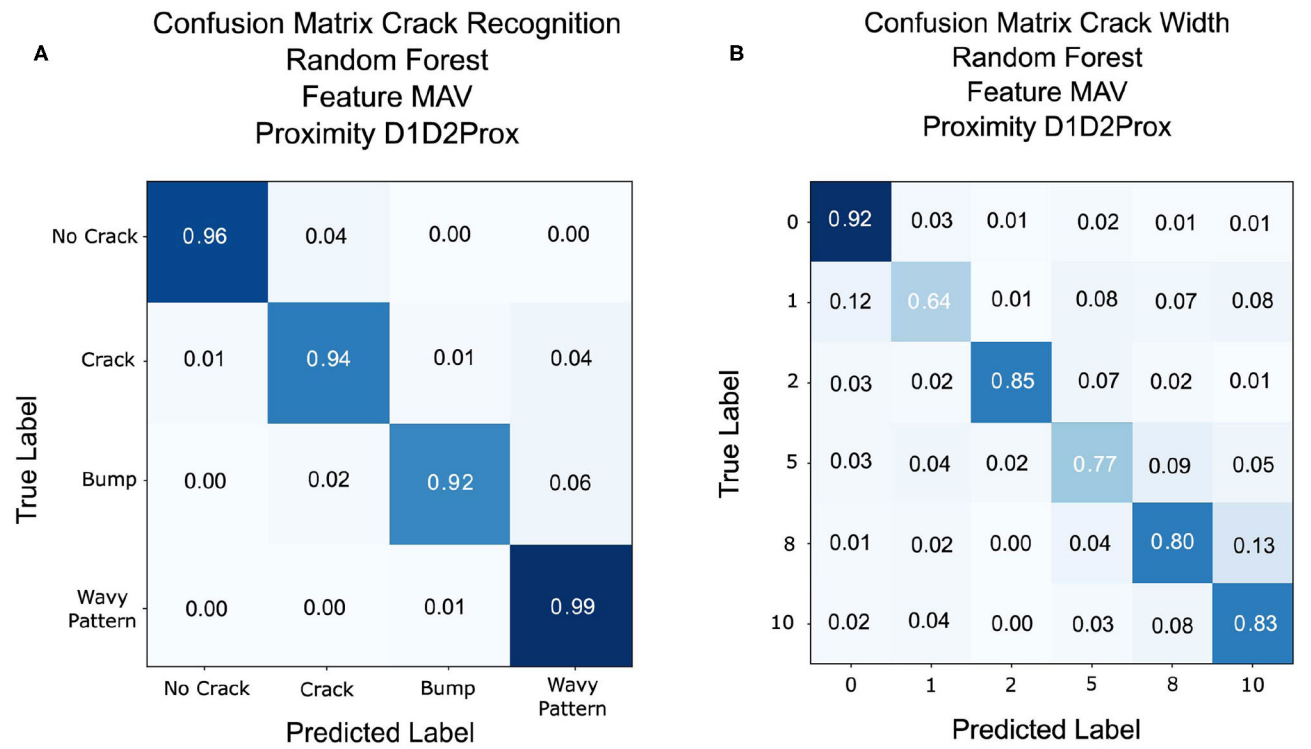

FIGURE 11 | (A) Confusion matrix result for crack recognition with Random Forest classification with MAV feature and left and right displacements combined with proximity data. **(B)** Confusion matrix result for crack width classification with Random Forest classification with MAV feature and left and right displacements combined with proximity data.

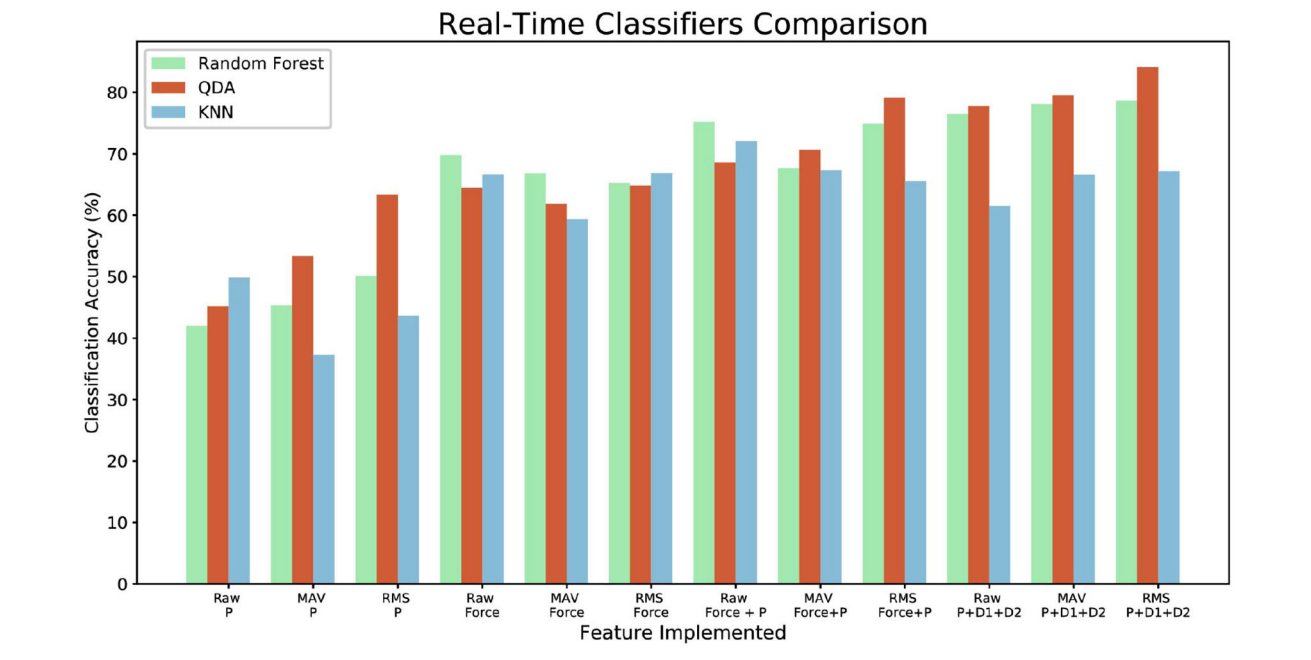

FIGURE 12 | Real-time classification comparison for Random Forest, QDA, and KNN for the complete set of implemented features.

TABLE 2 | Complete classification accuracy for crack width classification experiment with Random Forest classifier.

Implemented feature	Mean (%)	Standard deviation	Precision score (%)	Recall score (%)
RAW—Proximity	63.09	4.20	61.89	63.09
MAV—Proximity	61.54	3.02	62.00	61.54
RMS–Proximity	60.11	3.82	61.35	60.11
RAW—Deformation (D1, D2, D3)	48.33	5.05	45.63	48.33
MAV—Deformation (D1, D2, D3)	58.02	5.67	56.68	58.02
RMS—Deformation (D1, D2, D3)	58.13	6.05	56.71	58.14
RAW—Deformation + Proximity	73.97	3.55	74.47	73.97
MAV–Deformation + Proximity	79.81	4.46	81.27	79.81
RMS—Deformation + Proximity	79.55	4.39	81.02	79.55
RAW—Deformation (D1, D2) + Proximity	72.73	2.73	72.87	72.73
MAV—Deformation (D1, D2) + Proximity	80.29	3.84	81.22	80.29
RMS—Deformation (D1, D2) + Proximity	80.34	3.70	81.32	80.34

4.2. Crack Width Classification

The scope of the crack width classification experiment is to classify the width in millimeters (mm) of the fracture of the explored object. **Figure 9B** shows the comparison of accuracy, recall score, and precision score of the implemented classifier. Random Forest classifier achieves the best classification accuracy, followed by KNN. As described in the previous section, class data are close in proximity to each other, which is why KNN obtains good results for this experiment. **Figure 10B** and **Table 2** shows the complete results for the classification analysis. The lowest classification accuracy of 48.19% is obtained when classifying raw data with only deformation signals. Whereas, the best classification accuracy of ~80% is achieved when implementing MAV or RMS features with left and right deformation (D1 and D2) together with the proximity data. In this case, using only deformation or only proximity data is not sufficient to train the classifier. **Figure 11B** shows that the most difficult label to classify is the fracture of 1 mm width which can get mislabeled as a flat surface. This may be due to the fact that the fracture is so small that the left and right displacement are not big enough to trigger the recognition of the crack. Kruskal-Wallis results ($p < 0.001$) indicate statistically significant differences among results obtained when analyzing only deformation signals, only proximity data and both deformation and proximity data. In this case, instead of using a classifier, a regressor may be more appropriate to use since having a discrete class may not be the best solution when predicting the width of a fracture.

4.3. Real-Time Implementation

To further test the result of the classifier an online application was developed for the crack recognition analysis. During this experiment, it was found that the fiber optic cables position and their twisting influence the sensor data. Thus, an additional acquisition was necessary to obtain a model to use for the real-time classification. Offline models of Random Forest Classifier, KNN, and QDA were generated implementing the newly acquired data. The models were later used to predict the class of the data acquired in real-time while sliding the sensor over different surfaces. The software marks the start position of the detected crack and end position in relation to the Geomagic position. The same analysis, as the previously described offline classifier, was applied. Three continuous sliding movements were performed on each of the crack type surfaces as shown on the first row of **Figure 2**. Each movement was performed in a different section of the surface (top, center, bottom). The possible proximity data (only proximity, only deformation, deformation + proximity, and P + D1 + D2) and features (Raw Data, MAV, RMS) combination were investigated for each classifier, for a total of 432 classified movements. **Figure 12** shows the results of the real time classification accuracy. In this case, the KNN classifier achieves better results than the Random Forest. Increasing the number of classified sliding movements may reduce this difference. Having 3D printing and of regular shape objects may limit the training and testing but this will be addressed and improved in future analysis.

5. CONCLUSION AND FUTURE WORK

This work demonstrates how tactile and proximity sensing can be efficiently used to perform automatic crack detection. The proposed method uses machine learning techniques to detect the surface fractures and bumps of explored objects based on fiber optical proximity signals which are recorded during physical interaction between a custom-designed robotic finger and the remote environment. Experimental validation of the proposed method has shown that it is possible to achieve around 94% for crack detection and 80% for crack width classification accuracy. To achieve better results for the crack width classification, an alternative regressor may be more appropriate to use with respect to the implemented classifier. Real-time classification results, on three sliding movements, shows that it is possible to correctly

characterize the surface of the investigated object. During this experiment, it was found that the fiber optic cables position and their twisting influences the sensor data. Thus, additional analysis will be required to compensate for the change of flow of the data. In contrast to previous techniques, which rely on visual modality, the proposed approach, based on optical fibers, which may be more suitable for operation in extreme environments (such as nuclear facilities) where radiation damages electronic components of video cameras.

Future research will focus on integrating a multi-modal approach with visual patches and implementation of the proposed system on a teleoperated mobile manipulation system. We plan to demonstrate how automatic fracture characterization will be efficiently integrated with the mobile manipulator controller (Farkhatdinov and Ryu, 2008) and how the obtained tactile data can be visualized in a dedicated virtual reality-based human-operator interface (Omarali et al., 2020). Further studies will be performed on dimensionality reduction with principal component analysis which may increase the classification accuracy. Additional features, such as local min-max values, which may give a better comprehension of the data, will be analyzed. The implementation of an alternative bio-inspired ciliary force sensor will be investigated for small crack detection (Ribeiro et al., 2017).

DATA AVAILABILITY STATEMENT

The datasets generated for this study can be found in the figshare repository "Automatic Fracture Database" available at https://figshare.com/s/14deb00d874400e34d67.

AUTHOR CONTRIBUTIONS

FP developed the classification algorithm, experimental methods, performed the experiments, and wrote the draft of the paper. IF developed the experimental methods, contributed to the data analysis, and provided the input to the paper. KA advised on the experimental methods and corrected the paper. JK and SP assisted with the paper writing. All authors contributed to the article and approved the submitted version.

FUNDING

This work has been funded by the National Centre for Nuclear Robotics (UK EPSRC grant NCNR EP/R02572X/1), UK Innovate iGRASP project, by the European Community's Seventh Framework Programme grant agreement no. 610532 (SQUIRREL), and by the Alan Turing Fellowships to KA and IF.

REFERENCES

Ball, S. J., Holcomb, D. E., and Cetiner, S. (2012). *HTGR Measurements and Instrumentation Systems*. ORNL/TM-2012/107. Oak Ridge, TN: Oak Ridge National Laboratory.

Barhli, S., Saucedo-Mora, L., Jordan, M., Cinar, A., Reinhard, C., Mostafavi, M., et al. (2017). Synchrotron X-ray characterization of crack strain fields in polygranular graphite. *Carbon* 124, 357–371. doi: 10.1016/j.carbon.2017.08.075

Belgiu, M., and Drăguţ, L. (2016). Random forest in remote sensing: a review of applications and future directions. *ISPRS J. Photogramm. Rem. Sens.* 114, 24–31. doi: 10.1016/j.isprsjprs.2016.01.011

Berghmans, F., and Decreton, M. (1998). Evaluation of three different optical fibre temperature sensor types for application in gamma radiation environments. *IEEE Trans. Nucl. Sci.* 45, 1537–1542. doi: 10.1109/23.685236

Berghmans, F., Fernandez, A. F., Brichard, B., Vos, F., Decreton, M. C., Gusarov, A. I., et al. (1999). "Radiation hardness of fiber optic sensors for monitoring and remote handling applications in nuclear environments," in *Process Monitoring With Optical Fibers and Harsh Environment Sensors*, Vol. 3538 (Boston, MA: International Society for Optics and Photonics), 28–39. doi: 10.1117/12.335757

Berthold, J. W. III. (1994). "Overview of prototype fiber optic sensors for future application in nuclear environments," in *Optical Fibre Sensing and Systems in Nuclear Environments*, Vol. 2425 (Mol: Society of Photo-Optical Instrumentation Engineers), 74–83. doi: 10.1117/12.198644

Bogue, R. (2011). Robots in the nuclear industry: a review of technologies and applications. *Industrial Robot: An International Journal*.

Bogue, R. (2013). Radiation hardening and sensors for radioactive environments. *Sens. Rev.* 33, 191–196. doi: 10.1108/02602281311324636

Breiman, L. (2001). Random forests. *Mach. Learn.* 45, 5–32. doi: 10.1023/A:1010933404324

Chakraborty, J., Katunin, A., Klikowicz, P., and Salamak, M. (2019). Early crack detection of reinforced concrete structure using embedded sensors. *Sensors* 19:3879. doi: 10.3390/s19183879

Chen, F.-C., and Jahanshahi, M. R. (2017). NB-CNN: deep learning-based crack detection using convolutional neural network and naïve bayes data fusion. *IEEE Trans. Ind. Electron.* 65, 4392–4400. doi: 10.1109/TIE.2017.2764844

Drimus, A., Kootstra, G., Bilberg, A., and Kragic, D. (2014). Design of a flexible tactile sensor for classification of rigid and deformable objects. *Robot. Auton. Syst.* 62, 3–15. doi: 10.1016/j.robot.2012.07.021

Farkhatdinov, I., and Ryu, J.-H. (2008). "Teleoperation of multi-robot and multi-property systems," in *2008 6th IEEE International Conference on Industrial Informatics* (Daejeon: IEEE), 1453–1458. doi: 10.1109/INDIN.2008.4618333

Feng, D., Kaboli, M., and Cheng, G. (2018). Active prior tactile knowledge transfer for learning tactual properties of new objects. *Sensors* 18:634. doi: 10.3390/s18020634

Fishel, J. A., and Loeb, G. E. (2012). Bayesian exploration for intelligent identification of textures. *Front. Neurorobot.* 6:4. doi: 10.3389/fnbot.2012.00004

Hakonen, M., Piitulainen, H., and Visala, A. (2015). Current state of digital signal processing in myoelectric interfaces and related applications. *Biomed. Signal Process. Control* 18, 334–359. doi: 10.1016/j.bspc.2015.02.009

Hashemian, H. (2009). The state of the art in nuclear power plant instrumentation and control. *Int. J. Nucl. Energy Sci. Technol.* 4, 330–354. doi: 10.1504/IJNEST.2009.028597

Huang, K., Lancaster, P., Smith, J. R., and Chizeck, H. J. (2018). "Visionless tele-exploration of 3D moving objects," in *2018 IEEE International Conference on Robotics and Biomimetics (ROBIO)* (Kuala Lumpur: IEEE), 2238–2244. doi: 10.1109/ROBIO.2018.8665228

Huet, L. A., Rudnicki, J. W., and Hartmann, M. J. (2017). Tactile sensing with whiskers of various shapes: determining the three-dimensional location of object contact based on mechanical signals at the whisker base. *Soft Robot.* 4, 88–102. doi: 10.1089/soro.2016.0028

Iliopoulos, S., Aggelis, D., Pyl, L., Vantomme, J., Van Marcke, P., Coppens, E., et al. (2015). Detection and evaluation of cracks in the concrete buffer of the belgian nuclear waste container using combined ndt techniques. *Constr. Build. Mater.* 78, 369–378. doi: 10.1016/j.conbuildmat.2014.12.036

Inaudi, D., Glisic, B., Fakra, S., Billan, J., Redaelli, S., Perez, J. G., et al. (2001). Development of a displacement sensor for the CERN-LHC superconducting cryodipoles. *Meas. Sci. Technol.* 12:887. doi: 10.1088/0957-0233/12/7/324

Kaboli, M., and Cheng, G. (2018). Robust tactile descriptors for discriminating objects from textural properties via artificial robotic skin. *IEEE Trans. Robot.* 34, 985–1003. doi: 10.1109/TRO.2018.2830364

Kaboli, M., Walker, R., and Cheng, G. (2016). "Re-using prior tactile experience by robotic hands to discriminate in-hand objects via texture properties," in *2016 IEEE International Conference on Robotics and Automation (ICRA)* (Stockholm: IEEE), 2242–2247. doi: 10.1109/ICRA.2016.7487372

Kappassov, Z., Corrales, J.-A., and Perdereau, V. (2015). Tactile sensing in dexterous robot hands. *Robot. Auton. Syst.* 74, 195–220. doi: 10.1016/j.robot.2015.07.015

Kim, R., Park, C. H., Yoo, W. J., and Moon, J. H. (2017). Development and characterization of a fiber-optic monitoring system for the key environment variables of the spent nuclear fuel pool at a nuclear power plant. *Ann. Nucl. Energy* 99, 183–192. doi: 10.1016/j.anucene.2016.08.020

Konstantinova, J., Cotugno, G., Stilli, A., Noh, Y., and Althoefer, K. (2017). "Object classification using hybrid fiber optical force/proximity sensor," in *2017 IEEE Sensors* (Glasgow: IEEE), 1–3. doi: 10.1109/ICSENS.2017.8234057

Konstantinova, J., Stilli, A., Faragasso, A., and Althoefer, K. (2016). "Fingertip proximity sensor with realtime visual-based calibration," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (Daejeon: IEEE), 170–175. doi: 10.1109/IROS.2016.7759051

Lee, H.-K., Chang, S.-I., and Yoon, E. (2006). A flexible polymer tactile sensor: fabrication and modular expandability for large area deployment. *J. Microelectromech. Syst.* 15, 1681–1686. doi: 10.1109/JMEMS.2006.886021

Liu, H., Nguyen, K. C., Perdereau, V., Bimbo, J., Back, J., Godden, M., et al. (2015). Finger contact sensing and the application in dexterous hand manipulation. *Auton. Robots* 39, 25–41. doi: 10.1007/s10514-015-9425-4

Liu, H., Song, X., Bimbo, J., Seneviratne, L., and Althoefer, K. (2012). "Surface material recognition through haptic exploration using an intelligent contact sensing finger," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems* (Vilamoura: IEEE), 52–57. doi: 10.1109/IROS.2012.6385815

Lu, W., and Chu, F. (2011). Shaft crack identification based on vibration and AE signals. *Shock Vibr.* 18, 115–126. doi: 10.1155/2011/460178

Luo, S., Bimbo, J., Dahiya, R., and Liu, H. (2017). Robotic tactile perception of object properties: a review. *Mechatronics* 48, 54–67. doi: 10.1016/j.mechatronics.2017.11.002

Morita, Y., and Seguchi, T. (1983). Radiation resistance of nylon. *Denki Gakkai Zetsuen Zairyo Kenkyukai Shiryo* 83, 47–52.

Naragani, D., Sangid, M. D., Shade, P. A., Schuren, J. C., Sharma, H., Park, J.-S., et al. (2017). Investigation of fatigue crack initiation from a non-metallic inclusion via high energy X-ray diffraction microscopy. *Acta Mater.* 137, 71–84. doi: 10.1016/j.actamat.2017.07.027

Nicoletti, R., Cavalini, A. A., and Steffen, V. (2018). Detection of cracks in rotating shafts by using the combination resonances approach and the approximated entropy algorithm. *Shock Vibr.* 2018:4094631. doi: 10.1155/2018/4094631

Omarali, B., Denoun, B., Althoefer, K., Jamone, L., Valle, M., and Farkhatdinov, I. (2020). "Virtual reality based telerobotics framework with depth cameras," in *2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)* (Naples: IEEE). doi: 10.1109/RO-MAN47096.2020.9223445

Palermo, F., Cognolato, M., Gijsberts, A., Müller, H., Caputo, B., and Atzori, M. (2017). "Repeatability of grasp recognition for robotic hand prosthesis control based on sEMG data," in *2017 International Conference on Rehabilitation Robotics (ICORR)* (London: IEEE), 1154–1159. doi: 10.1109/ICORR.2017.8009405

Palermo, F., Konstantinova, J., Poslad, S., Althoefer, K., and Farkhatdinov, I. (2020). "Implementing tactile and proximity sensing for crack detection," in *2020 IEEE International Conference on Robotics and Automation (ICRA)* (Paris: IEEE). doi: 10.1109/ICRA40945.2020.9196936

Phéron, X., Girard, S., Boukenter, A., Brichard, B., Delepine-Lesoille, S., Bertrand, J., et al. (2012). High γ-ray dose radiation effects on the performances of brillouin scattering based optical fiber sensors. *Opt. Express* 20, 26978–26985. doi: 10.1364/OE.20.026978

Ribeiro, P., Khan, M. A., Alfadhel, A., Kosel, J., Franco, F., Cardoso, S., et al. (2017). Bioinspired ciliary force sensor for robotic platforms. *IEEE Robot. Autom. Lett.* 2, 971–976. doi: 10.1109/LRA.2017.2656249

Schmugge, S. J., Rice, L., Nguyen, N. R., Lindberg, J., Grizzi, R., Joffe, C., et al. (2016). "Detection of cracks in nuclear power plant using spatial-temporal grouping of local patches," in *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)* (Lake Placid, NY: IEEE), 1–7. doi: 10.1109/WACV.2016.7477601

Tomo, T. P., Regoli, M., Schmitz, A., Natale, L., Kristanto, H., Somlor, S., et al. (2018). A new silicone structure for uSkin—a soft, distributed, digital 3-axis skin sensor and its integration on the humanoid robot icub. *IEEE Robot. Autom. Lett.* 3, 2584–2591. doi: 10.1109/LRA.2018.2812915

Tomo, T. P., Schmitz, A., Wong, W. K., Kristanto, H., Somlor, S., Hwang, J., et al. (2017). Covering a robot fingertip with uSkin: a soft electronic skin with distributed 3-axis force sensitive elements for robot hands. *IEEE Robot. Autom. Lett.* 3, 124–131. doi: 10.1109/LRA.2017.2734965

Wong, R. D. P., Hellman, R. B., and Santos, V. J. (2014). "Haptic exploration of fingertip-sized geometric features using a multimodal tactile sensor," in *Next-Generation Robots and Systems*, Vol. 9116 [Baltimore: Society of Photo-Optical Instrumentation Engineers (SPIE)], 911605.

Wood, R. (2004). *Advanced Reactor Licensing: Experience With Digital I and C Technology in Evolutionary Plants*. Technical report, ORNL (US). US Department of Energy (United States).

Yao, Y., Tung, S.-T. E., and Glisic, B. (2014). Crack detection and characterization techniques—an overview. *Struct. Control Health Monit.* 21, 1387–1413. doi: 10.1002/stc.1655

Yuan, W., Zhu, C., Owens, A., Srinivasan, M. A., and Adelson, E. H. (2017). "Shape-independent hardness estimation using deep learning and a gelsight tactile sensor," in *2017 IEEE International Conference on Robotics and Automation (ICRA)* (Singapore: IEEE), 951–958. doi: 10.1109/ICRA.2017.7989116

Conflict of Interest: JK was employed by the company Ocado Technology.

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

frontiers in Robotics and AI

ORIGINAL RESEARCH
published: 18 June 2021
doi: 10.3389/frobt.2021.550644

A Holistic Approach to Human-Supervised Humanoid Robot Operations in Extreme Environments

*Murphy Wonsick[1], Philip Long[2], Aykut Özgün Önol[1], Maozhen Wang[1] and Taşkın Padır[1]**

[1]Institute for Experiential Robotics, Northeastern University, Boston, MA, United States, [2]Irish Manufacturing Research, National Science Park, Mullingar, Ireland

Nuclear energy will play a critical role in meeting clean energy targets worldwide. However, nuclear environments are dangerous for humans to operate in due to the presence of highly radioactive materials. Robots can help address this issue by allowing remote access to nuclear and other highly hazardous facilities under human supervision to perform inspection and maintenance tasks during normal operations, help with clean-up missions, and aid in decommissioning. This paper presents our research to help realize humanoid robots in supervisory roles in nuclear environments. Our research focuses on National Aeronautics and Space Administration (NASA's) humanoid robot, Valkyrie, in the areas of constrained manipulation and motion planning, increasing stability using support contact, dynamic non-prehensile manipulation, locomotion on deformable terrains, and human-in-the-loop control interfaces.

Keywords: humanoid robots, motion planning, supervised autonomy, nuclear, glovebox

OPEN ACCESS

Edited by:
Chie Takahashi,
University of Cambridge,
United Kingdom

Reviewed by:
Manuel Ferre,
Polytechnic University of Madrid,
Spain
William R. Hamel,
The University of Tennessee,
United States

****Correspondence:***
Taşkın Padır
t.padir@northeastern.edu

Specialty section:
This article was submitted to Robotic Control Systems, a section of the journal Frontiers in Robotics and AI

Received: *09 April 2020*
Accepted: *10 May 2021*
Published: *18 June 2021*

Citation:
Wonsick M, Long P, Önol AÖ, Wang M and Padır T (2021) A Holistic Approach to Human-Supervised Humanoid Robot Operations in Extreme Environments. Front. Robot. AI 8:550644. doi: 10.3389/frobt.2021.550644

1 INTRODUCTION

As the worldwide energy demand is expected to increase by 50% within the next 3 decades, nuclear energy will play a critical role in meeting clean energy targets worldwide. At the same time, many of the world's nuclear reactors are aging; from Japan to the United Kingdom to the United States, scientists, engineers and regulators are counting on new innovative technologies that will make decommissioning and clean-up missions safe for humans, environmentally-friendly and cost-effective. Furthermore, industrializing countries are investing in building new nuclear power plants to meet their growing energy demands.

Nuclear energy operations and nuclear disasters have great international impact with no boundaries. Ensuring safe, efficient, and productive operations of facilities and improving response to unplanned emergencies at any location around the globe is in the best interest of the international community. Moreover, the urgency and scale of the problems identified in high-consequence situations, such as the Fukushima (Japan) clean up and waste tank decommissioning in Savannah River Site (United States), require an interdisciplinary team of scientists, engineers, and technologists to solve similar yet sufficiently distinct challenges. For example, since 1989, the United States Department of Energy has spent over $250 billion of public funds on cleanup. The cleanup is less than half complete and the remaining mission scope is estimated to cost at least another $250 billion more over a 40–50- year period. Similarly, the Japanese government estimated the total costs for the Fukushima cleanup at $188 billion for the next 40 years.

It is now clear that robotics will play a key role in accelerating these cleanup timelines and reducing the costs, by addressing operational needs and challenges in nuclear facilities. Robotics technologies are needed to remotely access nuclear and other highly hazardous facilities under

human supervision to perform inspection and maintenance tasks during normal operations. During decommissioning, robots will become the eyes and hands of human operators in places no human has gone for more than 50 years. However, using robots in such environments is not without challenges, the Fukushima disaster has shown that conventional robots must be significantly modified, Nagatani et al. (2011b), to cope with highly radioactive environments. This is typically achieved by equipping the system with lead plates to protect electrical components, Nagatani et al. (2011a), which in-turn add weight and may impair functionality. Alternatively, the number of electronic components can be minimized and used in conjunction with specifically hardened parts, however high levels of radiation still can destroy such systems in a matter of days, Funakoshi (2016), Urabe and Stapczynski (2017), Yirmibeşoğlu et al. (2019). Nevertheless, robots have already been deployed in less active nuclear environments in commissioning and waste disposal task, Bogue (2011), and on-going research is demonstrating promising applications for human-supervised robotics in a range of different tasks in the future, Marturi et al. (2016). They will improve our ability to respond to and recover from unplanned events or operational emergencies in such critical and safety-significant applications. Nuclear environments are *dangerous* for humans to operate in due to the presence of highly radioactive materials. They are typically *distant* as the facilities separate these dangerous environments from humans by thick walls. And they present *daring* operational conditions by size and configuration with tight passages, debris accumulated over the years and cluttered internals.

Our research plan is motivated by the need for general purpose robots in routine and emergency operations in nuclear facilities for: 1) Disaster response and environmental clean-up, such as more than 175 waste tanks in Hanford (WA) or the F-canyon in Savannah River (SC). These processes rely heavily on accurate remote sampling and characterization before permanently grouting the facilities, by collecting samples from different spots for analyzing what and how much radioactive material remains. The operating environment is unknown and cluttered with vertical and horizontal piping, fallen debris and puddles and muck-like material on the floors (in the waste tanks). 2) Operational efficiencies—federal laws require nuclear facilities to develop and maintain emergency preparedness plans to protect the public. These emergencies include unusual events during normal operations to black swan events such as Fukushima, Chernobyl and Three Mile Island accidents. Furthermore, the aging workforce in the energy sector requires the adoption of technology to keep up with day-to-day operations. As a result, human-supervised robot assets with robust manipulation capabilities in these challenging environments and situations are needed. 3) Worker safety—Before any work can begin, human workers must enter a facility to characterize radioactive hazards, such as type of radiation, dose rates, and location of sources. This data is then used to determine the proper protective clothing and stay time limits for personnel. Replacing personnel with robots would be highly desirable in radioactive environments.

This paper discusses our approach to move towards utilizing humanoid robots in nuclear energy operations and nuclear disasters by furthering the development of human-supervised robot control and manipulation capabilities. **Section 2** presents a method capable of handling manipulation and motion planning in constrained environments, such as gloveboxes. **Section 3** presents work in utilizing support contacts to increase stability of a standing humanoid robot operating inside a glovebox. **Section 4** presents a method to plan dynamic non-prehensile manipulation behaviors in a highly-constrained environment, with focus on gloveboxes. **Section 5** presents a way to estimate *in-situ* deformable terrains in order to navigate in unknown environments. **Section 6** presents a human-in-the-loop user interface for operating humanoid robots.

2 CONSTRAINED MANIPULATION AND MOTION PLANNING

In hazardous environments it is crucial to perform manipulation tasks effectively. Additionally, such environments often provide constraints on the motions allowed, such as when operating through glovebox ports. In order to accomplish effective manipulation, different robot configurations should be evaluated. This is particularly important for redundant systems such as humanoid robots. A robot performance measure can be classified as local, e.g., manipulability (Yoshikawa, 1984) or global such as workspace analysis (Vahrenkamp and Asfour, 2015). Local indexes are advantageous as they provide a more generic solution and may be utilized in control frameworks without workspace knowledge. Hence, they can be used to choose a configuration based on a robot's inherent capability. However, local indexes study the system's kinematic transformations from configuration to task space, ignoring environmental constraints that have significant effects on the robot's admissible motions. In particular, for operations in hazardous environments, the workspace is often unstructured and uncontrolled. Moreover, unscheduled contacts may lead to catastrophic results. For these reasons, it is important to transmit an accurate measure of what the robot can or cannot do in its current pose to a remote operator/supervisor.

The following section recalls the work presented in (Long and Padır, 2018; Long and Padir, 2019; Long et al., 2019) in which a new measure called the constrained manipulability polytope (CMP) that considers the system's kinematic structure, including closed chains, composite sub-mechanisms, joint limits, and the presence of obstacles is developed.

2.1 Related Work

The manipulability ellipsoid first defined (Yoshikawa, 1984) measures the capabilities of a robot manipulator based on its kinematic structure. Extensions to include positional joint limits are proposed in (Tsai, 1986) using penalty functions and (Abdel-Malek et al., 2004) using augmented jacobian matrices. Robots with heterogeneous joint velocity limits are examined in (Lee, 1997), while dynamic constraints are considered (Bowling and Khatib, 2005; Zollo et al., 2008). For humanoid robots, improvements on local measures can be obtained by including the effects of contact while evaluating the dynamic manipulability

of a humanoid's center of mass (Gu et al., 2015; Azad et al., 2017). Moreover in (Azad et al., 2017) either joint torque and/or acceleration limits are accounted for by using a scaling matrix. In (Vahrenkamp et al., 2012, 2013) the manipulability ellipsoid is augmented to include environmental constraints. The authors include joint position limits and the detrimental effects of nearby obstacles using a spatial decomposition referred to as the hyperoctants approach. Manipulability polytopes (Kokkinis and Paden, 1989) provide a more elegant and indeed exact method for representing velocity limits in the Cartesian space. There are several examples where diverse constraints, defined by a set of inequality or equality equations, have been incorporated into polytopes. For instance, mobile robot toppling constraints have been integrated into the available wrench set for a cable-driven parallel robot in (Rasheed et al., 2018). Alternatively, friction constraints can be added after linearization (Caron et al., 2017).

2.2 Manipulability

Consider an n degree-of-freedom (DOF) manipulator in m dimensional space. Let $\boldsymbol{v}_n$ denote the twist at the end effector, comprising three translational and three angular velocities defined, respectively, as $\mathbf{v}$ and $\boldsymbol{\omega}$. $\boldsymbol{v}_n$ is obtained as

$$\boldsymbol{v}_n = \begin{bmatrix} \mathbf{v} \\ \boldsymbol{\omega} \end{bmatrix} = \mathbf{J}_n \dot{\mathbf{q}}, \tag{1}$$

where $\mathbf{J}_n \in \mathbb{R}^{6\times n}$ is the Jacobian matrix and $\dot{\mathbf{q}} = [\dot{q}_1, \dot{q}_2 \dots \dot{q}_n]^T$ is the joint velocity vector. An exact measure of the manipulator's capabilities can be obtained by studying the manipulability polytope in conjunction with the joint velocity limits. A polytope, $\mathcal{P}$ can be represented as the convex hull of its vertex set ($\mathcal{V}$-representation), i.e.,

$$\mathcal{P}^V = \left\{ \mathbf{x} : \mathbf{x} = \sum_{i=1}^{n} \alpha_i \mathbf{y}_i \middle| \alpha_i \geq 0, \sum_{i=1}^{n} \alpha_i = 1 \right\}, \tag{2}$$

where $\mathbf{y}_i$ denotes the i^{th} element of the vertex set and $\mathbf{x}$ is any point inside $\mathcal{P}$. Equivalently, $\mathcal{P}$ can be defined as the volume bounded by a finite number of half-spaces ($\mathcal{H}$-representation)

$$\mathcal{P}^H = \mathbf{A}\mathbf{x} \leq \mathbf{b}, \tag{3}$$

where $\mathbf{A}$ contains the half-spaces' normals and $\mathbf{b}$ is the shifted distance from the origin along the normal. Converting from $\mathcal{V}$ and $\mathcal{H}$ is possible, for example, using the double description[1] method (Fukuda and Prodon, 1996). The polytope representing joint velocities for an n-DOF robot, denoted by $\mathcal{Q}$, is written in $\mathcal{H}$-representation as

$$\mathcal{Q}^H = \begin{bmatrix} \mathbf{I}_n \\ -\mathbf{I}_n \end{bmatrix} \dot{\mathbf{q}} \leq \begin{bmatrix} \dot{\mathbf{q}}_{max} \\ -\dot{\mathbf{q}}_{min} \end{bmatrix}, \tag{4}$$

where $\mathbf{I}_n$ is the $n \times n$ identity matrix and $\dot{\mathbf{q}}_{max}$ and $\dot{\mathbf{q}}_{min}$ denote the robot's maximum and minimum joint velocities respectively. The equivalent polytope defined by its vertices is written as

$$\mathcal{Q}^V = \{ \dot{\mathbf{q}}_1^v, \quad \dot{\mathbf{q}}_2^v, \ \dots, \dot{\mathbf{q}}_{2^n}^v \}, \tag{5}$$

where $\dot{\mathbf{q}}_i^v$ denotes the i^{th} vertex of $\mathcal{Q}$. The convexity of a polytope is preserved under affine transformation, i.e., a linear transformation applied to $\mathcal{Q}^V$ is a convex combination of the same linear transformation applied to the vertices. Thus, a manipulability polytope (MP), denoted as $\mathcal{P}$, representing the Cartesian-space velocities can be obtained using the linear kinematic transform defined by one. $\mathcal{P}$'s vertex set representation is given as

$$\mathcal{P}^V = \{ \boldsymbol{v}_1^v \ \dots \boldsymbol{v}_{2^n}^v \} = \{ \mathbf{J}_n \dot{\mathbf{q}}_1^v \dots \mathbf{J}_n \dot{\mathbf{q}}_{2^n}^v \} \tag{6}$$

and its volume, denoted as w_p, can be used as an indicator of robot performance.

2.3 Constrained Manipulability

The manipulability polytope does not give a true picture of the robot's capabilities as they may be reduced due to environmental or joint limit constraints. Thus in our previous work (Long and Padır, 2018; Long and Padir, 2019), a method of considering obstacle and joint position limits is given. To do so the *kineostatic danger field* (Ragaglia et al., 2014) as an input which limits the maximum attainable velocity in the direction of a potential collision. The robot's velocity is reduced until the danger-field value is below a predefined threshold. The *kineostatic danger field* divides the robot's links into l control points (CPs) and the workspace into c cells. The danger field for the j^{th} $(j = 1, \dots, c)$ cell is calculated as

$$\phi_j = \max_{i=1\dots l} \left(\frac{1}{\left\| \mathbf{r}_i - \mathbf{r}_j \right\|} + \frac{||\mathbf{v}_i|| \cos\left(\angle \left(\mathbf{r}_i - \mathbf{r}_j, \mathbf{v}_i \right) \right)}{\left\| \mathbf{r}_i - \mathbf{r}_j \right\|^2} \right), \tag{7}$$

where $\mathbf{r}_j$ and $\mathbf{r}_i$ denote the position vector of the j^{th} cell and the robot's i^{th} CP, respectively. The translational velocity of point i is denoted by $\mathbf{v}_i$. To generate a set of inequality constraints, **Eq. 7** is re-defined as

$$\forall i \in \text{CP}, \phi_j \leq \frac{1}{\left\| \mathbf{r}_{ij} \right\|} + \frac{||\mathbf{v}_i|| \cos\left(\angle \left(\mathbf{r}_{ij}, \mathbf{v}_i \right) \right)}{\left\| \mathbf{r}_{ij} \right\|^2}, \tag{8}$$

where $\mathbf{r}_{ij} = \mathbf{r}_i - \mathbf{r}_j$. Substituting the dot product

$$\cos\left(\angle \left(\mathbf{r}_{ij}, \mathbf{v}_i \right) \right) = \frac{\mathbf{v}_i^T \mathbf{r}_{ij}}{||\mathbf{v}_i|| \left\| \mathbf{r}_{ij} \right\|}, \tag{9}$$

(8) becomes

$$\phi_j \leq \frac{1}{\left\| \mathbf{r}_{ij} \right\|} + \frac{\mathbf{v}_i^T \mathbf{r}_{ij}}{\left\| \mathbf{r}_{ij} \right\|^3}, \tag{10}$$

Finally, by introducing $\hat{r}_{ij}$ the normalized unit vector of $\mathbf{r}_{ij}$, **Eq. 10** becomes

$$\mathbf{v}_i^T \hat{\mathbf{r}}_{ij} \leq \phi_j \left\| \mathbf{r}_{ij} \right\|^2 - \left\| \mathbf{r}_{ij} \right\| \tag{11}$$

The robot's velocity is reduced until the danger field value at the obstacle location, denoted as o is below a threshold, i.e., a desired danger value. **Eq. 11** is re-written as

[1]We use the C++ wrapper for Fukuda's cdd library available here: https://github.com/vsamy/eigen-cdd, while the MATLAB© computations used the Multi-Parametric Toolbox 3.0 (Herceg et al., 2013).

$$\mathbf{v}_i^T \hat{\mathbf{r}}_{io} \leq \phi^d ||\mathbf{r}_{io}||^2 - ||\mathbf{r}_{io}||, \tag{12}$$

$\mathbf{r}_o$ is the obstacle's position vector with respect to the robot's fixed frame and ϕ^d denotes desired danger value. By introducing **Eq. 1**, the following expression is obtained in configuration space

$$\hat{\mathbf{r}}_{io}^T \mathbf{J}_i \dot{\mathbf{q}} \leq \phi^d \|\mathbf{r}_{io}\|^2 - \|\mathbf{r}_{io}\|, \tag{13}$$

where $\mathbf{J}_i \in \mathbb{R}^{3\times n}$ is the Jacobian matrix at point i. The Jacobian matrix at any CP will have n columns. However, if a joint does not contribute to the velocity at the i^{th} CP, the i^{th} column of contains only zeros. Hence, **Eq. 13** constrains the maximum velocity for the i^{th} CP in the direction toward the obstacle. Taking into account the l CPs leads to the following set of inequalities

$$\begin{bmatrix} \hat{\mathbf{r}}_{1o}^T \mathbf{J}_1 \\ \hat{\mathbf{r}}_{2o}^T \mathbf{J}_2 \\ \vdots \\ \hat{\mathbf{r}}_{lo}^T \mathbf{J}_l \end{bmatrix} \dot{q} \leq \begin{bmatrix} \phi^d \|\mathbf{r}_{1o}\|^2 - \|\mathbf{r}_{1o}\| \\ \phi^d \|\mathbf{r}_{2o}\|^2 - \|\mathbf{r}_{2o}\| \\ \vdots \\ \phi^d \|\mathbf{r}_{lo}\|^2 - \|\mathbf{r}_{lo}\| \end{bmatrix}, \tag{14}$$

rewritten, for the k^{th} obstacle, as

$$\mathbf{J}_{ok} \dot{\mathbf{q}} \leq \mathbf{b}_o. \tag{15}$$

Eq. 15 considers the reduced performance capabilities due to nearby obstacles. It is similarly convenient to consider the effects of joint limit proximity in the polytope before transformation to the Cartesian space, thus avoiding improper penalization due to redundancy (Tsai, 1986; Abdel-Malek et al., 2004). For the i^{th} joint, the penalization term is defined as

$$\psi_i^{max} = 1 - \left(\frac{\max(\overline{q}_i, q_i) - \overline{q}_i}{q_i^{max} - \overline{q}_i} \right)^k, \quad \psi_i^{min} = 1 - \left(\frac{\min(\overline{q}_i, q_i) - \overline{q}_i}{q_i^{min} - \overline{q}_i} \right)^k, \tag{16}$$

where $\overline{q}_i$ is given as $\overline{q}_i = \frac{1}{2}(q_i^{max} + q_i^{min})$, k is a positive integer, ψ_i^{max} varies from 1 to 0 as the i^{th} joint approaches its limit. **Eq. 4** is modified to consider the joint limits

$$\begin{bmatrix} \mathbf{I}_n \\ -\mathbf{I}_n \end{bmatrix} \dot{q} \leq \begin{bmatrix} \mathbf{\Psi}^{max} \dot{\mathbf{q}}_{max} \\ -\mathbf{\Psi}^{min} \dot{\mathbf{q}}_{min} \end{bmatrix}, \tag{17}$$

where $\mathbf{\Psi}^{max} = diag(\psi_1^{max} \dots \psi_n^{max})$ and $\mathbf{\Psi}^{min} = diag(\psi_1^{min} \dots \psi_n^{min})$.

By repeating **Eq. 15** for m obstacles and including the position constraints defined by **Eq. 17**, the following $\mathcal{H}$-representation of the joint-space polytope, denoted as $\mathcal{Q}^{H*}$, is obtained

$$\begin{bmatrix} \mathbf{J}_{o1} \\ \mathbf{J}_{o2} \\ \vdots \\ \mathbf{J}_{om} \\ \mathbf{I}_n \\ -\mathbf{I}_n \end{bmatrix} \dot{q} \leq \begin{bmatrix} \mathbf{b}_{o1} \\ \mathbf{b}_{o2} \\ \vdots \\ \mathbf{b}_{om} \\ \mathbf{\Psi}^{max} \dot{\mathbf{q}}_{max} \\ -\mathbf{\Psi}^{min} \dot{\mathbf{q}}_{min} \end{bmatrix}. \tag{18}$$

$\mathcal{Q}^{H*}$ can be converted to $\mathcal{V}$-representation using the double description method. The vertex form can then be transformed to the task space using **Eq. 1** and in doing so the CMP, denoted as $\mathcal{P}^*$ that characterizes the constrained task-space performance, is obtained. The volume of $\mathcal{P}^*$, denoted w_p^* measures the robot's velocity capacities while also considering joint position and velocity limits and the constraints imposed by the environment.

2.4 Applications of Constrained Manipulability Polytope for Humanoid Robots

In the following, applications for the performance measure are demonstrated using NASA's humanoid robot Valkyrie (Radford et al., 2015) interacting with a glovebox for nuclear decommissioning task.

2.4.1 Experiment 1: Right Arm Insertion

In the first experiment the right arm of Valkyrie is inserted into the glovebox. The glovebox is considered as an obstacle that reduces the manipulator's performance, as unwanted collision could be dangerous. **Figure 1** shows the right arm insertion task, demonstrating how the CMP changes with time. The initial reduction in manipulability is due to positional joint limits. As the right hand passes through the glovebox port the system experiences a reduction of velocity capacity due to the constrained space signifying that the hand cannot move quickly without increasing the likelihood of a collision. A partial recovery can be observed as the right hand is fully inserted, meaning the system can manipulate objects within the space.

2.4.2 Experiment 2: Dual Arm Insertion

Convex polytopes are geometric objects, thus can be combined through standard geometric operations. These combinations can be used to represent composite robotic chains both serial and parallel. In (Long and Padir, 2019), we have shown how the manipulability of mechanisms in series can be obtained from the *minkowski* sum of sub-mechanisms, while manipulability mechanisms in parallel can be obtained by a straightforward concatenation of inequality constraints.

In this experiment, we demonstrate the former, as Valkyrie inserts both arms into the glovebox. It is assumed that the two arms form a closed chain, in a scenario where the arms carry a common object or tool. The goal is to show how the CMP can be combined to obtain that of a closed chain system. To model the closed chain, we use the virtual object procedure, i.e., a rigid straight link extending from the left to the right hand (Long et al., 2015). **Figure 2** shows the motion associated with the dual-arm insertion task at four time instants and the CMP for the closed-chain system evaluated at the right-arm end effector. This is obtained by first calculating $\mathcal{P}_r$ and $\mathcal{P}_l$, then obtaining the intersection $\mathcal{P}_{r\cap l}$, while $\mathcal{P}_{r\cap l}^*$ is calculated in the same manner. In the third instant $\mathcal{P}_{r\cap l}^* = \varnothing$, as clearly it is impossible for the arms to enter through individual ports while holding a common object. In contrast, in the fourth instant, $\mathcal{P}_{r\cap l}^*$ is no longer empty demonstrating the ability to co-manipulate an object within the glovebox.

2.4.3 Experiment 3: Reachability Analysis

Finally, a reachability study/workspace analysis is presented in **Figure 3**. The environment is discretized into voxels. At each voxel, an optimization procedure obtains a feasible IK solution while trying to maximize the robot's distance to obstacles. The CMP is calculated in this configuration. The workspace

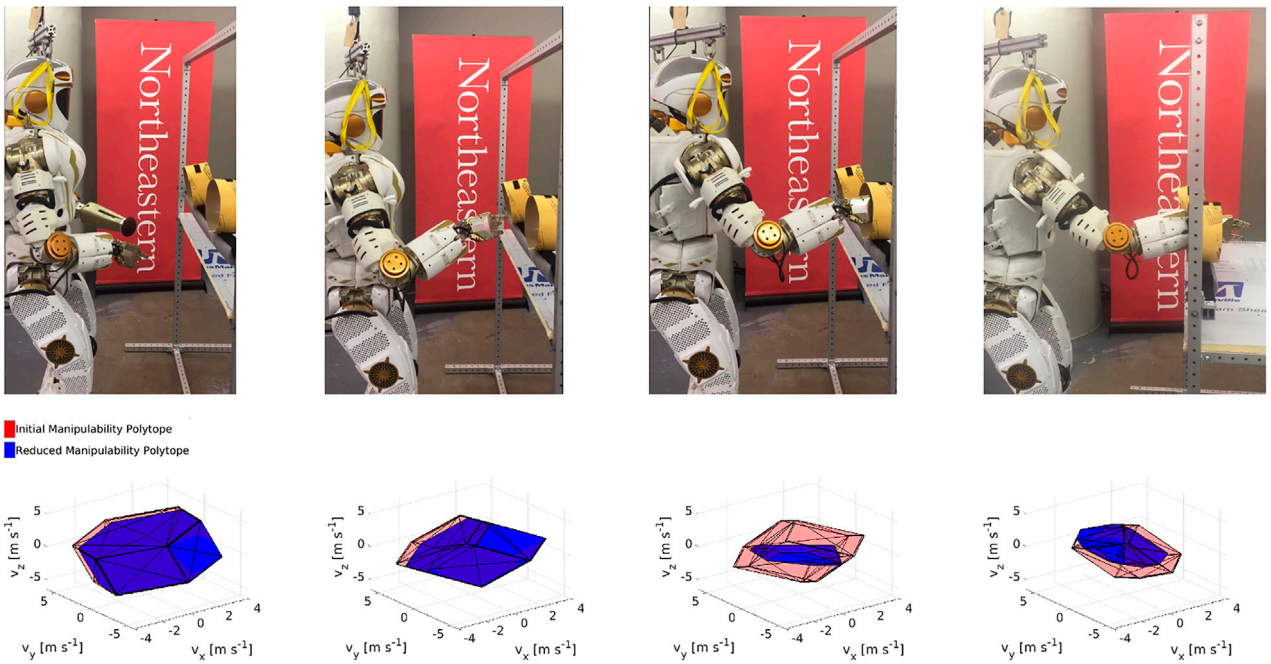

FIGURE 1 | Valkyrie inserting right arm into a glovebox and the evaluation of the manipulability polytope (red) and constrained manipulability polytope (blue) of the right arm. A video of the task is available here: https://www.youtube.com/watch?v=FzlhsLH5IPU.

discretization is shown in **Figure 3**. The voxel's color is defined by the volume of w_p^*, red implys high volume while blue implies empty set. **Figure 3B** shows the volume of $\mathcal{P}_r$ and $\mathcal{P}_l$ along the $x-axis$, i.e., along the centerline of the glovebox ports, while **Figure 3C** shows the reduced volume for $\mathcal{P}_r^*$ and $\mathcal{P}_l^*$. The increase in manipulator capacities can be observed as the arms align with the glovebox ports.

3 INCREASING STABILITY USING SUPPORT CONTACTS

As nuclear facilities reach the end of their life cycle they must be decommissioned in a safe and efficient manner. A particularly dangerous task is the decontamination of gloveboxes that have been previously used to manipulate radioactive material. Although a robotic system that is specifically designed for glovebox operations may be the best solution, humanoid robots are an attractive option since they can operate in a variety of environments and use tools that are designed for humans. While conducting operations within the glovebox, the constraints imposed by the ports, gloves, and the external structure, which effectively fix the arms at the entry points, must be considered. The inability to alter body configuration greatly diminishes the robot's capacity to take steps in arbitrary directions. This in turn leads to a danger of toppling during task execution as the system cannot easily change the support polygon's location.

Toppling occurs when the ZMP leaves the support polygon (SP) (Vukobratović and Borovac, 2004). If the SP cannot be displaced, alternative methods to maintain stability must be employed. For example, in (Rasheed et al., 2018), the ZMP position for a cable-driven mobile robot is modified online by a tension distribution algorithm. In (Khatib and Chung, 2014), it is shown that the SP size can be increased by using supplementary contact points. Similar to these approaches, we propose to exploit the contacts in the glovebox (i.e., leaning on the entry ports) in order to shift the ZMP towards the center of the SP while performing manipulation tasks.

This section presents our work on planning kinematic motions with support contacts without a predefined contact schedule that maintains the stability of the ZMP. To accomplish this, we model rigid-body contacts using complementarity constraints and solve a nonlinear constrained optimization for joint velocities and contact forces. Owing to the differentiable contact model, gradient-based optimization can reason about contacts between the robot arms and the glovebox ports. This optimization also respects constraints that ensure an object is grasped by the end effectors, the ZMP is in a safe region, and the deviation of the object's position from a desired position is acceptable. Furthermore, we present a null-space-based torque controller that prioritizes the stability, i.e., generating the support forces, and projects the torques needed for the manipulation task onto the null space of the support forces. The proposed methodology is tested through 2.5D, quasi-static simulations by considering a humanoid robot with two planar arms manipulating a relatively heavy object on an elevated plane representing the glovebox.

3.1 Related Work

For planning a motion with contact interactions, both the discrete contact events (e.g., making/breaking contacts at certain

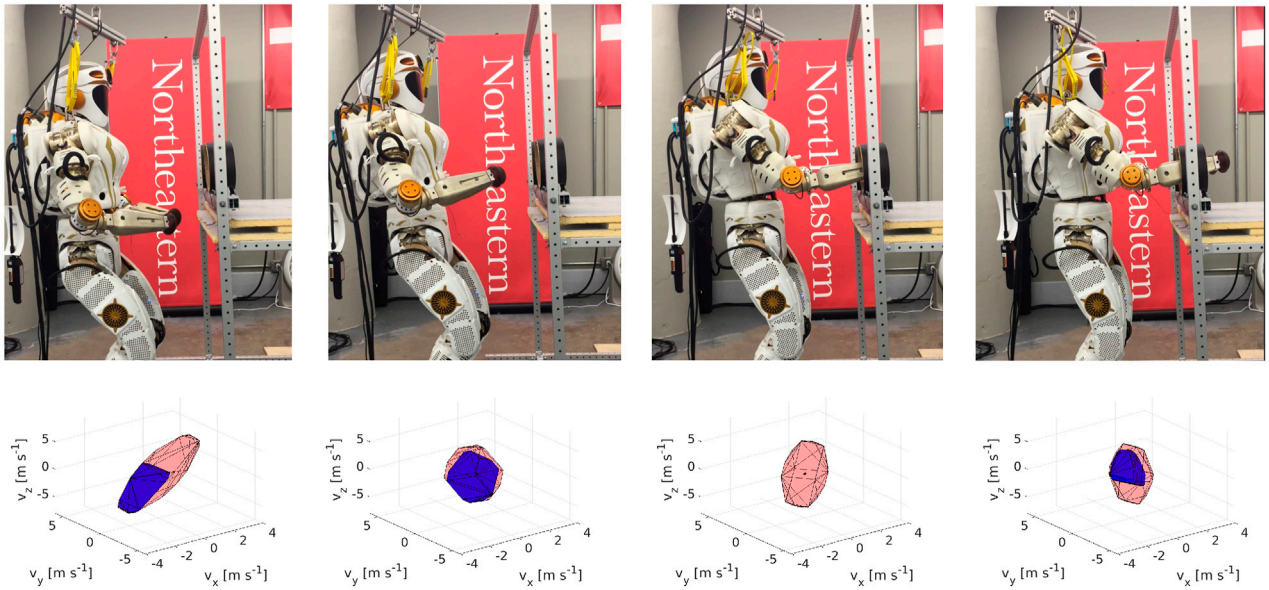

FIGURE 2 | Valkyrie inserting both arms into a glovebox shown at four timesteps along with the coordinating polytopes, $\mathcal{P}_{r\cap l}$ and $\mathcal{P}^*_{r\cap l}$ evaluated at the right end effector. A video of the task is available here: https://youtu.be/1Nouc4f_rlY.

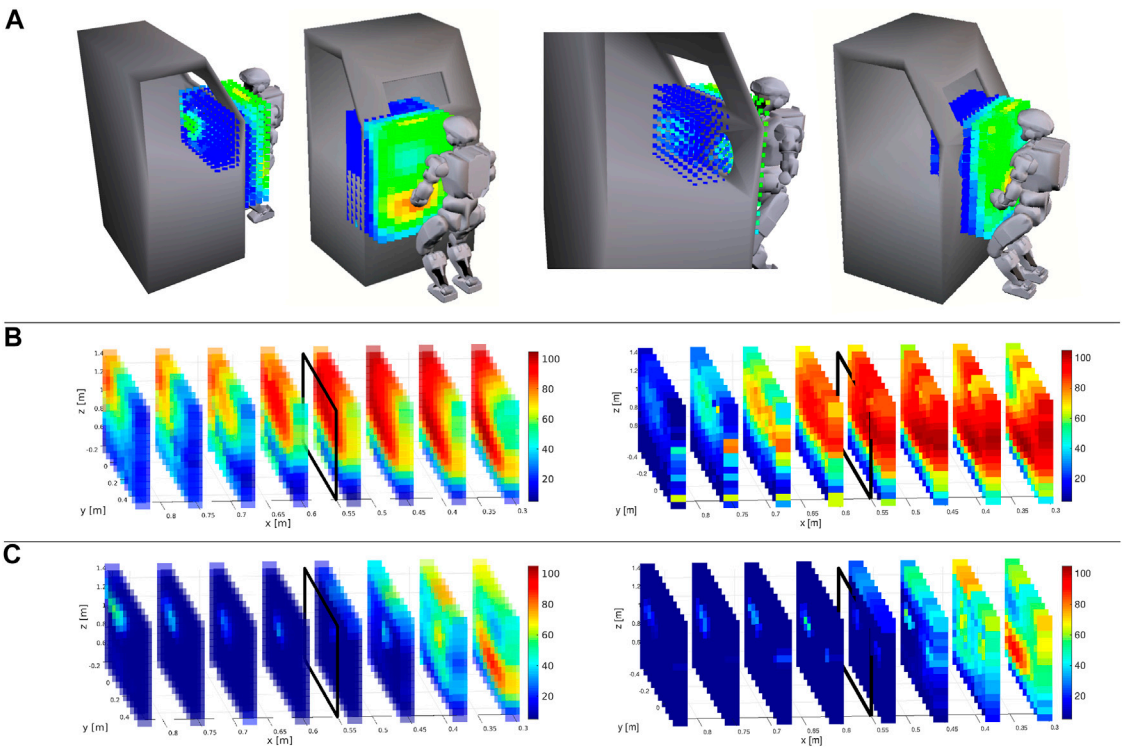

FIGURE 3 | The space is discretized into 3D voxels. **(A)** At each voxel, the IK solution is obtained for the left arm **(left images)** and the right arm **(right images)**. The corresponding volume of the CMP is calculated for each voxel, giving a good understanding of the robot's workspace. A video is available here: https://youtu.be/jc7X4WakdoE **(B)** The MP's volume w_p, i.e, for the left **(left image)** and the right **(right image)** end effectors. The black square shows the location of the glovebox front edge. **(C)** The CMP's volume w^*_p for the left **(left image)** and the right **(right image)** end effectors. High manipulability is possible far from the glovebox, the manipulability is extremely limited once either arm enters the glovebox.

locations) and the continuous variables (e.g., joint positions, contact forces, stability constraints) must be considered. One approach is to use a contact-before-motion planner such as those presented in (Hauser et al., 2005; Escande et al., 2013). In this case, first, a sequence of contacts at predefined locations is determined, then a continuous motion is planned subject to contact constraints. In contrast, in a motion-before-contact planner, the contacts are obtained as a result of the motion planning (Escande et al., 2013). Alternatively, contact-implicit motion planning (also known as, motion planning through contacts) can be used to plan for smooth motions and contact events at the same level.

In contact-implicit motion planning, a differentiable contact model is used to enable gradient-based optimization to reason about contacts. Complementarity constraints are widely used to model rigid-body contacts with friction, as proposed in (Stewart and Trinkle, 1996; Anitescu and Potra, 1997). (Yunt and Glocker, 2005; Posa et al., 2014) use complementarity constraints to model contacts in a trajectory optimization problem. The main idea here is to consider the contact-related parameters as additional optimization variables such that the contact events evolve along with continuous motion variables. Such an optimization problem can be solved locally through constrained nonlinear optimization algorithms such as sequential quadratic programming (Anitescu and Potra, 1997; Fletcher and Leyffer, 2004).

Once a kinematic motion with contact interactions is planned, it can be executed using a null-space-based controller. Park and Khatib (Park and Khatib, 2006) proposed a torque control framework for humanoid robots with multiple contacts and verified the method experimentally in (Park and Khatib, 2008). Moreover, they extended this method to a unified hierarchical whole-body control framework for humanoid robots in (Khatib et al., 2008). In this framework, tasks are hierarchically ranked. Thus, the torques required for a lower-priority task are projected onto the null-space of the Jacobian matrix associated with a higher-priority task. In (Henze et al., 2016a), a whole-body torque controller for humanoid robots is proposed that combines passivity-based balancing proposed in (Henze et al., 2016b) with a hierarchical null-space-based control that is similar to (Khatib et al., 2008).

3.2 Methodology

3.2.1 Static Equilibrium

For being balanced, the robot needs to be in static equilibrium (Vukobratović and Borovac, 2004). In this case, the static equilibrium of the system can be evaluated considering the following wrenches: the wrench due to the robot's mass, the wrenches at the end effectors due to the object wrench, and the wrenches at the support contact points. Henceforward, we enumerate the left and right arms as the first and second arms, respectively.

In the static equilibrium, the net force must be zero:

$$\sum \mathbf{f} = 0 \Rightarrow \mathbf{f}_r + m\mathbf{g} + \sum_{i=1}^{2} \mathbf{f}_{s_i} + \sum_{i=1}^{2} \mathbf{f}_{c_i} = 0, \tag{19}$$

where m is the total mass of the robot, $\mathbf{g} \triangleq [0, 0, -g]^T$ is the gravity vector, $\mathbf{f}_{s_i} \in \mathbb{R}^3$ is the force at the support point between the i^{th} arm and the glovebox port, $\mathbf{f}_{c_i} \in \mathbb{R}^3$ is the force at the contact point between the object and the i^{th} end effector, and $\mathbf{f}_r \in \mathbb{R}^3$ is the ground reaction force.

Additionally, the projection of the net moment, $\mathbf{M}$ onto the horizontal xy plane must be zero, i.e., $M^x = 0$ and $M^y = 0$:

$$\sum \mathbf{M}^H = 0 \Rightarrow (\mathbf{p}_r \times \mathbf{f}_r)^H + (\mathbf{p}_{CoM} \times m\mathbf{g})^H + \sum_{i=1}^{2} \left(\mathbf{p}_{s_i} \times \mathbf{f}_{s_i} + \mathbf{M}_{s,i}\right)^H + \sum_{i=1}^{2} \left(\mathbf{p}_{c_i} \times \mathbf{f}_{c_i} + \mathbf{M}_{c,i}\right)^H = 0, \tag{20}$$

where $\mathbf{a}^H$ denotes the horizontal projection of a vector $\mathbf{a}$, $\mathbf{p}_r, \mathbf{p}_{CoM}, \mathbf{p}_{s_i}, \mathbf{p}_{c_i} \in \mathbb{R}^3$ are the positions of the ground reaction force (i.e., the ZMP), the robot's center of mass (CoM), the support points on the glovebox ports and the contact points on the object, with respect to the world frame. $\mathbf{M}_{s,i}, \mathbf{M}_{c,i} \in \mathbb{R}^3$ are the moments at the support and grasp points. The position of the ZMP, $\mathbf{p}_r$, is obtained by solving (**Eq. 19**, **20**) simultaneously. In order to avoid toppling, the ZMP must lie in the support polygon (SP), namely, the convex hull of the robot's feet. The object wrench $\mathbf{h}_o \in \mathbb{R}^6$ can be obtained in terms of the wrenches applied by the end effectors as follows (Caccavale and Uchiyama, 2016):

$$\mathbf{h}_o = \begin{bmatrix} \mathbf{W}_{c_1} & \mathbf{W}_{c_2} \end{bmatrix} \begin{bmatrix} \mathbf{h}_{c_1} \\ \mathbf{h}_{c_2} \end{bmatrix} = \mathbf{W}\mathbf{h}_c, \tag{21}$$

using the wrench matrix $\mathbf{W}_{c_i} \in \mathbb{R}^{6\times6}$ that transforms the wrench at the i^{th} contact point, $\mathbf{h}_{c_i} \in \mathbb{R}^6$, to the wrench at the origin of the object frame, that is the center of the object in this case, and given by:

$$\mathbf{W}_{c_i} = \begin{bmatrix} \mathbf{I}_3 & 0_3 \\ -\widehat{\mathbf{r}}_{c_i} & \mathbf{I}_3 \end{bmatrix}, \tag{22}$$

where $\widehat{\mathbf{r}}_{c_i}$ is the skew-symmetric matrix representation of the vector from the i^{th} contact point c_i to the origin of the object frame, and $\mathbf{I}_3$ and 0_3 are 3×3 identity and zero matrices. Then, given the object wrench, the wrenches at the end effectors can be calculated from $\mathbf{h}_c = \mathbf{W}^+\mathbf{h}_o$, where $\mathbf{W}^+$ is the Moore-Penrose pseudo-inverse of the matrix $\mathbf{W} \in \mathbb{R}^{6\times12}$.

3.2.2 Motion Planning

In this work, we ignore the dynamic effects and investigate the quasi-static case for dual-arm manipulation of an object in a confined space, i.e., a glovebox. In the following, robot's joint positions and velocities are denoted by $\mathbf{q}$ and $\dot{\mathbf{q}}$, while those of the i^{th} arm are referred to as $\mathbf{q}_i$ and $\dot{\mathbf{q}}_i$. The objective is to preserve the robot's balance during the manipulation task. In other words, our goal is to find the joint positions that would keep the robot's ZMP in a safe region by leaning on the glovebox ports while simultaneously maintaining the manipulated object's desired position. For this purpose, we form a nonlinear constrained

optimization and solve it for the joint displacements and the support contact forces.

In order to take into account the rigid-body contacts between the robot and the glovebox ports, we use the following complementarity constraints, as in (Anitescu and Potra, 1997; Posa et al., 2014):

$$\phi(\mathbf{q}) \geq 0, \tag{23a}$$

$$\boldsymbol{\gamma} \geq 0, \tag{23b}$$

$$\boldsymbol{\gamma}^T \phi(\mathbf{q}) = 0, \tag{23c}$$

where $\phi(\mathbf{q}) \in \mathbb{R}^{n_p}$ is the vector of signed distance for n_p contact pairs, i.e., each pair comprises of a robot's link and a contact candidate in the environment; and $\boldsymbol{\gamma} \in \mathbb{R}^{n_p}$ is the vector of the magnitude of normal support force. (**Eq. 23a**) prevents any interpenetration, (**Eq. 23b**) ensures that the bodies can only push each other, and (**Eq. 23c**) allows force generation only when bodies are in contact. Thus, only one of these variables (either $\phi(\mathbf{q})$ or $\boldsymbol{\gamma}$) can be non-zero for a given time and contact pair. We relax the complementarity condition (**Eq. 23c**) by converting it into an inequality constraint through a slack variable and penalizing the slack variable in the cost so that potential numerical issues are mitigated, as described in (Fletcher and Leyffer, 2004; Manchester and Kuindersma, 2017). For numerical efficiency, the complementarity constraints, including the relaxation, are evaluated elementwise, i.e., separately for each contact pair.

As a result, the following optimization problem is solved for the joint displacements $\Delta\mathbf{q} \triangleq \mathbf{q}_{k+1} - \mathbf{q}_k$, the magnitudes of the normal contact forces at the support points $\boldsymbol{\gamma}$, and the slack variables $\mathbf{s}$:

$$\underset{\Delta\mathbf{q},\boldsymbol{\gamma},\mathbf{s}}{\text{minimize}} \quad w_1 \left\|\mathbf{p}_o^e\right\|^2 + w_2 \left\|\Delta\mathbf{q}\right\|^2 + w_3 \|s\|^2 \tag{24a}$$

subject to :

$$\phi(\mathbf{q}), \boldsymbol{\gamma}, \mathbf{s} \geq 0, \tag{24b}$$

$$\boldsymbol{\gamma}^T \phi(\mathbf{q}) \leq \mathbf{s}, \tag{24c}$$

$$c_g(\mathbf{q}) = 0, \tag{24d}$$

$$\left\|\mathbf{p}_r^d - \mathbf{p}_r\right\| \leq r_s, \tag{24f}$$

$$\left\|\mathbf{p}_o^e\right\| \leq r_o, \tag{24g}$$

where $\|\cdot\|$ is the Euclidean norm, w_i is the weight (a positive scalar) associated with the i^{th} term of the cost function, $\mathbf{p}_o^e$ is the deviation of the object's position from the desired position, $c_g(\mathbf{q}) = 0$ ensures that the end effectors are grasping the object, $\mathbf{p}_r^d$ is the desired position of the ZMP (i.e., the center of the SP), and r_s and r_o are the radii of the safe circle (SC) for the ZMP and the admissible sphere for the object position.

3.2.3 Torque Control

Using this optimization procedure, the robot configuration and the support forces' magnitude are obtained. Nevertheless, a torque controller is necessary to execute the planned motions.

The torques necessary to generate the desired object wrench $\boldsymbol{\tau}_h$ can be obtained as:

$$\boldsymbol{\tau}_h = \begin{bmatrix} \mathbf{J}_1^T & 0 \\ 0 & \mathbf{J}_2^T \end{bmatrix} \mathbf{W}^+ \mathbf{h}_o = \mathbf{J}^T \mathbf{W}^+ \mathbf{h}_o \tag{25}$$

$\mathbf{J}_i \in \mathbb{R}^{3\times4}$ is the kinematic Jacobian matrix that maps the joint velocities to the translational end-effector velocities for the i^{th} arm.

The support forces are oriented normal to the contacting robot geometry. Hence, using the contact angle β_i and the normal force magnitude $\boldsymbol{\gamma}$, the support force for the i^{th} arm can be calculated as:

$$\mathbf{f}_{s_i} = \begin{bmatrix} \gamma_i cos(\beta_i) & \gamma_i sin(\beta_i) & 0 \end{bmatrix}^T. \tag{26}$$

Similarly, the joint torques required to generate these forces, denoted as $\boldsymbol{\tau}_s$, can be calculated as described in (Park and Khatib, 2008). In our case, there is a maximum of two supports points at a given time, therefore:

$$\boldsymbol{\tau}_s = \begin{bmatrix} \mathbf{J}_{s_1}^T & 0 \\ 0 & \mathbf{J}_{s_2}^T \end{bmatrix} \begin{bmatrix} \mathbf{f}_{s_1} \\ \mathbf{f}_{s_2} \end{bmatrix} = \mathbf{J}_s^T \mathbf{f}_s. \tag{27}$$

$\mathbf{J}_{s_i} \in \mathbb{R}^{3\times4}$ is the Jacobian matrix that maps the joint velocities of the i^{th} arm to the translational velocities at the support point.

For the glovebox task, the support forces are crucial to maintain the stability of the robot, while generating the desired object wrench has a lower priority. Thus, we compose the joint torques, $\boldsymbol{\tau}$ such that the manipulation torques $\boldsymbol{\tau}_h$ are projected onto the null space of the stability torques:

$$\boldsymbol{\tau} = \boldsymbol{\tau}_s + \mathbf{N}_s \boldsymbol{\tau}_h, \tag{28}$$

where

$$\mathbf{N}_s = \mathbf{I}_{n_d} - \mathbf{J}_s^T \left(\mathbf{J}_s^{+T}\right) \tag{29}$$

is the null space projector of the support forces, and n_d is the DOF of the whole robot. Consequently, the resulting joint torques would generate the desired support forces to ensure the balance of the robot and create an object wrench using the redundancy of the robot.

3.3 Results

To test the proposed framework, we run simulation experiments in which a humanoid robot that has two planar 4-DOF arms with revolute joints manipulates a relatively heavy rigid bar on an elevated plane. The robot's arms pass through two ports representing the glovebox. We neglect the dynamics (i.e., velocities and accelerations) and assume point contacts without friction. The weights are selected as $w_1 = 10^3$, $w_2 = 10^2$, and $w_3 = 10^6$. The initial values for all the decision variables are zero. The radii of tolerance circles for the ZMP and the object position are selected as $r_s = 0.15$ m and $r_o = 0.1$ m. The masses of the robot and the object are 54 kg and 12 kg. The desired motion of the object is in $+y$ direction; hence, the desired object wrench to generate an acceleration in this direction is given as $\mathbf{h}_o = [0, 10, -117.72, 0, 0, 0]^T$.

We investigate the task of moving the object 40 cm forward on a straight path that consists of nine equally spaced waypoints. The results are depicted in **Figure 4**. Each step of the motion is indicated by a color from blue to red. In the initial configuration (indicated by blue), the robot grasps the object from both ends. During the simulation, the position of the ZMP is calculated with and without the effect of the support contacts on the glovebox

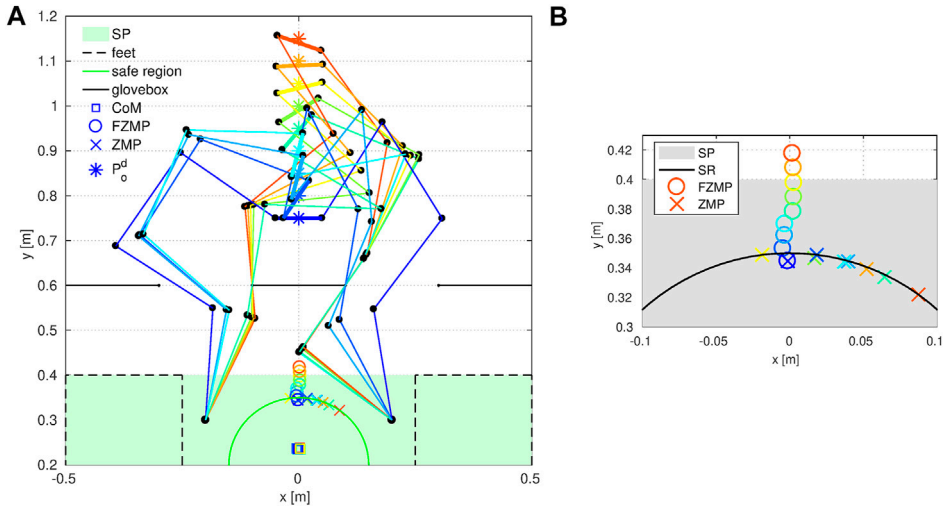

FIGURE 4 | (A) Planned motion for carrying an object on a straight path. **(B)** Zoomed in change of the ZMP and the FZMP (i.e., fictitious ZMP without supporting contacts). throughout the task.

frame. The latter is known as the fictitious ZMP (FZMP) (Vukobratović and Borovac, 2004) and may fall outside of the SP.

The results show that the contact-implicit motion planning method can increase the stability of the robot using support contacts while performing the manipulation task. The robot makes contacts with the glovebox to maintain its balance while moving the object on the desired path. As soon as the FZMP leaves the SR in the second step, the right arm makes a contact with the left end of the port to push the ZMP into the SR. As the object moves further away from the base, the contact angle is varied so that the magnitude of the support force in $-x$-direction is larger. This is required due to the circular shape of the SR. However, as object moves further from the base, simply changing the contact angle is no longer sufficient, thus the left arm also makes contact with the right end of the left port. As a result, the object is successfully transported along the desired path with a position error of 0.1 m (i.e., the allowed deviation) in each step after the initial configuration.

A zoomed-in version of the SP area is depicted in **Figure 4B** to show the change of the ZMP and the FZMP throughout the simulation. Even though the FZMP moves forward along with the object's position and leaves the SP eventually, the ZMP does not leave the SR owing to the support forces. It is also noted that, in Step 7 (indicated by yellow), the ZMP is more centralized with respect to the y axis compared to the other steps due to the symmetry of the support forces.

Figure 5 show the magnitudes of the support forces and the joint torques with respect to the distance between the object and the robot's base. The magnitudes of the support forces are much larger than the magnitude of the object wrench, and therefore the torques are much more affected by the support forces than the object wrench. This is why force and torque vs. distance characteristics are quite similar—i.e., the torque is dominated by the support forces (especially after Step 5). Moreover, the changes of $\|\mathbf{f}_s\|$ and $\|\boldsymbol{\tau}\|$ with the distance are almost linear, as one may anticipate. Apart from this, the magnitude of the support force on each arm is quite similar to each other in Step 7, as consistent with the observation regarding the more centralized ZMP in this step. Except for Steps 1 and 7, the magnitude of the support force on the right arm is always bigger than the one on the left arm, which shifts the ZMP in $+x$ direction. Such an unbalanced distribution of forces might be undesirable since higher joint torque limits would be required. Thus, enforcing a more uniform distribution of the support forces may be a future work.

4 DYNAMIC NON-PREHENSILE MANIPULATION

There is an increasing need to carry out decontamination and decommissioning tasks in safe and effective manner. A particularly dangerous task is glovebox decontamination and decommissioning that typically involves transporting debris and objects from the interior of the glovebox to an exit port, where they are bagged and removed (Long et al., 2018; Önol et al., 2018). Such tasks do not always require dexterous manipulation behaviors and instead simply require objects to be push from the interior to the exit port of a glovebox.

Contact-implicit trajectory optimization (CITO) is a promising method to generate contact-rich behaviors given only a high-level definition of a task. In this approach, a differentiable contact model is used to enable gradient-based optimization to reason about contacts such that discrete contact events and continuous trajectories are found simultaneously as a result of smooth optimization.

In this section, we present a CITO method based on a variable smooth contact model to plan dynamic non-prehensile manipulation behaviors for a 7-DOF robot arm in a highly-constrained environment. We demonstrate that the proposed

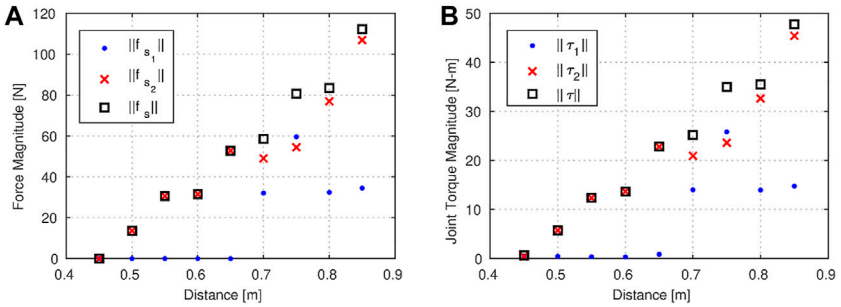

FIGURE 5 | The magnitudes of **(A)** the support forces vs. the distance of the object from the base, and **(B)** the joint torques vs. the distance of the object from the base.

method can solve complex tasks despite tight constraints imposed by the environment by exploiting the smooth virtual forces. Moreover, we experimentally verify that the physical inaccuracy introduced by the residual virtual forces is admissible and the motions found by this framework are realistic enough to be run on the hardware.

4.1 Related Work

Complementarity constraints are widely used to model rigid-body contacts in trajectory optimization (Yunt and Glocker, 2005; Posa et al., 2014; Gabiccini et al., 2018). This approach can find complex motions, but it typically suffers from poor convergence speed. Thus (Tassa et al., 2012; Mordatch et al., 2015; Mastalli et al., 2016; Neunert et al., 2016; Manchester and Kuindersma, 2017), use smoother fragments of the complementarity constraints. (Neunert et al., 2017; Giftthaler et al., 2017; Marcucci et al., 2017; Neunert et al., 2018), on the other hand, define contact forces as a smooth function of distance, i.e., a smooth contact model. Using such a contact model, highly-dynamic complex motions for a quadruped robot are planned and executed in real-time in (Neunert et al., 2018). However, it is difficult to tune smooth contact models (Carius et al., 2018), and the resulting motions may be physically inaccurate due to the non-physical contact forces that act from distance. In order to address these problems, we have recently proposed a variable smooth contact model (VSCM) (Önol et al., 2018) that injects virtual forces to the underactuated dynamics with frictional rigid-body contact mechanics, such that the states of the manipulator and the objects are coupled in a smooth way. Furthermore, the smoothness of the contact model is adjusted by optimization such that large virtual forces are permitted in the initial phases of optimization but vanish as the optimization converges. As a result, the VSCM improves the convergence of CITO without compromising the physical fidelity of resulting motions.

CITO has been used for animated characters (Mordatch et al., 2021a; Mordatch et al., 2012b) and in robotics (Tassa et al., 2012; Posa et al., 2014; Mordatch et al., 2015; Mastalli et al., 2016; Manchester and Kuindersma, 2017; Neunert et al., 2017; Carius et al., 2018; Neunert et al., 2018; Winkler et al., 2018). Although this method is task independent and can be generalized to both locomotion and manipulation problems, the majority of the related literature considers only the former. On the other hand, in (Mordatch et al., 2012a; Posa et al., 2014; Gabiccini et al., 2018), manipulation tasks are investigated but their analyses are either limited to a planar case or based on animated characters where physical fidelity is not critical. Recently (Önol et al., 2018, 2019, 2020; Sleiman et al., 2019), used CITO for non-prehensile manipulation tasks. Yet, they consider only tabletop pushing scenarios. Moreover, in general, experimental results in this domain are very limited, albeit with some notable exceptions (Mordatch et al., 2015; Mastalli et al., 2016; Neunert et al., 2017, 2018; Giftthaler et al., 2017; Carius et al., 2018; Winkler et al., 2018; Sleiman et al., 2019). Nonetheless, to the best of our knowledge, there is no experimental verification of CITO for constrained dynamic manipulation.

4.2 Methodology

4.2.1 Dynamic Model

The dynamics of an underactuated system consisting of an n_a-DOF manipulator and n_u-DOF objects that are subject to frictional rigid-body contacts and virtual forces is given by

$$\mathbf{M}(\mathbf{q})\ddot{\mathbf{q}} + \mathbf{c}(\mathbf{q}, \dot{\mathbf{q}}) = \mathbf{S}_a^T \boldsymbol{\tau} + \mathbf{J}_c^T(\mathbf{q})\boldsymbol{\lambda}_c + \mathbf{S}_u^T \boldsymbol{\lambda}_v, \tag{30}$$

where $\mathbf{q} \triangleq [\mathbf{q}_a^T, \mathbf{q}_u^T]^T \in \mathbb{R}^{n_a+n_u}$ is the configuration vector; $\mathbf{M}(\mathbf{q}) \in \mathbb{R}^{(n_a+n_u)\times(n_a+n_u)}$ is the mass matrix; $\mathbf{c}(\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^{n_a+n_u}$ is the bias term comprising of the Coriolis, centrifugal, and gravitational effects; $\mathbf{S}_a = [\mathbb{I}_{n_a\times n_a}\, 0_{n_a\times n_u}]$ is the selection matrix for the actuated DOF and $\mathbf{S}_u = [0_{n_u\times n_a}\, \mathbb{I}_{n_u\times n_u}]$ is the selection matrix for the unactuated DOF; $\boldsymbol{\tau} \in \mathbb{R}^{n_a}$ is the vector of generalized joint forces; $\boldsymbol{\lambda}_c \in \mathbb{R}^{6n_c}$ is the vector of generalized contact forces at n_c external contact points and $\mathbf{J}_c(\mathbf{q}) \in \mathbb{R}^{6n_c\times(n_a+n_u)}$ is the Jacobian matrix mapping the joint velocities to the Cartesian velocities at the contact points, and $\boldsymbol{\lambda}_v \in \mathbb{R}^{n_u}$ is the vector of generalized virtual forces on the unactuated DOF. For n_f free bodies in SE (**Eq. 3**) (e.g., objects), $n_u = 6n_f$. The state of the system is represented by $\mathbf{x} \triangleq [\mathbf{q}^T\, \dot{\mathbf{q}}^T]^T \in \mathbb{R}^n$ where $n = 2(n_a + n_u)$.

In this study, $\boldsymbol{\tau}$ is decomposed as $\boldsymbol{\tau} = \boldsymbol{\tau}_u + \tilde{\mathbf{c}}$, where $\tilde{c} \in \mathbb{R}^{n_a}$ is an estimation of the non-zero part of $\mathbf{S}_a^T \mathbf{c}(\mathbf{q}, \dot{\mathbf{q}})$ and $\boldsymbol{\tau}_u \in \mathbb{R}^{n_a}$ is the vector of control variables in terms of generalized joint forces. As a result, the control term $\boldsymbol{\tau}_u$ is linearly related to the joint accelerations in the absence of external contact.

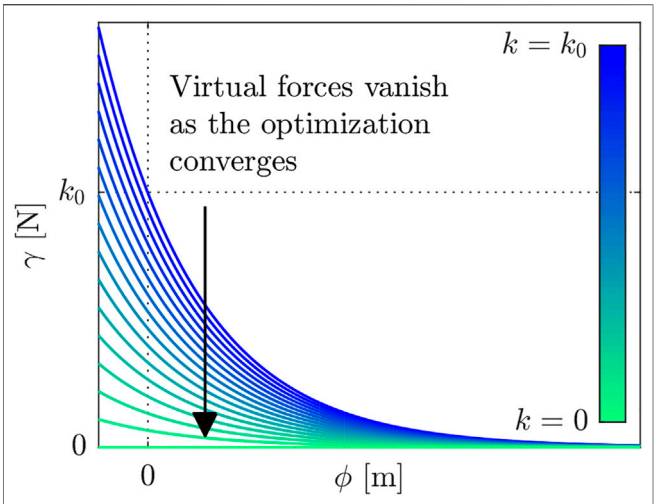

FIGURE 6 | The relationship between the virtual force and the distance. k_0 represents the initial value of the virtual stiffness k, and the arrow shows the evolution of the virtual forces throughout optimization.

4.2.2 Contact Model

The virtual forces generated by the contact model acts upon the unactuated DOF in addition to the external rigid-body contacts. Consequently, the robot's and objects' dynamics are related through the virtual forces. We assume an exponential relationship between the magnitude of the normal contact force γ and the signed distance between paired contact geometries ϕ, as depicted in **Figure 6**. While frictional forces are not considered in this contact model, the rigid-body contact mechanics [i.e., $\boldsymbol{\lambda}_c$ in (**Eq. 30**)] are frictional. Hence, the resulting motions include frictional contacts once the virtual forces vanish.

For the i^{th} contact candidate, the magnitude of the normal contact force is calculated from the virtual stiffness k_i and α_i that determines the curvature:

$$\gamma_i(\mathbf{x}) = k_i e^{-\alpha_i \phi_i(\mathbf{x})}. \tag{31}$$

This model is analogous to a spring model and (Marcucci et al., 2017) lists several reasons for not using damping (i.e., a velocity term) in such a contact model. The corresponding virtual force effective at the center of mass of the free body associated with the contact candidate $\boldsymbol{\lambda}_{v,i} \in \mathbb{R}^6$ is:

$$\boldsymbol{\lambda}_{v,i}(\mathbf{x}) = \gamma_i(\mathbf{x}) \begin{bmatrix} \mathbb{I}_3 \\ \widehat{\mathbf{l}}_i \end{bmatrix} \mathbf{n}_i(\mathbf{x}), \tag{32}$$

where $\mathbb{I}_3$ is 3×3 identity matrix; $\mathbf{l}_i$ is the vector between the end effector and the center of mass of the object that is associated with the contact candidate; $\widehat{\mathbf{l}}_i$ is the skew-symmetric matrix form of $\mathbf{l}_i$; and $\mathbf{n}_i \in \mathbb{R}^3$ is the unit vector that is normal to the contact surface on the object. Hence, the net virtual force on an object is the sum of the virtual forces associated with the contact candidates on that object.

In the variable smooth contact model, the virtual stiffness k for each time step and contact pair is a decision variable of optimization and initialized with a large value such that there is a non-zero virtual force on each contact candidate. Nonetheless, the virtual forces are penalized as an integrated cost, so that they vanish as the optimization converges, see **Figure 6**. This approach helps to discover contact candidates that are initially distant.

4.2.3 Trajectory Optimization

In this study, the optimal control problem is transcribed into a finite-dimensional nonlinear constrained optimization by assuming constant control inputs over N discretization intervals. Final cost terms penalize the deviations of the objects' poses from desired poses, $\mathbf{p}_o^e$ and $\boldsymbol{\theta}_o^e$. Integrated cost terms are defined in terms of the velocities $\dot{\mathbf{x}}$ and the virtual forces $\boldsymbol{\gamma}$. As a result, the final and integrated components of the cost (c_F and c_I) are calculated in terms of the weights $w_{1,\ldots,4}$, the control sampling period t_c by:

$$c_F = w_1 \left\| \mathbf{p}_o^e \right\|^2 + w_2 \left\| \boldsymbol{\theta}_o^e \right\|^2, \tag{33a}$$

$$c_I = t_c \sum_{i=1}^{N} \left(w_3 \left\| \dot{\mathbf{x}}_i \right\|^2 + w_4 \left\| \boldsymbol{\gamma}_i \right\|^2 \right). \tag{33b}$$

The following optimization problem is solved by a sequential quadratic programming (SQP) algorithm by rolling out the dynamics:

$$\underset{\boldsymbol{\tau}_{u,1,\ldots,N},\, \mathbf{k}_{1,\ldots,N}}{\text{minimize}} \quad c_F + c_I \tag{34a}$$

$$\text{subject to :} \quad \boldsymbol{\tau}_{u,L} \le \boldsymbol{\tau}_{u,1,\ldots,N} \le \boldsymbol{\tau}_{u,U},\ 0 \le \mathbf{k}_{1,\ldots,N} \le \mathbf{k}_0. \tag{34b}$$

The lower and upper bounds for the control variables $\boldsymbol{\tau}_{u,L}$ and $\boldsymbol{\tau}_{u,U}$ are determined from the torque limits of the robot. However, it is noted that the bias in the torque decomposition, $\widehat{\mathbf{c}}$, is not considered while setting the torque limits. The virtual stiffness variables are bounded above by their initial values $\mathbf{k}_0$, which is selected as a large value to facilitate convergence.

4.3 Experiments

4.3.1 Application: Non-Prehensile Manipulation in a Glovebox

Our overall objective is to enable a robot to carry out such manipulation tasks with only high-level commands such as desired object poses. **Figure 7D** shows Sawyer, a 7-DOF robot arm, and a mock-up glovebox environment. Non-prehensile manipulation is advantageous in this case, as the highly-constrained environment means that grasp configurations are difficult to obtain. Thus, we consider non-prehensile manipulation tasks in a glovebox.

The proposed method is tested in three different scenarios of increasing complexity: 1) pushing an object on a table, **Figure 7A-C** 2) pushing an object in a glovebox, **Figure 7D-F**, and 3) ejecting an unreachable object from the glovebox by exploiting physical interactions in the environment, **Figure 7G-I**. In the first case, there is a (red) box on a table and the task is to move it 20 cm along the x axis, see **Figure 7A** for the reference frame. In the second case, the task is to move the object 10 cm along the $-y$ axis in the glovebox, **Figure 7D**. In the last case, two boxes that are placed next to each other are considered, as shown in **Figure 7G**, and the task is to eject the one that is further away from the robot (i.e., the blue box) from the glovebox. In other words, this task

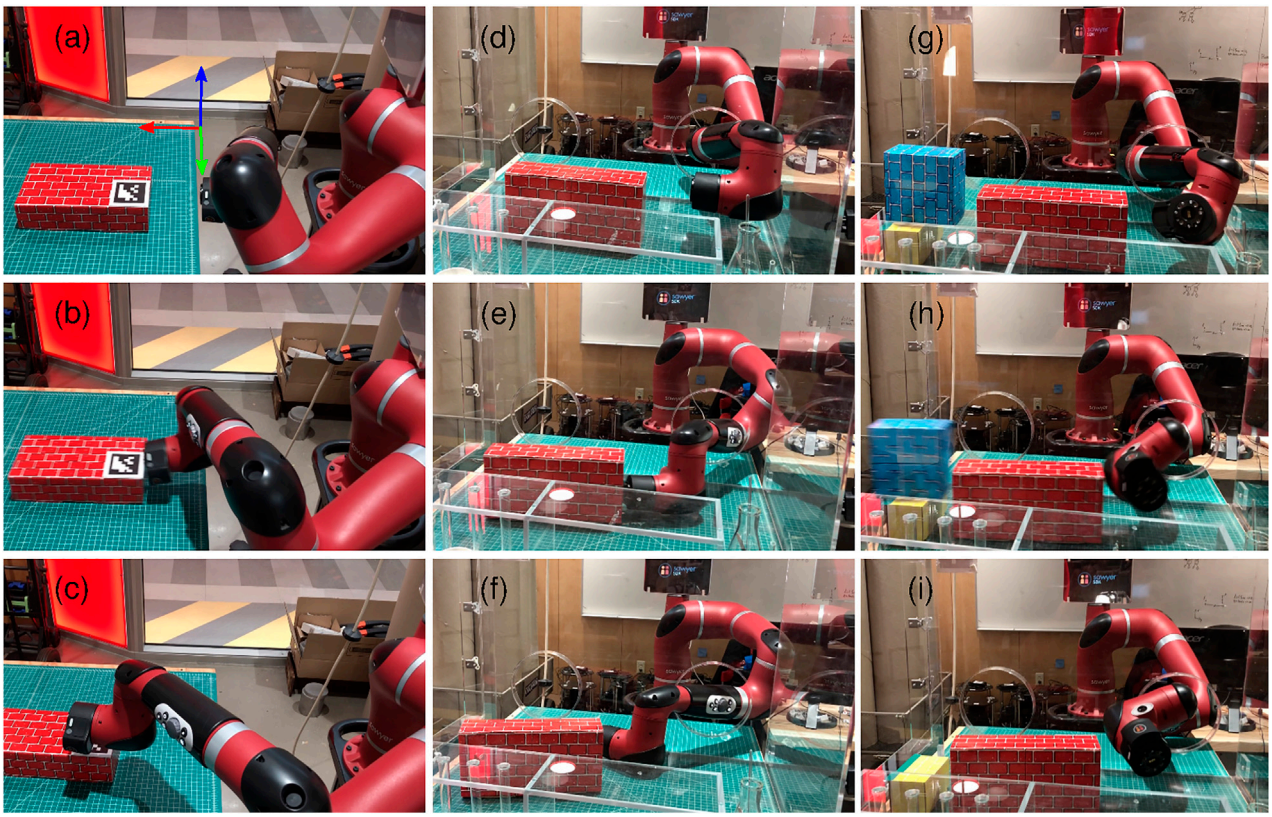

FIGURE 7 | Snapshots from the hardware experiments: pushing an object on a table **(A–C)**, pushing an object within a glovebox **(D–F)**, manipulating an unreachable object by exploiting inter-object contacts **(G–I)**.

requires moving the blue box at least 15 cm along the $-y$ axis so that it will leave the glovebox through the exit port. In all cases, the desired rotation is zero.

4.3.2 Experimental Setup

In the experiments, standard and commercially-available hardware[2] is used in order to facilitate reproducibility. The dynamics is modeled using MuJoCo physics engine (Todorov et al., 2012) with time steps of 5 ms while the control sampling period is 50 ms. The SQP solver SNOPT (Gill et al., 2005) is used to solve the optimization problem. As an interface between MuJoCo and SNOPT, IFOPT (Winkler et al., 2018) is employed. The planned position, velocity, and acceleration trajectories are interpolated with 10 ms steps and executed on the robot by using the built-in inverse dynamics feed-forward controller of Sawyer. For detecting the poses of the objects and the glovebox through the head camera of Sawyer, AprilTag 2 algorithm (Wang and Olson, 2016; Malyuta, 2017) is used.

For all cases, the weights are $w_1 = 10^3$, $w_2 = 10^3$, $w_3 = 1$, and $w_4 = 1$. The initial trajectory is set at zero torque values. The initial value and the upper bound for the virtual stiffness is 5 N/m for the red box and 1 N/m for the blue box, since the blue box is lighter than the red box. α in **Eq. 31** is selected such that $\gamma = k_0 \times 10^{-2}$ for each contact candidate in the initial configuration. Namely, the optimization is started with a trivial initial guess in which the robot stands still for the whole simulation, and there is no heuristic regarding the contact interactions for any task.

4.4 Results

Figure 7 demonstrates initial, intermediate, and final snapshots from the experiments. **Table 1** shows the position and orientation deviations for the object ($\|\mathbf{p}_o^e\|$ and $\|\boldsymbol{\theta}_o^e\|$) for the simulation and hardware experiments and the discrepancy between them. Additionally, the physical inaccuracy caused by the residual virtual forces $\psi = t_c \sum_i \|\boldsymbol{\gamma}_i\|$ is shown for the simulations. For the last case, only $\delta\|\mathbf{p}_o^e\|$ and $\delta\|\boldsymbol{\theta}_o^e\|$ for the red box are shown since the blue box is ejected (i.e., could not be tracked) and there is no desired pose for the red box.

Despite the trivial initial guess and no additional tuning for different tasks, the proposed method is capable of finding a motion that successfully completes each task in simulation. That is, the object is moved to within 1-cm radius of the desired position while the change of orientation is negligible

[2]Information regarding the glovebox (https://www.belart.com/bel-art-h50026-0000-sidentry-glove-box-30-x-24-x-24.html): and the objects (http://www.melissaanddoug.com/deluxe-jumbo-cardboard-blocks—40-pieces/2784.html): are available online.

TABLE 1 | Numerical results from simulation and hardware experiments for all cases.

Task	Simulation			Experiment		Discrepancy	
	ψ [N-s]	$\|\mathbf{p}_o^e\|$ [m]	$\|\theta_o^e\|$ [rad]	$\|\mathbf{p}_o^e\|$ [m]	$\|\theta_o^e\|$ [rad]	$\|\mathbf{p}_o^e\|$ [m]	$\|\theta_o^e\|$ [rad]
1	0.8847	0.0025	0.2772	0.0495	0.0112	0.0470	0.2660
2	1.0698	0.0085	0.1222	0.0777	0.3404	0.0692	0.2182
3	0.1466	N/A		N/A		0.0336	0.1562

(i.e., smaller than 15°). It is noteworthy that in the second and third cases, the glovebox port imposes tight constraints on the motion and the workspace of the robot, yet still our method can handle this without any additional penalties or constraints for collisions with the glovebox or tuning. Moreover, the last task requires the blue box to have a high velocity when the contact between it and the red box is broken because when the red box collides with the yellow box that is under the blue box, the red box cannot apply a force on the blue box anymore and the blue box would still be in the glovebox. Thus, using a non-dynamic planner or running the resulting motions through a position controller without velocities and accelerations would not work in this case.

A more detailed numerical analysis of the experiments is given in the following. In the simulation for Task 1, the box is moved 19.8 cm along the x axis and 0.7 cm along the y axis; whereas, in the hardware experiment, the box is pushed only 15 cm along the x axis and 0.5 cm along the y axis. For Task 2, the displacements of the box along the $-y$ and x axes are 9.8 and 0.8 cm in the simulation. In the experiment, the box is moved 8 cm along the $-y$ axis, which is satisfactory, but also 7 cm along the $-x$ axis due to the relatively large rotation (ca. 20°) about the z axis. In Task 3, the blue box is ejected from the glovebox, namely the task is completed, both in the simulation and the hardware experiment. Since the final position of the blue box could not be detected in the experiment, only the final positions of the red box are compared here. In the simulation, the red box is moved 10.8 cm along the $-y$ axis and 0.4 cm along the $-x$ axis, and these quantities are 7.5 and 0.3 cm for the hardware experiment.

On average, the position and orientation discrepancies between the simulation and hardware experiments are 5 cm and 12°. Such differences can be deemed reasonable since we playback the planned trajectories on the robot using a naive joint-level controller, i.e., without a closed-loop controller that tracks and compensates for the deviation of the object's trajectory from the planned one. The deviations of the executed motions from the planned motions are expected considering errors caused by modeling and perception. The main goal of this study is to show that the proposed method can solve for complex tasks by exploiting the smooth virtual forces and the residual non-physical forces do not hinder the task performance.

5 *IN-SITU* TERRAIN CLASSIFICATION AND ESTIMATION

Robust locomotion on deformable terrains is necessary for biped humanoid robots to perform tasks effectively in unstructured environments. The knowledge of deformable terrain properties, particularly the stiffness, has major implications in modeling the robot walking dynamics, which is the key to achieve stable gait patterns. Prior studies on walking stabilization chose to model such walking dynamics using pre-identified stiffness or damping constants (Wittmann et al., 2016; Wu et al., 2018; Hashimoto et al., 2012). However, it is unlikely for robots to access such terrain properties in advance when deployed to unknown environments.

In this section, we present an *in-situ* ground classification and estimation method that can be used to improve the stability of the robot while traversing unknown terrain, utilizing NASA's humanoid robot Valkyrie (Radford et al., 2015). The terrain estimation works in two steps: i) The robot tries to identify the terrain type from a database. If the terrain is recognized, all needed data can be retrieved and used by the controller. ii) If the terrain is classified as an unknown type, the robot then estimates its stiffness by using Bernstein-Goriatchkin (Ding et al., 2013; Caurin and Tschichold-Gurman, 1994) pressure-sinkage model. The estimated stiffness can then be fed to stabilizers such as the one proposed in (Hashimoto et al., 2012).

5.1 Related Work

Our study on robot foot-terrain interaction is inspired by (Skonieczny et al., 2014), where the interaction between soil and single wheel is analyzed using optical flow techniques. Computer vision techniques have been widely used for terrain classification in the past (Weiss et al., 2008; Brandão et al., 2016). However, due to the poor lighting conditions in the outer space, it is desirable to augment vision-based techniques with a terrain classification approach that relies on physical contacts between the robot foot and the terrain. We thus aim at providing a "sense-of-walking" to the robot by using on-board sensors. In (Walas et al., 2016; Otte et al., 2016), ankle mounted force/torque sensors and accelerometers are used, respectively, to achieve terrain classification. Our approach is comparable to (Otte et al., 2016) as Recurrent Neural Networks (RNNs) are used but they differ in the aspect that we perform terrain classification with a bipedal robot while (Otte et al., 2016) uses a wheeled mobile robot. To describe terrains' properties under pressure, various pressure-sinkage models have been studied in the past (Komizunai et al., 2010; Ding et al., 2013). There is no common opinion on which model is better than others. We choose Bernstein-Goriatchkin model considering it is one of the most commonly used models and it is relatively easy to implement. The method we developed for terrain estimation can be viewed as an extension of (Will Bosworth1 and Hogan, 2016), where spring

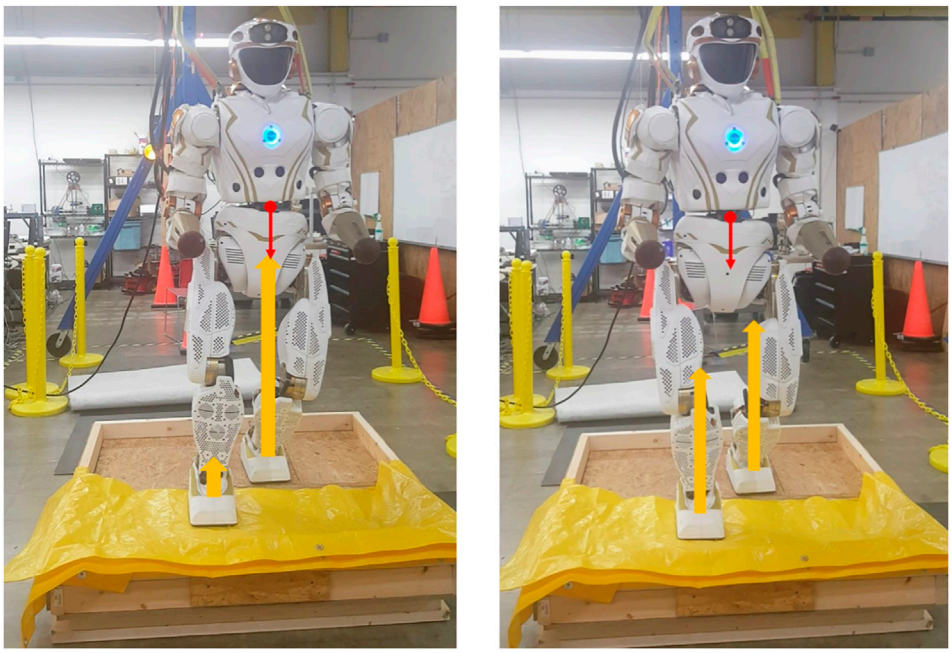

FIGURE 8 | Valkyrie is performing the touch-land motion. The red arrow represents CoM and yellow arrow represents support force on each foot.

model is used to estimate the ground stiffness by measuring the force and leg displacement during the interaction between a quadrupedal robot, Super Mini Cheetah (SMC), and the ground. Since SMC has spherical shaped feet, point contacts are used when modeling the foot-terrain interaction. This is not applicable for bipedal robots with flat soles because the support force on sole is distributed unevenly. To overcome this challenge (Caurin and Tschichold-Gurman, 1994), develops a special sensor system consisting of multiple sensor cells to measure the force distribution. Instead of using additional sensors, our approach achieves estimation for stiffness using a single load cell which provides only a net force measurement.

5.2 Terrain Classification

To study the foot-ground interaction, we designed a set of experiments involving Valkyrie interacting with four types of terrains, including wood mulch, rubber mulch, mason sand, and gym foam mats. A custom sandbox, as seen in **Figure 8**, was built for Valkyrie to stand in and perform a set of pre-designed motions. Half of the sandbox is covered by a wood plate, where Valkyrie can stand stably, and the other half is filled with specified deformable ground materials to be tested. The sandbox is not used when testing on the gym foam mat, which Valkyrie can stand on directly. The data for terrain classification and estimation is collected by commanding Valkyrie to perform a touch-land motion.

Valkyrie starts from standing on the wood plate of the sandbox and then moves the right foot onto the terrain. During this motion, Valkyrie slowly moves the center of mass (CoM) from the left foot to the center of the two feet (**Figure 8**). The support force on the right foot increases from zero to approximately 650 N, which is about half the weight of Valkyrie, in this process. Valkyrie then moves the right foot back onto the wood plate and the motion is finished. The right foot force and torque changes during this motion are measured by the load cell. Meanwhile, the right foot sinkage in the terrain is also calculated using joint positions.

One method to calculate the foot sinkage, δ, is to use the kinematic chain of the robot. Transformation matrix from coordinate frame *RightFoot1* to *RightFoot2* is denoted with ${}^{R1}T_{R2}$ representing the transformation of the right foot at two different poses. *RightFoot1* corresponds to the foot pose when the foot makes contact with the terrain. Since Valkyrie keeps its foot flat during the swing, it is reasonable to take *RightFoot1* as the fixed ground frame. *RightFoot2* is the frame when the foot sank in the soft terrain. There could be many different choices for *RightFoot2*. Using the transformation ${}^{R1}T_{R2}$, we can attain the position for each specific point of the sole relative to ground as ${}^{R1}T_{R2} \cdot P$, where $P = [x, y, 0, 1]^T$ is the coordinates in *RightFoot2* frame. The sinkage δ_{xy} can then be represented as the third element of ${}^{R1}T_{R2} \cdot P$, denoted as ${}^{R1}T_{R2} \cdot P \cdot P(3)$. To get ${}^{R1}T_{R2}$, we introduce a third coordinate frame *LeftFoot*. Since Valkyrie's left foot is stationary during the experiments (right foot taking the step), we take *LeftFoot* as the reference coordinate frame. Using the robot's kinematic chain, we can get ${}^{L}T_{R1}$, ${}^{L}T_{R2}$, and thus ${}^{R1}T_{R2}$.

We use RNNs, particularly Long Short-Term Memory (LSTM) network, to classify the terrain. Comparing to other classification methods such as Support Vector Machine (SVM), the use of RNNs gains us two advantages. First, our method can be applied to a raw data stream rather than accumulated data over time, thus, we can achieve terrain classification in real-time. Second, it has been shown that RNNs have better performance

than SVM (Otte et al., 2016) and we achieved 95% accuracy on average during experiments.

Fifteen features are selected as input for the LSTM network, which are: right foot force in *X, Y, Z* directions, right foot torque in *X, Y, Z* directions, pitch/roll angular displacement/acceleration of the right ankle, pitch/roll angular torque of the right ankle, and the right foot sinkage, represented by the first three elements on third row of the transformation ${}^{R1}T_{R2}$. One data set is collected when Valkyrie performs one touch-land motion for one specific terrain. In total, we collected 30 data sets for the foam mat, 30 data sets for the wood mulch, 34 data sets for the rubber mulch, and 42 data sets for the sand. 90% of the data sets (118 data sets) are used for training and the rest (18 data sets) are used for testing. The LSTM network thus has 15 input neurons, one hidden layer with 64 hidden blocks and one linear output layer with four neurons. A softmax function is used to determine the terrain type. For each experiment run, the model is trained with a fixed learning rate of 0.0005 for 16 epochs. To track the performance of our model, the prediction accuracy is calculated using the test data set after each training epoch. Out of the five experiment runs, we achieved 100% prediction accuracy for four runs and 89% accuracy for the remaining one.

5.3 Terrain Estimation

Although the RNN based classification method works well, there are chances that Valkyrie may encounter with terrains of unknown type. To handle such situations, a method that can estimate an unknown terrain's stiffness, which governs the foot sinkage behavior under load, is desired.

5.3.1 Terrain Stiffness Model

To describe the terrain's stiffness, the well-known pressure-sinkage model developed by Bernstein-Goriatchkin (Caurin and Tschichold-Gurman, 1994; Ding et al., 2013) is used:

$$\sigma = k \cdot \delta^n \tag{35}$$

where $\sigma(Pa)$ is the normal pressure, $k(Pa/m^n)$ and $n(non-dimensional)$ are the terrain parameters and $\delta(m)$ is the foot sinkage in terrain. The two terrain parameters, k and n, are to be estimated by using data from Valkyrie's on-board sensors. The challenge here is that Valkyrie uses only one 6-axis load cell mounted at the ankle to measure support force acting on the foot, which cannot provide a detailed force distribution upon sole. Ideally, Bernstein-Goriatchkin model can be easily implemented to determine the terrain stiffness when the terrain undergoes an even and flat sinkage, as seen in **Figure 9A**. However, when a bipedal robot walks on deformable terrains, the sinkage is usually non-uniform, thus the contact surface between the foot and the ground could be oblique, as seen in **Figure 9B**.

In this case, it is impractical to identify the sinkage value corresponding to the pressure calculated from the measured force on ankle. To handle this mismatch, we calculate the net force acting on the sole resulting from different levels of sinkage using the following equation

$$F = \iint_A \sigma(x,y)\, dA = \int_{y_1}^{y_2} \int_{x_1}^{x_2} k \cdot \delta_{xy}^n \, dxdy \tag{36}$$

where F is the net force acting on the foot, δ_{xy} is the local foot sinkage at (x, y) expressed in the foot frame. The contact area between the foot and the ground is defined by $(x1, x2)$ and $(y1, y2)$, which are also expressed in the foot frame. These same parameters are used to calculate the torque values for verification, thus we pick the projection of the ankle in the foot frame as the origin. The contact area for Valkyrie is then defined as (−0.09 m, 0.18 m) along the x-axis and (-0.08 m, 0.08 m) along the y-axis based on the physical dimensions of the robot. Since F can be measured directly, using **Eq. 36**, we obtain k and n once we know the sinkage at different positions of the foot, i.e., δ_{xy}. The sinkage δ_{xy} can be calculated using the same frame transformation method described in **Section 3**. **Eq. 36** can then be written as

$$F = \int_{y_1}^{y_2} \int_{x_1}^{x_2} k \cdot ({}^{F1}T_{F2} \cdot P(3))^n dxdy \tag{37}$$

Theoretically, if two sets of F and ${}^{F1}T_{F2}$ for the sinking foot of different states are provided, we can solve for k and n using **Eq. 37**. It is worth mentioning that in **Eq. 37**, both transformation ${}^{F1}T_{F2}$ and support force F can be calculated or measured in real-time. Contact range x and y are pre-defined based on the foot dimensions. Therefore, this method can estimate the unknown terrain's stiffness *in-situ*.

5.3.2 Experiments and Results

To test the proposed estimation method, we run four experiments: two with gym foam mats (Gym Foam Mat one and two), Rubber Mulch, and Mason Sand. The same touch-land motion as described in **Section 5.2** is performed by Valkyrie to collect data for terrain estimation. The support force acting on Valkyrie's right foot and the frame transformation ${}^{F1}T_{F2}$ are recorded during each touch-land motion. Nonlinear Least-Squares Data Fitting (Coleman and Li, 1996) is then used to solve for parameters k and n. **Table 2** shows the estimated parameters for four tested terrains. The pressure and sinkage relationships of the four terrains with the estimated parameters are depicted in **Figure 10**. We intuitively realize that the estimation results are not valid for the sand. This is because the sand is known to be stiffer than both foam and rubber mulch but **Figure 10** shows that the sand is estimated to be the softest. The reason of inaccurate estimation for sand is that the sand demonstrates a much smaller deformation under the same pressure compared with the form and the rubber mulch. During the experiments, the foam mat and the rubber mulch have a deformation between 20 and 25 mm while the sand has a deformation between 5 and 10 mm, which is a section that is too short to do regression for solving k and n.

To validate the estimation results for the foam mats and the rubber mulch, we compare ankle torques measured from the load cell with ankle torques calculated using the estimated parameters. We choose torque comparison as the validation approach for one practical reason: many bipedal robots walking algorithms keep the robot balanced by tuning torques at the ankle (Kim et al., 2007; Kang et al., 2012). We thus posit that a set of terrain parameters that can correctly predict torque changes can be useful in walking algorithms for keeping robots stable on

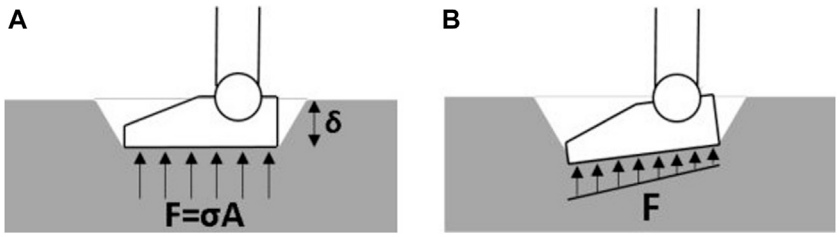

FIGURE 9 | Contact between the foot and ground terrain. **(A)** Even sinkage when the contact is flat. **(B)** Non-even sinkage when the contact is oblique.

TABLE 2 | Estimated stiffness parameters *k* and *n* for tested terrains using Bernstein-Goriatchkin Model.

Terrain	k (10^3Pa/m^n)	N	Validity
Gym foam mat 1	1167.861	1.367	Good
Gym foam mat 2	609.296	1.212	Good
Rubber mulch	1738.93	1.45	Good
Mason sand	421146.385	3.00	Poor

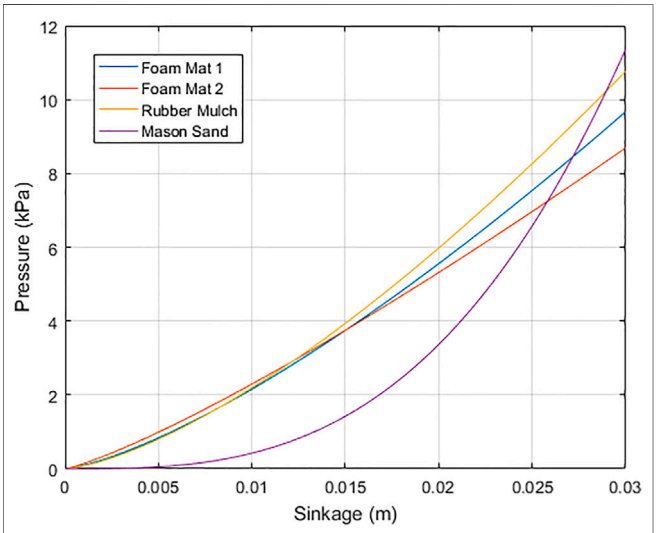

FIGURE 10 | Relationship of pressure and sinkage for different terrains using estimated parameters.

deformable terrain. When Valkyrie's foot lands on the terrain, the torque measured at the ankle should be equal to the torque caused by the force acting on the sole around the ankle axis. Considering the torque caused by gravity, we have.

$$\tau_{ankle} = \tau_f + \tau_g \tag{38}$$

where τ_{ankle} is the net torque measured at the ankle; τ_g is the torque due to gravity, which equals to the measured torque at the ankle during the swing; τ_f is the torque due to support force resulting from terrain deformation. Since we now have parameter *k* and *n* for the terrain, using **Eq. 36**, the pitch and roll torque can then be calculated as

$$\tau_{pitch} = \int_{y_1}^{y_2}\int_{x_1}^{x_2} k\delta_{xy}^n x\, dxdy,\ \tau_{roll} = \int_{y_1}^{y_2}\int_{x_1}^{x_2} k\delta_{xy}^n y\, dxdy \tag{39}$$

Figure 11 shows the comparison between the calculated foot pitch and roll torques using estimated parameters and measured foot pitch and roll torques from the ankle load cell. The blue curves represent the measured results while the red curves stand for the estimated results. The plots show that the estimated parameters can predict the foot pitch torque well. As for the foot roll torque, the estimated results match the trend of torque changes with an offset of 1.5 Nm on average. One possible source of this offset is the friction between the foot and the terrain. The friction coefficient between the foot and the terrain is currently not estimated and is planned as the future work.

To numerically examine how well the estimated parameters predict the foot pitch torque, we calculate the Normalized Root-Mean-Square Error (NRMSE) between the estimated pitch torques and the measured pitch torques. The Gym Foam Mat one and two show a NRMSE of 15.9 and 15.7%, respectively, which is the equivalent to an accuracy of 84.1 and 84.3%. While the Rubber Mulch showed an NRMSE of 24.2%, which is an accuracy of 75.8%. Since the estimated parameters from the Gym Foam Mat one and two and Rubber Mulch experiments can well predict foot pitch torque (an average accuracy of 80%) and roll torque (same trend but with an offset), we assert that the proposed method can be used to measure the stiffness of an unknown deformable terrain. Future work will focus on improving the estimation accuracy, including estimation of friction coefficient for a thorough knowledge of unknown terrains.

6 USER INTERFACE

Supervisory-control interfaces are an important aspect to robots in extreme environments as you often want a human-in-the-loop to help evaluate and make decisions for critical objectives of a task at hand. In any supervisory-control interface, it is important to find the balance of control between robot and human. Too much control or autonomy on the robot side leaves the human unable to potentially provide valuable feedback to the robot. While too much control on the human side can leave the operator overwhelmed with information and decision making, leaving

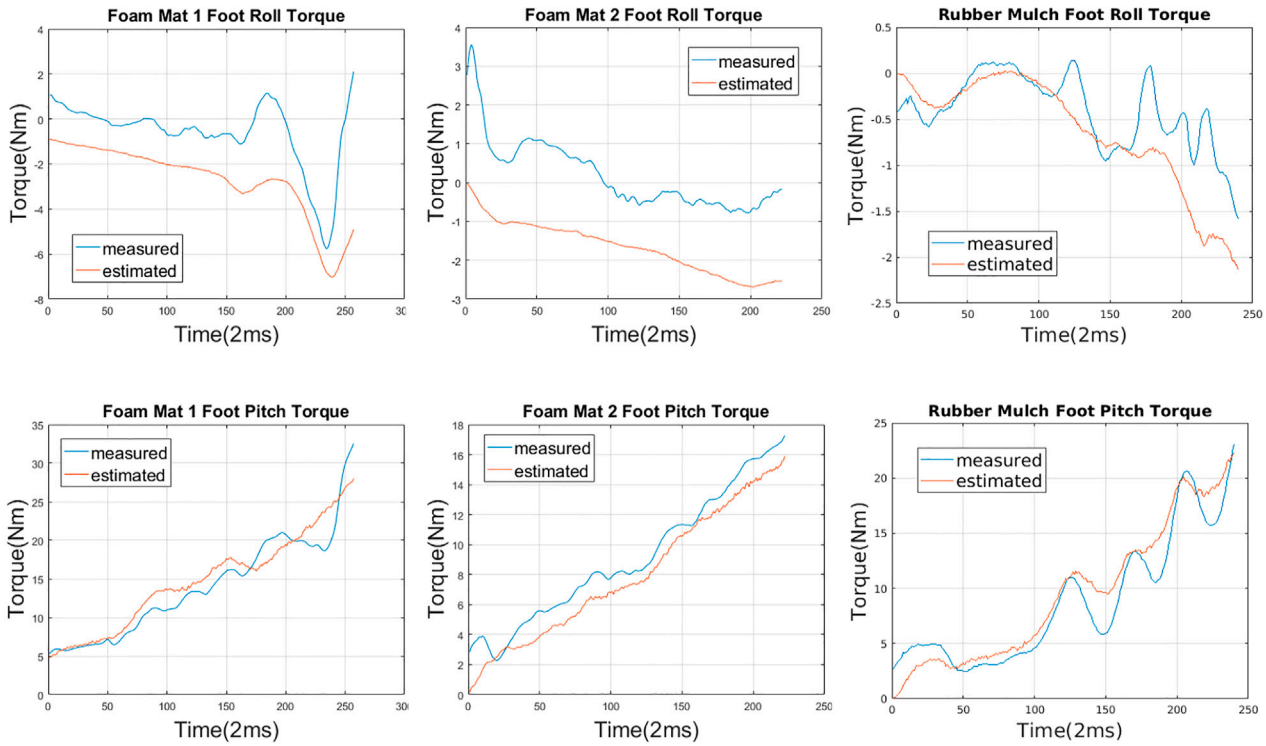

FIGURE 11 | Foot pitch and roll torque comparison. Blue curves stand for measured torques and red curves represent the estimated torques.

less focus for the critical elements of a given task. Additionally, too much control on the human side tends to increase the interface complexity and therefore increase the amount of required operator training to utilize the interface.

In this section, we present our work on a computer-based, supervisory-control interface for NASA's humanoid robot Valkyrie (Radford et al., 2015), that aims to balance the supervisory-control to allow for relatively untrained operators. Our interface was originally designed to accomplish the tasks put forward by the Space Robotics Challenge, a challenge using a simulated Valkyrie robot in a virtual, mock-up Mars environment with tasks such as turning valves, moving equipment, and inspection. However, our interface was purposely designed to be generic enough to accomplish a wide variety of tasks and not require tasks to be pre-programmed into the interface and is therefore extendable to a wide variety of domains.

6.1 Related Work

The DARPA Robotics Challenge (DRC) was a two-year challenge with the goal of accelerating progress in human-supervised robotic technology for disaster-response operations. The DRC helped support development for a variety of robotic research, with one key area being the design of human-robot interaction with focus on supervisory control methodologies (Yanco et al., 2015; Norton et al., 2017). The DRC allowed for a variety of supervisory-control interfaces for humanoid robots to be developed furthering the state-of-the-art. However, most of the developed interfaces required multiple well-trained operators in order to successfully complete tasks (Norton et al., 2017). Several interfaces aimed to reduce the amount of input required on the human operator, but still allowed operators to take over low-level commands of the robot when needed and provided a wide range of robot sensor data (Stentz et al., 2015; DeDonato et al., 2017).

6.2 Interface

Our Valkyrie interface was designed to mimic several of the humanoid interfaces designed for the DRC finals, but with a reduced number of capabilities, inputs, and sensor data to allow for both a single operator and to reduce operator training time. The goal was to create a comparable interface to ones used at the DRC, but user-friendly enough to allow for non-experts to quickly learn and operate.

Our interface, as seen in **Figure 12**, utilizes RViz, a 3D visualization tool for ROS, and custom Qt Widgets, also known as Panels in RViz. RViz was chosen as the interface backbone since ROS is widely used in the robotics community making it familiar to many and for its built-in data visualization tools. RViz was also used for several of the interfaces designed for DRC (DeDonato et al., 2015; Kohlbrecher et al., 2015; Stentz et al., 2015). The interface can be broken down into four custom panels: Arm Motion Planner, Step Planner, Neck Interface, and Hand Interface. Additionally, there are three built-in RViz panels: the main visualization window, a display panel, and a camera feed panel.

The main window consists of the built in 3D data visualization of RViz and for our specific interface includes a point-cloud from Valkyrie's Multisense SL sensor located in the head, a robot model of Valkyrie's current joint state and planned joint state

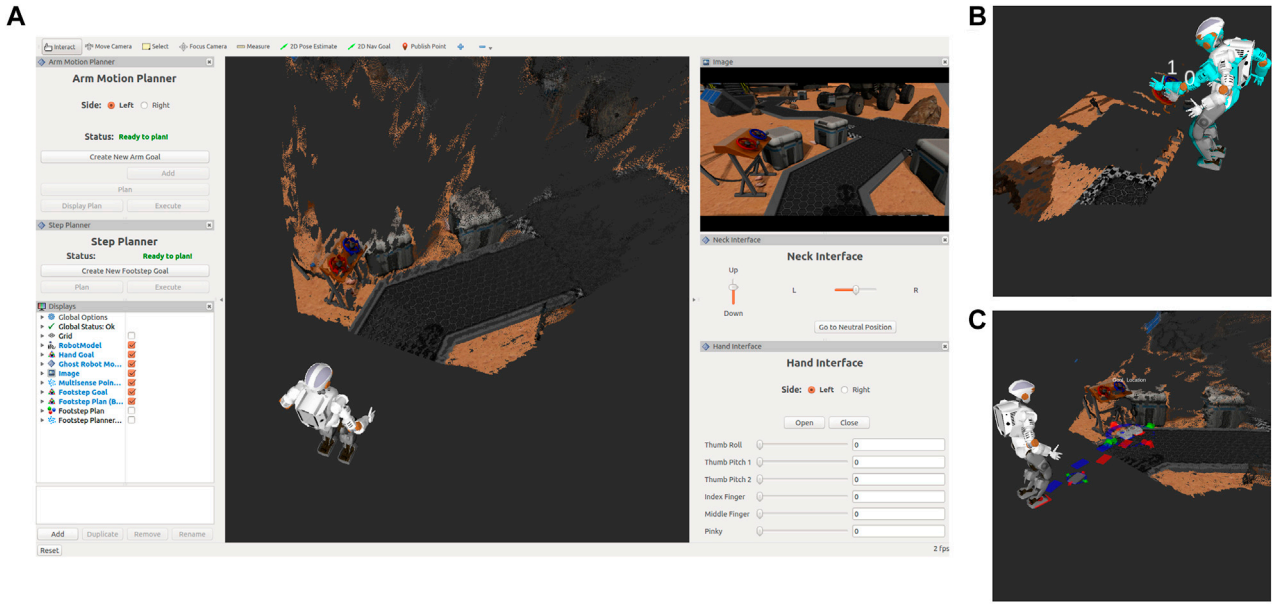

FIGURE 12 | Valkyrie interface.

from the arm motion planner, and interactive markers to interface with the arm motion and step planners. The display panel allows users to select what is actively visualized in the main window and the camera feed panel for our interface is a camera feed from one of the stereo cameras in the Multisense SL.

6.2.1 Arm Motion Planner Panel

The Arm Motion Planner panel utilizes a whole body inverse kinematics solver that finds motion plans based on a set of cost and constraints (Long et al., 2016, 2017). Cost and constraints can be categorized as either kinematic, collision avoidance, or ZMP. To reduce the complexity of the interface, operators only need to define the desired end-effector positions and a predefined set of cost and constraints are used, i.e., balancing constraint or velocity cost, etc. To define the end-effector positions, operators can add waypoints for the end-effector to traverse though. This is accomplished through the use of interactive markers that represent a waypoint and are visually represented as a single end-effector. Operators are able to translate and rotate the waypoint by using arrows and circular scroll bars that surround the marker visual. The waypoints are ordered to allow for a variety of paths. A plan can be requested after the waypoints have been placed in a satisfactory manner. Operators are provided a planning indicator and are informed when the planning is complete and whether the planner was successful in finding a plan or unsuccessful. If a plan returns unsuccessful, then the returned plan either violated one of the default set of costs or it was unable to traverse through all the placed waypoints. Operators are notified of the reason the planned returned unsuccessful and informed on how to improve the waypoints to results in a successful plan. For example, a waypoint may be out of the robot's workspace and need to be moved closer to the robot. After a plan is returned, the operator can visualize the robot's movement through the entire plan before allowing the robot to execute it to ensure that the plan both provides the desired outcome and is collision free. **Figure 12B** demonstrates an example of placed waypoints and the final robot configuration of the returned plan. After visualizing the plan, operators can request the robot execute it.

6.2.2 Step Planner Panel

The Step Planner panel allows operators to designate a goal position for the robot to navigate to. It works in a similar strategy as the Arm Motion Planner panel, except rather than multiple waypoints, there is only a single waypoint that represents the final goal pose for feet with a visual of a pair of feet. Additionally, the interactive marker that represents the waypoint only allows for 2D interaction on a plane. After this single waypoint is placed, operators can request a plan which is generated using an A* search-based footstep planner. Same as the Arm Motion Planner panel, operators are provided a planning indicator and informed when the plan returns. As soon as a successful plan is found the planner will return. However, the planner will infinitely plan until a successful plan is found, therefore, there is a 10 s timeout in place to force a return and the operator is notified of this timeout and to try again. As soon as a successful plan is returned, colored interactive markers that represent each footstep in the returned plan are visualized, red for the right foot and blue for the left foot. These interactive markers can be adjusted by the operator by clicking on the footstep marker to activate the interactive marker controls and then adjusting the footstep in the same method as the waypoints, the interactive marker controls can also be switched off by clicking the footstep marker. **Figure 12C** shows an example of a placed goal marker, a returned footstep plan, and selected footstep in adjustment mode. When finished

making adjustments, operators are then able to send the plan to the robot for execution.

6.2.3 Neck Interface and Hand Interface Panels

The Neck Interface and Hand Interface panels both work in the same manner with a series of joint sliders and buttons with some predefined joint positions. The Neck Interface panel allows operators to move the head of the robot right-left and up-down with sliders and a predefined neutral position, where the robot is looking straight forward. The Hand Interface panel has sliders for each individual finger to open-close along with two predefined open hand and closed fist grasp. Sliders were added over additional grasp types to allow the operator to position the hand at any intermediate pose if needed.

7 CONCLUSION

Humanoid robots are equipped with the necessary combination of location and manipulation capacities to fulfill a variety of roles in man-made hazardous environments. However, there are numerous challenges before this potential can be realized. In this paper, we presented our research to help further the development of utilizing humanoid robots in supervisory roles for extreme environments, with emphasis on operation in nuclear facilities.

We first examined ways to improve manipulation within constrained environments, with a focus on glovebox operation, by looking into three different areas. First, we aimed to better understand a robot's feasible workspace in constrained environments by developing a method called the constrained manipulability polytope (CMP) that considers both the robot's capabilities and the constraints imposed by an environment. The CMP methodology assists in both autonomous motion planning and human-in-the-loop teleoperation by providing feedback of the robot's workspace to the operator and indeed by providing a representation of both robot joint and Cartesian space constraints as a bounded limits in tool motion. Second, we investigated using contact supports to increase stability and we presented our theoretical methodology to accomplish contact-implicit motion planning. The objective of this work was to increase the robot's workspace by using support contacts to reach areas that are not within the robot's workspace without support. Third, we considered using non-prehensile manipulation over dexterous manipulation and presented a methodology capable of pushing objects in a constrained environment. The purpose of this work was to add the capability of moving objects around that might not be within reach of the robot or easily manipulated with dexterous hands.

We also detailed methods to improve locomotion for humanoid robots by developing a technique to accomplish *in-situ* terrain classification and estimation that can be used to better inform a walking controller. Finally, we presented our generic humanoid robot interface that allows for robot operation from operators with limited training and/or knowledge of robotics. When operating in high-risk environments, such as nuclear facilities, it will be necessary to keep humans-in-the-loop especially individuals with expertise in the environment the robot is operating in. Future work should consider further developing all these areas and combining them into one robotic platform with a supervisory-control interface for a single operator. Human supervised robots will be a key role in the future operation, maintenance, and decommissioning of nuclear facilities.

DATA AVAILABILITY STATEMENT

The original contributions presented in the study are included in the article/Supplementary Material, further inquiries can be directed to the corresponding author.

AUTHOR CONTRIBUTIONS

PL authored the Constrained Manipulation and Motion Planning section. AO authored the Increasing Stability Using Support Contacts section and Dynamic Non-Prehensile Manipulation section. MaW authored the *in-Situ* Terrain Classification and Estimation. MuW authored the Supervisory-Control Interface and compiled the manuscript. TP supervised the work.

FUNDING

This research is supported by the Department of Energy under Award Number DE-EM0004482, by the National Aeronautics and Space Administration under Grant No. NNX16AC48A issued through the Science and Technology Mission Directorate and by the National Science Foundation under Award No. 1451427, 1944453.

REFERENCES

Abdel-Malek, K., Yu, W., and Yang, J. (2004). Placement of robot manipulators to maximize dexterity. *Int. Journal of Robotics and Automation* 19, 6–14. doi:10.2316/journal.206.2004.1.206-2029

Anitescu, M., and Potra, F. A. (1997). Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems. *Nonlinear Dynamics* 14, 231–247. doi:10.1023/a:1008292328909

Azad, M., Babič, J., and Mistry, M. (2017). Dynamic manipulability of the center of mass: A tool to study, analyse and measure physical ability of robots. In IEEE Int. Conf. on Robotics and Automation. IEEE, 3484–3490.

Bogue, R. (2011). Robots in the nuclear industry: a review of technologies and applications. *Industrial Robot: An International Journal.*

Bowling, A., and Khatib, O. (2005). The dynamic capability equations: a new tool for analyzing robotic manipulator performance. *IEEE Trans. Robot.* 21, 115–123. doi:10.1109/tro.2004.837243

Brandão, M., Shiguematsu, Y. M., Hashimoto, K., and Takanishi, A. (2016). Material recognition cnns and hierarchical planning for biped robot locomotion on slippery terrain. In 2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids). 81–88. doi:10.1109/HUMANOIDS.2016.7803258

Caccavale, F., and Uchiyama, M. (2016). Cooperative manipulation. In Springer Handbook of Robotics. Springer, 989–1006. doi:10.1007/978-3-319-32552-1_39

Carius, J., Ranftl, R., Koltun, V., and Hutter, M. (2018). Trajectory optimization with implicit hard contacts. *IEEE Robot. Autom. Lett.* 3, 3316–3323. doi:10.1109/lra.2018.2852785

Caron, S., Pham, Q.-C., and Nakamura, Y. (2017). Zmp support areas for multicontact mobility under frictional constraints. *IEEE Trans. Robot.* 33, 67–80. doi:10.1109/tro.2016.2623338

Caurin, G., and Tschichold-Gurman, N. (1994). The development of a robot terrain interaction system for walking machines. *Proceedings of the 1994 IEEE International Conference on Robotics and Automation* 2, 1013–1018. doi:10.1109/ROBOT.1994.351224

Coleman, T. F., and Li, Y. (1996). An interior trust region approach for nonlinear minimization subject to bounds. *SIAM J. Optim.* 6, 418–445. doi:10.1137/0806023

DeDonato, M., Dimitrov, V., Du, R., Giovacchini, R., Knoedler, K., Long, X., et al. (2015). Human-in-the-loop control of a humanoid robot for disaster response: A report from the darpa robotics challenge trials. *J. Field Robotics* 32, 275–292. doi:10.1002/rob.21567

DeDonato, M., Polido, F., Knoedler, K., Babu, B. P. W., Banerjee, N., Bove, C. P., et al. (2017). Team wpi-cmu: Achieving reliable humanoid behavior in the darpa robotics challenge. *Journal of Field Robotics*. doi:10.1002/rob.21685

Ding, L., Gao, H., Deng, Z., Song, J., Liu, Y., Liu, G., et al. (2013). Foot-terrain interaction mechanics for legged robots: Modeling and experimental validation. *The International Journal of Robotics Research* 32, 1585–1606. doi:10.1177/0278364913498122

Escande, A., Kheddar, A., and Miossec, S. (2013). Planning contact points for humanoid robots. *Robotics and Autonomous Systems* 61, 428–442. doi:10.1016/j.robot.2013.01.008

Fletcher*, R., and Leyffer,‡, S. (2004). Solving mathematical programs with complementarity constraints as nonlinear programs. *Optimization Methods and Software* 19, 15–40. doi:10.1080/10556780410001654241

Fukuda, K., and Prodon, A. (1996). Double description method revisited. In Combinatorics and computer science. Springer, 91–111. doi:10.1007/3-540-61576-8_77

Funakoshi, A. S. M. (2016). *Fukushima's ground zero: No place for man or robot*. Reuters.

Gabiccini, M., Artoni, A., Pannocchia, G., and Gillis, J. (2018). A computational framework for environment-aware robotic manipulation planning. In Robotics Research. Springer, 363–385. doi:10.1007/978-3-319-60916-4_21

Giftthaler, M., Neunert, M., Stäuble, M., Buchli, J., and Diehl, M. (2017). *A family of iterative gauss-newton shooting methods for nonlinear optimal control* (arXiv preprint arXiv:1711.11006).

Gill, P. E., Murray, W., and Saunders, M. A. (2005). SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Rev.* 47, 99–131. doi:10.1137/s0036144504446096

Gu, Y., Lee, C. G., and Yao, B. (2015). Feasible center of mass dynamic manipulability of humanoid robots. In Int. Conf. on Robotics and Automation. IEEE, 5082–5087.

Hashimoto, K., Kang, H.-j., Nakamura, M., Falotico, E., Lim, H.-o., Takanishi, A., et al. (2012). Realization of biped walking on soft ground with stabilization control based on gait analysis. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems. 2064–2069. doi:10.1109/IROS.2012.6385684

Hauser, K., Bretl, T., and Latombe, J.-C. (2005). Non-gaited humanoid locomotion planning. In Humanoid Robots, 2005 5th IEEE-RAS International Conference on IEEE), 7–12.

Henze, B., Dietrich, A., and Ott, C. (2016a). An approach to combine balancing with hierarchical whole-body control for legged humanoid robots. *IEEE Robot. Autom. Lett.* 1, 700–707. doi:10.1109/lra.2015.2512933

Henze, B., Roa, M. A., and Ott, C. (2016b). Passivity-based whole-body balancing for torque-controlled humanoid robots in multi-contact scenarios. *The International Journal of Robotics Research* 35, 1522–1543. doi:10.1177/0278364916653815

Herceg, M., Kvasnica, M., Jones, C., and Morari, M. (2013). Multi-Parametric Toolbox 3.0. In Proc. of the European Control Conf. Switzerland: Zürich, 502–510.

Kang, H.-j., Nishikawa, K., Falotico, E., Lim, H.-o., Takanishi, A.H., et al. (2012). Biped walking stabilization on soft ground based on gait analysis. In 2012 4th IEEE RAS EMBS International Conference on Biomedical Robotics and Biomechatronics. BioRob, 669–674. doi:10.1109/BioRob.2012.6290870

Khatib, O., and Chung, S.-Y. (2014). Suprapeds: Humanoid contact-supported locomotion for 3d unstructured environments. In Robotics and Automation (ICRA), 2014 IEEE International Conference on. IEEE, 2319–2325.

Khatib, O., Sentis, L., and Park, J.-H. (2008). *European Robotics Symposium 2008*. Springer, 303–312. A unified framework for whole-body humanoid robot control with multiple constraints and contacts.

Kim, J.-Y., Park, I.-W., and Oh, J.-H. (2007). Walking control algorithm of biped humanoid robot on uneven and inclined floor. *J. Intell. Robot. Syst.* 48, 457–484. doi:10.1007/s10846-006-9107-8

Kohlbrecher, S., Romay, A., Stumpf, A., Gupta, A., Von Stryk, O., Bacim, F., et al. (2015). Human-robot Teaming for Rescue Missions: Team ViGIR's Approach to the 2013 DARPA Robotics Challenge Trials. *J. Field Robotics* 32, 352–377. doi:10.1002/rob.21558

Kokkinis, T., and Paden, B. (1989). Kinetostatic performance limits of cooperating robot manipulators using force-velocity polytopes. *ASME Winter Annual Meeting*, 151–155.

Komizunai, S., Konno, A., Abiko, S., and Uchiyama, M. (2010). Development of a static sinkage model for a biped robot on loose soil. In 2010 IEEE/SICE International Symposium on System Integration. 61–66. doi:10.1109/SII.2010.5708302

Lee, J. (1997). A study on the manipulability measures for robot manipulators. Int. Conf. on Intelligent Robots and Systems, *(IEEE)* 3, 1458–1465.

Long, P., Keleştemur, T., Önol, A. Ö., and Padir, T. (2019). optimization-based human-in-the-loop manipulation using joint space polytopes. In 2019 International Conference on Robotics and Automation (ICRA). IEEE, 204–210.

Long, P., Khalil, W., and Caro, S. (2015). Kinematic and dynamic analysis of lower-mobility cooperative arms. *Robotica* 33, 1813–1834. doi:10.1017/s0263574714001039

Long, P., and Padır, T. (2018). Evaluating robot manipulability in constrained environments by velocity polytope reduction. In 2018 IEEE-RAS 16th Int. Conf. on Humanoid Robots.

Long, P., and Padir, T. (2019). Constrained manipulability for humanoid robots using velocity polytopes. *International Journal of Humanoid Robotics*.

Long, X., Long, P., Onol, A. O., and Padır, T. (2018). *Integrating risk in humanoid robot control for applications in the nuclear industry* (arXiv preprint arXiv:1807.04814).

Long, X., Wonsick, M., Dimitrov, V., and Padır, T. (2016). Task-oriented planning algorithm for humanoid robots based on a foot repositionable inverse kinematics engine. In 2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids). IEEE, 1114–1120.

Long, X., Wonsick, M., Dimitrov, V., and Padır, T. (2017). IEEE), 4452–4459. Anytime multi-task motion planning for humanoid robots. 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS).

Malyuta, D. (2017). *Guidance, Navigation, Control and Mission Logic for Quadrotor Full-cycle Autonomy. Master thesis.* USA: Jet Propulsion Laboratory, 4800 Oak Grove Drive, Pasadena, CA 91109.

Manchester, Z., and Kuindersma, S. (2017). Variational contact-implicit trajectory optimization. In Proceedings of the International Symposium on Robotics Research (ISRR).

Marcucci, T., Gabiccini, M., and Artoni, A. (2017). A two-stage trajectory optimization strategy for articulated bodies with unscheduled contact sequences. *IEEE Robot. Autom. Lett.* 2, 104–111. doi:10.1109/lra.2016.2547024

Marturi, N., Rastegarpanah, A., Takahashi, C., Adjigble, M., Stolkin, R., Zurek, S., et al. (2016). Towards advanced robotic manipulation for nuclear decommissioning: A pilot study on tele-operation and autonomy. In 2016 International Conference on Robotics and Automation for Humanitarian Applications (RAHA). IEEE, 1–8.

Mastalli, C., Havoutis, I., Focchi, M., Caldwell, D. G., and Semini, C. (2016). Hierarchical planning of dynamic movements without scheduled contact sequences. IEEE International Conference on Robotics and Automation (ICRA) *(IEEE)*, 4636–4641.

Mordatch, I., Lowrey, K., and Todorov, E. (2015). Ensemble-cio: Full-body dynamic motion planning that transfers to physical humanoids. In 2015

IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 5307–5314.

Mordatch, I., Popović, Z., and Todorov, E. (2012a). Contact-invariant optimization for hand manipulation. In Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation. Eurographics Association), 137–144.

Mordatch, I., Todorov, E., and Popović, Z. (2012b). Discovery of complex behaviors through contact-invariant optimization. *ACM Transactions on Graphics (TOG)* 31, 43. doi:10.1145/2185520.2185539

Nagatani, K., Kiribayashi, S., Okada, Y., Otake, K., Yoshida, K., Tadokoro, S., et al. (2011a). Gamma-ray irradiation test of electric components of rescue mobile robot quince. In 2011 IEEE International Symposium on Safety, Security, and Rescue Robotics. IEEE, 56–60.

Nagatani, K., Kiribayashi, S., Okada, Y., Tadokoro, S., Nishimura, T., Yoshida, T., et al. (2011b). Redesign of rescue mobile robot quince. In 2011 IEEE International Symposium on Safety, Security, and Rescue Robotics. IEEE, 13–18.

Neunert, M., Farshidian, F., and Buchli, J. (2016). Efficient whole-body trajectory optimization using contact constraint relaxation. In 2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids). IEEE, 43–48.

Neunert, M., Farshidian, F., Winkler, A. W., and Buchli, J. (2017). Trajectory optimization through contacts and automatic gait discovery for quadrupeds. *IEEE Robot. Autom. Lett.* 2, 1502–1509. doi:10.1109/lra.2017.2665685

Neunert, M., Stäuble, M., Giftthaler, M., Bellicoso, C. D., Carius, J., Gehring, C., et al. (2018). Whole-body nonlinear model predictive control through contacts for quadrupeds. *IEEE Robot. Autom. Lett.* 3, 1458–1465. doi:10.1109/lra.2018.2800124

Norton, A., Ober, W., Baraniecki, L., McCann, E., Scholtz, J., Shane, D., et al. (2017). Analysis of human-robot interaction at the DARPA Robotics Challenge Finals. *The International Journal of Robotics Research* 36, 483–513. doi:10.1177/0278364916688254

Önol, A. Ö., Corcodel, R., Long, P., and Padır, T. (2020). Tuning-free contact-implicit trajectory optimization. In 2020 IEEE International Conference on Robotics and Automation (ICRA). IEEE.

Önol, A. Ö., Long, P., and Padır, T. (2018). A comparative analysis of contact models in trajectory optimization for manipulation. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE.

Önol, A. O., Long, P., and Padır, T. (2018). *Using contact to increase robot performance for glovebox d&d tasks* (arXiv preprint arXiv:1807.04198).

Önol, A. Ö., Long, P., and Padır, T. (2019). Contact-implicit trajectory optimization based on a variable smooth contact model and successive convexification. In 2019 IEEE International Conference on Robotics and Automation (ICRA). IEEE, 2447–2453.

Otte, S., Weiss, C., Scherer, T., and Zell, A. (2016). Recurrent neural networks for fast and robust vibration-based ground classification on mobile robots. *IEEE International Conference on Robotics and Automation (ICRA)*, 5603–5608. doi:10.1109/ICRA.2016.7487778

Park, J., and Khatib, O. (2006). Contact consistent control framework for humanoid robots. In Robotics and Automation, ICRA 2006. Proceedings 2006 IEEE International Conference on. IEEE, 1963–1969.

Park, J., and Khatib, O. (2008). Robot multiple contact control. *Robotica* 26, 667–677. doi:10.1017/s0263574708004281

Posa, M., Cantu, C., and Tedrake, R. (2014). A direct method for trajectory optimization of rigid bodies through contact. *The International Journal of Robotics Research* 33, 69–81. doi:10.1177/0278364913506757

Radford, N. A., Strawser, P., Hambuchen, K., Mehling, J. S., Verdeyen, W. K., Donnan, A. S., et al. (2015). Valkyrie: NASA's First Bipedal Humanoid Robot. *J. Field Robotics* 32, 397–419. doi:10.1002/rob.21560

Ragaglia, M., Bascetta, L., Rocco, P., and Zanchettin, A. M. (2014). Integration of perception, control and injury knowledge for safe human-robot interaction. In IEEE Int. Conf. on Robotics and Automation. IEEE, 1196–1202.

Rasheed, T., Long, P., Marquez-Gamez, D., and Caro, S. (2018). Tension distribution algorithm for planar mobile cable-driven parallel robots. In Cable-Driven Parallel Robots. Springer, 268–279. doi:10.1007/978-3-319-61431-1_23

Skonieczny, K., Moreland, S. J., Asnani, V. M., Creager, C. M., Inotsume, H., and Wettergreen, D. S. (2014). Visualizing and analyzing machine-soil interactions using computer vision. *J. Field Robotics* 31, 820–836. doi:10.1002/rob.21510

Sleiman, J.-P., Carius, J., Grandia, R., Wermelinger, M., and Hutter, M. (2019). Contact-implicit trajectory optimization for dynamic object manipulation. In 2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE.

Stentz, A., Herman, H., Kelly, A., Meyhofer, E., Haynes, G. C., Stager, D., et al. (2015). Chimp, the cmu highly intelligent mobile platform. *J. Field Robotics* 32, 209–228. doi:10.1002/rob.21569

Stewart, D. E., and Trinkle, J. C. (1996). An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction. *Int. J. Numer. Meth. Engng.* 39, 2673–2691. doi:10.1002/(sici)1097-0207(19960815)39:15<2673::aid-nme972>3.0.co;2-i

Tassa, Y., Erez, T., and Todorov, E. (2012). and stabilization of complex behaviors through online trajectory optimization. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE), 4906–4913.

Todorov, E., Erez, T., and Tassa, Y. (2012). MuJoCo: A physics engine for model-based control. In 2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 5026–5033.

Tsai, M.-J. (1986). *Workspace geometric characterization and manipulability of industrial robots*. Ph.D. thesis, The Ohio State University.

Urabe, E., and Stapczynski, S. (2017). *Fukushima fuel-removal quest leaves trail of dead robots*.

Vahrenkamp, N., and Asfour, T. (2015). Representing the robot's workspace through constrained manipulability analysis. *Auton. Robot.* 38, 17–30. doi:10.1007/s10514-014-9394-z

Vahrenkamp, N., Asfour, T., and Dillmann, R. (2013). Robot placement based on reachability inversion. In 2013 IEEE Int. Conf. on Robotics and Automation. IEEE, 1970–1975.

Vahrenkamp, N., Asfour, T., Metta, G., Sandini, G., and Dillmann, R. (2012). Manipulability analysis. In IEEE-RAS Int. Conf. on Humanoid Robots. IEEE, 568–573.

Vukobratović, M., and Borovac, B. (2004). Zero-moment point—thirty five years of its life. *International journal of humanoid robotics* 1, 157–173.

Walas, K., Kanoulas, D., and Kryczka, P. (2016). Terrain classification and locomotion parameters adaptation for humanoid robots using force/torque sensing. In 2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids). 133–140. doi:10.1109/HUMANOIDS.2016.7803265

Wang, J., and Olson, E. (2016). AprilTag 2: Efficient and robust fiducial detection. In 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 4193–4198. doi:10.1109/IROS.2016.7759617

Weiss, C., Tamimi, H., and Zell, A. (2008). A combination of vision- and vibration-based terrain classification. In 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems. 2204–2209. doi:10.1109/IROS.2008.4650678

Will Bosworth1, S. K., Jonas, W., and Hogan, N. (2016). Robot locomotion on hard and soft ground: measuring stability and ground properties *in-situ*. In IEEE International Conference on Robotics and Automation (ICRA). 3582–3589.

Winkler, A. W., Bellicoso, C. D., Hutter, M., and Buchli, J. (2018). Gait and trajectory optimization for legged systems through phase-based end-effector parameterization. *IEEE Robot. Autom. Lett.* 3, 1560–1567. doi:10.1109/lra.2018.2798285

Wittmann, R., Hildebrandt, A.-C., Wahrmann, D., Sygulla, F., Rixen, D., and Buschmann, T. (2016). Model-based predictive bipedal walking stabilization. In 2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids). IEEE, 718–724.

Wu, Y., Yao, D., and Xiao, X. (2018). The effects of ground compliance on flexible planar passive biped dynamic walking. *J. Mech. Sci. Technol.* 32, 1793–1804. doi:10.1007/s12206-018-0336-0

Yanco, H. A., Norton, A., Ober, W., Shane, D., Skinner, A., and Vice, J. (2015). Analysis of human-robot interaction at the darpa robotics challenge trials. *J. Field Robotics* 32, 420–444. doi:10.1002/rob.21568

Yirmibeşoğlu, OD, Oshiro, T, Olson, G, Palmer, C, and Mengüç, Y (2019). Evaluation of 3d printed soft robots in radiation environments and

comparison with molded counterparts. *Front. Robot. AI* 6, 40. doi:10.3389/frobt.2019.00040

Yoshikawa, T. (1984). Analysis and control of robot manipulators with redundancy. In Robotics research: the first Int. symp. Cambridge, MA): Mit Press, 735–747.

Yunt, K., and Glocker, C. (2005). Trajectory optimization of mechanical hybrid systems using SUMT. In 9th IEEE International Workshop on Advanced Motion Control. IEEE, 665–671.

Zollo, L., Accoto, D., Torchiani, F., Formica, D., and Guglielmelli, E. (2008). Design of a planar robotic machine for neuro-rehabilitation. In IEEE Int. Conf. on Robotics and Automation. IEEE, 2031–2036.

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Advantages of publishing in Frontiers

OPEN ACCESS
Articles are free to read for greatest visibility and readership

FAST PUBLICATION
Around 90 days from submission to decision

HIGH QUALITY PEER-REVIEW
Rigorous, collaborative, and constructive peer-review

TRANSPARENT PEER-REVIEW
Editors and reviewers acknowledged by name on published articles

Frontiers
Avenue du Tribunal-Fédéral 34
1005 Lausanne | Switzerland

Visit us: www.frontiersin.org
Contact us: frontiersin.org/about/contact

REPRODUCIBILITY OF RESEARCH
Support open data and methods to enhance research reproducibility

DIGITAL PUBLISHING
Articles designed for optimal readership across devices

FOLLOW US
@frontiersin

IMPACT METRICS
Advanced article metrics track visibility across digital media

EXTENSIVE PROMOTION
Marketing and promotion of impactful research

LOOP RESEARCH NETWORK
Our network increases your article's readership