



A Serverless Tool for Platform Agnostic Computational Experiment Management

Gregory Kiar^{1,2*}, Shawn T. Brown¹, Tristan Glatard^{3†} and Alan C. Evans^{1,2,4†}

¹ Montreal Neurological Institute, McGill University, Montreal, QC, Canada, ² Department of Biomedical Engineering, McGill University, Montreal, QC, Canada, ³ Department of Computer Science, Concordia University, Montreal, QC, Canada, ⁴ Department of Neurology and Neurosurgery, McGill University, Montreal, QC, Canada

OPEN ACCESS

Edited by:

Andrew P. Davison,
Institut des Neurosciences
Paris-Saclay (Neuro-PSI), France

Reviewed by:

Fidel Santamaria,
The University of Texas at
San Antonio, United States
Werner Van Geit,
École Polytechnique Fédérale
de Lausanne, Switzerland
Daniel Wheeler,
National Institute of Standards
and Technology (NIST), United States

*Correspondence:

Gregory Kiar
gregory.kiar@mail.mcgill.ca

[†] These authors have contributed
equally to this work

Received: 31 August 2018

Accepted: 15 February 2019

Published: 05 March 2019

Citation:

Kiar G, Brown ST, Glatard T and
Evans AC (2019) A Serverless Tool
for Platform Agnostic Computational
Experiment Management.
Front. Neuroinform. 13:12.
doi: 10.3389/fninf.2019.00012

Neuroscience has been carried into the domain of big data and high performance computing (HPC) on the backs of initiatives in data collection and an increasingly compute-intensive tools. While managing HPC experiments requires considerable technical acumen, platforms, and standards have been developed to ease this burden on scientists. While web-portals make resources widely accessible, data organizations such as the Brain Imaging Data Structure and tool description languages such as Boutiques provide researchers with a foothold to tackle these problems using their own datasets, pipelines, and environments. While these standards lower the barrier to adoption of HPC and cloud systems for neuroscience applications, they still require the consolidation of disparate domain-specific knowledge. We present Clowdr, a lightweight tool to launch experiments on HPC systems and clouds, record rich execution records, and enable the accessible sharing and re-launch of experimental summaries and results. Clowdr uniquely sits between web platforms and bare-metal applications for experiment management by preserving the flexibility of do-it-yourself solutions while providing a low barrier for developing, deploying and disseminating neuroscientific analysis.

Keywords: cloud computing, reproducibility and tools, high performance computing, microservice, interactive visualization

INTRODUCTION

The increasing adoption of distributed computing, including cloud and high-performance computing (HPC), has played a crucial role in the expansive growth of neuroscience. With an emphasis on big-data analytics, collecting large datasets such as the Consortium for Reliability and Reproducibility (Zuo et al., 2014), UK-Biobank (Sudlow et al., 2015), and Human Connectome Project (Van Essen et al., 2013) is becoming increasingly popular and necessary. While these datasets provide the opportunity for unprecedented insight into human brain function, their size makes non-automated analysis impractical.

At the backbone of science is the necessity that claims are reproducible. The reproducibility of findings has entered the spotlight as a key question of interest, and has been explored extensively in psychology (Open Science Collaboration, 2015), neuroimaging (Eklund et al., 2016; Bowring et al., 2018), and other domains (Baker, 2016; Miłkowski et al., 2018). Computational experiments must be re-executable as a critical condition for reproducibility, and this bare minimum requirement becomes increasingly challenging with larger datasets and more complex

analyses. While sharing all code and data involved may appear a compelling solution, this is often inadequate for achieving re-runnability or reproducibility of the presented findings and models (Miłkowski et al., 2018). When re-executable applications fail to reproduce findings, there is a gray area where the source of errors are often unknown and may be linked to misinterpretation of data, computing resources or undocumented execution details, rather than scientific meaning.

As a result, new tools and standards have emerged to aid in producing more reusable datasets and tools, and thereby more reproducible science. The Brain Imaging Data Structure (BIDS) (Gorgolewski et al., 2016) and the associated BIDS apps (Gorgolewski et al., 2017) prescribe a standard for sharing and accessing datasets, and therefore, increasing the accessibility and impact of both datasets and tools. This standard includes the definition of file organization on disk, as well as key-value pairs of metadata information in JavaScript Object Notation (JSON) files, and assigns specific meaning to command-line arguments to be used when processing these datasets. The Boutiques framework (Glatard et al., 2018) provides a standard for software documentation in a machine-interpretable way, allowing the automation of tool execution and evaluation. These descriptions fully encapsulate the runtime instructions for a given tool in JSON-files, and are appropriate for a majority of command-line applications. Software containerization initiatives such as Docker (Merkel, 2014) and Singularity (Kurtzer et al., 2017) facilitate execution consistently across arbitrary computing environments with minimal burden on the user.

Web-platforms such as OpenNeuro (Poldrack et al., 2013), LONI Pipeline (Rex et al., 2003), and CBRAIN (Sherif et al., 2014) simplify the analysis process further by providing an accessible way to construct neuroscience experiments on commonly used tools and uploaded-datasets. These systems deploy tools on HPC environments and record detailed execution information so that scientists can keep accurate records and debug their workflows. Tools such as LONI's provenance manager (Dinov et al., 2010), Reprozip (Chirigati et al., 2016), and ReCAP (Hasham et al., 2018) capture system-level properties such as system resources consumed and files accessed, where tools supporting the Neuroimaging Data Model (NIDM) (Sochat and Nichols, 2016), a neuroimaging-specific provenance model based on W3C-PROV (Missier et al., 2013), capture information about the domain-specific transformations applied to the data of interest.

The initiatives enumerated above have synergistic relationships, where each solves a small but significant piece of the larger puzzle that is computational and scientific reproducibility and replicability. However, the learning curve associated with adopting each of these technologies is considerable, leveraging them in an impactful way is difficult, and certain applications may benefit from different approaches so these learning curves may have to be paid multiple times. For instance, interoperability is mainly valuable in contexts which there is a large variety of datasets or tools, and provenance may be of importance to identify the impact of an underlying system on a processing or modeling task. We present Clowdr, a lightweight tool which ties these approaches together so that researchers can minimize the learning burden and lower the

barrier to develop, perform, and disseminate reproducible, interoperable and provenance-rich neuroscience experiments.

EMERGENT TECHNOLOGIES IN REPRODUCIBLE NEUROSCIENCE

Conducting reproducible analyses in neuroscience requires many complementary facets, building on technologies which are commonly adopted as *de facto* standards.

Data and Code Interoperability

Due in part to its simplicity and active public development community, BIDS (Gorgolewski et al., 2016) has become an increasingly prominent data organization format in neuroimaging. This standard makes use of the Nifti file format (Cox et al., 2004) and human-readable JSON files to capture both imaging and subject-specific metadata. An important benefit of this data organization is the ability to launch data processing pipelines in the form of BIDS applications (Gorgolewski et al., 2017), which expose a consistent set of instructions compatible with the data organization. Together, these complementary standards are suitable for performing a large variety of neuroimaging experiments. In contexts where tools have heterogeneous interfaces, or data organizations are custom-built for a particular context, the Boutiques (Glatard et al., 2018, 2019) framework allows the rich description of a pipeline such that tool execution, validation, and exploration can be automated. These descriptors include the command-line structure to be populated as well as rich parameter descriptions and interactions, such as mutually exclusivity or dependence, such that complicated data interactions required or forbidden by the tool can be accounted for.

Software Virtualization

While virtual machines have long been used for deploying analysis pipelines with complex dependencies in heterogeneous environments, software containers have recently emerged as lighter-weight alternatives suitable for transient data processing applications. Docker (Merkel, 2014) provides this functionality across all major host operating systems, but is often not supported by HPC centers due to security vulnerabilities (Bui, 2015; Combe et al., 2016). Singularity (Kurtzer et al., 2017) addresses the security risks of Docker, but currently only supports Linux operating systems, filling the niche of containerization on shared computing resources. A detailed comparison in the context of medical imaging can be found in Matelsky et al. (2018).

Workflow Engines

Custom scientific pipelines can be composed in Python with Nipype (Gorgolewski et al., 2011), Dask (Rocklin, 2015), Pegasus (Deelman et al., 2015), Toil (Vivian et al., 2017), or several other tools which facilitate the modular interaction of complex independent processing stages. While the underlying tasks in Nipype, Dask, and Toil are defined in Python, Pegasus uses a Domain Specific Language (DSL) for representing tasks, increasing the barrier for defining tasks but ultimately

increasing their portability. While Nipype is a widely used tool in neuroimaging and has many readily-defined interfaces available for researchers, the others require non-insignificant development to describe interfaces for common neuroimaging applications such as FSL (Jenkinson et al., 2012) or MRtrix (Tournier and Calamante, 2012). PSOM (Bellec et al., 2012) and Scipipe (Lampa et al., 2018) are functionally similar to Nipype but have been developed for GNU Octave/MATLAB and Golang, respectively. Several domains have more specialized tools which accomplish similar feats in their area of interest. These include Pypet (Meyer and Obermayer, 2016), Neuromanager (Stockton and Santamaria, 2015), Arachne (Aleksin et al., 2017), and others which facilitate the automation of modeling and simulation workflows using tools such as Neuron (Hines and Carnevale, 2001). For an in depth look at other tools in this space please refer to Meyer and Obermayer (2016) and (Stockton and Santamaria, 2015). Each of these tools enables the construction of dependency graphs between pipeline components, and allow the deployment either to cluster scheduling software, multiple processing threads, and in some cases computing clouds. These tools primarily function through programmatic interfaces, though LONI pipeline (Rex et al., 2003), OpenMOLE (Reuillon et al., 2013), and Galaxy (Goecks et al., 2010) provide both DSL and graphical user interfaces. Many of these tools also embed provenance capture, fault-tolerance features, and data tracking to avoid recomputations across similar executions. While each of these tools is a powerful and attractive option for creating workflows, they remain complex and potentially overkill when launching atomic single-step analyses, prebuilt pipelines, or analysis software developed in a different language than the workflow engine of choice.

Provenance

Building on the W3-PROV (Missier et al., 2013) standard for data provenance metadata put forth by the World Wide Web Consortium, NIDM (Sochat and Nichols, 2016) is a standard which represents a processing and data provenance graph specific to neuroimaging analyses. While this standard is machine-interpretable and interoperable-by-design, supporting it currently requires tight integration with analysis pipelines. In LONI pipeline, a provenance model exists which includes detailed records of data use and file lifecycle (Dinov et al., 2010), which is designed to inform data consumers what types of analyses can be and have been performed with the data in question; this tool is tightly coupled with the LONI pipeline ecosystem. The ReCAP (Hasham et al., 2018) project has been developed to evaluate the resource consumption of arbitrary pipelines on the cloud and can aid in cloud-instance optimization. While this tool has potential for a large impact in designing both cost effective and scalable analyses, there is considerable overhead as it manages executions through a persistent server and workflow engine. While various other libraries exist to monitor some piece of data or compute provenance, Reprozip (Chirigati et al., 2016) is perhaps the most exciting as it uniquely captures records of all files accessed and created throughout an

execution, which allows for the creation of rich file dependency graphs. The limitation of this technique is that it requires data of interest to be written to disk, as opposed to managed in memory, which may not always be the case in some applications.

Web Platforms

Increasing the portability and accessibility of launching large scale analyses, web platforms such as CBRAIN (Sherif et al., 2014), LONI pipeline (Rex et al., 2003), and OpenNeuro (Poldrack et al., 2013) provide science-as-a-service where users can upload and process their data on distant computing resources. Additionally, these platforms provide an accessible and immediate way to share the results produced from experiments with collaborators or the public. These tools provide incredible value to the community and allow the deployment of production-level pipelines from the web, but they are not suitable for prototyping analyses or developing pipelines, and it is cumbersome to run these services on a lab's own resources. In addition, monolithic Web interfaces are only suitable for a certain type of use-cases and high-level users, while developers or computer-savvy users prefer to rely on modular command-line tools and libraries.

THE CLOWDR MICROTOOL

While the technologies enumerated above are essential pieces toward reproducible neuroscience, they are largely isolated from one another and place a large burden on researchers who wish to adopt all of these best practices. Clowdr leverages these tools to increase the deployability, provenance capture, and shareability of experiments. In summary, Clowdr:

- (i) is tightly based on Boutiques and is BIDS-aware, supporting both arbitrary pipelines and providing an accessible endpoint for neuroimaging;
- (ii) executes both bare-metal workflows and Docker or Singularity virtualized pipelines through Boutiques on local, HPC, and cloud resources;
- (iii) supports the parallelized batch deployment and redeployment of pipelines constructed with workflow-engines, while being agnostic to programming language and construct;
- (iv) captures system-level provenance information (i.e., CPU and RAM usage), supports Reprozip, and internal provenance captured by arbitrary pipelines such as NIDM; and
- (v) supports the deployment of both development- and production-level tools without an active server, and provides a web-report for exploring and sharing experiments.

A typical workflow using Clowdr is summarized in **Figure 1**. While a Clowdr experiment follows the same workflow as traditional experiments, beginning with tool and data curation through prototyping, deployment, and exploration, there are several considerable benefits provided by Clowdr

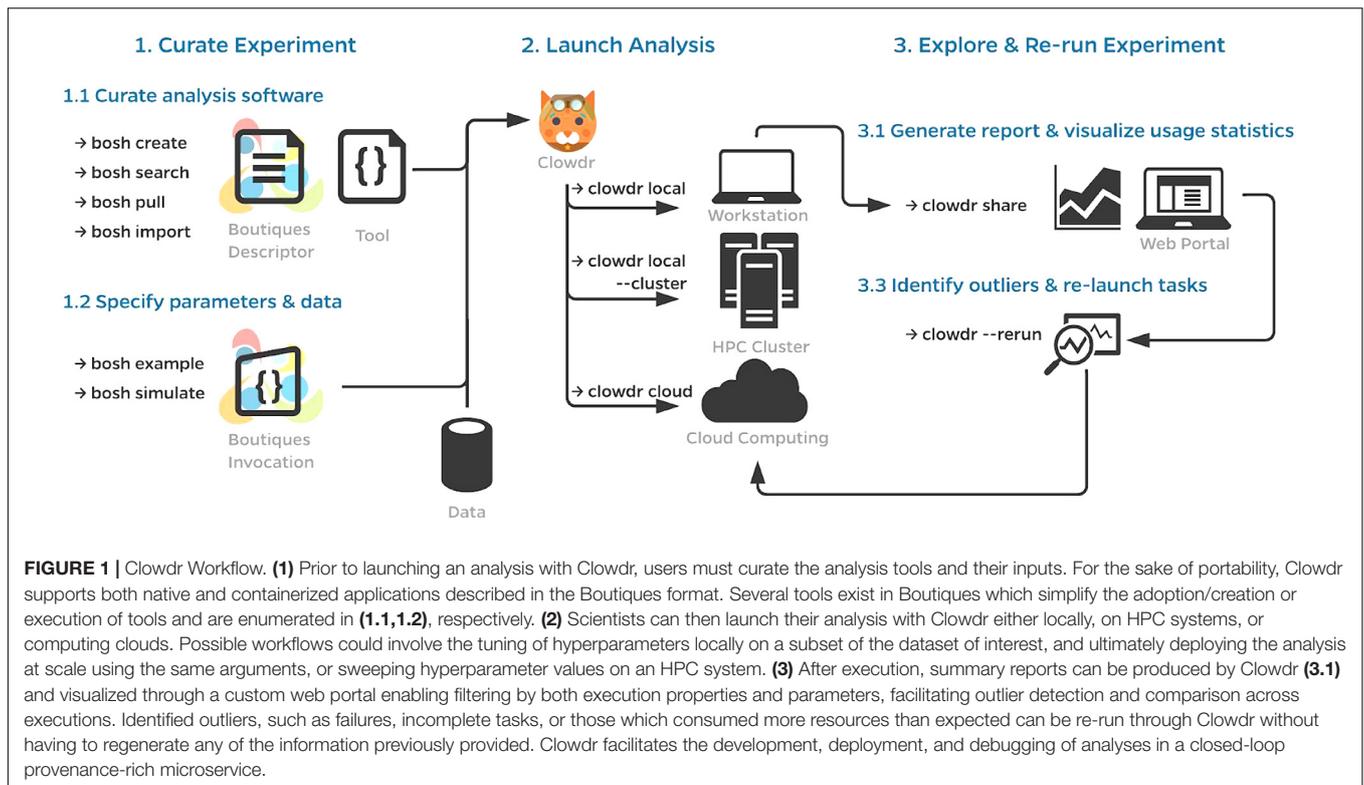


FIGURE 1 | Clowdr Workflow. **(1)** Prior to launching an analysis with Clowdr, users must curate the analysis tools and their inputs. For the sake of portability, Clowdr supports both native and containerized applications described in the Boutiques format. Several tools exist in Boutiques which simplify the adoption/creation or execution of tools and are enumerated in **(1.1,1.2)**, respectively. **(2)** Scientists can then launch their analysis with Clowdr either locally, on HPC systems, or computing clouds. Possible workflows could involve the tuning of hyperparameters locally on a subset of the dataset of interest, and ultimately deploying the analysis at scale using the same arguments, or sweeping hyperparameter values on an HPC system. **(3)** After execution, summary reports can be produced by Clowdr **(3.1)** and visualized through a custom web portal enabling filtering by both execution properties and parameters, facilitating outlier detection and comparison across executions. Identified outliers, such as failures, incomplete tasks, or those which consumed more resources than expected can be re-run through Clowdr without having to regenerate any of the information previously provided. Clowdr facilitates the development, deployment, and debugging of analyses in a closed-loop provenance-rich microservice.

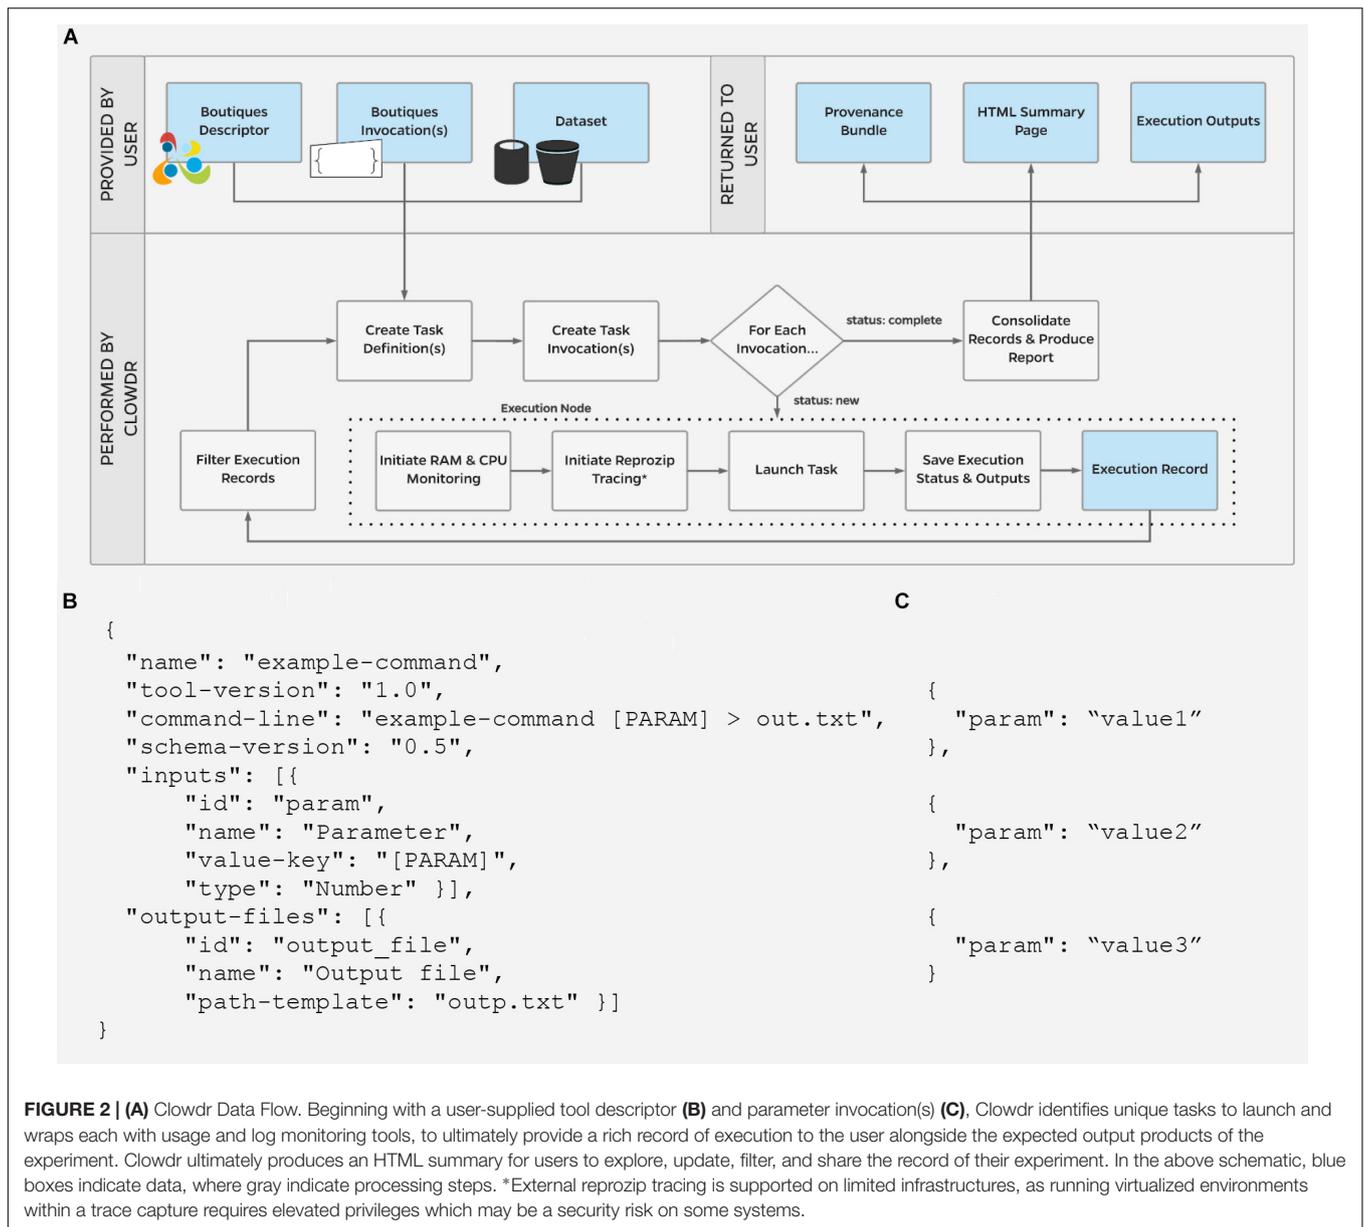
over traditional approaches. In particular, Clowdr is based on the rich Boutiques framework for tool description and execution, ensuring that documentation, parameter definitions, and real-world parameter values accompany the tool at all times. Clowdr also treats all computing systems the same, from the users perspective, so transitioning from local development of analyses to at-scale systems is seamless, which minimizes errors made during this transition. Clowdr also provides a visualization portal for exploring executions and filtering either based on parameter values or runtime statistics, allowing for quality control of the execution in addition to commonly used quality control of processed derivatives themselves.

Figure 2 shows the execution lifecycle within Clowdr. Starting from user-provided Boutiques descriptor (B) and invocation(s) (C), and access to any required datasets, Clowdr begins by identifying a list of tasks to launch. For a new experiment, tasks are identified in one of three main ways: (1) a one:one mapping from a list of invocations, (2) a one:many mapping from a single invocation in which parameter(s) have been specified for sweeping during execution, or a BIDS-specific, and (3) one:many mapping from a single invocation for a BIDS app, which will iterate upon the participant- and session-label fields, and described in the BIDS app specification (Gorgolewski et al., 2017). Experiments can be re-run, and determine the task-list based on whether a full, failure-only, or incomplete-only re-execution is desired. Once the task-list is determined, Clowdr creates an independent invocation which explicitly defines the arguments used in each task.

At this stage, Clowdr distributes tasks to the Cloud system or local cluster scheduler being used for deployment. Presently Clowdr supports the SLURM scheduler and Amazon Web Services (AWS) cloud through their Batch service with adoption of more platforms ongoing. Each task is launched through a Clowdr-wrapper, which initializes CPU and RAM monitoring and triggers Rezip tracing prior to launching the analysis itself. Rezip tracing has limited support in conjunction with containerized analyses on HPC systems due to potential security issues. Rezip is built upon the Linux command “ptrace,” which traces processes to monitor or control them. To eliminate the potential risk of using this tool, it is common for systems to disallow the tracing of administrator-level processes. The requirement of limited administrator privileges by Singularity (for interacting with the daemon) makes encapsulating these tools within the restricted ptrace scope not possible on many shared systems. For more information on the specific conditions in which these technologies can be made to interoperate please view the GitHub repository for this manuscript, linked below.

Upon completion of the analysis, Clowdr bundles the system monitored records, standard output and error, exit status, and any other information collected by either the tool itself or the Boutiques runtime engine, and concludes its execution. Once the experiment has begun, Clowdr provides the user with the Clowdr provenance directory which will be updated automatically as executions progress.

The researcher can monitor the provenance directory using the Clowdr share portal (**Figure 3**), which provides



a web interface summarizing the task executions. Once the analysis concludes, the figures on this web page and the associated metadata can be saved and serve as a record of the experiment either for evaluation or dissemination alongside published results.

The Clowdr package is open-source on GitHub (Kiar and Glatard, 2019), and installable through the Python Package Index.

PERFORMING EXPERIMENTS WITH CLOWDR

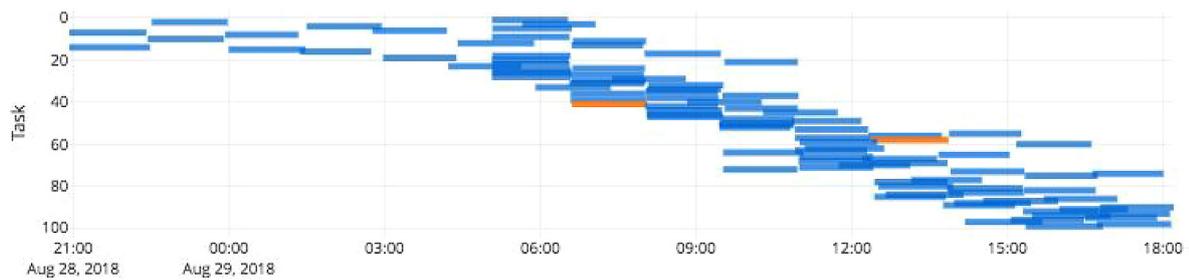
Here, we explore an experiment in which we used Clowdr to process the Human Connectome Project (HCP) (Van Essen et al., 2013) dataset with a structural and functional

connectome estimation pipeline, ndmg (Kiar et al., 2018a,b). The records of this experiment, and materials and instructions that can be used to reproduce a similar analysis with Clowdr using the publicly available DS114 BIDS dataset (Poldrack et al., 2013) and the example BIDS application (Gorgolewski et al., 2017) can be found on Github at: <https://github.com/clowdr/clowdr-paper>. Specific packages and their versions for both experiments can be found at the end of this manuscript.

As summarized above, performing an analysis with Clowdr requires the creation of a Boutiques descriptor summarizing the pipeline of interest, an invocation containing the parameters to supply to this pipeline on execution, and curation of the data to be processed. There are several utilities in Boutiques which aid in this setup process, including to automatically generate a

Statistics				Invocations		
FILTER ROWS						
<input type="checkbox"/>	Task	Tool	Duration (s)	Max RAM (MB)	▲ Max CPU (%)	Exit Code
<input type="checkbox"/>	58	ndmg				Incomplete
<input type="checkbox"/>	0	ndmg				Incomplete
<input checked="" type="checkbox"/>	80	ndmg	9.95317149162292	189.88671875	55.33	0
<input checked="" type="checkbox"/>	46	ndmg	5225.12778949737	8700.62178125	198.38	0
<input type="checkbox"/>	18	ndmg	5302.22937703132	8318.92759375	198.43	0
<input checked="" type="checkbox"/>	61	ndmg	5370.41310644149	7760.80071875	198.52	0

Experiment Timeline



Usage Stats

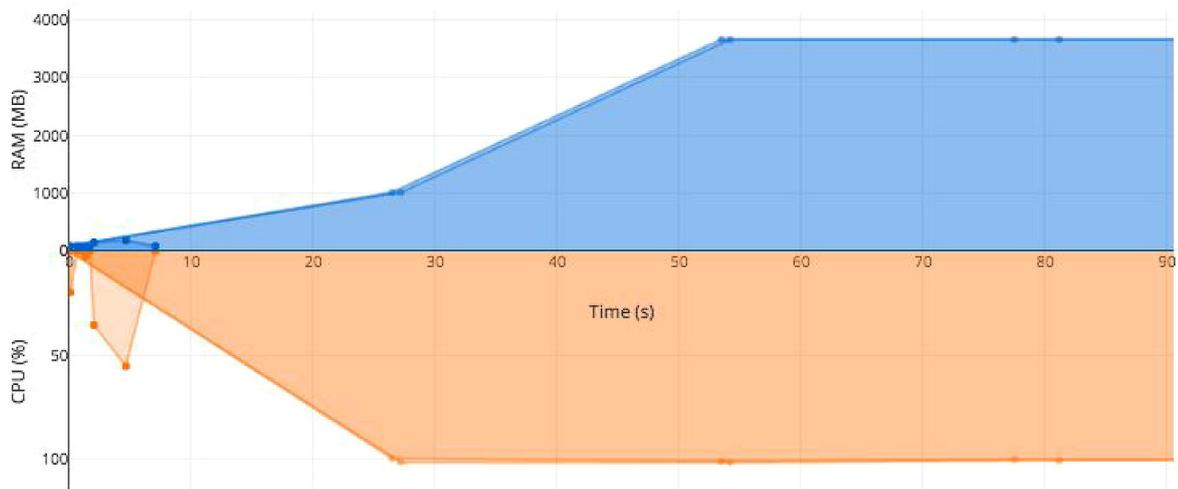


FIGURE 3 | Cloudr Experiment Viewer. Experiments launched with Cloudr can be monitored and both progress and runtime statistics explored. The page is produced using Plotly Dash to produce highly interactive plots and tables, enabling rich filtering, rescaling, and exploration of executions. The table can be toggled to present summary statistics about experiment execution or invocation parameters identifying parameters used for each task in the experiment. The subsequent Gantt plot shows the timeline of executed tasks in the experiment, where those selected for visualization in the usage plot below are highlighted. The final plot in this view shows the memory and processing footprint throughout all selected tasks. Selection and filtering may be done by value in the tables or selection in the task timeline. In this example, several tasks did not complete and one appeared to exit after 10 s erroneously. The Cloudr portal enables quick identification of these outliers, and the table view can be switched to identify more information such as parameters pertaining to the executions of interest. For more information about the pipeline being executed in particular, please see Kiar et al. (2018b).

descriptor and sample inputs from a tool, and can be explored in the associated documentation (Glatard et al., 2019).

Clowdr experiments can be launched locally, on cluster, or submitted to cloud resources. In each case, invocations and task definitions are created locally, and then the jobs are run serially, submitted to a cluster queue, or pushed to cloud storage and called remotely. The commands used in Clowdr to launch these commands are the local, local with the cluster switch, and cloud modes, respectively. Upon completion of each tasks, summary files created by Clowdr can be either inspected manually or consolidated and visualized in the web with the Clowdr share command (Figure 3).

The share tool, launchable on any computer with access to the experiment, creates a lightweight web service displaying summary statistics and invocation information from the experiment, including memory usage, task duration, launch order, and log information. The visualizations provided are filterable and sortable, enabling users to interrogate and identify outliers in their experiment, explore potential sources of failure, and effectively profile the analysis pipeline in use. The modified figures can be downloaded from this interface, serving as accessible records of execution.

In the example above, the HCP dataset has been processed using a pipeline performing image denoising, registration, model fitting, and connectivity estimation, all of which are commonly used processing steps in neuroimaging. For more information on this pipeline, please see Kiar et al. (2018b).

In this experiment the table has been filtered to show several tasks which appeared spurious in their execution compared to the others. We can see that several tasks failed to complete and one appeared to terminate in significantly less time than the others. After identifying these tasks and exploring the time series' to see at what stage of processing the job failed (in this case, immediately), we can investigate parameter selections used in each and attempt the re-execution of these jobs using the local or cloud command with the rerun switch in Clowdr. Clowdr provides a layer of quality control on executions, in addition to that which is regularly performed by researchers on their datasets, which provides immediate value when identifying task failures which otherwise may be difficult to identify, especially in cases which intermediate and terminal derivatives are written to the same location, which can often be the case with transformations estimated by registration pipelines, for example.

While the share tool currently requires maintaining an active server, the plots can be exported statically and it is in the development roadmap to enable exporting the entire web page as static files, as discussed here: <https://github.com/plotly/dash/issues/266>. Since the record created by Clowdr is stored in the machine-readable and JSON format, researchers can easily extract their records and integrate it into other interfaces that suit their application.

DISCUSSION

Clowdr addresses several barriers to performing reproducible neuroscience. Clowdr experiments consist of enclosed

computational environments, versioned-controlled Boutiques-described tools with explicit usage parameters, rich execution history, and can be re-executed or distributed with minimal effort. Clowdr provides an accessible interface for initially running analyses locally, and translating them seamlessly to HPC environments. The rich record keeping provided with Clowdr is system-agnostic resulting in uniformly interpretable summaries of execution. As a Python library, Clowdr can be used as a module in a larger platform, or directly as a command-line tool.

Clowdr uniquely packages an executable tool summary, parameters, and results together, in a language- and tool-agnostic way, and therefore, greatly increases the transparency and shareability of experiments. Importantly, this adds clarity to experimental failures and documents the hyper-parameter tuning process of experiments, which has been historically largely undocumented in literature (Reunanen, 2003).

There are several axes upon which the value of Clowdr can be discussed. In particular: lines of code written, time spent, and the ultimate re-runnability of analyses. While these remain subjective areas for comparison, we can conceptually consider a workflow dependent on Clowdr to those constructed with traditional scripting, workflow engines (WEs), and software-as-a-service (SaaS) platforms.

Where Clowdr has been built upon tools and standards to provide users with a series of single-commands for launching and managing analyses, accomplishing a similar result with traditional scripting would take considerably more lines of code and time. Similarly, where command-line execution may be similar in complexity to tools developed with WEs, their integration within tools requires substantial development and is only practical in cases for which there is a WE written in the same language as the underlying application. SaaS platforms provide a similar type of abstraction to Clowdr, where tools are treated as black-box objects, but come with the added overhead of maintaining complex database architectures, often complex integration of tools, and primarily restrict access through web-based interfaces which leads to reduced flexibility for the user.

The clear benefit of Clowdr is in the simplicity it provides for identifying outliers or failed tasks and either re-launching specific subsets of an analysis or the entire experiment. Clowdr records and visualizes detailed logging information about executions and the specific instructions which were used, which isn't guaranteed in either traditional scripting or WE-based applications. To replicate this feature across these systems, tools which (1) record execution instructions, (2) identify parameters used for parallelization, (3) produce summary plots, and (4) reconstruct and (5) re-execute instructions would require development.

While SaaS platforms often contain these features, an additional limitation of large platforms is that they are often designed for consumers of widely adopted tools consumers rather than tool developers. Clowdr fills the void between these types of pipeline deployments by providing a programmatic tool-independent method for managing job submission and collecting provenance across multiple architectures and enabling the rapid prototyping of analyses.

Several immediate applications of the provenance information captured by Clowdr include the benchmarking of tools, and

resource optimization during the selection of cloud resources, as was done in Hasham et al. (2018). While the value of comparing provenance records has not been demonstrated here, other studies such as (Salari et al., 2018) have demonstrated the efficacy of leveraging provenance information to identify sources of variability or instability within pipelines.

Future work includes adopting a W3C-PROV compatible format for Clowdr provenance records, increasing the machine-readability and interoperability of these records with other standards such as NIDM. Integrating the reports produced by Clowdr with a system such as Datalad would allow for record versioning and more strictly enforce the complete reporting of experiments. Clowdr will also continually be extended with greater testing and support for more HPC schedulers, clouds, and provenance capture models.

TOOLS AND VERSIONS

The following is a list of tools and data used in this manuscript, and their respective versions. The architecture and analysis presented for the Clowdr package corresponds to version 0.1.0. The key Python packages and specific versions tested are: boutiques (version 0.5.12), boto3 (1.7.81), botocore (1.10.81), slurmpy (0.0.7), psutil (5.4.7), pandas (0.23.4), plotly (3.1.1), and plotly dash (0.24.1), including dash-core-components (0.27.1), dash-html-components (0.11.0), dash-renderer (0.13.0), dash-table-experiments (0.6.0), and flask (0.12.2). Executions were tested locally using Docker (17.12.0-ce), and on Compute Canada's Cedar high performance cluster using Singularity (2.5.1-dist). The Docker container used for ndmg can be

found on Docker hub as neurodata/m3r-release (0.0.5), which contains ndmg (0.1.0-f). The Singularity container used was pulled and dynamically created from this Docker hub endpoint. The dataset use was a subset of the HCP 1200 collection (Van Essen et al., 2013).

AUTHOR CONTRIBUTIONS

GK designed and developed the tools, experiments, and figures, and wrote the majority of the manuscript. SB supported the design and development processes, and edited the manuscript and provided valuable feedback. TG provided insight and contributed to the design and development of the tools and experiments, and contributed to writing the manuscript. AE edited the manuscript and provided valuable feedback. TG and AE jointly supervised this project.

FUNDING

Funding for this work was provided by CFREF/HBHL (Canada First Research Excellence Fund/Healthy Brains for Healthy Lives) and the Natural Sciences and Engineering Research Council of Canada (CGSD3-519497-2018).

ACKNOWLEDGMENTS

The authors would like to thank Pierre Rioux and Valerie Hayot-Sasson for their insight and many helpful discussions.

REFERENCES

- Aleksin, S. G., Zheng, K., Rusakov, D. A., and Savtchenko, L. P. (2017). ARACHNE: a neural-neuroglial network builder with remotely controlled parallel computing. *PLoS Comput. Biol.* 13:e1005467. doi: 10.1371/journal.pcbi.1005467
- Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. *Nature* 533, 452–454. doi: 10.1038/533452a
- Bellec, P., Lavoie-Courchesne, S., Dickinson, P., Lerch, J. P., Zijdenbos, A. P., and Evans, A. C. (2012). The pipeline system for Octave and Matlab (PSOM): a lightweight scripting framework and execution engine for scientific workflows. *Front. Neuroinform.* 6:7. doi: 10.3389/fninf.2012.00007
- Bowring, A., Maumet, C., and Nichols, T. (2018). Exploring the impact of analysis software on task fMRI results. *bioRxiv* [Preprint]. doi: 10.1101/285585
- Bui, T. (2015). *Analysis of Docker Security*. Available at: <https://arxiv.org/abs/1501.02967v1>.
- Chirigati, F., Rampin, R., Shasha, D., and Freire, J. (2016). "ReproZip: computational reproducibility with ease," in *Proceedings of the International Conference on Management of Data*, San Francisco, CA, 2085–2088. doi: 10.1145/2882903.2899401
- Combe, T., Martin, A., and Pietro, R. D. (2016). To Docker or not to Docker: a security perspective. *IEEE Cloud Comput.* 3, 54–62. doi: 10.1109/MCC.2016.100
- Cox, R. W., Ashburner, J., Breman, H., Fissell, K., Haselgrove, C., Holmes, C. J., et al. (2004). A (Sort of) new image data format standard: NIfTI-1: WE 150. *Neuroimage* 22:e1440.
- Deelman, E., Vahi, K., Juve, G., Rynge, M., Callaghan, S., Maechling, P., et al. (2015). Pegasus, a workflow management system for science automation. *Future Gen. Comput. Syst.* 46, 17–35. doi: 10.1016/j.future.2014.10.008
- Dinov, I., Lozev, K., Petrosyan, P., Liu, Z., Eggert, P., Pierce, J., et al. (2010). Neuroimaging study designs, computational analyses and data provenance using the LONI pipeline. *PLoS One* 5:e13070. doi: 10.1371/journal.pone.0013070
- Eklund, A., Nichols, T. E., and Knutsson, H. (2016). Cluster failure: why fMRI inferences for spatial extent have inflated false-positive rates. *Proc. Natl. Acad. Sci. U.S.A.* 113, 7900–7905. doi: 10.1073/pnas.1602004113
- Glatard, T., Kiar, G., Aumentado-Armstrong, T., Beck, N., Bellec, P., Bernard, R., et al. (2018). Boutiques: a flexible framework to integrate command-line applications in computing platforms. *Gigascience* 7:giy016. doi: 10.1093/gigascience/giy016
- Glatard, T., Kiar, G., Benderoff, E., Tristan, A. A., Charrière, J., ValHayot, Simon-dube, et al. (2019). boutiques/boutiques: release 0.5.19 (version 0.5.19). *Zenodo*. doi: 10.5281/zenodo.2574166
- Goecks, J., Nekrutenko, A., Taylor, J., and Galaxy Team (2010). Galaxy: a comprehensive approach for supporting accessible, reproducible, and transparent computational research in the life sciences. *Genome Biol.* 11:R86. doi: 10.1186/gb-2010-11-8-r86
- Gorgolewski, K., Burns, C. D., Madison, C., Clark, D., Halchenko, Y. O., Waskom, M. L., et al. (2011). Nipype: a flexible, lightweight and extensible neuroimaging data processing framework in python. *Front. Neuroinform.* 5:13. doi: 10.3389/fninf.2011.00013
- Gorgolewski, K. J., Alfaro-Almagro, F., Auer, T., Bellec, P., Capotã, M., Chakravarty, M. M., et al. (2017). BIDS apps: improving ease of use, accessibility, and reproducibility of neuroimaging data analysis methods. *PLoS Comput. Biol.* 13:e1005209. doi: 10.1371/journal.pcbi.1005209
- Gorgolewski, K. J., Auer, T., Calhoun, V. D., Craddock, R. C., Das, S., Duff, E. P., et al. (2016). The brain imaging data structure, a format for organizing

- and describing outputs of neuroimaging experiments. *Sci. Data* 3:160044. doi: 10.1038/sdata.2016.44
- Hasham, K., Munir, K., and McClatchey, R. (2018). Cloud infrastructure provenance collection and management to reproduce scientific workflows execution. *Future Gen. Comput. Syst.* 86, 799–820. doi: 10.1016/j.future.2017.07.015
- Hines, M. L., and Carnevale, N. T. (2001). NEURON: a tool for neuroscientists. *Neuroscientist* 7, 123–135. doi: 10.1177/107385840100700207
- Jenkinson, M., Beckmann, C. F., Behrens, T. E. J., Woolrich, M. W., and Smith, S. M. (2012). FSL. *Neuroimage* 62, 782–790. doi: 10.1016/j.neuroimage.2011.09.015
- Kiar, G., Bridgeford, E., Roncal, W. G., Chandrashekhar, V., Mhembere, D., Burns, R., et al. (2018a). neurodata/ndmg: stable ndmg-DWI pipeline release (version v0.1.0). *Zenodo*. doi: 10.5281/zenodo.1161284
- Kiar, G., Bridgeford, E., Roncal, W. G., Chandrashekhar, V., Mhembere, D., Ryman, S., et al. (2018b). A high-throughput pipeline identifies robust connectomes but troublesome variability. *bioRxiv* [Preprint]. doi: 10.1101/188706
- Kiar, G., and Glatard, T. (2019). clowdr/clowdr: clowdr version 0.1.2 (version 0.1.2). *Zenodo*. doi: 10.5281/zenodo.2537168
- Kurtzer, G. M., Sochat, V., and Bauer, M. W. (2017). Singularity: scientific containers for mobility of compute. *PLoS One* 12:e0177459. doi: 10.1371/journal.pone.0177459
- Lampa, S., Dahlö, M., Alvarsson, J., and Spjuth, O. (2018). SciPipe - a workflow library for agile development of complex and dynamic bioinformatics pipelines. *bioRxiv* [Preprint]. doi: 10.1101/380808
- Matelsky, J., Kiar, G., Johnson, E., Rivera, C., Toma, M., and Gray-Roncal, W. (2018). Container-based clinical solutions for portable and reproducible image analysis. *J. Digit. Imaging* 31, 315–320. doi: 10.1007/s10278-018-0089-4
- Merkel, D. (2014). Docker: lightweight linux containers for consistent development and deployment. *Linux J.* 2014:2.
- Meyer, R., and Obermayer, K. (2016). pypet: a python toolkit for data management of parameter explorations. *Front. Neuroinform.* 10:38. doi: 10.3389/fninf.2016.00038
- Milkowski, M., Hensel, W. M., and Hohol, M. (2018). Replicability or reproducibility? On the replication crisis in computational neuroscience and sharing only relevant detail. *J. Comput. Neurosci.* 45, 163–172. doi: 10.1007/s10827-018-0702-z
- Missier, P., Belhajjame, K., and Cheney, J. (2013). “The W3C PROV family of specifications for modelling provenance metadata,” in *Proceedings of the 16th International Conference on Extending Database Technology*, Genoa, 773–776. doi: 10.1145/2452376.2452478
- Open Science Collaboration (2015). PSYCHOLOGY estimating the reproducibility of psychological science. *Science* 349:aac4716. doi: 10.1126/science.aac4716
- Poldrack, R. A., Barch, D. M., Mitchell, J. P., Wager, T. D., Wagner, A. D., Devlin, J. T., et al. (2013). Toward open sharing of task-based fMRI data: the OpenfMRI project. *Front. Neuroinform.* 7:12. doi: 10.3389/fninf.2013.00012
- Reuillon, R., Leclaire, M., and Rey-Coyrehourcq, S. (2013). OpenMOLE, a workflow engine specifically tailored for the distributed exploration of simulation models. *Future Gen. Comput. Syst.* 29, 1981–1990. doi: 10.1016/j.future.2013.05.003
- Reunanen, J. (2003). Overfitting in making comparisons between variable selection methods. *J. Mach. Learn. Res.* 3, 1371–1382.
- Rex, D. E., Ma, J. Q., and Toga, A. W. (2003). The LONI pipeline processing environment. *Neuroimage* 19, 1033–1048. doi: 10.1016/S1053-8119(03)00185-X
- Rocklin, M. (2015). “Dask: parallel computation with blocked algorithms and task scheduling,” in *Proceedings of the 14th Python in Science Conference*, Austin, TX. doi: 10.25080/Majora-7b98e3ed-013
- Salari, A., Scaria, L., Kiar, G., and Glatard, T. (2018). Numerical error propagation in the HCP structural pre-processing pipelines. *Organ. Hum. Brain Mapp.* 2052.
- Sherif, T., Rioux, P., Rousseau, M. E., Kassis, N., Beck, N., Adalat, R., et al. (2014). CBRAIN: a web-based, distributed computing platform for collaborative neuroimaging research. *Front. Neuroinform.* 8:54. doi: 10.3389/fninf.2014.00054
- Sochat, V., and Nichols, B. N. (2016). The neuroimaging data model (n.d.) API. *Gigascience* (Suppl. 1), 23–24. doi: 10.1186/s13742-016-0147-0-u
- Stockton, D. B., and Santamaria, F. (2015). NeuroManager: a workflow analysis based simulation management engine for computational neuroscience. *Front. Neuroinform.* 9:24. doi: 10.3389/fninf.2015.00024
- Sudlow, C., Gallacher, J., Allen, N., Beral, V., Burton, P., Danesh, J., et al. (2015). UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Med.* 12:e1001779. doi: 10.1371/journal.pmed.1001779
- Tournier, J. D., and Calamante, F. (2012). MRtrix: diffusion tractography in crossing fiber regions. *Int. J. Imaging Syst. Technol.* 22, 53–66. doi: 10.1016/j.nicl.2018.02.015
- Van Essen, D. C., Smith, S. M., Barch, D. M., Behrens, T. E., Yacoub, E., Ugurbil, K., et al. (2013). The WU-Minn human connectome project: an overview. *Neuroimage* 80, 62–79. doi: 10.1016/j.neuroimage.2013.05.041
- Vivian, J., Rao, A. A., Nothaft, F. A., Ketchum, C., Armstrong, J., Novak, A., et al. (2017). Toil enables reproducible, open source, big biomedical data analyses. *Nat. Biotechnol.* 35, 314–316. doi: 10.1038/nbt.3772
- Zuo, X.-N., Anderson, J. S., Bellec, P., Birn, R. M., Biswal, B. B., Blautzik, J., et al. (2014). An open science resource for establishing reliability and reproducibility in functional connectomics. *Sci. Data* 1:140049. doi: 10.1038/sdata.2014.49

Conflict of Interest Statement: 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.

Copyright © 2019 Kiar, Brown, Glatard and Evans. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.