



Graph Representation Forecasting of Patient's Medical Conditions: Toward a Digital Twin

Pietro Barbiero ^{*†}, Ramon Viñas Torné [†] and Pietro Lió [†]

Department of Computer Science and Technology, University of Cambridge, Cambridge, United Kingdom

Objective: Modern medicine needs to shift from a wait and react, curative discipline to a preventative, interdisciplinary science aiming at providing personalized, systemic, and precise treatment plans to patients. To this purpose, we propose a “digital twin” of patients modeling the human body as a whole and providing a panoramic view over individuals' conditions.

Methods: We propose a general framework that composes advanced artificial intelligence (AI) approaches and integrates mathematical modeling in order to provide a panoramic view over current and future pathophysiological conditions. Our modular architecture is based on a graph neural network (GNN) forecasting clinically relevant endpoints (such as blood pressure) and a generative adversarial network (GAN) providing a proof of concept of transcriptomic integrability.

Results: We tested our digital twin model on two simulated clinical case studies combining information at organ, tissue, and cellular level. We provided a panoramic overview over current and future patient's conditions by monitoring and forecasting clinically relevant endpoints representing the evolution of patient's vital parameters using the GNN model. We showed how to use the GAN to generate multi-tissue expression data for blood and lung to find associations between cytokines conditioned on the expression of genes in the renin–angiotensin pathway. Our approach was to detect inflammatory cytokines, which are known to have effects on blood pressure and have previously been associated with SARS-CoV-2 infection (e.g., CXCR6, XCL1, and others).

Significance: The graph representation of a computational patient has potential to solve important technological challenges in integrating multiscale computational modeling with AI. We believe that this work represents a step forward toward next-generation devices for precision and predictive medicine.

Keywords: digital twin, generative adversarial networks, monitoring, graph representation learning, precision medicine

OPEN ACCESS

Edited by:

Wei Lan,
Guangxi University, China

Reviewed by:

Cheng Liang,
Shandong Normal University, China
Ruiqing Zheng,
Central South University, China

*Correspondence:

Pietro Barbiero
pb737@cam.ac.uk

[†]These authors have contributed
equally to this work

Specialty section:

This article was submitted to
Computational Genomics,
a section of the journal
Frontiers in Genetics

Received: 13 January 2021

Accepted: 24 June 2021

Published: 16 September 2021

Citation:

Barbiero P, Viñas Torné R and Lió P
(2021) Graph Representation
Forecasting of Patient's Medical
Conditions: Toward a Digital Twin.
Front. Genet. 12:652907.
doi: 10.3389/fgene.2021.652907

1. INTRODUCTION

Modern medicine is shifting from a wait and react, curative discipline to a preventative, interdisciplinary science aiming at providing personalized, systemic, and precise treatment plans to patients. Systems and network medicine are rapidly emerging in medical research providing new paradigms to address.

In the next decades, precision and predictive medicine will have a pivotal role in revolutionizing the healthcare system making it more flexible and efficient. Precision and predictive medicine are challenging research fields as they need to deal with the complexity of the human body (Ginsburg and Willard, 2009; Naylor and Chen, 2010). Precision requires integrating large amount of observations at individual and population levels simultaneously. These measures need to be taken at different scales, from genome to clinical and family history and at systemic levels, i.e., considering multiple tissues and organs. In the last years, systems and network medicine have introduced a variety of novel approaches with the aim of integrating and gaining knowledge on the human body. We have no capacity to integrate such disparate information into equation-based models but we can use machine learning and, in particular, deep learning methods to achieve this integration goal.

The primary objective of this work is exploring challenges and opportunities in modeling the human body as a whole, providing a panoramic view over individuals' conditions. To this aim, we propose a proof of concept of a "digital twin," i.e., a virtual prototype of patients mirroring the underlying biological system (Gelernter, 1993; Laubenbacher et al., 2021) combining information at organ, tissue, and cellular level. Existing prominent examples of digital twins in healthcare include "the artificial pancreas" (Brown et al., 2019; Kovatchev, 2019), pediatric cardiac digital twins (Gutierrez et al., 2019; Shang et al., 2019), and diabetes models (Eddy and Schlessinger, 2003). However, all these examples focus on just one single aspect of the human body due to its extreme complexity. As a result, they are not suitable to provide a holistic overview over the whole human body. We believe that recent graph representation approaches could overcome digital twin's limitations scaling across all the variety of body signals at different levels, making possible a revolution in healthcare. This work provides a first proof of concept providing the first elements for a novel class of machine-learning-assisted tools that scale to medical device deployment and run time monitoring and verification. By fusing ideas from systems medicine with scientific computing and machine learning, our software integrates and automates the analysis of vital parameters models under large uncertainty. A high degree of automation could transform how we use models in the scientific and medical discovery cycle and open up for a next-generation of powerful medical devices for probing the inner workings of full body in well-being and disease conditions.

The proposed architecture combines the qualities of generative and (Goodfellow et al., 2014) graph-based models (Scarselli et al., 2008) (see **Figure 1**). On the one hand, the generative model can be used to produce synthetic data under different biological states that might not be observed in reality. By augmenting the set of explorable states of the underlying biological system, the generative model may be employed for the simulation of extremely rare clinical scenarios representing precarious conditions, which might be difficult to analyze otherwise (Yi et al., 2019). In clinical contexts, this means that physicians will be able to set up personalized experiments in a virtual environment representing their patients in a very detailed and realistic way. On the other hand, the graph model

represents the actual digital twin, providing a general and flexible framework to run probabilistic simulations. A panoramic view of individuals' conditions is provided by the final network configuration that combines information at organ, tissue, and cellular level. Cross-modal signals are also supported by the most recent graph learning frameworks, thus allowing the combination of different data sources, both structured and unstructured, real or simulated by generative methods. Finally, by relying upon flexible and modular architectures, our "digital twin" model can be conveniently deployed in dedicated hardware modules paving the way for a next-generation of medical devices.

2. DESIGN OF A BIOMEDICAL DIGITAL TWIN

The birth of the term "digital twin" could be the NASA's Apollo program where one spacecraft was launched into the outer space, while a "twin" spacecraft remained on earth to mirror flight conditions. Digital twin has been defined as "an integrated multiphysics, multiscale, probabilistic simulation of a vehicle or system that uses the best available physical models, sensor updates, fleet history, etc., to mirror the life of its flying twin" (Shafto et al., 2010; Grieves, 2015). The digital twin is a virtual prototype; the analysis of its digital life cycle provides information to understand a product's functionality, manufacturing, behavior, and usage prior to building it. Here, the meaning of digital twin is slightly different: there is no product to be built, instead experimenting therapies on a digital twin will be cost-effective and will provide us with a rigorous testbed to conduct medical interventions. Within this framework, the artificial intelligence model could enable the prediction of disease trajectories before the insurgence of symptoms. The personal medical digital twin could also represent a pragmatic way for the cyber-physical fusion, as a new approach to support biomedical engineering design. In our vision, a composable AI architecture could enable the development of automatic analysis and verification techniques that are key to translational medicine.

Our digital twin consists of a modular AI-aided system that can be used to model the human body as a whole and to forecast the evolution of pathophysiological conditions (see **Figure 2**). The first module is based on a graph neural network (GNN) forecasting clinically relevant endpoints (such as blood pressure), while the second one is represented by a generative adversarial network (GAN) providing a proof of concept of multi-omic integrability.

2.1. The Effectiveness of GNNs and GANs in Biomedical Signal Analysis

The lack of interpretability of deep learning models has been one of the most significant barriers preventing their application in healthcare. Such models exhibit great capacity (Hornik, 1991) but understanding their behavior and following their decision-making process is not trivial (Castelvecchi, 2016). There is a growing body of literature focusing on interpretable artificial intelligence and interpretable deep learning aiming at developing white box models or at explaining black box ones (Das and Rad,

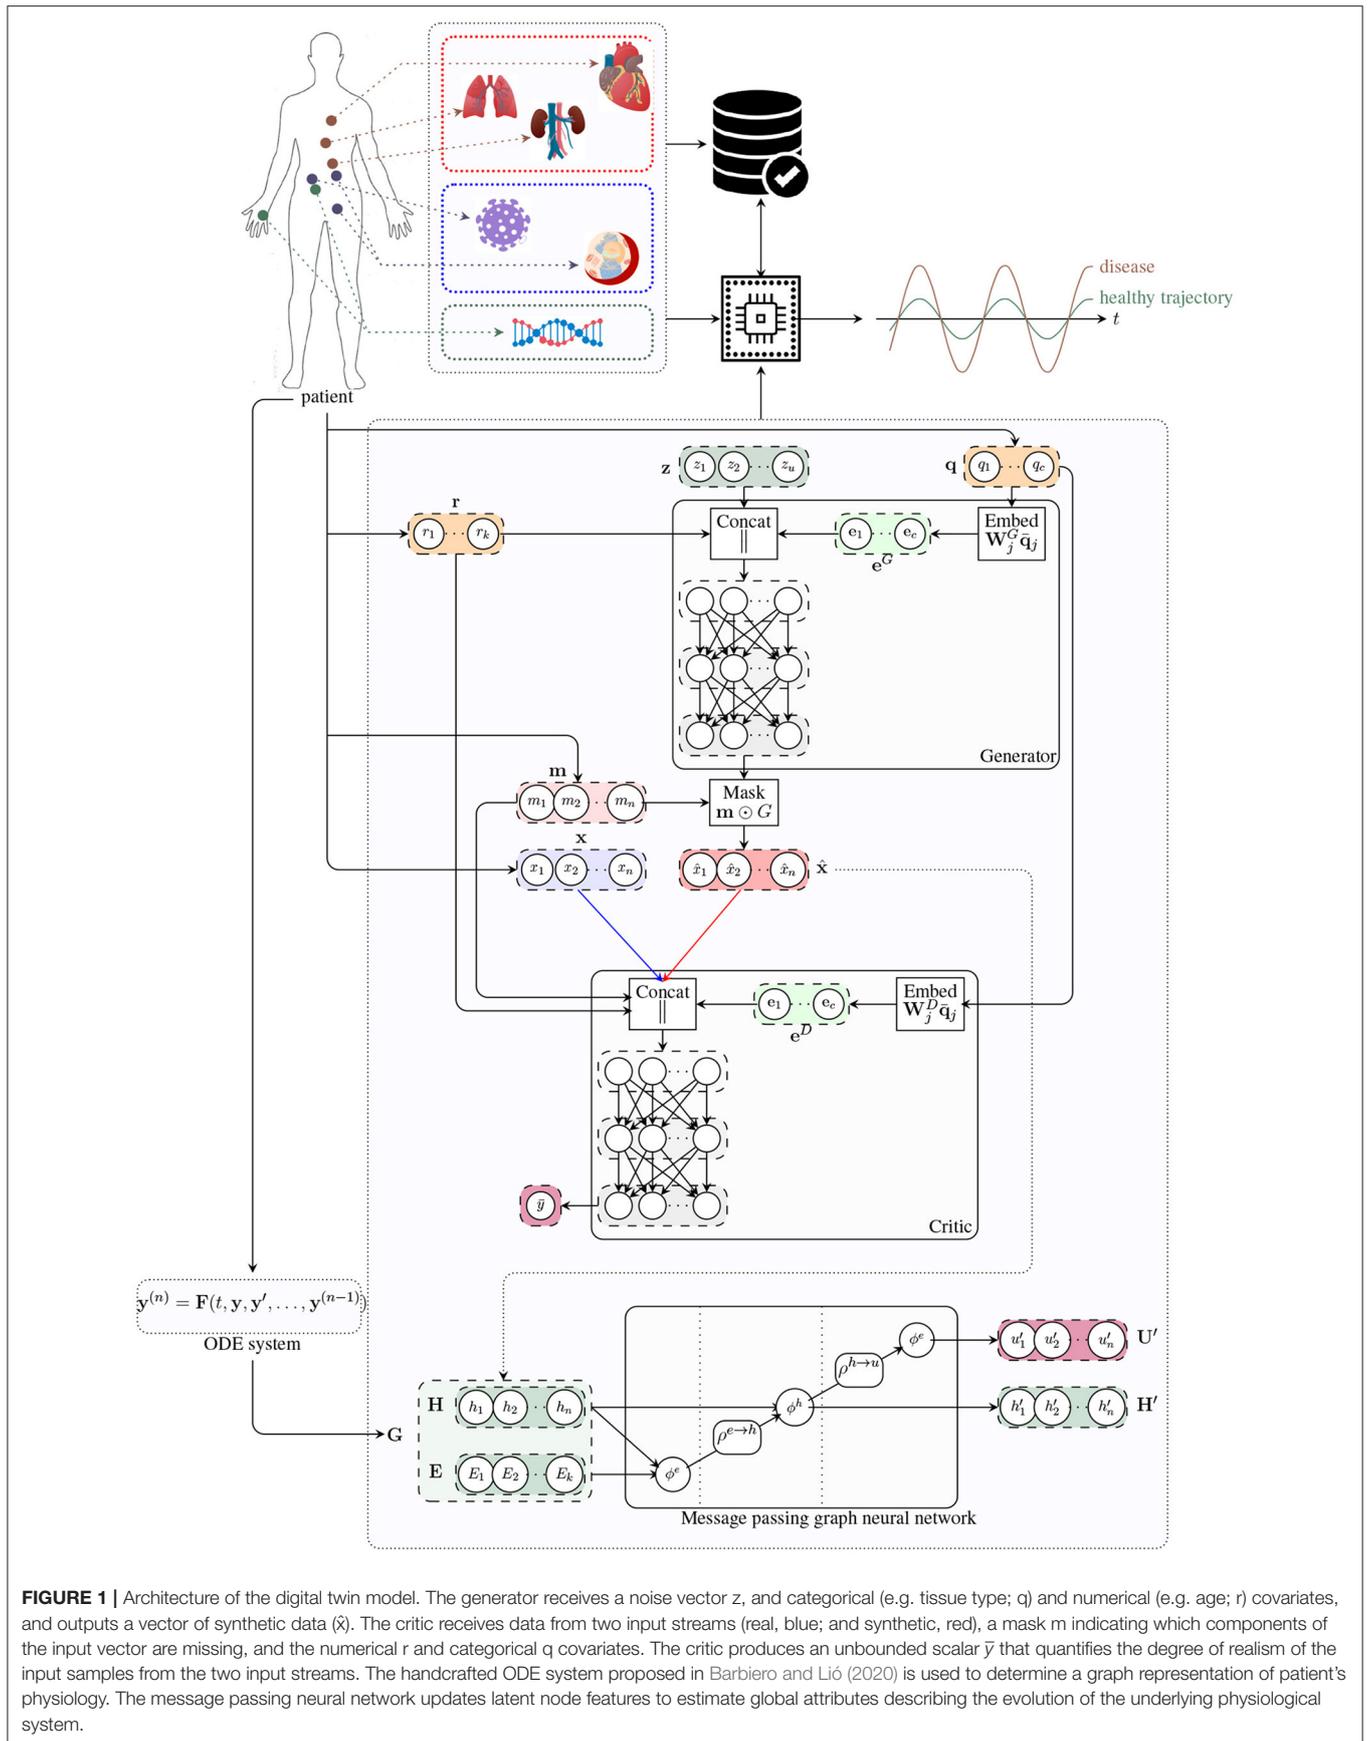
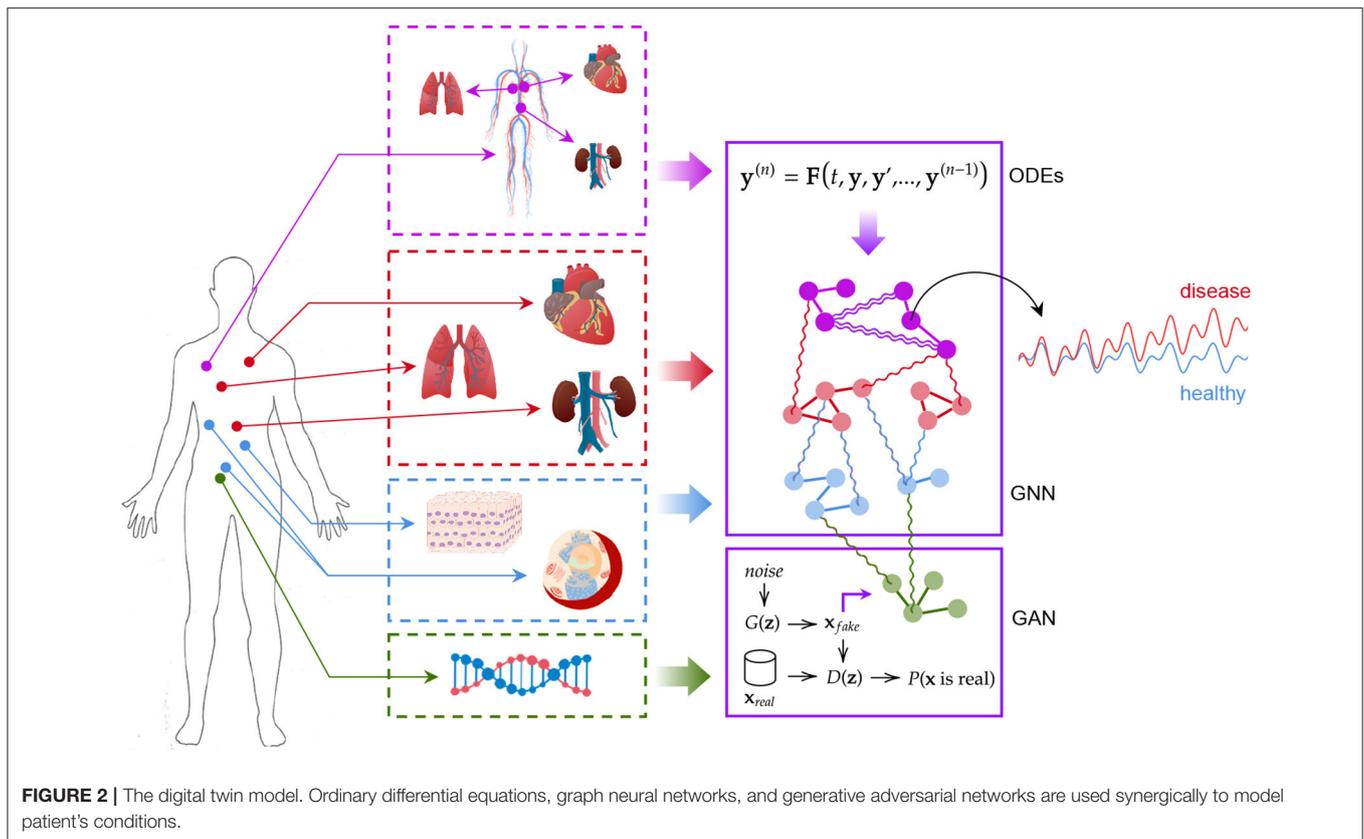


FIGURE 1 | Architecture of the digital twin model. The generator receives a noise vector z , and categorical (e.g. tissue type; q) and numerical (e.g. age; r) covariates, and outputs a vector of synthetic data (\hat{x}). The critic receives data from two input streams (real, blue; and synthetic, red), a mask m indicating which components of the input vector are missing, and the numerical r and categorical q covariates. The critic produces an unbounded scalar \bar{y} that quantifies the degree of realism of the input samples from the two input streams. The handcrafted ODE system proposed in Barbiero and Lió (2020) is used to determine a graph representation of patient's physiology. The message passing neural network updates latent node features to estimate global attributes describing the evolution of the underlying physiological system.



2020). Among such techniques, GNNs have started drawing the attention of both research and industry communities (Bronstein et al., 2017; Zhou et al., 2018). Such models are much more interpretable with respect to other neural approaches thanks to their graph structure, which is quite easy to understand from a human standpoint, and a few studies have already shown how graph networks can be effectively employed in biology and healthcare (Zitnik et al., 2018; Gysi et al., 2020).

Several properties of graph and generative adversarial neural networks make them suitable for medical data analysis. (1) *Non-linearity*: Both GNNs and GANs are able to detect non-linear patterns, which is of key interest as most systems are inherently non-linear in nature. Examples in medicine include heart rate dynamics, pulmonary functions, vascular structure, and gait dynamics. There is often a loss of non-linearity and multiscale fractal in aging and disease conditions (Goldberger et al., 2002). (2) *Interpretability*: Graph-based models are much easier to interpret with respect to other neural approaches thanks to their structure. The possibility of interpreting the behavior of models and the reason for their predictions is pivotal if not critical in many fields including clinical practice. (3) *Non-Euclidean geometry*: As a unique non-Euclidean data structure for machine learning, graphs can be used to model a variety of biological systems at different scales. Tissue and organ distributions could be modeled as graph models where each node or the graph contain time-dependent signals, similarly for pressure and electric sensors positioned at various parts of the body.

Lymphatic vessels can also be modeled as a network where lymph nodes are vertices. At lower scale, cell arrangements in tissues form particular manifolds; proteins and genes are organized in regulatory networks; other examples are cytoskeleton and organelles (mitochondria networks). Additionally, diseases could be seen as nodes in a graph where edges represent comorbidity or underlying polygenic causes. (4) *Modularity*: A key property of GNNs is modularity, which allows to learn independent mechanisms that can be reused in several parts of the graph. Modularity facilitates scalability and allows to model dynamic properties of graphs. (5) *Cross-modality*: Both GNNs and GANs can learn how to combine structured and unstructured data sources, spanning different levels of biological complexity. This is particularly relevant when integrating signals at different levels of biological scale such as DNA methylation and functional magnetic resonance imaging (fMRI) data. (6) *Generative*: Both GNNs and GANs can learn how to generate new data preserving the statistical properties of the training set. This could be used to compare statistics at individual level with those at specific groups identified with stratification analysis or at general population levels. (7) *Multiscale*: The graph representation has the capability of integrating granular information organized as networks at different layers of biological complexity. This allows to recognize patterns in higher-order structures such as motifs, pathways, tissues (as compositions of cells), organs (as composition of tissues), processes and apparatus (as composition of organs), and stratification (as composition of individuals). (8) *Spectral*

density: Together with spatial properties, GNN are amenable to frequency domain analysis. This allows to investigate network motifs, substructures, and periodical patterns at network levels.

2.2. Graph Neural Model

Graphs are mathematical structures that are used to model a set of objects (nodes) and their mutual relationships (edges) (Bollobás, 2013). Graphs are employed in a variety of research areas as they provide a general and flexible data structure for modeling real-world systems (Lieberman et al., 2005; Zhou et al., 2018; Rakocevic et al., 2019; Bica et al., 2020). GNNs are deep learning-based models working on the graph domain (Scarselli et al., 2008; Battaglia et al., 2018; Wu et al., 2020). Their properties have been recently drawn the attention of the artificial intelligence research community given their high interpretability (Lecue, 2019; Huang et al., 2020). The combination of graph theory and neural network elements have made GNNs one of the most promising tools to analyze complex systems in the graph domain. From neural networks, GNNs inherit a data-driven approach associated with a multi-layer architecture, which is the key to extract hierarchical patterns from data. However, unlike other deep-learning models, GNNs exploit additional features from graph theory and other mathematical disciplines. The main advantage with respect to other machine learning models relies in their extremely flexible and interpretable architecture. Once defined, the main endpoints of a system together with their mutual relationships directly induce a corresponding graph representation, which can be easily interpreted from a human standpoint. The abstract graph representation can be handcrafted, when the complexity of the underlying system allows it, or even automatically induced from data using generative approaches (Li et al., 2018). Hybrid techniques may also be explored taking advantage of generative algorithms for handling system complexity and human modeling to customize the most relevant endpoints. The design of GNNs is based on two basic principles, flexibility, and composability. GNNs support different graph structures as well as flexible representations of global, node, and edge attributes, customizable according to specific demands of tasks.

2.2.1. Stratification of Human Body Layers in a GNN

GNNs natively allow the design of complex systems using a modular approach. First, the complexity of the human body is broken up by developing independent subsystems representing genomic alterations, biological pathways, and organ physiology. Each subsystem can be represented as a different node or a network of nodes in a GNN, while inter-process signals can be reframed as message passing operations supporting multiscale ripple effects. Homogeneous subsystems can be aggregated into layers according to their characteristics. Our digital patient model is composed of four biological layers: the transcriptomic layer, the cellular layer, the organ layer, and the exposomic layer. Other layers can be easily implemented.

2.2.1.1. Transcriptomic Layer

The transcriptomic layer operates on the set of RNA transcripts produced by the genome at a particular time. Currently, RNA

sequencing (RNA-seq) can measure RNA abundance across the entire genome with high resolution. The resulting high-throughput gene expression data can be used to uncover disease mechanisms (Emilsson et al., 2008; Cookson et al., 2009; Gamazon et al., 2018), propose novel drug targets (Evans and Relling, 2004; Sirota et al., 2011), provide a basis for comparative genomics (Colbran et al., 2019), and address a wide range of fundamental biological problems.

In this work, we study the crosstalk between tissues in the organ layer (see **Figure 1**) through the communicome, e.g., communication factors in blood (Ray et al., 2007). Specifically, we analyze to what extent the expression of genes involved in the renin-angiotensin system (RAS) can be explained by genes from signaling and receptor pathways, including the chemokine, TNE, and TGF- β pathways. We further develop a transcriptomics generative model based on a generative adversarial network (Goodfellow et al., 2014) and simulate the effects of SARS-CoV-2 infection by conditioning on high expression of ACE2 in the lung, kidney, and pancreas.

2.2.1.2. Cellular Layer

The cellular layer involves biological processes affecting individual cells from metabolism and protein synthesis to replication and motility. In this study, we focus on modeling the RAS, one of the main biological pathways regulating blood pressure and closely related to SARS-CoV-2 infectivity. Hence, it represents a suitable case study to demonstrate the flexibility and expressiveness of our GNN-based approach. The RAS is a hormone system regulating vasoconstriction and inflammatory response (Fountain and Lappin, 2019). The key hormone of the system is the peptide angiotensin II (ANG-II) generated from the decapeptide angiotensin I by the angiotensin-converting enzyme (ACE). ANG II promotes vasoconstriction, hypertension, inflammation, and fibrosis by activating the ANG-II type 1 receptor (AT1R) (Kuba et al., 2010; Gironacci et al., 2011). Glucose concentration, ACE inhibitor treatments, and viral infections binding to ACE2, such as SARS-CoV-2, can all have a significant impact on the RAS. A high glucose concentration may determine chronic hypertensive conditions. Reducing ANG II production with ACE inhibitors increases vasodilation and vasoprotection effects stimulated by the overproduction of AT2R and ANG-(1-7) (Zaman et al., 2002). Viral infections such as SARS-CoV-2 may also have an impact on RAS, as the virus binds to ACE2 in order to gain entry into the host cell. This results in an altered ACE2 activity and concentration, possibly leading to hypertension and inflammatory response (South et al., 2020).

2.2.1.3. Organ Layer

The organ layer comprises group of tissues with similar functions (organs) and complex networks of cooperating organs. Given the nature of the multi-factorial disease under study, we limited the organ layer to the circulatory system and a physiological representation of a few organs (Barbiero and Lió, 2020): heart, lungs, and kidneys. The heart model includes four compartments known as chambers (Neal and Bassingthwaite, 2007). Deoxygenated blood collected from the superior and

inferior venae cavae flows into the right atrium. When the right atrium contracts, the blood is pumped through the tricuspid valve into the right ventricle. From the right ventricle, the blood is pumped into the pulmonary trunk through the pulmonary valve flowing toward the lungs where carbon dioxide is exchanged for oxygen. The pulmonary circulation is composed of five vascular segments: proximal and distal pulmonary artery, small arteries, capillaries, and veins. Oxygenated blood collects into the left atrium via the pulmonary veins. From there, it flows into the left ventricle through the mitral valve and it is pumped into the aorta through the aortic valve for systemic circulation, providing oxygen and nutrients to body cells for metabolism in exchange for carbon dioxide and waste products. The mean arterial blood pressure is controlled by baroreceptors, special sensory neurons excited by a stretch in the carotid sinus and aortic arch vessels. They relay sensory information regarding blood pressure changes to the central nervous system where it is processed and utilized primarily in autonomic reflexes, regulating short-term blood pressure.

2.2.1.4. Exposomic Layer

The exposome refers to the totality of exposure individuals experience from conception until death and its impact on chronic and acute diseases (Wild, 2005). Toxicants, dietary regimens, treatments, physical exercise, posture, and lifestyle habits are possible exposures taking part to individual's well-being or disease condition. All such environmental factors are deeply coupled among themselves but also with individuals influencing the effects of new or present exposures. The exposome is intrinsically co-dependent on a person's genetics, epigenetics, health status, and physiology. For instance, regular exposure to pollution may lead to the outbreak of a lung carcinoma, which in turn may call for clinical intervention. In this work, we consider four types of exposures: dietary habits, physical activity, therapeutic treatments, and viral infections.

2.2.2. Inter-process Signals and Clinical Endpoints

One of the main advantages of using GNN-based models relies in that inter-process and multiscale communications can be natively implemented using message passing. In a GNN, each biological entity can be represented as a node, while the relationship between two entities can be modeled using directional edges. Signals exchanged between nodes are implemented using message functions ϕ^h (see Equation 1), which are used to update the hidden states of nodes. Such state transition will then have an impact on messages exchanged at the following time steps. Another strength of GNN models consists of the possibility of supervising the evolution of the underlying system by using the readout functions ϕ^u . Hence, the endpoints of multi-factorial diseases can be directly controlled by checking the output of readout functions in critical nodes. The resulting GNN model will combine a simple and modular design with a versatile structure accommodating for complex multiscale systems where clinical endpoints can be easily monitored and forecast in real time.

2.3. Generative Adversarial Model

One way of studying probability distributions is by means of generative models, which describe the random phenomenon in terms of the joint probability distribution of observed and target variables (Jebara, 2012). Generative adversarial networks (GANs) are a framework for estimating generative models via an adversarial process (Goodfellow et al., 2014). They are often described as a two-player game in which both players are encouraged to improve. One player, the *generator*, creates samples that are intended to be indistinguishable from the ones coming from a given data distribution. The other player, the *discriminator*, learns to determine whether samples come from the *fake* distribution (*fake* samples) or the *real* data distribution (*real* samples). **Figure 3** shows the basic idea of generative adversarial networks. With respect to other generative models, they provide a general and flexible framework for the analysis of joint probability distributions. The architecture itself allows a fine control of the data generation process and a high level of customization, making them suitable for a variety of experimental scenarios.

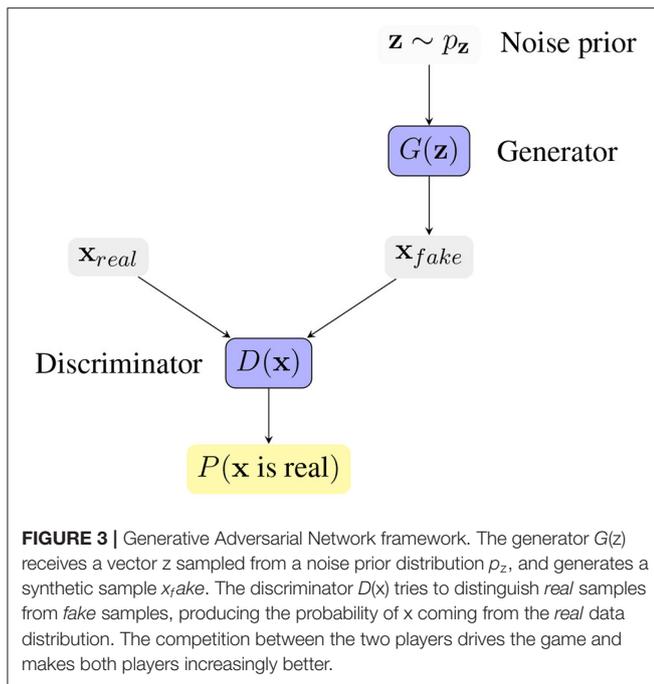
2.3.1. Crosstalk Between Tissue Types

The activity of biological systems is determined by internal factors, determined by intrinsic and functional properties, and by external factors shaping the interconnections between different systems. Chemical and molecular events, like oxygenation or protein phosphorylation, are often the vehicles of biological signals' transduction. A chain of biochemical events forms a signaling pathway whose activation may give rise to a biochemical cascade of events affecting the organism at different levels. In complex organisms, several signal transduction pathways communicate and react reciprocally generating biological crosstalks. Crosstalks have been widely characterized and observed in a variety of biological processes from micro- to macroscale from genomics (Poyton and McEwen, 1996; Du et al., 2015), to internal and external cell activity (Geiger et al., 2001; Li et al., 2016), and even between tissues (Lengyel et al., 2018). In particular, receptors and signaling factors from the chemokine, TNF, and TGF- β pathways are known to take an active role in tissue communication as well as inflammatory-associated diseases (e.g., cardiovascular diseases affecting that the heart and the stiffness of blood vessels). Here, we develop a generative model based on a generative adversarial network to produce synthetic transcriptomics data describing the ripple effects of a viral infection on crosstalks between different tissues. The aim is to demonstrate how generative approaches can be used both to reproduce and enhance the set of observable states of a patient allowing for a deeper understanding of complex biological processes.

3. RESULTS

3.1. Clinical Case Studies

In Barbiero and Lió (2020), the authors proposed a computational tool for running simulations integrating a variety of mechanistic and phenomenological models describing



the human body with ordinary differential equations (ODEs). This computational framework is hereby used to generate two clinical case studies. The main difference of the proposed approach with respect to the computational tool proposed in Barbiero and Lió (2020) consists of a different modeling approach based on state-of-the-art AI models instead of ODEs.

The first scenario consists of an elderly patient experiencing hypertension and type 2 diabetes with diabetic nephropathy. Her lifestyle is mainly sedentary and her diet is rich in carbohydrates. The patient needs a therapeutic plan for the treatment of her hypertension. The task for the clinician is to personalize the therapy assigning a proper daily dosage of benazepril. This case study is used to show how the digital patient model can be employed to simulate the evolution over time of clinical endpoints under a set of possible therapeutic plans and to choose the best option.

In the second scenario, the same patient is seeking medical help for a mild flu caused by a SARS-CoV infection. For this case study, the model can be used to constantly monitor and forecast clinical endpoints to prevent complications threatening patient's life. The decreased oxygenation caused by flu may have detrimental effects on both heart and brain activities indeed. Studies have reported that SARS-CoV infections can activate the blood clotting pathway by impairing left heart pumping performance, which results in a blood back up in the lungs and in an increased blood pressure. High blood pressure can reduce blood vessel's compliance decreasing blood and oxygen flows and leading to a higher risk of developing systemic conditions. For this reason, heparin-based therapies have been recommended to prevent clot formation or tissue plasminogen activator (tPA) (Sardu et al., 2020; Tang et al., 2020). Although some variation in blood pressure throughout the day is normal,

a high blood pressure variability is associated with a higher risk of cardiovascular disease (O'Rourke and Nichols, 2005; Mitchell et al., 2010; Wen et al., 2015; Clark et al., 2019; Bangalore et al., 2020) and all-cause mortality (Tao et al., 2017; Kim et al., 2018). Clogged arteries, fibrosis, and strokes caused by blood pressure spikes are among the main complications threatening patient's life and calling for the foremost necessity for treatment. Hence, blood pressure is one of the most relevant clinical endpoints that need to be constantly monitored in real time and accurately forecast.

3.2. Forecasting Clinical Scenarios

3.2.1. Dataset

Our digital twin model is hereby used to actively monitor and forecast the endpoints highlighted in the two clinical case studies. First, the computational system described in Barbiero and Lió (2020) based on ordinary differential equations (ODEs) is used to generate a time series of clinical endpoints for each differential equation with a window size of $\tau = 500$ time steps (Barbiero and Lió, 2020). Time series are collected, randomly shuffled, and stacked in a dataset. Each item of the collection is randomly assigned either to a training ($n_{train} = 3,200$), validation ($n_{val} = 800$), or test set ($n_{test} = 1,000$).

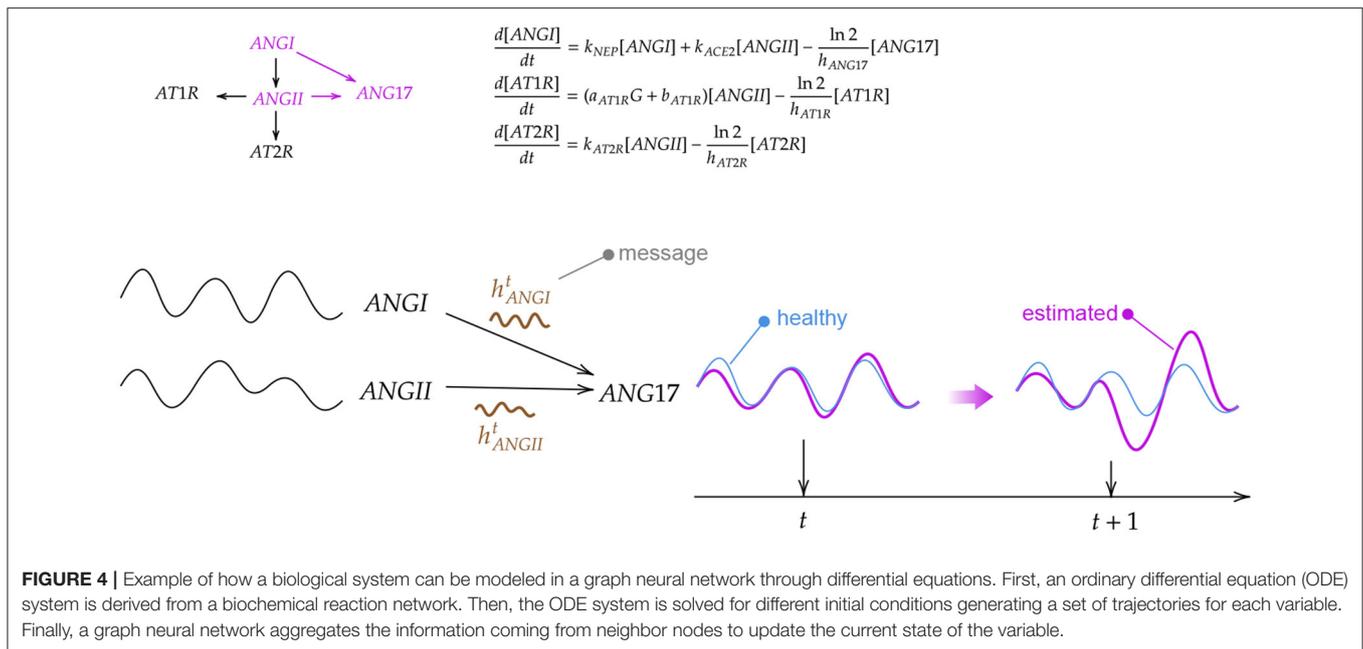
3.2.2. Training

The graph model is derived from the structure of the ODE system, thus leveraging human knowledge (an example is shown in Figure 4). Nodes correspond to variables represented by the differential equations in Barbiero and Lió (2020) while edges follow the underlying relationships. In a GNN-based model, each node learns a latent representation of the state using the messages received from its neighborhood. Hence, the rigid mathematical structure of the ODE system is relaxed in our model as such structure can be learned directly from data. The learning process lasts for $\eta = 50$ epochs with a learning rate of $\epsilon = 0.01$. Once trained and validated, the model is used to generate a bundle of possible trajectories for the elements of the test set. As a result, the model estimates a 95% confidence interval of the evolution of each variable over time.

3.2.3. Results

Providing a complete overview of the clinical state of a patient is not trivial. Focusing just on one endpoint might be misleading. On the contrary, a global vision comprising pathophysiological conditions is required in order to provide a clear and effective overview where organs and physiological systems can be monitored as a whole. One of the most effective approaches consists of applying a dimensionality reduction technique (Van Der Maaten et al., 2009) condensing the information of each organ and projecting forecasts in a lower-dimensional space.

Figure 5 shows an overview of the clinical state of the heart in a two-dimensional projected phase space. For each clinical case study, a GNN-based model is used to simulate a therapeutic intervention and its impact on blood pressure in heart chambers (right and left atrium and ventricle). In order to provide an overview of heart conditions, we



projected the predicted trajectories using principle component analysis (PCA) (Pearson, 1901). The interpretation of both pictures is straightforward. The first one shows the effect of a therapeutic intervention comprising an increased physical exercise, a reduced amount of calorie intake, and the subscription of a daily dosage of benazepril (5 mg). The predicted result of the prescription (green density reporting the 95% CI of the trajectories) reveals an overall reduction of blood pressure mean and variability in heart chambers. This results in a reduced risk of developing severe cardiovascular conditions with detrimental ripple effects for the whole system.

The second figure reports the simulation corresponding to the second case study. The same patient is seeking medical help to treat the first symptoms of a SARS-CoV-2 infection. The first simulation (red density) shows the long-term impact on heart blood pressure of an untreated viral infection. In this case, blood pressure spikes may cause irreparable damages to blood vessel walls, reducing their compliance, and impairing their capacity for adaptation to different environmental conditions. A synergic therapy including both benazepril (5 mg/day) and intravenous injection of heparin (5,000 U/ml) may have a beneficial effect on blood pressure mean and variability (orange density). On the one hand, benazepril lowers blood pressure by inhibiting ACE activity in cleaving ANG-I and producing ANG-II, which is the key RAS regulator of blood pressure. On the other hand, heparin is used to prevent and dissolve blood clots (Sardu et al., 2020; Tang et al., 2020). The treatment has an indirect impact on blood pressure by making blood less dense, reducing clotting formation, and lowering inflammation.

A lower-dimensional representation of an organ or system as a whole could be interesting to get a rapid and clear overview of the long-term impact of a disease or a therapeutic intervention. Nonetheless, bundle of predicted trajectories can be visualized

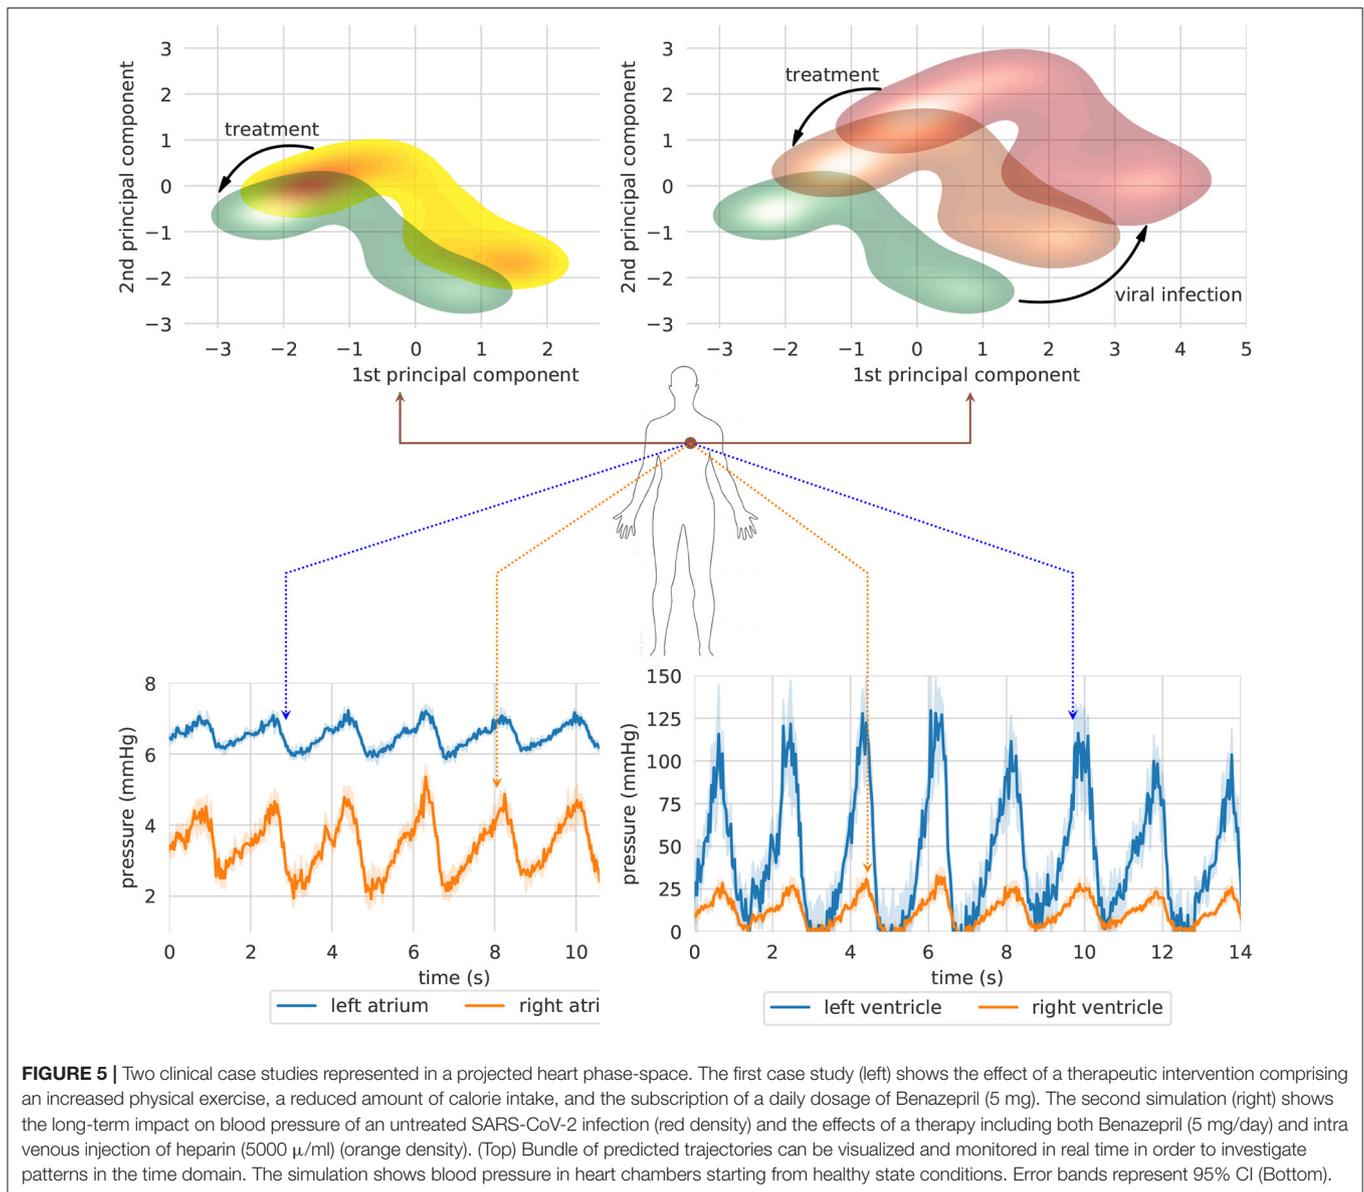
and monitored individually in real time when needed in order to investigate patterns in the time domain. **Figure 5** shows an example where blood pressure trajectories in heart chambers are predicted in real time starting from a healthy state condition (green density). In some cases, this representation in the time domain might be closer to common clinical approaches, thus providing a more conventional visualization tool for monitoring clinical endpoints in real time.

3.3. Transcriptomics Analysis of the Crosstalk Between Tissue Types

We hypothesize that the communication factors in blood might be playing an important role in the development of the SARS-CoV-2 infection by facilitating the spread of the virus in the human body. Here, we study whether the expression of genes involved in the RAS can be explained by genes that take part of the communicate in blood. This analysis might shed light on whether it is sensible to model the crosstalk between tissue types with a GNN where tissue nodes communicate with each other through whole blood.

3.3.1. Dataset

We leverage data from the Genotype-Tissue Expression (GTEx) project (v8), a resource that has generated a comprehensive collection of human transcriptome data in a diverse set of tissues (Aguet et al., 2019). The dataset contains 15,201 RNA-Seq samples collected from 49 tissues of 838 unique donors. We select genes based on expression thresholds of ≥ 0.1 TPM in $\geq 20\%$ of samples and ≥ 6 reads in $\geq 20\%$ of samples. We normalize the read counts between samples using the trimmed mean of M-values (TMM) normalization method (Robinson and Oshlack, 2010) and we inverse normal transform the expression values for each gene. From all the donors, we select those that have gene expression measurements for whole blood, yielding



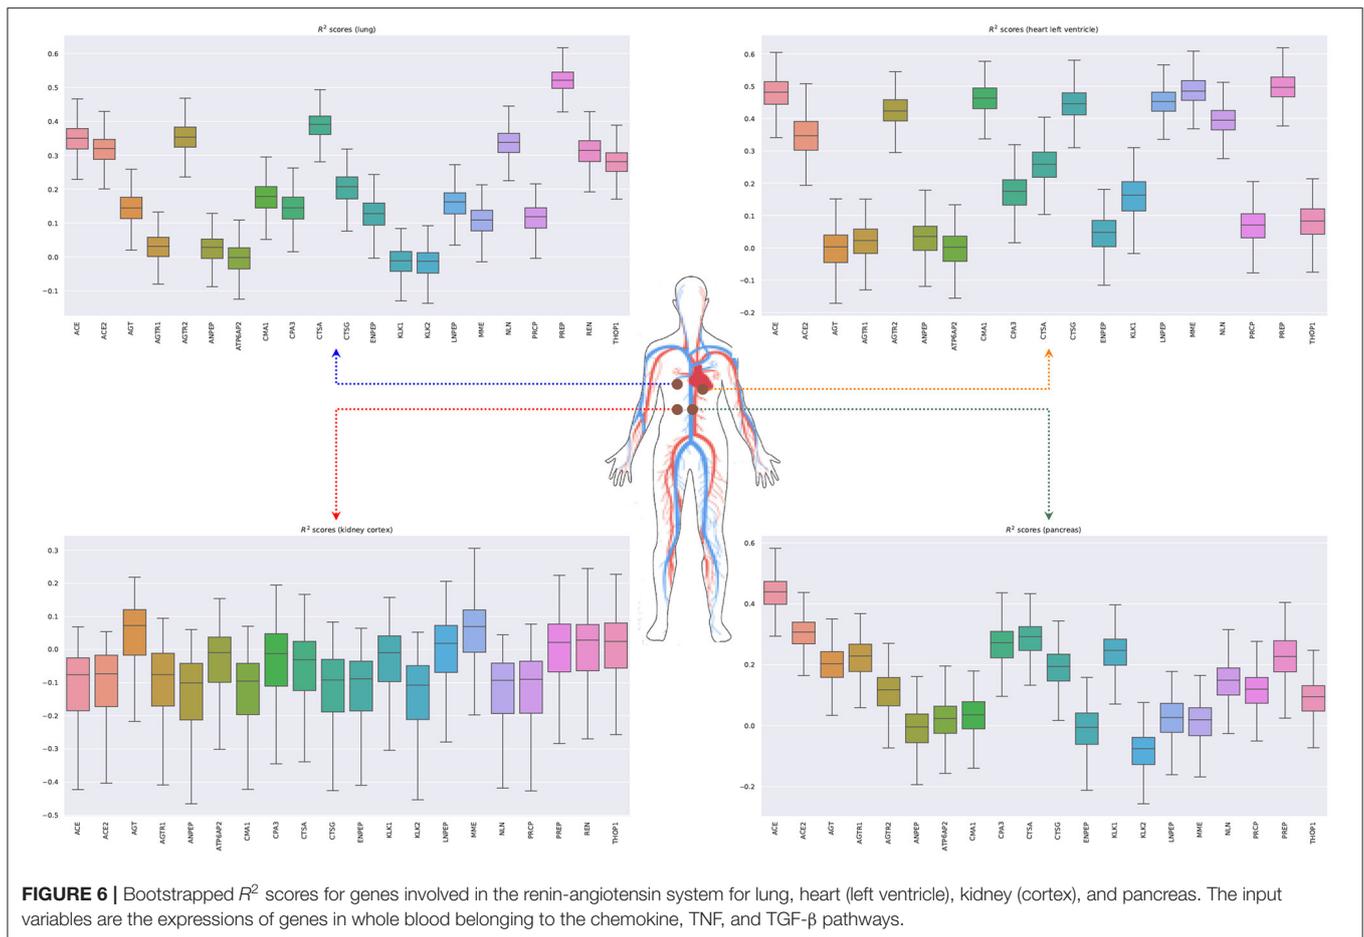
670 unique individuals. We then match the patients' whole blood samples with the corresponding measurements in lung (418), cortex of kidney (62), pancreas (257), and left ventricle of heart (324). Finally, we use the KEGG pathway database (Kanehisa et al., 2010) to select genes from the RAS (*hsa04614*), chemokine (*hsa04062*), TNF (*hsa04668*), and TGF- β (*hsa04350*) pathways.

3.3.2. Results

Figure 6 shows the bootstrapped R^2 scores for each gene in the RAS pathway in different tissue types. To compute the bootstrapped scores, we sampled donors with replacement (sample size: 75% of the total observations), trained the ridge regression model (Equation 5.3) on the sampled data, and evaluated the performance on the remaining out of bag (OOB) observations. Appendix 3 in **Supplementary Material** shows the held-out performances for different regularization strengths. We

repeated this process 1,000 times to obtain a distribution of R^2 scores for each gene. Our results show that the expression of some genes in the ACE2 pathway can be partially explained by signaling genes from whole blood. Notably, the associations for the kidney (cortex) are weaker or non-existent, potentially because the data are limited for this tissue (62 samples) or because the biological associations are indeed small. Overall, these results suggest that signaling pathways such as TNF, TGF- β , and chemokine might be playing an important role in the development of the SARS-CoV-2 infection.

We next model the expression of cytokines and receptors from whole blood (TNF, TGF- β , and chemokine pathways) as a function of cytokines from other tissue types (lung, kidney, heart, and pancreas) (Lijnen et al., 2003; Elmarakby et al., 2007; Rudemiller and Crowley, 2017). **Figure 7** shows the bootstrapped R^2 scores for the top 20 cytokines (chemokine pathway) for each



target tissue type. These results illustrate the associations between cytokines in blood and other tissue types, which facilitate tissue communication and crosstalks.

3.4. Generative Model for Transcriptomics Data

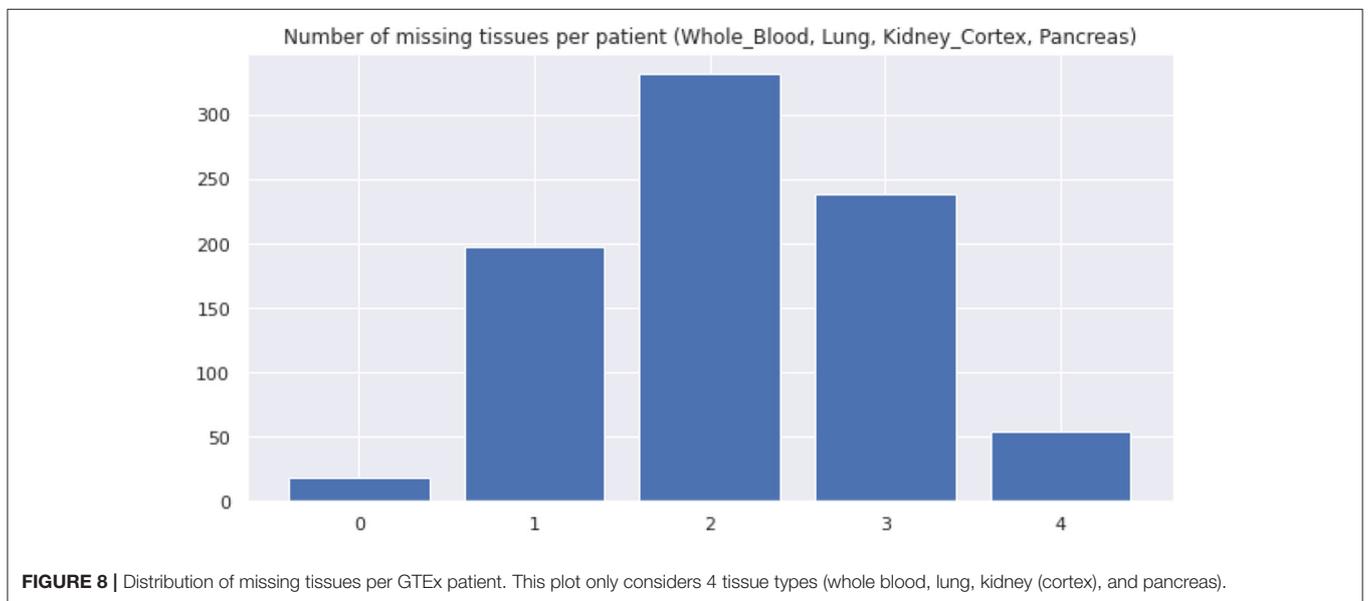
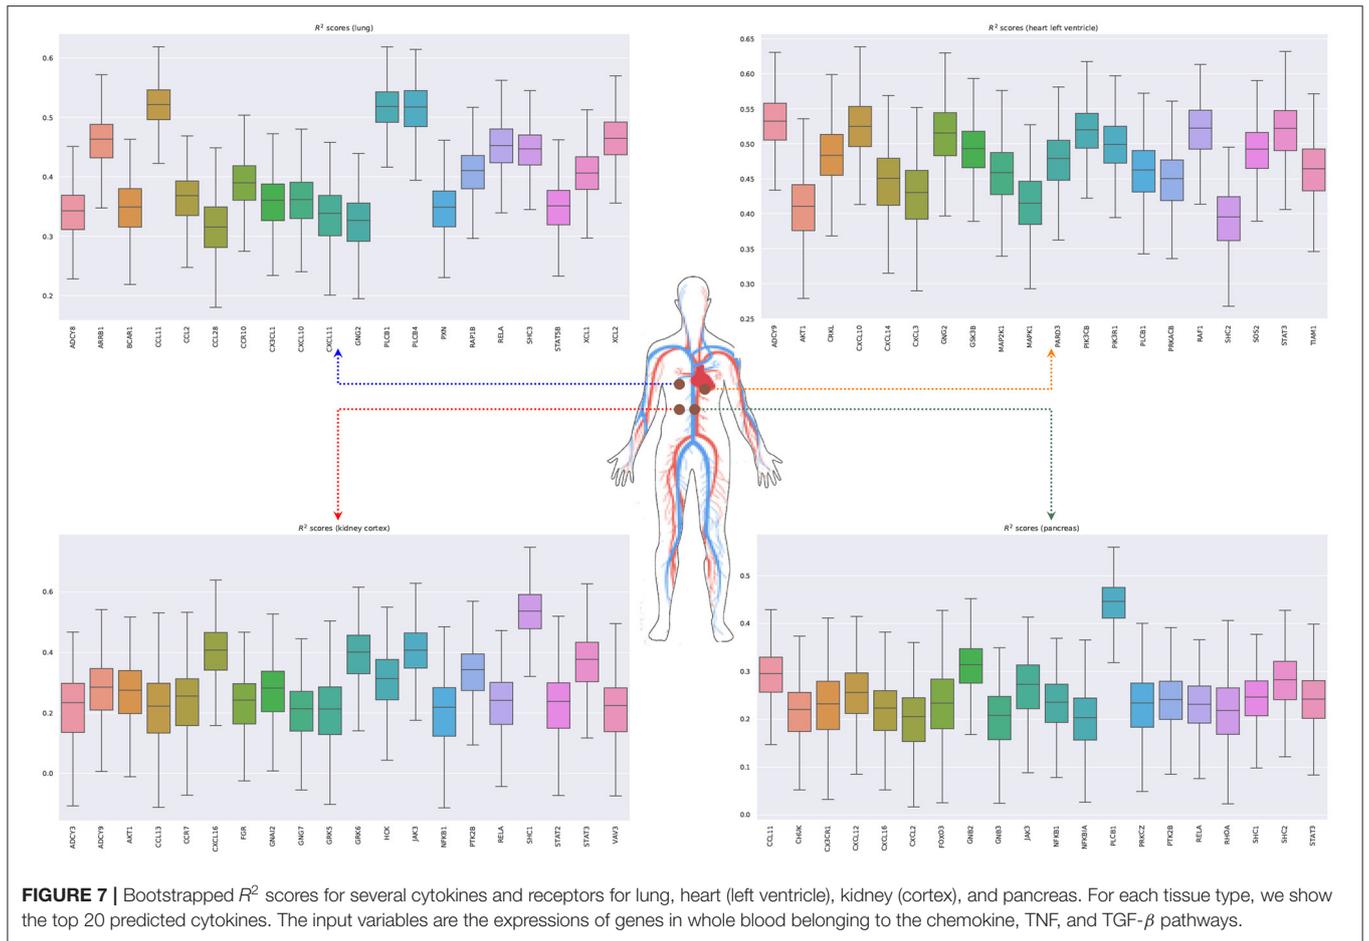
The generative model is here used to produce synthetic transcriptomics data. By conditioning on high expression of ACE2 in the lung, kidney, and pancreas, we aim to simulate the effects of SARS-CoV-2 infection in the expression of genes involved in communication and signaling pathways such as TNE, TGF- β , and chemokines. These pathways are implicated in many physiological and pathological processes including the regulation of blood pressure and inflammatory processes, and have been hypothesized to play a central role in SARS-CoV-2 infection (Garvin et al., 2020). For this analysis, we use data from the GTEx project previously described. In Appendix 2 in **Supplementary Material**, we analyze the held-out performance for different architectures of the generator and critic and describe all the training details.

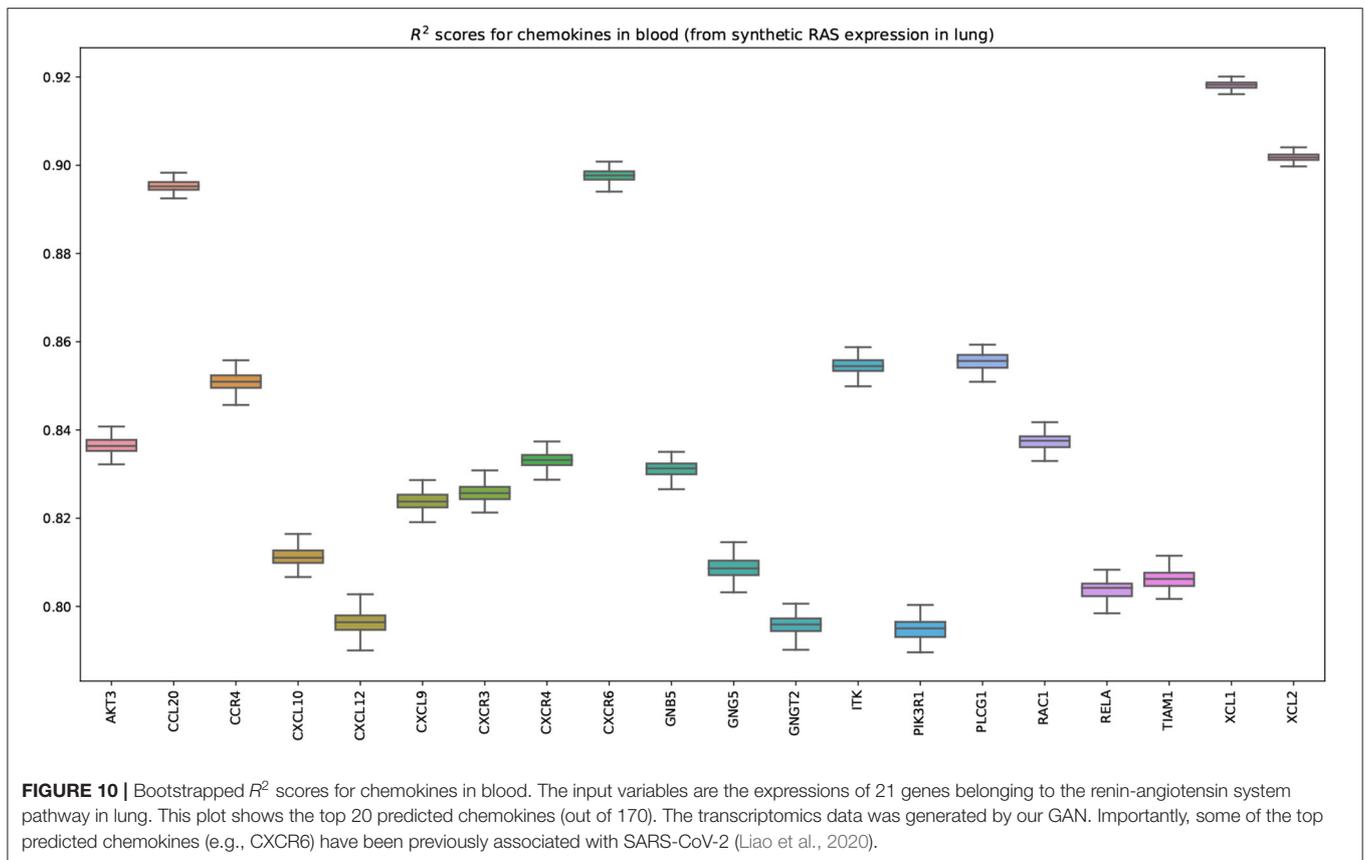
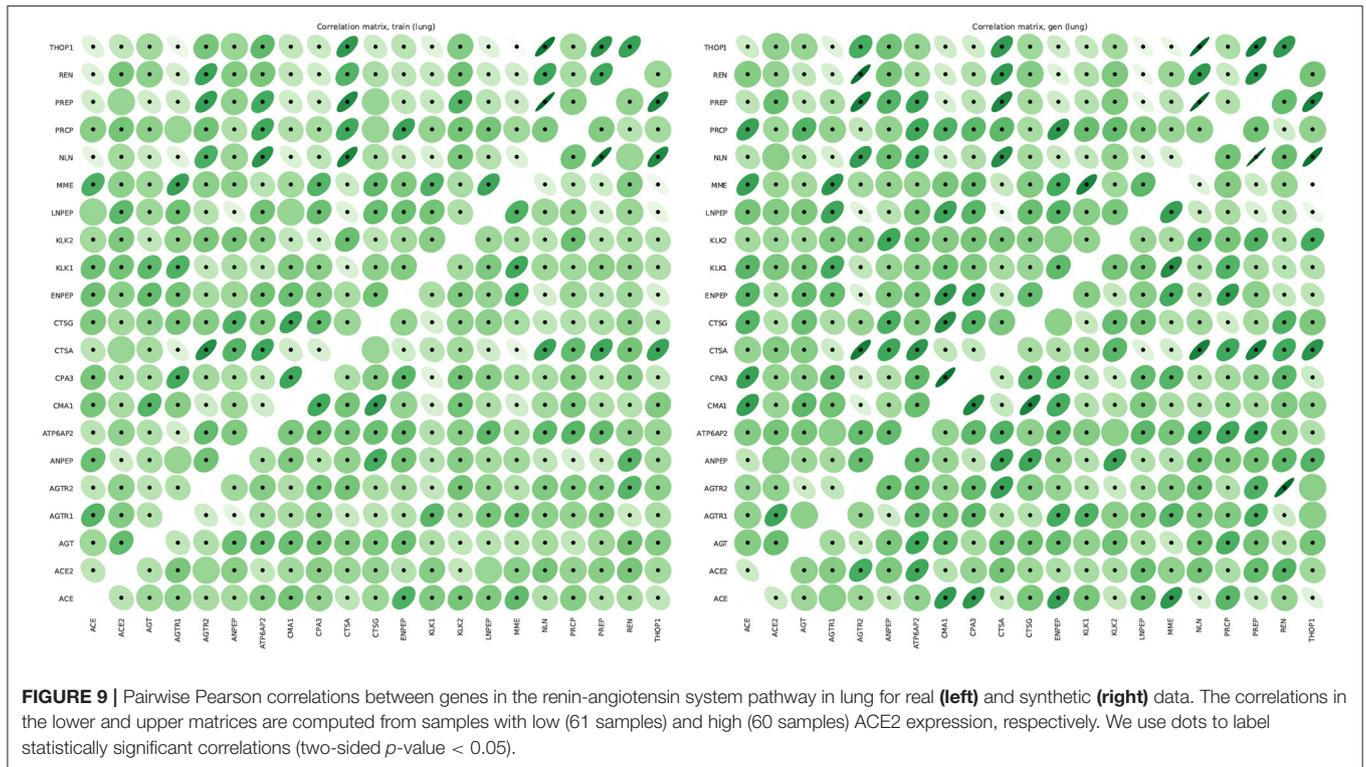
Real datasets often lack transcriptomic measurements that account for multiple tissue types jointly. For example, out of 838 GTEx donors, only 257 of them present joint observations

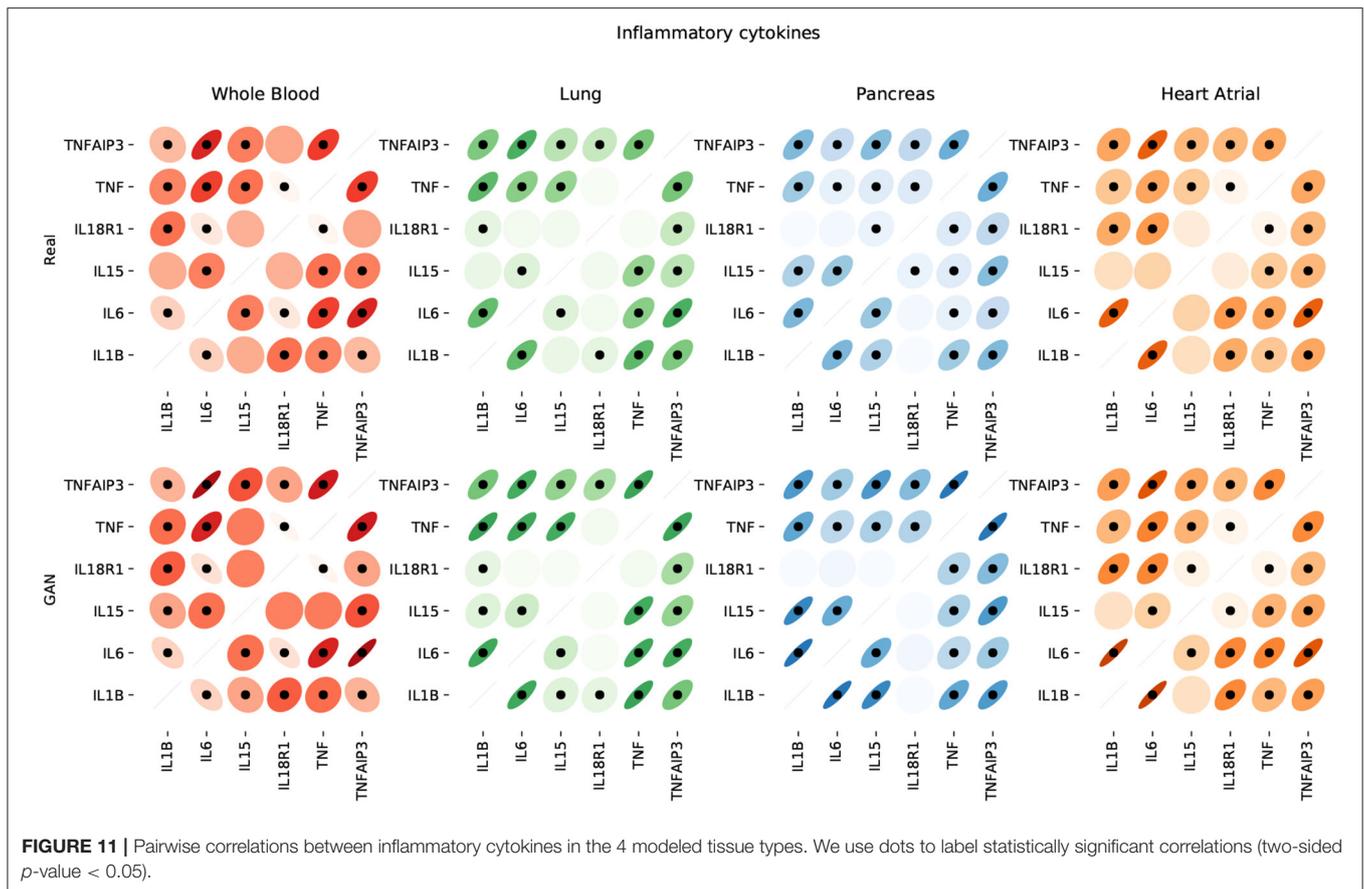
for pancreas and whole blood (**Figure 8** shows the distribution of missing tissues per patient). Importantly, our model allows to sample gene expression data for synthetic patients in every modeled tissue type and without any missing values, facilitating the cross-tissue analysis of gene expression.

3.4.1. Results

Figure 9 shows that the pairwise correlations between genes in the ACE2 pathway (lung) are well-preserved in the synthetic data. We observe that some genes in the RAS pathway (CTSA, AGTR2, NLN, and PREP) that can be relatively well-explained as a function of blood signaling factors (see **Figure 6**) are simultaneously correlated with ACE2. This suggests that these genes could be playing an important role in the spread of SARS-CoV-2 in our body through blood. Next, we use the GAN to generate multi-tissue expression data for blood and lung, and fit a linear model to predict the expression of 170 chemokines in blood as a function of the expression of 21 genes in the renin-angiotensin pathway from lung. **Figure 10** shows the R^2 scores for the top 20 chemokines. We find that some of the top predicted chemokines (e.g., CXCR6 and XCL1) have previously been associated with SARS-CoV-2 infection (Kusnadi et al., 2020; Liao et al., 2020). Additionally, our GAN captures associations between inflammatory cytokines, which are







known to have effects on blood pressure (Groth et al., 2014). **Figure 11** illustrates the real and synthetic pairwise correlations for 6 inflammatory cytokines in the 4 modeled tissue types. Finally, **Figure 12** shows that it is also possible to sample data for synthetic patients conditioned on different levels of ACE2 expression in lung.

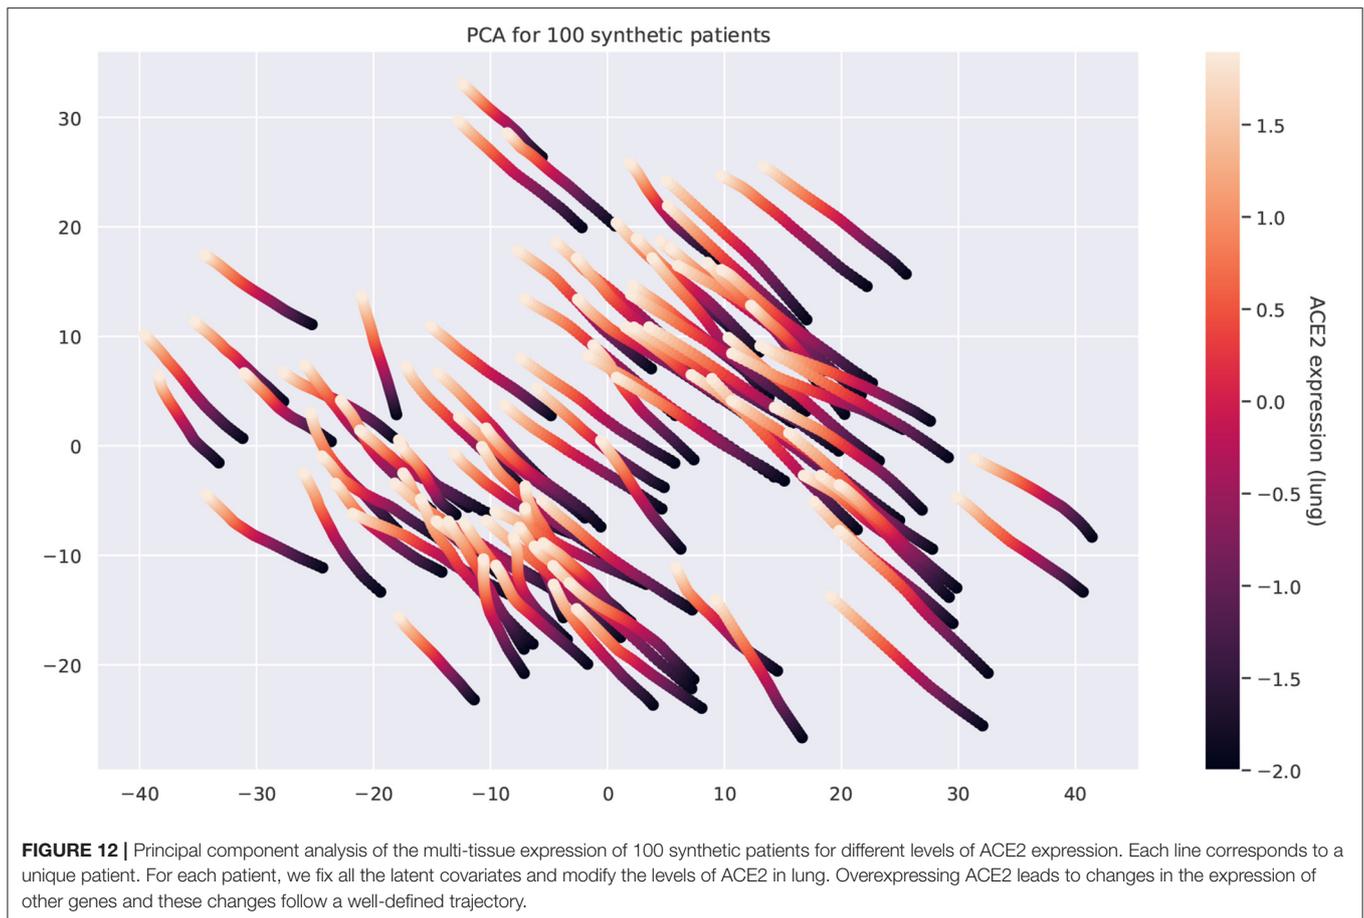
4. DISCUSSION

In this work, we presented an interpretable digital twin model providing an holistic view over patients' conditions. We tested our proof of concept on two clinical case studies combining information at organ, tissue, and cellular level showing the potential of our framework in clinical practice. We demonstrate the feasibility of representing and integrating physiological models and molecular information using GNNs and generative adversarial networks. This composite approach provides modularity and scalability across layers of biomedical data, it is amenable of a battery of modeling approaches, and generates integrated predictions that translate into patients trajectories. We have assimilated our product to a digital twin of the patient.

4.1. Technological Perspectives

4.1.1. Digital Twin Deployment

Mechanistic computational modeling and machine learning should be considered together when building innovative healthcare solutions. Building a puzzle is often an example of participatory activity. Clinicians, mechanistic computational modeling and machine learning researchers, data policy makers, and public and private sectors could build a puzzle (i.e., the healthcare) together and they should first develop a shared vision about what is the puzzle. Our vision is to consider a co-simulation (say doctor checkup visits vs. computational experiments) of the two twins to allow co-verification. From a theoretical computer science perspective, this could open the direction of an interplay between AI and verification/synthesis and the use of reachability analysis to identify constraints over the well-being and disease system state space. Although different architectures seem suitable (e.g., only GNNs, only GANs, VAEs, etc.), our design has important advantages: the GNN could provide a physical mapping of the human body (in the same way a tube map or bus route is a map of a city); GANs could be specialized on processing molecular information or they could operate cross-modal operations such as omic-omic, omic-clinical, and clinical-clinical.



4.1.2. A Modular Approach

The models presented in this work (GAN and GNN) are independent of each other. On the one hand, the main goal of the GNN model is to forecast various patient's conditions based on real or synthetic data, integrating information that spans multiple layers of the human body. On the other hand, the GAN model is able to generate data under different states, effectively enriching the space of pathophysiological conditions and endowing the digital twin with the ability to simulate the effects of counterfactual events. The independence of these two models enables a modular framework wherein each module can be trained separately on a distinct data modality. Importantly, these modules can be composed and reused through transfer learning. In this work, we have shown how computational models can be used to generate synthetic training data representing physiological conditions. Following the same principles, each module of a complex architecture could be pre-trained on synthetic simulations, refined using data obtained from horizontal population studies, and finally personalized according to clinical health records.

4.1.3. Next-Generation Datasets

The GAN and the GNN models can be interconnected in a synergistic way. In order to train the GNN effectively, it is

necessary to have access to heterogeneous, paired data modalities (from different layers: genomic, transcriptomic, cellular, organ, exposomic, etc.) collected from a comprehensive collection of patients and encompassing a wide variety of conditions. However, to the best of our knowledge, to this date no such dataset exists. This is mainly because collecting paired, multilayer data from patients is expensive and entails important ethical and privacy concerns (Jobin et al., 2019; Mittelstadt, 2019). To address this issue, our GAN framework can synthesize data at multiple layers conditioned on the patient's conditions (e.g., diabetic, hypertension, etc.) and clinical information (e.g., heart rate, blood pressure, age, sex, ethnicity, exercise, nutrition, etc.). This synthetic data can be used to train the GNN and impute missing data modalities of real patients. Yet, the lack of real data from patients remains the key limitation for the introduction of our framework in clinical practice.

4.1.4. Explainable AI

The lack of interpretability of deep learning models has been one of the most significant barriers preventing their application in healthcare. Such models exhibit great capacity (Hornik, 1991) but understanding their behavior and following their decision-making process is not trivial (Castelvecchi, 2016). There is a growing body of literature focusing on interpretable artificial

intelligence and interpretable deep learning aiming at developing white box models or at explaining black box ones (Das and Rad, 2020). Among such techniques, GNNs have started drawing the attention of both research and industry communities (Bronstein et al., 2017; Zhou et al., 2018). Such models are much more interpretable with respect to other neural approaches thanks to their graph structure, which is quite easy to understand from a human standpoint and a few studies have already shown how graph networks can be effectively employed in biology and healthcare (Zitnik et al., 2018; Gysi et al., 2020).

4.2. Advantages, Limitations, and Visions

4.2.1. Toward Precision and Predictive Medicine

The future of medicine is already bound to AI (Topol, 2019). Technological innovations are completely changing medicine perspectives expanding its horizons and moving toward an holistic view of human beings. The destiny of the whole healthcare system depends on this radical paradigm shift. Embracing AI innovations is just a technological prerequisite, and the first step toward a total transformation of how medicine currently works is delivered and perceived by patients. Thinking that AI will just and mainly improve clinical decision making is wrong. AI may actually open the doors to completely new ways of investigating the human body as a whole. The core and ultimate purpose of health will be developing preventative and personalized pathways to well-being rather than delivering treatments. The future foreseen is that AI will assist medicine in improving diagnosis and devising novel therapeutic strategies to deliver more effective solutions. The current healthcare revolution will not take back all the past technological advances, but it will show them under a new light.

4.2.2. Patient's Benefit

A meaningful quote about twins is the following: being a twin is like being born with a best friend. The data integration will make a better portrait of patient's condition trajectories but will require data inter-operability and data security. Technology is often not neutral, but transformed to be biased in one way or another (Ellul et al., 1954). Individuals can have different unforeseen readings and usage of new technologies. It may increase both user vulnerability and user empowerment. The vulnerability is the combination of exposure to the variety of personal medical data and the coping capabilities of users that could be different between young and mature people, as young are usually quicker in incorporating a new technology into everyday life. The user is empowered if he/she acquires awareness and control of his/her condition and context. A common example are online (website and blogs) initiatives such as patientslikeme that allow the user to search and make up his/her mind about a disease (Wicks et al., 2010). Instead, the user disempowerment depends on the lack of technical knowledge of how mechanisms work; this is even enhanced in black box techniques such as deep learning.

4.2.3. Training Clinicians

We believe that improving both data integration and predictability will provide physicians with improved medical

decisions support systems and a decrease in both costs, through the evaluation of best therapies, and errors. A limitation is the poor interpretability and explainability in deep learning architectures. This limitation will also greatly affect the training of the new clinicians on AI technologies. There are growing efforts to make neural networks more interpretable in order to keep the human (doctors and patients) in the loop. The interpretability could be improved by using parallel mechanistic computational modeling and simulations (Milanesi et al., 2009; Bartocci and Lió, 2016), model extraction libraries (see, for instance, Kazhdan et al., 2020), and visual inference tools (Bodnar et al., 2020). This tool could also be complemented by clinical decision support systems such as Müller and Lio (2020). The complexly structured and multilevel comorbidity and frailty patterns of most diseases describe a highly dynamical system and are, therefore, challenging current medical therapies.

4.2.4. AI for Evidence-Based Medicine

From a clinical standpoint, AI will support a plethora of different tasks from medical check up to personalized intervention strategies to contrast ripple effects or to promote healthy habits. In non-acute states, predictive inference will propose prevention plans for comorbidity management, particularly in presence of multiple therapies (Rivera, 2020). Increasingly large amount of personal data will be collected to feed modular machine learning (ML) models organized to address specific and personalized medical issues. Clinical endpoints will be constantly monitored, shared, and compared in order to answer relevant research questions and to deliver the best possible service. A deeper understanding and practice of modeling in medicine will produce better investigation of complex biological processes, and even new ideas and better feedback into medicine. Modeling-based approaches combined with data-driven ML techniques will progressively provide models with higher degree of interpretability and generalization ability (Barbiero et al., 2020a), which will make evidence-based medicine even more accessible intensifying the involvement of patients in the decision-making process. AI simulations forecasting the evolution of clinical endpoints over time will also reshape clinical guidelines (Rivera, 2020), which will no longer be based just on *horizontal* population studies. Cross-modality data will be collected for each patient and machine learning models will be used to predict a bundle of possible trajectories representing the future states of the patient allowing for personalized prescriptions, surgical planning, and medical interventions.

4.2.5. Social Impact

Ethical repercussions will also be huge (Jobin et al., 2019; Mittelstadt, 2019). The transition will call for deeper trans-disciplinary research and a substantial technological innovation in a variety of research and social areas. Here, education will play a key role in changing lifestyle habits and the way health is perceived, communicated, and delivered (Yu et al., 2017). For each individual, both healthcare systems and private companies will collect, save, and eventually exploit an enormous amount of personal data. Providing an effective, stable, and unified juridical overview is critical on this matter (Panch et al., 2019).

4.2.6. Next-Generation Medical Devices

AI will change the leading *vehicle* of medicine. The demand for AI-powered and internet of things (IoT) devices is increasing worldwide. The future equipment for precision medicine will likely required to be cheap and extremely modular, but more importantly it needs to be deployable in dedicated hardware to be distributed in larger markets. Our digital twin model aims at providing the first example of a novel class of AI-assisted tools for precision and predictive medicine. Our framework is designed to scale to medical device deployment and run time monitoring and verification combining ideas from systems medicine with scientific computing and machine learning. The integration of interpretable AI models in clinical devices may lead to a deep transformation in healthcare paving the way for a next generation of tools for precision medicine probing the inner workings of full body in well-being and disease conditions.

5. METHODS

5.1. Graph Neural Network

5.1.1. Graph Network Blocks

The GNN framework proposed by Battaglia et al. (2018) is based on modules called graph network blocks (GN blocks) representing the core computation units of a GNN. Multiple GN blocks can be composed of or even combined with other neural networks to generate complex architectures. A GNN can be defined as a 3-tuple $G = (\mathbf{u}, H, E)$. $H = \{\mathbf{h}_i\}_{i=1:N^v}$ is the node set where the feature of each node is denoted by \mathbf{h}_i . $E = \{(e_k, r_k, s_k)\}$ is the edge set where each node is represented by its features e_k , the receiver node r_k , and the sender node s_k . \mathbf{u} denotes a set of global attributes representing the state of the underlying system. Each GN block consists of three update functions, ϕ , and three aggregation functions, ρ :

$$\begin{aligned} \mathbf{e}'_k &= \phi^e(\mathbf{e}_k, \mathbf{h}_{r_k}, \mathbf{h}_{s_k}, \mathbf{u}) & \bar{\mathbf{e}}'_i &= \rho^{e \rightarrow h}(E'_i) \\ \mathbf{h}'_i &= \phi^h(\bar{\mathbf{e}}'_i, \mathbf{h}_i, \mathbf{u}) & \bar{\mathbf{e}}' &= \rho^{e \rightarrow u}(E') \\ \mathbf{u}' &= \phi^u(\mathbf{e}', \mathbf{h}', \mathbf{u}) & \bar{\mathbf{h}}' &= \rho^{h \rightarrow u}(H') \end{aligned} \quad (1)$$

where $E'_i = \{(\mathbf{e}'_k, r_k, s_k)\}$, $H' = \{(\mathbf{h}'_i)\}_{i=1:N^v}$, and $E' = \bigcup_i E'_i = \{(\mathbf{e}'_k, r_k, s_k)\}_{k=1:N^e}$. In order to train a GN block in full, six computation steps are required, alternating the update and aggregation functions. For each edge, E'_i is computed through the update function ϕ^e . The result is then aggregated by means of the function $\rho^{e \rightarrow v}$. The output $\bar{\mathbf{e}}'_i$ corresponds to an edge update and it is employed to update node representations \mathbf{h}'_i by means of ϕ^h . $\rho^{e \rightarrow u}$ and $\rho^{h \rightarrow u}$ perform aggregation steps generating $\bar{\mathbf{e}}'$ and $\bar{\mathbf{h}}'$ from edge and node updates, respectively. Global attributes represented by \mathbf{u}' are computed leveraging the information from $\bar{\mathbf{e}}'$, $\bar{\mathbf{h}}'$, and \mathbf{u} via the function ϕ^u . The learning process of each GN block may be independent or co-dependent with other blocks. Constraints may apply on edges, information flows, or global attributes, depending on the application. In this work, we are just interested in the evaluation of global attributes to monitor clinical endpoints and we did not apply any learning constraint, even if in clinical practice may still be of great interest. Given a set of labels for global attributes $\mathbf{t} = \{t_i\}_{i=1:N^v}$ and

the corresponding predictions provided by the GN block $\hat{\mathbf{u}}' = \{\hat{u}'_i\}_{i=1:N^v}$ representing the evolution of the underlying biological system, we aim at minimizing the following objective function:

$$\min_{\theta} \frac{1}{N^v} \sum_{i=1}^{N^v} (t_i - \hat{u}'_i)^2 \quad (2)$$

where θ is the set of model's parameters.

5.1.2. Assessing Prediction Uncertainty

The aim of developing a digital patient model is to provide an accurate estimation of the trajectory of a patient by forecasting clinically relevant endpoints. In such a context, quantifying model uncertainty is critical. One of the most established techniques relies upon the use of dropout (Srivastava et al., 2014) at test time, as a Bayesian approximation, without sacrificing either computational complexity or test performance (Gal and Ghahramani, 2016b). In this framework, the first two moments of the predictive distribution q performing T stochastic forward passes for a sample \mathbf{x}^* with label \mathbf{y}^* can be estimated as (Gal and Ghahramani, 2016a):

$$\begin{aligned} \mathbb{E}_{q(\mathbf{y}^*, \mathbf{x}^*)}(\mathbf{y}^*) &\approx \frac{1}{T} \sum_{i=1}^T \hat{\mathbf{y}}^*(\mathbf{x}^*, W_1^t, \dots, W_L^t) & (3) \\ \text{Var}_{q(\mathbf{y}^*, \mathbf{x}^*)}(\mathbf{y}^*) &\approx \tau^{-1} I_D \\ &+ \frac{1}{T} \sum_{i=1}^T \hat{\mathbf{y}}^*(\mathbf{x}^*, W_1^t, \dots, W_L^t)^T \hat{\mathbf{y}}^*(\mathbf{x}^*, W_1^t, \dots, W_L^t) \\ &- \mathbb{E}_{q(\mathbf{y}^*, \mathbf{x}^*)}(\mathbf{y}^*)^T \mathbb{E}_{q(\mathbf{y}^*, \mathbf{x}^*)}(\mathbf{y}^*) & (4) \end{aligned}$$

where $\hat{\mathbf{y}}^*$ is the predicted label, $\{W_i\}_{i=1}^L$ is a set of random variables representing the weights of a neural network with L layers, I_D is an identity matrix, D is the number of output units of the neural network, and τ is a precision hyper-parameter. The method has also been generalized to convolutional (Gal and Ghahramani, 2015) and recurrent networks (Gal and Ghahramani, 2016c).

Here, we show how such technique can be used to quantify the uncertainty of a GNN by generating a predictive distribution of the trajectories representing the future states of the patient. Let x_1^*, \dots, x_k^* be a sequence of real values representing a clinical endpoint measured at $1, \dots, k$ time steps. Let f^t be a stochastic model that takes a sequence x_1^*, \dots, x_k^* as input and it outputs a prediction $\hat{y}^* \in \mathbb{R}$. We are interested in estimating a predictive distribution of the trajectories of the variable x over the next $k + 1, \dots, k + h$ time steps. To this aim, we can use an iterative algorithm by generating one trajectory at a time. The first prediction \hat{y}_{k+1}^* can be generated as:

$$\hat{y}_{k+1}^{*,t} = f^t(x_1^*, \dots, x_k^*) \quad (5)$$

By using the obtained prediction and sliding the time window one time step further, we can generate the first prediction for the second time step $k + 2$:

$$\hat{y}_{k+2}^{*,t} = f^t(x_2^*, \dots, x_k^*, \hat{y}_{k+1}^{*,t}) \quad (6)$$

The procedure can be repeated for $k + h$ time steps to generate a single trajectory. Model uncertainty can be assessed building multiple trajectories by performing T stochastic forward passes. The resulting algorithm is equivalent to a Monte Carlo sampling as proven by Gal and Ghahramani (2016b). In our GNN model, the approach we just described can be easily applied for each node in order to assess the uncertainty of clinical endpoints.

5.2. Generative Adversarial Network

Consider a dataset $\mathcal{D} = \{(\mathbf{x}, \mathbf{m}, \mathbf{r}, \mathbf{q})\}$ of samples from an unknown distribution $\mathbb{P}_{\mathbf{x}, \mathbf{m}, \mathbf{r}, \mathbf{q}}$, where $\mathbf{x} \in \mathbb{R}^{t \times n}$ represents a matrix of n gene expression values in t tissues; $\mathbf{m} \in \{0, 1\}^t$ is a mask vector indicating whether the expression of each tissue has been measured for the given patient; and $\mathbf{r} \in \mathbb{R}^k$ and $\mathbf{q} \in \mathbb{N}^c$ are vectors of k quantitative covariates (e.g., age) and c categorical (e.g., gender), respectively. Our goal is to produce realistic gene expression samples by modeling the conditional probability distribution $\mathbb{P}(\mathbf{X} = \mathbf{x} | \mathbf{M} = \mathbf{m}, \mathbf{R} = \mathbf{r}, \mathbf{Q} = \mathbf{q})$, where \mathbf{r} includes the expression of ACE2 in different tissues (e.g., lung, kidney, and pancreas). By modeling this distribution, we can sample data for different conditions and quantify the uncertainty of the generated expression values.

To address this problem, we extend the model proposed in Viñas et al. (2021) to simultaneously account for t tissue types from the same donor. In particular, our method builds on a Wasserstein GAN with gradient penalty (WGAN-GP) (Arjovsky et al., 2017; Gulrajani et al., 2017). Similar to Generative Adversarial Networks (GAN) (Goodfellow et al., 2014), WGAN-GPs estimate a generative model via an adversarial process driven by the competition between two players, the *generator* and the *critic*.

The generator aims at producing samples from the conditional $\mathbb{P}(\mathbf{X} | \mathbf{M}, \mathbf{R}, \mathbf{Q})$. Formally, we define the generator as a function $G_\theta: \mathbb{R}^u \times \mathbb{R}^k \times \mathbb{N}^c \rightarrow \mathbb{R}^{t \times n}$ parameterized by θ that generates gene expression values $\hat{\mathbf{x}}$ as follows:

$$\hat{\mathbf{x}} = \mathbf{m} \odot G_\theta(\mathbf{z}, \mathbf{r}, \mathbf{q}) \quad (7)$$

where $\mathbf{z} \in \mathbb{R}^u$ is a vector sampled from a fixed noise distribution $\mathbb{P}_{\mathbf{z}}$ and u is a user-definable hyperparameter. We apply the mask \mathbf{m} element-wise to match the distribution of missing tissues of the training dataset.

The critic takes gene expression samples \mathbf{x} from two input streams (the generator and the data distribution) and attempts to distinguish the true input source. Formally, the critic is a function $D_\omega: \mathbb{R}^{t \times n} \times \{0, 1\}^t \times \mathbb{R}^k \times \mathbb{N}^c \rightarrow \mathbb{R}$ parameterized by ω that we define as follows:

$$\bar{y} = D_\omega(\bar{\mathbf{x}}, \mathbf{m}, \mathbf{r}, \mathbf{q})$$

where the output \bar{y} is an unbounded scalar that quantifies the degree of realism of an input sample $\bar{\mathbf{x}}$ given the covariates \mathbf{r} and \mathbf{q} (e.g., high values correspond to real samples and low values correspond to fake samples). When the expression of a certain tissue t is unavailable for a given patient, we set the unobserved

values of tissue t in $\bar{\mathbf{x}}$ to 0 and the t -th component of the mask \mathbf{m} to 0.

We optimize the generator and the critic adversarially. Following (Arjovsky et al., 2017), we train the generator G_θ and the critic D_ω to solve the following minimax game based on the Wasserstein distance:

$$\begin{aligned} \min_{\theta} \max_{\omega} \quad & \mathbb{E}_{\mathbf{x}, \mathbf{m}, \mathbf{r}, \mathbf{q} \sim \mathbb{P}_{\mathbf{x}, \mathbf{m}, \mathbf{r}, \mathbf{q}}} \left[D_\omega(\mathbf{x}, \mathbf{m}, \mathbf{r}, \mathbf{q}) - \mathbb{E}_{\mathbf{z} \sim \mathbb{P}_{\mathbf{z}}} [D_\omega(\hat{\mathbf{x}}, \mathbf{m}, \mathbf{r}, \mathbf{q})] \right] \\ \text{subject to} \quad & \|D_\omega(\mathbf{x}_i, \mathbf{m}, \mathbf{r}, \mathbf{q}) - D_\omega(\mathbf{x}_j, \mathbf{m}, \mathbf{r}, \mathbf{q})\| \leq \|\mathbf{x}_i - \mathbf{x}_j\| \\ & \forall \mathbf{x}_i, \mathbf{x}_j \in \mathbb{R}^{t \times n}, \mathbf{m} \in \{0, 1\}^t, \mathbf{r} \in \mathbb{R}^k, \mathbf{q} \in \mathbb{N}^c \end{aligned} \quad (8)$$

where $\hat{\mathbf{x}}$ is defined as in Equation (7) and the constraint enforces a soft version of the 1-Lipschitz constraint (e.g., the norm of the critic's gradient with respect to \mathbf{x} must be at most 1 everywhere).

Let $\{(\mathbf{x}_i, \mathbf{m}_i, \mathbf{r}_i, \mathbf{q}_i)\}_{i=1}^b$ be a mini-batch of b independent samples from the training dataset \mathcal{D} . Let $\{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_k\}$ be a set of k vectors sampled independently from the noise distribution $\mathbb{P}_{\mathbf{z}}$ and let us define the synthetic samples corresponding to the mini-batch as $\hat{\mathbf{x}}_i = \mathbf{m}_i \odot G_\theta(\mathbf{z}_i, \mathbf{r}_i, \mathbf{q}_i)$ for each i in $[1, 2, \dots, k]$. We solve the minimax problem described in Equation (8) by interleaving mini-batch gradient updates for the generator and the critic, optimizing the following problems:

$$\begin{aligned} \text{Generator:} \quad & \min_{\theta} \quad -\frac{1}{k} \sum_{i=1}^k D_\omega(\hat{\mathbf{x}}_i, \mathbf{m}_i, \mathbf{r}_i, \mathbf{q}_i) \\ \text{Critic:} \quad & \min_{\omega} \quad \frac{1}{k} \sum_{i=1}^k D_\omega(\hat{\mathbf{x}}_i, \mathbf{m}_i, \mathbf{r}_i, \mathbf{q}_i) - D_\omega(\mathbf{x}_i, \mathbf{m}_i, \mathbf{r}_i, \mathbf{q}_i) \\ & + \frac{\lambda}{k} \sum_{i=1}^k (\|\nabla_{\hat{\mathbf{x}}_i} D_\omega(\hat{\mathbf{x}}_i, \mathbf{m}_i, \mathbf{r}_i, \mathbf{q}_i)\|_2 - 1)^2 \end{aligned} \quad (9)$$

where λ is a user-definable hyperparameter and each $\tilde{\mathbf{x}}_i$ is a random point along the straight line that connects \mathbf{x}_i and $\hat{\mathbf{x}}_i$, that is, $\tilde{\mathbf{x}}_i = \alpha_i \mathbf{x}_i + (1 - \alpha_i) \hat{\mathbf{x}}_i$ with $\alpha_i \sim \mathcal{U}(0, 1)$. Intuitively, since enforcing the 1-Lipschitz constraint everywhere is intractable (see Equation 8), the second term of the critic problem is a relaxed version of the constraint that penalizes the gradient norm along points in the straight lines that connect real and synthetic samples (Gulrajani et al., 2017).

5.2.1. Architecture

Figure 1 shows the architecture of both players. The generator G receives a noise vector \mathbf{z} as input (green box) as well as sample covariates \mathbf{r} and \mathbf{q} (orange boxes) and produces a vector $\hat{\mathbf{x}}$ of synthetic expression values (red box). The critic D takes either a real gene expression sample \mathbf{x} (blue box) or a synthetic sample $\hat{\mathbf{x}}$ (red box), in addition to sample covariates \mathbf{r} and \mathbf{q} , and attempts to distinguish whether the input sample is real or fake. For both players, we use word embeddings (Mikolov et al., 2013) to model the sample covariates (light green boxes), a distinctive feature that allows to learn distributed, dense representations for the different tissue types and, more generally, for all the categorical covariates $\mathbf{q} \in \mathbb{N}^c$.

Formally, let q_j be a categorical covariate (e.g., tissue type) with vocabulary size v_j , that is, $q_j \in \{1, 2, \dots, v_j\}$, where each value in the vocabulary $\{1, 2, \dots, v_j\}$ represents a different category (e.g., whole blood or kidney). Let $\bar{\mathbf{q}}_j \in \{0, 1\}^{v_j}$ be a one-hot vector such that $\bar{q}_{jk} = 1$ if $q_j = k$ and $\bar{q}_{jk} = 0$ otherwise. Let d_j be the dimensionality of the embeddings for covariate j . We obtain a vector of embeddings $\mathbf{e}_j \in \mathbb{R}^{d_j}$ as follows:

$$\mathbf{e}_j = \mathbf{W}_j \bar{\mathbf{q}}_j$$

where each $\mathbf{W}_j \in \mathbb{R}^{d_j \times v_j}$ is a matrix of learnable weights. Essentially, this operation describes a lookup search in a dictionary with v_j entries, where each entry contains a learnable d_j -dimensional vector of embeddings that characterizes each of the possible values that q_j can take. To obtain a global collection of embeddings \mathbf{e} , we concatenate all the vectors \mathbf{e}_j for each categorical covariate j :

$$\mathbf{e} = \left\|_{j=1}^c \mathbf{e}_j \right\|$$

where c is the number of categorical covariates and $\|$ represents the concatenation operator. We then use the learnable embeddings \mathbf{e} in downstream tasks.

In terms of the player's architecture, we model both the generator G_θ and critic D_ω as neural networks that leverage independent instances \mathbf{e}^G and \mathbf{e}^D of the categorical embeddings for their corresponding downstream tasks. Specifically, we model the two players as follows:

$$G_\theta(\mathbf{z}, \mathbf{r}, \mathbf{q}) = \text{MLP}(\mathbf{z} \parallel \mathbf{r} \parallel \mathbf{e}^G) \quad D_\omega(\bar{\mathbf{x}}, \mathbf{m}, \mathbf{r}, \mathbf{q}) = \text{MLP}(\bar{\mathbf{x}} \parallel \mathbf{m} \parallel \mathbf{r} \parallel \mathbf{e}^D)$$

where MLP denotes a multilayer perceptron.

5.3. Ridge Regression

We model the expression of genes from the renin-angiotensin system in lung, kidney, pancreas, and heart as a function of genes in the chemokine, TNF, and TGF- β pathways in blood. Let $\mathbf{Y} = (Y_1, \dots, Y_n)^\top$ and $\mathbf{X} = (X_1, \dots, X_m)^\top$ be multivariate random variables representing the expression of the n genes in the renin-angiotensin system and the m genes in the signaling pathways, respectively. Our model is based on ridge regression (Hoerl and Kennard, 1970):

$$\mathbf{Y} = \mathbf{X}\mathbf{W} + \boldsymbol{\epsilon}$$

where $\mathbf{W} \in \mathbb{R}^{m \times n}$ is a matrix of learnable weights and $\boldsymbol{\epsilon} \in \mathbb{R}^n$ are the residuals. We optimize the following objective:

$$\min_{\mathbf{W}} \|\mathbf{Y} - \mathbf{X}\mathbf{W}\|_2^2 + \alpha \|\mathbf{W}\|_2^2$$

where α is a hyperparameter that controls the regularization strength. Alternative non-linear models such as support vector machines, Gaussian processes, and random forests did not improve our cross-validation scores.

DATA AVAILABILITY STATEMENT

We leverage data from the Genotype-Tissue Expression (GTEx) project (v8), a resource that has generated a comprehensive collection of human transcriptome data in a diverse set of tissues (Aguet et al., 2019). We use the KEGG pathway database (Kanehisa et al., 2010) to select genes from the renin-angiotensin system (*hsa04614*), chemokine (*hsa04062*), TNF (*hsa04668*), and TGF- β (*hsa04350*) pathways.

The computational ODE-based system described in Barbiero and Lió (2020) is used to generate a time series for each differential equation with a window size of $\tau = 500$ time steps (Barbiero and Lio, 2020). Time series are collected, randomly shuffled, and stacked in a dataset. Each item of the collection is randomly assigned either to a training ($n_{train} = 3200$), validation ($n_{val} = 800$), or test set ($n_{test} = 1000$).

CODE AVAILABILITY

All the code for the experiments has been implemented in Python 3, relying upon open-source libraries (Pedregosa et al., 2011; Abadi et al., 2016; Wang et al., 2019). All the experiments have been run on the same machine: Intel® Core™ i7-8750H 6-Core Processor at 2.20 GHz equipped with 8 GB RAM. To enable code reuse, the Python code for the mathematical models including parameter values and documentation is freely available under GNU Public License from a GitHub repository¹ (Barbiero et al., 2020b). Unless required by applicable law or agreed to in writing, software is distributed on an “as is” basis, without warranties or conditions of any kind, either express or implied.

AUTHOR CONTRIBUTIONS

PB and RV performed the experiments. All authors were involved in conceiving the approach, co-wrote the paper, designing the experiments, reviewing and discussing the data, and commented on the manuscript.

FUNDING

This project has received funding from “la Caixa” Foundation (ID 100010434), under agreement 667 LCF/BQ/EU19/11710059, and from the Engineering and Physical Sciences 668 Research Council under grant agreement No EP/R513180/1 (R.V. EPSRC DTG 2018/19).

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fgene.2021.652907/full#supplementary-material>

¹<https://github.com/pietrobarbiero/digital-patient>

REFERENCES

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., et al. (2016). “Tensorflow: a system for large-scale machine learning,” in *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)* (Savannah, GA), 265–283.
- Aguet, F., Barbeira, A. N., Bonazzola, R., Brown, A., Castel, S. E., Jo, B., et al. (2019). The gtx consortium atlas of genetic regulatory effects across human tissues. *bioRxiv [Preprint]*. doi: 10.1101/787903
- Arjovsky, M., Chintala, S., and Bottou, L. (2017). Wasserstein GAN. *arXiv [Preprint]*. arXiv:1701.07875.
- Bangalore, S., Maron, D. J., O’Brien, S. M., Fleg, J. L., Kretov, E. I., Briguori, C., et al. (2020). Management of coronary disease in patients with advanced kidney disease. *N. Engl. J. Med.* 382, 1608–1618. doi: 10.1056/nejmoa1915925
- Barbiero, P., and Lió, P. (2020). The computational patient has diabetes and a covid. *arXiv [Preprint]*. arXiv:2006.06435.
- Barbiero, P., and Lio, P. (2020). Pietrobarbiero/computational-patient: absolutno. doi: 10.5281/zenodo.4030228
- Barbiero, P., Squillero, G., and Tonda, A. (2020a). Modeling generalization in machine learning: A methodological and computational study. *arXiv [Preprint]*. arXiv:2006.15680.
- Barbiero, P., Torne, R. V., and Lio, P. (2020b). pietrobarbiero/digital-patient: absolutno. doi: 10.5281/zenodo.4030220
- Bartocci, E., and Lió, P. (2016). Computational modeling, formal analysis, and tools for systems biology. *PL S Computat. Biol.* 12:e1004591. doi: 10.1371/journal.pcbi.1004591
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., et al. (2018). Relational inductive biases, deep learning, and graph networks. *arXiv [Preprint]*. arXiv:1806.01261.
- Bica, I., Andrés-Terré, H., Cvejic, A., and Lió, P. (2020). Unsupervised generative and graph representation learning for modelling cell differentiation. *Sci. Rep.* 10, 1–13. doi: 10.1038/s41598-020-66166-8
- Bodnar, C., Cangea, C., and Lió, P. (2020). Deep graph mapper: seeing graphs through the neural lens. *arXiv [Preprint]*. arXiv:2002.03864.
- Bollobás, B. (2013). *Modern graph Theory*, Vol. 184. New York, NY: Springer Science & Business Media.
- Bronstein, M. M., Bruna, J., LeCun, Y., Szlam, A., and Vandergheynst, P. (2017). Geometric deep learning: going beyond euclidean data. *IEEE Signal Process. Mag.* 34, 18–42. doi: 10.1109/msp.2017.2693418
- Brown, S. A., Kovatchev, B. P., Raghinaru, D., Lum, J. W., Buckingham, B. A., Kudva, Y. C., et al. (2019). Six-month randomized, multicenter trial of closed-loop control in type 1 diabetes. *N. Engl. J. Med.* 381, 1707–1717. doi: 10.1056/NEJMoa1907863
- Castelvecchi, D. (2016). Can we open the black box of AI? *Nat. News* 538:20. doi: 10.1038/538020a
- Clark, D., Nicholls, S. J., John, J. S., Elshazly, M. B., Ahmed, H. M., Khraishah, H., et al. (2019). Visit-to-visit blood pressure variability, coronary atheroma progression, and clinical outcomes. *JAMA Cardiol.* 4:437. doi: 10.1001/jamacardio.2019.0751
- Colbran, L. L., Gamazon, E. R., Zhou, D., Evans, P., Cox, N. J., and Capra, J. A. (2019). Inferred divergent gene regulation in archaic hominins reveals potential phenotypic differences. *Nat. Ecol. Evol.* 3, 1598–1606. doi: 10.1038/s41559-019-0996-x
- Cookson, W., Liang, L., Abecasis, G., Moffatt, M., and Lathrop, M. (2009). Mapping complex disease traits with global gene expression. *Nat. Rev. Genet.* 10, 184–194. doi: 10.1038/nrg2537
- Das, A., and Rad, P. (2020). Opportunities and challenges in explainable artificial intelligence (XAI): a survey. *arXiv [Preprint]*. arXiv:2006.11371.
- Du, J., Johnson, L. M., Jacobsen, S. E., and Patel, D. J. (2015). Dna methylation pathways and their crosstalk with histone methylation. *Nat. Rev. Mol. Cell Biol.* 16, 519–532. doi: 10.1038/nrm4043
- Eddy, D. M., and Schlessinger, L. (2003). Archimedes: a trial-validated model of diabetes. *Diabetes Care* 26, 3093–3101. doi: 10.2337/diacare.26.11.3093
- Ellul, J., Ellul, J., Ellul, J., Juriste, P., and Ellul, J. (1954). *La technique ou l'enjeu du siècle*. A. Colin Paris.
- Elmarakby, A. A., Quigley, J. E., Olearczyk, J. J., Sridhar, A., Cook, A. K., Insko, E. W., et al. (2007). Chemokine receptor 2b inhibition provides renal protection in angiotensin ii-salt hypertension. *Hypertension* 50, 1069–1076. doi: 10.1161/HYPERTENSIONAHA.107.098806
- Emilsson, V., Thorleifsson, G., Zhang, B., Leonardson, A. S., Zink, F., Zhu, J., et al. (2008). Genetics of gene expression and its effect on disease. *Nature* 452, 423–428. doi: 10.1038/nature06758
- Evans, W. E., and Relling, M. V. (2004). Moving towards individualized medicine with pharmacogenomics. *Nature* 429, 464–468. doi: 10.1038/nature02626
- Fountain, J. H., and Lappin, S. L. (2019). “Physiology, renin angiotensin system,” in *StatPearls*, (StatPearls Publishing).
- Gal, Y., and Ghahramani, Z. (2015). Bayesian convolutional neural networks with bernoulli approximate variational inference. *arXiv [Preprint]*. arXiv:1506.02158.
- Gal, Y., and Ghahramani, Z. (2016a). Dropout as a bayesian approximation: appendix 20. *arXiv [Preprint]*. arxiv:1506.02157.
- Gal, Y., and Ghahramani, Z. (2016b). “Dropout as a bayesian approximation: representing model uncertainty in deep learning,” in *International Conference on Machine Learning*, 1050–1059.
- Gal, Y., and Ghahramani, Z. (2016c). “A theoretically grounded application of dropout in recurrent neural networks,” in *Advances in Neural Information Processing Systems*, 1019–1027.
- Gamazon, E. R., Segre, A. V., van de Bunt, M., Wen, X., Xi, H. S., Hormozdiari, F., et al. (2018). Using an atlas of gene regulation across 44 human tissues to inform complex disease-and trait-associated variation. *Nat. Genet.* 50, 956–967. doi: 10.1038/s41588-018-0154-4
- Garvin, M. R., Alvarez, C., Miller, J. I., Prates, E. T., Walker, A. M., Amos, B. K., et al. (2020). A mechanistic model and therapeutic interventions for covid-19 involving a ras-mediated bradykinin storm. *eLife* 9:e59177. doi: 10.7554/eLife.59177.sa2
- Geiger, B., Bershadsky, A., Pankov, R., and Yamada, K. M. (2001). Transmembrane crosstalk between the extracellular matrix and the cytoskeleton. *Nat. Rev. Mol. Cell Biol.* 2, 793–805. doi: 10.1038/35099066
- Gelernter, D. (1993). *Mirror Worlds: Or: The Day Software Puts the Universe in a Shoebox. How It Will Happen and What It Will Mean*. Oxford University Press.
- Ginsburg, G. S., and Willard, H. F. (2009). Genomic and personalized medicine: foundations and applications. *Transl. Res.* 154, 277–287. doi: 10.1016/j.trsl.2009.09.005
- Gironacci, M. M., Adamo, H. P., Corradi, G., Santos, R. A., Ortiz, P., and Carretero, O. A. (2011). Angiotensin (1-7) induces mas receptor internalization. *Hypertension* 58, 176–181. doi: 10.1161/HYPERTENSIONAHA.111.173344
- Goldberger, A. L., Amaral, L. A., Hausdorff, J. M., Ivanov, P. C., Peng, C.-K., and Stanley, H. E. (2002). Fractal dynamics in physiology: alterations with disease and aging. *Proc. Natl. Acad. Sci. U.S.A.* 99, 2466–2472. doi: 10.1073/pnas.012579499
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). “Generative adversarial nets,” in *Advances in Neural Information Processing Systems*, 2672–2680.
- Grieves, M. (2015). “Digital twin: manufacturing excellence through virtual factory replication,” *Whitepaper*.
- Groth, A., Vrugt, B., Brock, M., Speich, R., Ulrich, S., and Huber, L. C. (2014). Inflammatory cytokines in pulmonary hypertension. *Respir. Res.* 15, 1–10. doi: 10.1186/1465-9921-15-47
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., and Courville, A. (2017). Improved training of Wasserstein GANs. *arXiv [Preprint]*. arXiv: 1704.00028.
- Gutierrez, N. G., Mathew, M., McCrindle, B. W., Tran, J. S., Kahn, A. M., Burns, J. C., et al. (2019). Hemodynamic variables in aneurysms are associated with thrombotic risk in children with Kawasaki disease. *Int. J. Cardiol.* 281, 15–21. doi: 10.1016/j.ijcard.2019.01.092
- Gysi, D. M., Valle, Í. D., Zitnik, M., Ameli, A., Gan, X., Varol, O., et al. (2020). Network medicine framework for identifying drug repurposing opportunities for covid-19. *arXiv [Preprint]*. arXiv:2004.07229.
- Hoerl, A. E., and Kennard, R. W. (1970). Ridge regression: biased estimation for nonorthogonal problems. *Technometrics* 12, 55–67. doi: 10.1080/00401706.1970.10488634
- Hornik, K. (1991). Approximation capabilities of multilayer feedforward networks. *Neural Netw.* 4, 251–257. doi: 10.1016/0893-6080(91)90009-T

- Huang, Q., Yamada, M., Tian, Y., Singh, D., Yin, D., and Chang, Y. (2020). Graphlime: local interpretable model explanations for graph neural networks. *arXiv [Preprint]*. arXiv:2001.06216.
- Jebara, T. (2012). *Machine Learning: Discriminative and Generative*, Vol. 755. Springer Science & Business Media.
- Jobin, A., Ienca, M., and Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nat. Mach. Intell.* 1, 389–399. doi: 10.1038/s42256-019-0088-2
- Kanehisa, M., Goto, S., Furumichi, M., Tanabe, M., and Hirakawa, M. (2010). Keg for representation and analysis of molecular networks involving diseases and drugs. *Nucleic Acids Res.* 38, D355–D360. doi: 10.1093/nar/gkp896
- Kazhdan, D., Shams, Z., and Liò, P. (2020). Marleme: a multi-agent reinforcement learning model extraction library. *arXiv [Preprint]*. arXiv:2004.07928.
- Kim, B. J., Kwon, S. U., Wajsbrot, D., Koo, J., Park, J. M., and Jeffers, B. W. (2018). Relationship of inter-individual blood pressure variability and the risk for recurrent stroke. *J. Am. Heart Assoc.* 7:e009480. doi: 10.1161/jaha.118.009480
- Kovatchev, B. (2019). A century of diabetes technology: signals, models, and artificial pancreas control. *Trends Endocrinol. Metab.* 30, 432–444. doi: 10.1016/j.tem.2019.04.008
- Kuba, K., Imai, Y., Ohto-Nakanishi, T., and Penninger, J. M. (2010). Trilogy of ace2: a peptidase in the renin–angiotensin system, a sars receptor, and a partner for amino acid transporters. *Pharmacol. Therapeut.* 128, 119–128. doi: 10.1016/j.pharmthera.2010.06.003
- Kusnadi, A., Ramírez-Suástegui, C., Fajardo, V., Chee, S. J., Meckiff, B. J., Simon, H., et al. (2020). Severely ill covid-19 patients display augmented functional properties in sars-cov-2-reactive cd8+ t cells. *bioRxiv [Preprint]*.
- Laubenbacher, R., Sluka, J. P., and Glazier, J. A. (2021). Using digital twins in viral infection. *Science* 371, 1105–1106. doi: 10.1126/science.abf370
- Lecue, F. (2019). On the role of knowledge graphs in explainable AI. *Semantic Web* 1, 1–11. doi: 10.3233/SW-190374
- Lengyel, E., Makowski, L., DiGiovanni, J., and Kolonin, M. G. (2018). Cancer as a matter of fat: the crosstalk between adipose tissue and tumors. *Trends Cancer* 4, 374–384. doi: 10.1016/j.trecan.2018.03.004
- Li, M., Gao, P., and Zhang, J. (2016). Crosstalk between autophagy and apoptosis: potential and emerging therapeutic targets for cardiac diseases. *Int. J. Mol. Sci.* 17:332. doi: 10.3390/ijms17030332
- Li, Y., Vinyals, O., Dyer, C., Pascanu, R., and Battaglia, P. (2018). Learning deep generative models of graphs. *arXiv [Preprint]*. arXiv:1803.03324.
- Liao, M., Liu, Y., Yuan, J., Wen, Y., Xu, G., Zhao, J., et al. (2020). Single-cell landscape of bronchoalveolar immune cells in patients with covid-19. *Nat. Med.* 26, 842–844. doi: 10.1038/s41591-020-0901-9
- Lieberman, E., Hauert, C., and Nowak, M. A. (2005). Evolutionary dynamics on graphs. *Nature* 433, 312–316. doi: 10.1038/nature03204
- Lijnen, P. J., Petrov, V. V., and Fagard, R. H. (2003). Association between transforming growth factor- β and hypertension.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G., and Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *arXiv [Preprint]*. arXiv: 1310.4546.
- Milanesi, L., Romano, P., Castellani, G., Remondini, D., and Liò, P. (2009). Trends in modeling biomedical complex systems. *BMC Bioinformatics* 10:11. doi: 10.1186/1471-2105-10-s12-i1
- Mitchell, G. F., Hwang, S.-J., Vasan, R. S., Larson, M. G., Pencina, M. J., Hamburg, N. M., et al. (2010). Arterial stiffness and cardiovascular events. *Circulation* 121, 505–511. doi: 10.1161/circulationaha.109.886655
- Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nat. Mach. Intell.* 1, 501–507. doi: 10.1038/s42256-019-0114-4
- Müller, T. T., and Lio, P. (2020). PECLIDES neuro: a personalisable clinical decision support system for neurological diseases. *Front. Artif. Intell.* 3:23. doi: 10.3389/frai.2020.00023
- Naylor, S., and Chen, J. Y. (2010). Unraveling human complexity and disease with systems biology and personalized medicine. *Pers. Med.* 7, 275–289. doi: 10.2217/pme.10.16
- Neal, M. L., and Bassingthwaite, J. B. (2007). Subject-specific model estimation of cardiac output and blood volume during hemorrhage. *Cardiovasc. Eng.* 7, 97–120. doi: 10.1007/s10558-007-9035-7
- O'Rourke, M. F., and Nichols, W. W. (2005). Aortic diameter, aortic stiffness, and wave reflection increase with age and isolated systolic hypertension. *Hypertension* 45, 652–658. doi: 10.1161/01.hyp.0000153793.84859.b8
- Panch, T., Mattie, H., and Celi, L. A. (2019). The “inconvenient truth” about AI in healthcare. *NPJ Digit. Med.* 2, 1–3. doi: 10.1038/s41746-019-0155-4
- Pearson, K. (1901). On lines of closes fit to system of points in space, London, E Dinb. *Dublin Philos. Mag. J. Sci.* 2, 559–572. doi: 10.1080/14786440109462720
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830.
- Poyton, R. O., and McEwen, J. E. (1996). Crosstalk between nuclear and mitochondrial genomes. *Annu. Rev. Biochem.* 65, 563–607. doi: 10.1146/annurev.bi.65.070196.003023
- Rakocevic, G., Semenyuk, V., Lee, W.-P., Spencer, J., Browning, J., Johnson, I. J., et al. (2019). Fast and accurate genomic analyses using genome graphs. *Nat. Genet.* 51, 354–362. doi: 10.1038/s41588-018-0316-4
- Ray, S., Britschgi, M., Herbert, C., Takeda-Uchimura, Y., Boxer, A., Blennow, K., et al. (2007). Classification and prediction of clinical Alzheimer's diagnosis based on plasma signaling proteins. *Nat. Med.* 13, 1359–1362. doi: 10.1038/nm1653
- Rivera, S. (2020). Guidelines for clinical trial protocols for interventions involving artificial intelligence: the spirit-AI extension. *BMJ* 26, 1351–1363. doi: 10.1038/s41591-020-1037-7
- Robinson, M. D., and Oshlack, A. (2010). A scaling normalization method for differential expression analysis of RNA-Seq data. *Genome Biol.* 11, 1–9. doi: 10.1186/gb-2010-11-3-r25
- Rudemiller, N. P., and Crowley, S. D. (2017). The role of chemokines in hypertension and consequent target organ damage. *Pharmacol. Res.* 119, 404–411. doi: 10.1016/j.phrs.2017.02.026
- Sardu, C., Gambardella, J., Morelli, M. B., Wang, X., Marfella, R., and Santulli, G. (2020). Hypertension, thrombosis, kidney failure, and diabetes: Is COVID-19 an endothelial disease? a comprehensive evaluation of clinical and basic evidence. *J. Clin. Med.* 9:1417. doi: 10.3390/jcm9051417
- Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. (2008). The graph neural network model. *IEEE Trans. Neural Netw.* 20, 61–80. doi: 10.1109/TNN.2008.2005605
- Shafto, M., Conroy, M., Doyle, R., Glaessgen, E., Kemp, C., LeMoigne, J., et al. (2010). *Modeling, Simulation, Information Technology and Processing Roadmap*. National Aeronautics and Space Administration.
- Shang, J. K., Esmaily, M., Verma, A., Reinhartz, O., Figliola, R. S., Hsia, T.-Y., et al. (2019). Patient-specific multiscale modeling of the assisted bidirectional glenn. *Ann. Thorac. Surg.* 107, 1232–1239. doi: 10.1016/j.athoracsur.2018.10.024
- Sirota, M., Dudley, J. T., Kim, J., Chiang, A. P., Morgan, A. A., Sweet-Cordero, A., et al. (2011). Discovery and preclinical validation of drug indications using compendia of public gene expression data. *Sci. Transl. Med.* 3:96ra77. doi: 10.1126/scitranslmed.3001318
- South, A. M., Tomlinson, L., Edmonston, D., Hiremath, S., and Sparks, M. A. (2020). Controversies of renin–angiotensin system inhibition during the covid-19 pandemic. *Nat. Rev. Nephrol.* 16, 305–307. doi: 10.1038/s41581-020-0279-4
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15, 1929–1958.
- Tang, N., Bai, H., Chen, X., Gong, J., Li, D., and Sun, Z. (2020). Anticoagulant treatment is associated with decreased mortality in severe coronavirus disease 2019 patients with coagulopathy. *J. Thromb. Haemost.* 18, 1094–1099. doi: 10.1111/jth.14817
- Tao, Y., Xu, J., Song, B., Xie, X., Gu, H., Liu, Q., et al. (2017). Short-term blood pressure variability and long-term blood pressure variability: which one is a reliable predictor for recurrent stroke. *Journal of Human Hypertension* 31, 568–573. doi: 10.1038/jhh.2017.32
- Topol, E. (2019). *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again*. Hachette.
- Van Der Maaten, L., Postma, E., and Van den Herik, J. (2009). Dimensionality reduction: a comparative. *J. Mach. Learn. Res.* 10:13.
- Viñas, R., Andres-Terre, H., Lio, P., and Bryson, K. (2021). Adversarial generation of gene expression data. *Bioinformatics*. doi: 10.1093/bioinformatics/ctab035. [Epub ahead of print].
- Wang, M., Yu, L., Zheng, D., Gan, Q., Gai, Y., Ye, Z., et al. (2019). Deep graph library: towards efficient and scalable deep learning on graphs. *arXiv [Preprint]*. arXiv:1909.01315.

- Wen, W., Luo, R., Tang, X., Tang, L., Huang, H. X., Wen, X., et al. (2015). Age-related progression of arterial stiffness and its elevated positive association with blood pressure in healthy people. *Atherosclerosis* 238, 147–152. doi: 10.1016/j.atherosclerosis.2014.10.089
- Wicks, P., Massagli, M., Frost, J., Brownstein, C., Okun, S., Vaughan, T., et al. (2010). Sharing health data for better outcomes on patientslikeme. *J. Med. Internet Res.* 12:e19. doi: 10.2196/jmir.1549
- Wild, C. P. (2005). Complementing the genome with an “exposome”: the outstanding challenge of environmental exposure measurement in molecular epidemiology. *Cancer Epidemiol. Biomarkers Prev.* 14, 1847–1850. doi: 10.1158/1055-9965.EPI-05-0456
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., and Philip, S. Y. (2020). A comprehensive survey on graph neural networks. *IEEE Trans. Neural Netw. Learn. Syst.*
- Yi, X., Walia, E., and Babyn, P. (2019). Generative adversarial network in medical imaging: a review. *Med. Image Anal.* 58:101552. doi: 10.1016/j.media.2019.101552
- Yu, H., Miao, C., Leung, C., and White, T. J. (2017). Towards AI-powered personalization in MOOC learning. *NPJ Science of Learning* 2, 1–5. doi: 10.1038/s41539-017-0016-3
- Zaman, M. A., Oparil, S., and Calhoun, D. A. (2002). Drugs targeting the renin-angiotensin-aldosterone system. *Nat. Rev. Drug Discov.* 1, 621–636. doi: 10.1038/nrd873
- Zhou, J., Cui, G., Zhang, Z., Yang, C., Liu, Z., Wang, L., et al. (2018). Graph neural networks: a review of methods and applications. *arXiv [Preprint]. arXiv:1812.08434.*
- Zitnik, M., Agrawal, M., and Leskovec, J. (2018). Modeling polypharmacy side effects with graph convolutional networks. *Bioinformatics* 34, i457–i466. doi: 10.1093/bioinformatics/bty294

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.

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.

Copyright © 2021 Barbiero, Viñas Torné and Lió. 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.