GRAPH REPRESENTATION FORECASTING OF PATIENT'S MEDICAL CONDITIONS: TOWARDS A DIGITAL TWIN

A PREPRINT

Pietro Barbiero

Department of Computer Science and Technology
University of Cambridge
Cambridge, UK, CB3 0FD
pb737@cam.ac.uk

Ramon Viñas Torné

Department of Computer Science and Technology
University of Cambridge
Cambridge, UK, CB3 0FD
rv340@cam.ac.uk

Pietro Lió

Department of Computer Science and Technology
University of Cambridge
Cambridge, UK, CB3 0FD
pl219@cam.ac.uk

September 19, 2020

ABSTRACT

Objective: Modern medicine needs to shift from a wait and react, curative discipline to a preventative, interdisciplinary science aiming at providing personalised, systemic and precise treatment plans to patients. The aim of this work is to present how the integration of machine learning approaches with mechanistic computational modelling could yield a reliable infrastructure to run probabilistic simulations where the entire organism is considered as a whole.

Methods: We propose a general framework that composes advanced AI approaches and integrates mathematical modelling in order to provide a panoramic view over current and future physiological conditions. The proposed architecture is based on a graph neural network (GNNs) forecasting clinically relevant endpoints (such as blood pressure) and a generative adversarial network (GANs) providing a proof of concept of transcriptomic integrability.

Results: We show the results of the investigation of pathological effects of overexpression of ACE2 across different signalling pathways in multiple tissues on cardiovascular functions. We provide a proof of concept of integrating a large set of composable clinical models using molecular data to drive local and global clinical parameters and derive future trajectories representing the evolution of the physiological state of the patient.

Significance: We argue that the graph representation of a computational patient has potential to solve important technological challenges in integrating multiscale computational modelling with AI. We believe that this work represents a step forward towards a healthcare digital twin.

Keywords Precision medicine · Graph representation · Generative adversarial networks

1 Introduction

Large part of humanity cannot afford basic medical care. Where medicine has reached high level of biological and engineering technological complexity, costs have grown so much to reduce the access to a substantial percentage of the population. The target is very clear: we need healthcare to be precise and personalised at a reduced cost per capita in order to have 100% access [1]. A meaningful way to reach this target is to design a precise and predictive medicine. Precision will require the integration of large amount of observations at individual and other observations at population levels. These measures will be taken at different scales, from genome to clinical and family history and at systemic

levels, i.e. considering multiple tissues and organs. 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 predictive capability will allow us to shift from a wait and react, curative modus to a proactive modus of maintaining of the well-being condition or reverse disease status before more damage spreading (emergence of comorbidities and frailty).

Precision and predictable properties will therefore implicitly include personalised and preventative properties. The predictive capacity would act in lowering the overall and general cost of healthcare if the data integration step (leading to precision) would come at no extra cost. We believe that recent graph representation approaches [2] could scale across all the variety of body signals at different levels and could sum all the above mentioned properties making possible a revolution in healthcare.

The development of reliable systems accurately forecasting physiological conditions of patients is one of the primary objectives of precision medicine [3]. The underlying complexity of human physiology makes this research field extremely challenging [4, 5]. On one side, such complexity requires the development of models with very high capacity. On the other hand, the internal activity of several physiological systems are actually independent from others. For instance, the periodic contraction myocardiocytes does not depend on the pancreatic release of insulin. Nay, the mutual influence of several physiological processes is not hard-coded in their behaviour, but it is usually carried out by specialised transmitters, generating sophisticated signalling pathways. Besides, such signals can be transmitted from micro to macroscale systems and vice versa, affecting the whole organism at different levels [6, 7]. From a computational perspective, this means that each biological system could be modelled independently in the first place, provided that it receives expected signals from other systems [8].

This work proposes a new class of black-box machine-learning assisted methods for model analysis that scale to medical device deployment and run time monitoring and verification. By fusing ideas from systems medicine with scientific computing and machine learning we have developed a software that 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 a generative model [9] with a graph-based representation [10] of pathophysiological conditions (see Figure 1). On one side, the generative model can be used to produce synthetic data under different biological states that might not even 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 analyse otherwise [11]. In clinical contexts, this means that physicians will be able to set up personalised experiments in a virtual environment representing their patients in a very detailed and realistic way. On the other side, the graph model represents a virtual prototype of patients, a sort of "digital twin" mirroring the actual multiscale biological system [12], thus providing a general and flexible framework to run probabilistic simulations. The intrinsic characteristics of graph models make them suitable for the analysis of complex systems, while still providing highly interpretable results [13]. Graph are not just interpretable, but the network itself can be induced from data or even handcrafted. Researchers may take advantage of generative models for graphs to find optimal network configurations [14] or formalise mathematical properties in form of differential equations or logical constraints with the purpose of describing the underlying system [15]. A panoramic view of individuals' conditions is provided by the final network configuration which 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. By relying upon flexible and modular architectures, our "digital twin" model can be conveniently deployed in dedicated hardware modules which might also be composed in a hierarchical fashion according to clinical needs.

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" [16, 17]. The Digital twin is a virtual prototype; the analysis of its digital life cycle provides information to understand a product's functionality, manufacturing, behaviour 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 best practice to be used on the biological twin. Here an 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.

This work proposes a modular approach which can be used to model the evolution of pathophysiological conditions. The first module is based on a graph neural network (GNNs) forecasting clinically relevant endpoints (such as blood pressure) while the second one is represented by a generative adversarial network (GANs) providing a proof of concept of omic integrability. For experimental simulations we use and expand a clinical case study presented in [18] as a reference. In [18], authors introduced a handcrafted dynamical system of equations representing some of the main physiological processes involved in cardiovascular and infection diseases. The computational model was developed to simulate systemic ripple effects caused by viral infections targeting the renin-angiotensin-aldosterone system (RAS) in patients with comorbid conditions. In such a context, the most relevant endpoints requiring constant monitoring include blood pressure, oxygenation, and insulin levels (for diabetic patients). In patients with multi-factorial diseases a very large set of factors may cause an irreparable impairment of such a delicate physiological equilibrium. Genomic traits, ageing, therapies, and lifestyle routines like physical exercise or dietary habits may all have an impact on keeping patients in a healthy state. Recording and keeping track of the stratification of all such factors is problematic by itself, but it is critical to get the whole picture. However, the actual challenge consists in exploiting the acquired information to forecast the evolution of clinical endpoints over time. Depending on clinical needs and patient's comorbidities, it may be worth accounting for just a subset of such endpoints. Thus a sparse modelling approach whose components can be combined or activated on the fly may have a great potential in clinical practice.

2.1 The effectiveness of GNNs and GANs in biomedical signal analysis

Several properties of graph and generative adversarial neural networks make them suitable for medical data analysis:

- **Non-linearity:** detection of non linear patterns is of key interest as most systems are inherently nonlinear 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 ageing and disease conditions [19].
- **Interpretability:** the possibility of interpreting the behaviour of models and the reason for their predictions is pivotal if not critical for in many fields including clinical practice. Thanks to their structure, graph-based models are much easier to interpret with respect to other neural approaches.
- **Non-Euclidean geometry:** tissue and organ distributions could be modelled 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 modelled as a network where lymph nodes are vertices. At lower scale, cell arrangements in tissues form particular manifolds; proteins and genes are organised 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.
- **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.
- **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 fMRI data.
- **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.
- **Multiscale:** the graph representation has the capability of integrating granular information organised as networks at different layers of biological complexity. This allows to recognise 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), stratification (as composition of individuals).
- **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 which are used to model a set of objects (nodes) and their mutual relationships (edges) [20]. Graphs are employed in a variety of research areas as they provide a general and flexible data structure for

modelling real-world systems [21, 22, 23, 24]. Graph neural networks (GNN) are deep learning-based models working on the graph domain [10, 25, 26]. Their properties have been recently drawn the attention of the artificial intelligence research community given their high interpretability and as the only non-Euclidean models available in machine learning [27, 13]. The combination of graph theory and neural network elements have made GNNs one of the most promising tools to analyse 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 [14]. Hybrid techniques may also be explored taking advantage of generative algorithms for handling system complexity and human modelling to customise 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 Graph network blocks

The GNN framework proposed by Battaglia et al. [25] is based on modules called graph network blocks (GN blocks) representing the core computation units of a GNN. Multiple GN blocks can be composed or even combined with other neural networks to generate complex architectures. A graph neural network 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 $\mathbf{h}_i$. $E = \{(\mathbf{e}_k, r_k, s_k)\}$ is the edge set where each node is represented by its features $\mathbf{e}_k$, 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 :

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

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 minimising the following objective function:

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

where θ is the set of model's parameters.

2.2.2 Graph layers

GNNs natively allow the design of complex systems using a modular approach. First, the complexity 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 set 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 layers: the transcriptomic layer,

Figure 1: Architecture of the digital twin model. The generator receives a noise vector $\mathbf{z}$, and categorical (e.g. tissue type; $\mathbf{q}$) and numerical (e.g. age; $\mathbf{r}$) covariates, and outputs a vector of synthetic data ($\hat{\mathbf{x}}$). The critic receives data from two input streams (real, blue; and synthetic, red), a mask $\mathbf{m}$ indicating which components of the input vector are missing, and the numerical $\mathbf{r}$ and categorical $\mathbf{q}$ 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 [18] 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.

the cellular layer, the organ layer, and the exposomic layer. Other layers that bring information on other omics or body sensor network i.e. a collection of networked sensors that can be used to monitor human physiological signals, could be similarly implemented.

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 [28, 29, 30], propose novel drug targets [31, 32], provide a basis for comparative genomics [33], 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 [34]. Specifically, we analyse to what extent the expression of genes involved in the renin-angiotensin system can be explained by genes from signalling and receptor pathways, including the chemokine, TNF, and TGF- β pathways. We further develop a transcriptomics generative model based on a generative adversarial network [35] and simulate the effects of SARS-CoV-2 infection by conditioning on high expression of ACE2 in the lung, kidney, and pancreas.

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 modelling the Renin-Angiotensin System (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 renin-angiotensin system is a hormone system regulating vasoconstriction and inflammatory response [36]. 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) [37, 38]. 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) [39]. 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 [40].

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 [18]: heart, lungs, and kidneys. The heart model includes four compartments known as chambers [41]. 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 towards 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 utilised primarily in autonomic reflexes, regulating short-term blood pressure.

Exposomic layer The exposome refers to the totality of exposure individuals experience from conception until death and its impact on chronic and acute diseases [42]. Toxicants, dietary regimens, treatments, physical exercise, posture, lifestyle habits, all of them 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.

2.2.3 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 Eq. 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 in 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.

2.2.4 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 [43] at test time, as a Bayesian approximation, without sacrificing either computational complexity or test performance [44]. 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 [45]:

$$\mathbb{E}_{q(\mathbf{y}^*, \mathbf{x}^*)}(\mathbf{y}^*) \approx \frac{1}{T} \sum_{i=1}^T \hat{\mathbf{y}}^*(\mathbf{x}^*, W_1^t, \dots, W_L^t) \quad (3)$$

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

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 generalised to convolutional [46] and recurrent networks [47].

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 which 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 et al. [44]. 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.

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 [48]. Generative adversarial networks (GANs) are a framework for estimating generative models via an adversarial process [35]. 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 2 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 customisation, making them suitable for a variety of experimental scenarios.

Figure 2: Generative Adversarial Network framework. The generator $G(\mathbf{z})$ receives a vector $\mathbf{z}$ sampled from a noise prior distribution $p_{\mathbf{z}}$, and generates a synthetic sample $\mathbf{x}_{fake}$. The discriminator $D(\mathbf{x})$ tries to distinguish *real* samples from *fake* samples, producing the probability of $\mathbf{x}$ coming from the *real* data distribution. The competition between the two players drives the game and makes both players increasingly better.

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 signalling 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 characterised and observed in a variety of biological processes from micro to macroscale from genomics [49, 50], to internal and external cell activity [51, 52], and even between tissues [53]. 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.

Model. 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 vector of a patient's gene expression values in t tissues; n is the number of genes; $\mathbf{m} \in \{0, 1\}^t$ 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 modelling 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 modelling this distribution, we can sample data for different conditions and quantify the uncertainty of the generated expression values.

Our method builds on a Wasserstein GAN with gradient penalty (WGAN-GP) [54, 55]. Similar to Generative Adversarial Networks (GAN) [35], 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 $G_{\theta} : \mathbb{R}^u \times \mathbb{R}^k \times \mathbb{N}^c \rightarrow \mathbb{R}^{t \times n}$ parametrised 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 $D_\omega : \mathbb{R}^{t \times n} \times \{0, 1\}^t \times \mathbb{R}^k \times \mathbb{N}^c \rightarrow \mathbb{R}$ parametrised 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 is unavailable for a given patient, we impute the unobserved values of $\bar{\mathbf{x}}$ with zeroes and set the corresponding value of the mask $\mathbf{m}$ to zero.

We optimise the generator and the critic adversarially. Following [54], we train the generator G_θ and the critic D_ω to solve the following minimax game based on the Wasserstein distance:

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

where $\hat{\mathbf{x}}$ is defined as in Equation 7 and the constraint enforces the critic D_ω to be 1-Lipschitz, that is, 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, optimising the following problems:

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

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 penalises the gradient norm along points in the straight lines that connect real and synthetic samples [55].

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 [56] 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 characterises 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 $\parallel$ 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 and discriminator D 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.

3 Simulations

3.1 Clinical case studies

The clinical case studies used for the simulations are derived from [18]. 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 personalise the therapy according 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 a 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) [57, 58]. 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 [59, 60, 61, 62, 63] and all-cause mortality [64, 65]. 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 which need to be constantly monitored in real time and accurately forecast.

3.2 Forecasting clinical scenarios and interpreting GNN simulations

The proposed GNN-based model presented in Section 2 is hereby used to actively monitor and forecast the endpoints highlighted in the two clinical case studies. First, the computational ODE-based system described in [18] is used to generate a time series for each differential equation with a window size of $\tau = 500$ time steps [66]. 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 ($n_{test} = 1000$).

The graph model is derived from the structure of ODE system, thus leveraging human knowledge. Nodes correspond to variables represented by the differential equations in [18] 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 neighbourhood. 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 elements of the test set using the procedure described in Section 2.2.4. As a result, the model estimates a 95% confidence interval of the evolution of each variable over time.

Providing a complete overview of the clinical state of a patient is not trivial. Arguably, 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 in applying a dimensionality reduction technique [67] condensing the information of each organ and projecting forecasts in a lower-dimensional space. Figure 3 shows an overview of the clinical state of the heart in a two-dimensional 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, [68]). 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

Figure 3: Two clinical case studies represented in a projected heart phase-space. The first case study (left figure) 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 (5mg). The second simulation (right figure) shows the long-term impact on blood pressure of an untreated SARS-CoV-2 infection (red density) and the effects of a therapy including both Benazepril (5 mg/day) and intra venous injection of heparin (5000 U/ml) (orange density). **[top figures]** Bundle of predicted trajectories can be visualised and monitored in real time in order to investigate patterns in the time domain. The simulation shows blood pressure in heart chambers starting from healthy state conditions. Error bands represent 95% CI. **[bottom figures]**

a daily dosage of Benazepril (5mg). 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, instead, 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 intra venous injection of heparin (5000 U/ml) may have a beneficial effect on blood pressure mean and variability (orange density). On one side, 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 [57, 58]. 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 visualised and monitored individually in real time when needed in order to investigate patterns in the time domain. Figure 3 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 visualisation tool for monitoring clinical endpoints in real-time.

3.3 Transcriptomics analysis of the crosstalk between tissue-types

We hypothesise 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 renin-angiotensin system 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.

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 [69]. 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 normalise the read counts between samples using the trimmed mean of M-values (TMM) normalisation method [70] 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 [71] to select genes from the renin-angiotensin system (*hsa04614*), chemokine (*hsa04062*), TNF (*hsa04668*), and TGF- β (*hsa04350*) pathways.

Model. 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 signalling pathways, respectively. Our model is based on ridge regression [72]:

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

where $\mathbf{W} \in \mathbb{R}^{m \times n}$ is a matrix of learnable weights and $\epsilon \in \mathbb{R}^n$ are the residuals. We optimise 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 regularisation strength. We also tried non-linear models such as support vector machines, Gaussian processes, and random forests, but they were not significantly better than ridge regression according to our cross-validation scores.

Results. Figure 4 show the bootstrapped R^2 scores for each gene in the renin-angiotensin system pathway in different tissue types. These results show that the expression of some genes in the ACE2 pathway can be partially explained by signalling genes from whole blood. Notably, the associations for the kidney (cortex) are weaker or nonexistent, potentially because the data is limited for this tissue (62 samples) or because the biological associations are indeed small. Overall, these results suggest that signalling pathways such as TNF, TGF- β , and chemokine might be playing an important role in the development of the SARS-CoV-2 infection.

Figure 4: Bootstrapped R^2 scores for genes involved in the renin-angiotensin system for lung, heart (left ventricle), kidney (cortex), and pancreas. The input variables are the expressions of genes in whole blood belonging to the chemokine, TNF, and TFG- β pathways.

Figure 5: Pairwise correlations between genes in the renin-angiotensin system pathway in lung for real (left) and synthetic (right) data. The correlations in the lower and upper matrices are computed from samples with low (60 samples) and high (61 samples) ACE2 expression, respectively.

3.4 Generative model for transcriptomics data

The generative model developed in Section 2.3 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 communicome and signalling pathways such as TNF, TGF- β and chemokines. These pathways are implicated in many physiological and pathological processes including the regulation of blood pressure and inflammatory processes, and has been hypothesised to play a central role in SARS-CoV-2 infection [73, 74].

Dataset. We use data from the GTEx project and process it as described in Section 3.3.

Results. Figure 5 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 renin-angiotensin system pathway (CTSA, AGTR2, NLN, and PREP) that can be relatively well explained as a function of blood signalling factors (see Figure 4) are simultaneously correlated with ACE2. This suggests that these genes might be playing an important role in the spread of SARS-CoV-2 in our body through blood. Figure 6 shows that it is possible to sample data for synthetic patients conditioned on different levels of ACE2 expression in lung.

3.5 Software

All the code for the experiments has been implemented in Python 3, relying upon open-source libraries [75, 76, 77]. All the experiments have been run on the same machine: Intel® Core™ i7-8750H 6-Core Processor at 2.20 GHz equipped with 8 GiB 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 [78]. 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.

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

Figure 6: Principal component analysis of the multi-tissue expression of 100 synthetic patients for different levels of ACE2 expression. Each line corresponds to a unique patient. For each patient, we fix all the latent covariates and modify the levels of ACE2 in lung. Overexpressing ACE2 leads to changes in the expression of other genes and these changes follow a well-defined trajectory.

4 Discussion and conclusion

4.1 Interplay between GAN and GNN models

The models presented in this work (GAN, GNN) are independent of each other. On one hand, the main goal of the GNN model (see Section 2.2) 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 (see Section 2.3) 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 personalised according to clinical health records.

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 heterogenous, paired data modalities (from different layers: genomic, transcriptomic, cellular, organ, exposomic, ...) 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 [79, 80]. To address this issue, our GAN framework can synthesise data at multiple layers conditioned on the patient's conditions (e.g. diabetic, hypertension, ...) and clinical information (e.g. heart rate, blood pressure, age, sex, ethnicity, exercise, nutrition, ...). This synthetic data can be used to train the GNN and impute missing data modalities of real patients.

4.2 Advantages, limitations and visions

The promise of artificial intelligence in medicine is to provide composite, panoramic views of individuals' medical data; to improve decision making; to avoid errors such as misdiagnosis and unnecessary procedures; to help in the ordering and interpretation of appropriate tests; and to recommend treatment [81]. The future of medicine is already bound to AI. Technological innovations are completely changing medicine perspectives expanding its horizons and moving towards 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, the first step towards 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 personalised pathways to well-being rather than delivering treatments. This does not mean that medicine will no longer be connected to illness. On the contrary, 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.

Ethical repercussions will also be huge [79, 80]. The transition has begun, but it will call for deeper trans-disciplinary research and a substantial technological innovation in a variety of research areas. Education will also play a key role in changing lifestyle habits and the way health is perceived, communicated, and delivered [82]. From a clinical standpoint, AI will support a plethora of different tasks from medical check up to personalised 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 [83].

Increasingly large amount of personal data will be collected to feed modular machine learning (ML) models organised to address specific and personalised 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 modelling in medicine will produce better investigation of complex biological processes, and even new ideas and better feedback into medicine. Modelling-based approaches combined with data-driven ML techniques will progressively provide models with higher degree of interpretability and generalisation ability [84] which will make evidence-based medicine even more accessible intensifying the involvement of patients in the decision making process. Besides, 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 [85].

AI simulations forecasting the evolution of clinical endpoints over time will also reshape clinical guidelines [83] 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 personalised prescriptions, surgical planning and medical interventions. Finally, AI will change the leading *vehicle* of medicine. The demand for AI-powered and IoT devices is increasing worldwide. The future medical equipment 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.

The most important question around the integration of AI in medicine is about the benefit for the patient. 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 [86]. 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 which 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* which allow the user to search and make up his/her mind about a disease [87]. 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. The second question is about impact on how clinicians work or are trained. 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 limitations 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 in parallel mechanistic computational modelling and simulations ([88, 89]), model extraction libraries (see for instance [90]), and visual inference tools [91]. This tool could also be complemented by clinical decision support systems such as [92].

The complexly structured and multi level comorbidity and frailty patterns of most diseases describe a highly dynamical system and are, therefore, challenging current medical therapies.

Mechanistic computational modelling 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 modelling and machine learning researchers, data policy makers, 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 specialised on processing molecular information or they could operate cross modal operations such as omic-omic, omic-clinical, clinical-clinical.

In this work we demonstrate the feasibility of representing and integrating physiological models and molecular information using graph neural networks 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 which translate into patients trajectories. We have assimilated our product to a digital twin of the patient.

5 Acknowledgement

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 848077. This project has received funding from "la Caixa" Foundation (ID 100010434), under agreement LCF/BQ/EU19/11710059.

References

- [1] Nenad Tomašev, Julien Cornebise, Frank Hutter, Shakir Mohamed, Angela Picciariello, Bec Connelly, Danielle CM Belgrave, Daphne Ezer, Fanny Cachat van der Haert, Frank Mugisha, et al. Ai for social good: unlocking the opportunity for positive impact. *Nature Communications*, 11(1):1–6, 2020.
- [2] Michael M. Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: Going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, July 2017.
- [3] Denis Noble. The rise of computational biology. *Nature Reviews Molecular Cell Biology*, 3(6):459–463, 2002.
- [4] Geoffrey S Ginsburg and Huntington F Willard. Genomic and personalized medicine: foundations and applications. *Translational research*, 154(6):277–287, 2009.
- [5] Stephen Naylor and Jake Y Chen. Unraveling human complexity and disease with systems biology and personalized medicine. *Personalized medicine*, 7(3):275–289, 2010.
- [6] Olaf Wolkenhauer, Charles Auffray, Olivier Brass, Jean Clairambault, Andreas Deutsch, Dirk Drasdo, Francesco Gervasio, Luigi Preziosi, Philip Maini, Anna Marciniak-Czochra, et al. Enabling multiscale modeling in systems medicine. *Genome medicine*, 6(3):21, 2014.
- [7] Kipp W Johnson, Khader Shameer, Benjamin S Glicksberg, Ben Readhead, Partho P Sengupta, Johan LM Björkegren, Jason C Kovacic, and Joel T Dudley. Enabling precision cardiology through multiscale biology and systems medicine. *JACC: Basic to Translational Science*, 2(3):311–327, 2017.
- [8] Mark AJ Chaplain. Multiscale mathematical modelling in biology and medicine. *IMA Journal of Applied Mathematics*, 76(3):371–388, 2011.
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [10] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE Transactions on Neural Networks*, 20(1):61–80, 2008.
- [11] Xin Yi, Ekta Walia, and Paul Babyn. Generative adversarial network in medical imaging: A review. *Medical image analysis*, 58:101552, 2019.

- [12] David Gelernter. *Mirror worlds: or: the day software puts the universe in a shoebox... How it will happen and what it will mean*. Oxford University Press, 1993.
- [13] Qiang Huang, Makoto Yamada, Yuan Tian, Dinesh Singh, Dawei Yin, and Yi Chang. Graphlime: Local interpretable model explanations for graph neural networks. *arXiv preprint arXiv:2001.06216*, 2020.
- [14] Yujia Li, Oriol Vinyals, Chris Dyer, Razvan Pascanu, and Peter Battaglia. Learning deep generative models of graphs. *arXiv preprint arXiv:1803.03324*, 2018.
- [15] Giuseppe Marra, Francesco Giannini, Michelangelo Diligenti, and Marco Gori. Lyrics: A general interface layer to integrate logic inference and deep learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 283–298. Springer, 2019.
- [16] Mike Shafto, M Conroy, R Doyle, E Glaessgen, C Kemp, J LeMoigne, and L Wang. Modeling, simulation, information technology and processing roadmap. *National Aeronautics and Space Administration*, 05 2010.
- [17] Michael Grieves. Digital twin: Manufacturing excellence through virtual factory replication. *Whitepaper*, 03 2015.
- [18] Pietro Barbiero and Pietro Lió. The computational patient has diabetes and a covid. *arXiv preprint arXiv:2006.06435*, 2020.
- [19] Ary L Goldberger, Luis AN Amaral, Jeffrey M Hausdorff, Plamen Ch Ivanov, C-K Peng, and H Eugene Stanley. Fractal dynamics in physiology: alterations with disease and aging. *Proceedings of the national academy of sciences*, 99(suppl 1):2466–2472, 2002.
- [20] Béla Bollobás. *Modern graph theory*, volume 184. Springer Science & Business Media, 2013.
- [21] Jie Zhou, Ganqu Cui, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *arXiv preprint arXiv:1812.08434*, 2018.
- [22] Erez Lieberman, Christoph Hauert, and Martin A Nowak. Evolutionary dynamics on graphs. *Nature*, 433(7023):312–316, 2005.
- [23] Goran Rakocovic, Vladimir Semenyuk, Wan-Ping Lee, James Spencer, John Browning, Ivan J Johnson, Vladan Arsenijevic, Jelena Nadj, Kaushik Ghose, Maria C Suci, et al. Fast and accurate genomic analyses using genome graphs. *Nature genetics*, 51(2):354–362, 2019.
- [24] Ioana Bica, Helena Andrés-Terré, Ana Cvejic, and Pietro Lió. Unsupervised generative and graph representation learning for modelling cell differentiation. *Scientific reports*, 10(1):1–13, 2020.
- [25] Peter W Battaglia, Jessica B Hamrick, Victor Bapst, Alvaro Sanchez-Gonzalez, Vinicius Zambaldi, Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, et al. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*, 2018.
- [26] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems*, 2020.
- [27] Freddy Lecue. On the role of knowledge graphs in explainable ai. *Semantic Web*, (Preprint):1–11, 2019.
- [28] William Cookson, Liming Liang, Gonçalo Abecasis, Miriam Moffatt, and Mark Lathrop. Mapping complex disease traits with global gene expression. *Nature Reviews Genetics*, 10(3):184–194, 2009.
- [29] Valur Emilsson, Gudmar Thorleifsson, Bin Zhang, Amy S Leonardson, Florian Zink, Jun Zhu, Sonia Carlson, Agnar Helgason, G Bragi Walters, Steinunn Gunnarsdottir, et al. Genetics of gene expression and its effect on disease. *Nature*, 452(7186):423–428, 2008.
- [30] Eric R Gamazon, Ayellet V Segrè, Martijn van de Bunt, Xiaoquan Wen, Hualin S Xi, Farhad Hormozdiari, Halit Ongen, Anuar Konkashbaev, Eske M Derks, François Aguet, et al. Using an atlas of gene regulation across 44 human tissues to inform complex disease-and trait-associated variation. *Nature genetics*, 50(7):956–967, 2018.
- [31] Marina Sirota, Joel T Dudley, Jeewon Kim, Annie P Chiang, Alex A Morgan, Alejandro Sweet-Cordero, Julien Sage, and Atul J Butte. Discovery and preclinical validation of drug indications using compendia of public gene expression data. *Science translational medicine*, 3(96):96ra77–96ra77, 2011.
- [32] William E Evans and Mary V Relling. Moving towards individualized medicine with pharmacogenomics. *Nature*, 429(6990):464–468, 2004.
- [33] Laura L Colbran, Eric R Gamazon, Dan Zhou, Patrick Evans, Nancy J Cox, and John A Capra. Inferred divergent gene regulation in archaic hominins reveals potential phenotypic differences. *Nature ecology & evolution*, 3(11):1598–1606, 2019.

- [34] Sandip Ray, Markus Britschgi, Charles Herbert, Yoshiko Takeda-Uchimura, Adam Boxer, Kaj Blennow, Leah F Friedman, Douglas R Galasko, Marek Jutel, Anna Karydas, Jeffrey A Kaye, Jerzy Leszek, Bruce L Miller, Lennart Minthon, Joseph F Quinn, Gil D Rabinovici, William H Robinson, Marwan N Sabbagh, Yuen T So, D Larry Sparks, Massimo Tabaton, Jared Tinklenberg, Jerome A Yesavage, Robert Tibshirani, and Tony Wyss-Coray. Classification and prediction of clinical alzheimer’s diagnosis based on plasma signaling proteins. *Nature Medicine*, 13(11):1359–1362, October 2007.
- [35] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems 27*, pages 2672–2680. Curran Associates, Inc., 2014.
- [36] John H Fountain and Sarah L Lappin. Physiology, renin angiotensin system. In *StatPearls [Internet]*. StatPearls Publishing, 2019.
- [37] Keiji Kuba, Yumiko Imai, Takayo Ohto-Nakanishi, and Josef M Penninger. Trilogy of ace2: A peptidase in the renin–angiotensin system, a sars receptor, and a partner for amino acid transporters. *Pharmacology & therapeutics*, 128(1):119–128, 2010.
- [38] Mariela M Gironacci, Hugo P Adamo, Gerardo Corradi, Robson A Santos, Pablo Ortiz, and Oscar A Carretero. Angiotensin (1-7) induces mas receptor internalization. *Hypertension*, 58(2):176–181, 2011.
- [39] Mohammad Amin Zaman, Suzanne Oparil, and David A Calhoun. Drugs targeting the renin–angiotensin–aldosterone system. *Nature reviews Drug discovery*, 1(8):621–636, 2002.
- [40] Andrew M South, Laurie Tomlinson, Daniel Edmonston, Swapnil Hiremath, and Matthew A Sparks. Controversies of renin–angiotensin system inhibition during the covid-19 pandemic. *Nature Reviews Nephrology*, pages 1–3, 2020.
- [41] Maxwell Lewis Neal and James B Bassingthwaite. Subject-specific model estimation of cardiac output and blood volume during hemorrhage. *Cardiovascular engineering*, 7(3):97–120, 2007.
- [42] Christopher Paul Wild. Complementing the genome with an “exposome”: the outstanding challenge of environmental exposure measurement in molecular epidemiology, 2005.
- [43] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1):1929–1958, 2014.
- [44] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- [45] Y Gal and Z Ghahramani. Dropout as a bayesian approximation: Appendix 20 (2016). URL <http://arxiv.org/abs/1506.02157>, 1506.
- [46] Yarin Gal and Zoubin Ghahramani. Bayesian convolutional neural networks with bernoulli approximate variational inference. *arXiv preprint arXiv:1506.02158*, 2015.
- [47] Yarin Gal and Zoubin Ghahramani. A theoretically grounded application of dropout in recurrent neural networks. In *Advances in neural information processing systems*, pages 1019–1027, 2016.
- [48] Tony Jebara. *Machine learning: discriminative and generative*, volume 755. Springer Science & Business Media, 2012.
- [49] Jiamu Du, Lianna M Johnson, Steven E Jacobsen, and Dinshaw J Patel. Dna methylation pathways and their crosstalk with histone methylation. *Nature reviews Molecular cell biology*, 16(9):519–532, 2015.
- [50] Robert O Poyton and Joan E McEwen. Crosstalk between nuclear and mitochondrial genomes. *Annual review of biochemistry*, 65(1):563–607, 1996.
- [51] Meng Li, Ping Gao, and Junping Zhang. Crosstalk between autophagy and apoptosis: potential and emerging therapeutic targets for cardiac diseases. *International journal of molecular sciences*, 17(3):332, 2016.
- [52] Benjamin Geiger, Alexander Bershadsky, Roumen Pankov, and Kenneth M Yamada. Transmembrane crosstalk between the extracellular matrix and the cytoskeleton. *Nature reviews Molecular cell biology*, 2(11):793–805, 2001.
- [53] Ernst Lengyel, Liza Makowski, John DiGiovanni, and Mikhail G Kolonin. Cancer as a matter of fat: the crosstalk between adipose tissue and tumors. *Trends in cancer*, 4(5):374–384, 2018.
- [54] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein GAN. *arXiv e-prints*, page arXiv:1701.07875, January 2017.

- [55] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans, 2017.
- [56] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg Corrado, and Jeffrey Dean. Distributed representations of words and phrases and their compositionality, 2013.
- [57] Celestino Sardu, Jessica Gambardella, Marco Bruno Morelli, Xujun Wang, Raffaele Marfella, and Gaetano Santulli. Hypertension, thrombosis, kidney failure, and diabetes: Is COVID-19 an endothelial disease? a comprehensive evaluation of clinical and basic evidence. *Journal of Clinical Medicine*, 9(5):1417, May 2020.
- [58] Ning Tang, Huan Bai, Xing Chen, Jiale Gong, Dengju Li, and Ziyong Sun. Anticoagulant treatment is associated with decreased mortality in severe coronavirus disease 2019 patients with coagulopathy. *Journal of Thrombosis and Haemostasis*, 18(5):1094–1099, April 2020.
- [59] Wen Wen, Rong Luo, Xiaojing Tang, Lan Tang, Hunter X. Huang, Xiaoyan Wen, Shan Hu, and Bin Peng. Age-related progression of arterial stiffness and its elevated positive association with blood pressure in healthy people. *Atherosclerosis*, 238(1):147–152, January 2015.
- [60] Michael F. O’Rourke and Wilmer W. Nichols. Aortic diameter, aortic stiffness, and wave reflection increase with age and isolated systolic hypertension. *Hypertension*, 45(4):652–658, April 2005.
- [61] Gary F. Mitchell, Shih-Jen Hwang, Ramachandran S. Vasan, Martin G. Larson, Michael J. Pencina, Naomi M. Hamburg, Joseph A. Vita, Daniel Levy, and Emelia J. Benjamin. Arterial stiffness and cardiovascular events. *Circulation*, 121(4):505–511, February 2010.
- [62] Sripal Bangalore, David J. Maron, Sean M. O’Brien, Jerome L. Fleg, Evgeny I. Kretov, Carlo Briguori, Upendra Kaul, Harmony R. Reynolds, Tomasz Mazurek, Mandeep S. Sidhu, Jeffrey S. Berger, Roy O. Mathew, Olga Bockeria, Samuel Broderick, Radoslaw Pracon, Charles A. Herzog, Zhen Huang, Gregg W. Stone, William E. Boden, Jonathan D. Newman, Ziad A. Ali, Daniel B. Mark, John A. Spertus, Karen P. Alexander, Bernard R. Chaitman, Glenn M. Chertow, and Judith S. Hochman. Management of coronary disease in patients with advanced kidney disease. *New England Journal of Medicine*, 382(17):1608–1618, April 2020.
- [63] Donald Clark, Stephen J. Nicholls, Julie St John, Mohamed B. Elshazly, Haitham M. Ahmed, Haitham Khraishah, Steven E. Nissen, and Rishi Puri. Visit-to-visit blood pressure variability, coronary atheroma progression, and clinical outcomes. *JAMA Cardiology*, 4(5):437, May 2019.
- [64] Bum Joon Kim, Sun U. Kwon, Dalia Wajsbrot, Jaseong Koo, Jong Moo Park, and Barrett W. Jeffers. Relationship of inter-individual blood pressure variability and the risk for recurrent stroke. *Journal of the American Heart Association*, 7(24), December 2018.
- [65] Y Tao, J Xu, B Song, X Xie, H Gu, Q Liu, L Zhao, Y Wang, Y Xu, and Y Wang. 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(9):568–573, April 2017.
- [66] Pietro Barbiero and Pietro Liò. pietrobarbiero/computational-patient: Absolutno, Sep 2020.
- [67] Laurens Van Der Maaten, Eric Postma, and Jaap Van den Herik. Dimensionality reduction: a comparative. *J Mach Learn Res*, 10(66-71):13, 2009.
- [68] K Pearson. On lines of closes fit to system of points in space, london, e dinb. *Dublin Philos. Mag. J. Sci*, 2:559–572, 1901.
- [69] François Aguet, Alvaro N Barbeira, Rodrigo Bonazzola, Andrew Brown, Stephane E Castel, Brian Jo, Silva Kasela, Sarah Kim-Hellmuth, Yanyu Liang, Meritxell Oliva, et al. The gtex consortium atlas of genetic regulatory effects across human tissues. *BioRxiv*, page 787903, 2019.
- [70] Mark D Robinson and Alicia Oshlack. A scaling normalization method for differential expression analysis of rna-seq data. *Genome biology*, 11(3):1–9, 2010.
- [71] Minoru Kanehisa, Susumu Goto, Miho Furumichi, Mao Tanabe, and Mika Hirakawa. Kegg for representation and analysis of molecular networks involving diseases and drugs. *Nucleic acids research*, 38(suppl_1):D355–D360, 2010.
- [72] Arthur E Hoerl and Robert W Kennard. Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1):55–67, 1970.
- [73] Hui Xiao, Krzysztof Bartoszek, and Pietro Lio’. Multi-omic analysis of signalling factors in inflammatory comorbidities. *BMC Bioinformatics*, 19(S15), November 2018.
- [74] Michael R Garvin, Christiane Alvarez, J Izaak Miller, Erica T Prates, Angelica M Walker, B Kirtley Amos, Alan E Mast, Amy Justice, Bruce Aronow, and Daniel Jacobson. A mechanistic model and therapeutic interventions for covid-19 involving a ras-mediated bradykinin storm. *Elife*, 9:e59177, 2020.

- [75] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*, pages 265–283, 2016.
- [76] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830, 2011.
- [77] Minjie Wang, Lingfan Yu, Da Zheng, Quan Gan, Yu Gai, Zihao Ye, Mufei Li, Jinjing Zhou, Qi Huang, Chao Ma, et al. Deep graph library: Towards efficient and scalable deep learning on graphs. *arXiv preprint arXiv:1909.01315*, 2019.
- [78] Pietro Barbiero, Ramon V. Torné, and Pietro Liò. *pietrobarbiero/digital-patient: Absolutno*, Sep 2020.
- [79] Anna Jobin, Marcello Ienca, and Effy Vayena. The global landscape of ai ethics guidelines. *Nature Machine Intelligence*, 1(9):389–399, 2019.
- [80] Brent Mittelstadt. Principles alone cannot guarantee ethical ai. *Nature Machine Intelligence*, pages 1–7, 2019.
- [81] Eric Topol. *Deep medicine: how artificial intelligence can make healthcare human again*. Hachette UK, 2019.
- [82] Han Yu, Chunyan Miao, Cyril Leung, and Timothy John White. Towards ai-powered personalization in mooc learning. *npj Science of Learning*, 2(1):1–5, 2017.
- [83] S Rivera. Guidelines for clinical trial protocols for interventions involving artificial intelligence: the spirit-ai extension. *BMJ*. <https://doi.org/10.1038/s41591-020-1037-7>, 2020.
- [84] Pietro Barbiero, Giovanni Squillero, and Alberto Tonda. Modeling generalization in machine learning: A methodological and computational study. *arXiv preprint arXiv:2006.15680*, 2020.
- [85] Trishan Panch, Heather Mattie, and Leo Anthony Celi. The “inconvenient truth” about ai in healthcare. *Npj Digital Medicine*, 2(1):1–3, 2019.
- [86] Jacques Ellul, Jacques Ellul, Jacques Ellul, Politologue Juriste, and Jacques Ellul. *La technique ou l’enjeu du siècle*. A. Colin Paris, 1954.
- [87] Paul Wicks, Michael Massagli, Jeana Frost, Catherine Brownstein, Sally Okun, Timothy Vaughan, Richard Bradley, and James Heywood. Sharing health data for better outcomes on patientslikeme. *Journal of medical Internet research*, 12(2):e19, 2010.
- [88] Ezio Bartocci and Pietro Liò. Computational modeling, formal analysis, and tools for systems biology. *PLOS Computational Biology*, 12(1):e1004591, January 2016.
- [89] Luciano Milanese, Paolo Romano, Gastone Castellani, Daniel Remondini, and Pietro Liò. Trends in modeling biomedical complex systems. *BMC Bioinformatics*, 10(S12), October 2009.
- [90] Dmitry Kazhdan, Zohreh Shams, and Pietro Liò. Marleme: A multi-agent reinforcement learning model extraction library. *arXiv preprint arXiv:2004.07928*, 2020.
- [91] Cristian Bodnar, Cătălina Cangea, and Pietro Liò. Deep graph mapper: Seeing graphs through the neural lens, 2020.
- [92] Tamara T. Müller and Pietro Lio. PECLIDES neuro: A personalisable clinical decision support system for neurological diseases. *Frontiers in Artificial Intelligence*, 3, April 2020.