



HAL
open science

Machine learning and biophysical models: how to benefit each other?

Nicolas Duchateau, Oscar Camara

► **To cite this version:**

Nicolas Duchateau, Oscar Camara. Machine learning and biophysical models: how to benefit each other?. *Reduced Order Models for the Biomechanics of Living Organs*, Elsevier, pp.147-164, 2023, 10.1016/B978-0-32-389967-3.00009-3 . hal-04212035

HAL Id: hal-04212035

<https://hal.science/hal-04212035v1>

Submitted on 25 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Machine learning and biophysical models: how to benefit each other?

Nicolas Duchateau^{1,2}, Oscar Camara³

¹ Univ Lyon, Université Claude Bernard Lyon 1, INSA-Lyon, CNRS, Inserm, CREATIS UMR 5220, U1294, F-69621, LYON, France

² Institut Universitaire de France (IUF)

³ BCN MedTech, Department of Information and Communication Technologies, Universitat Pompeu Fabra, Barcelona, Spain

Word count: around 8.700 words, abstract and refs included.

Abbreviations

AI	Artificial intelligence
ML	Machine learning
BM	Biophysical models
PINN	Physics-informed neural networks
ROM	Reduced order model
FOM	Full order model
PCA	Principal component analysis

Abstract

Biophysical models and machine learning may be perceived as rather different entities, or on the contrary as very related forms of modelling. In this chapter, we precisely develop on the latter idea to provide a didactic and up-to-date overview of some major research tracks where these two fields can collaborate and benefit each other. We specifically articulate contents around two complementary points-of-view on the potential benefits of one field to the other. For biophysical modelling, we focused on accelerating computations, estimating unobservable parameters, and examining complex outputs; for machine learning, we laid stress on adding physiologically-relevant knowledge, and generating synthetic data for training and validation. Along this review, we detail specific questions of relevance with examples majoritarly in the context of computational cardiology, which is our field of interest, and encourage further interaction between these two areas of active research.

Keywords

Machine learning; Simulations; Biophysical models; Synthetic data; Reduced-order models

Introduction

At first sight, machine learning and biophysical modelling could be seen as rather antinomic entities, which continue the long-lasting distinction between experimentalists and modellers, or even within the modelling field between phenomenological and detailed biophysical models. However, both types of methodologies are based on the same principles of creating simplified and controlled versions of real observations (e.g., machine learning from large datasets) or physical phenomena. More concretely, a model stands as an abstract representation that mimics a phenomenon in a realistic way, being governed by a set of parameters and state variables. For example, a 3D dynamic biomechanical model of the heart aims to reproduce the motion and deformation of the heart along the cardiac cycle, represented by a 3D dynamic volume

that evolves over time. The state variables¹ of the model are the equation unknowns that need to be solved (e.g., displacement in each mesh node). In addition, several parameters¹ that control the behaviour of such equations (for example, stiffness or contractility constants, among many others) are also present. In a similar manner, a neural network classifier is a machine learning model that aims to approximate a complex relationship between input data and output labels, to achieve the best accuracy in a given task (e.g., classification), once a metric has been defined (e.g., correlation). Its state variables consist of the weights given to the links between the different neurons, and the (hyper)parameters on top of them serve to condition the model (for example, the number of neurons in each layer, the importance given to some regularisation, etc.).

Both biophysical and machine learning models need to be optimised; model parameters are updated to best reproduce a given behaviour, either defined according to prior knowledge or from patient data. In consequence, model design is strictly dependent on the question to be answered, keeping in mind the principle of parsimony (i.e., Occam's razor), which recommends opting for the simplest path to solve a problem. In both cases, and although more complex models may be tempting, simpler models have clear advantages regarding optimisation (several simulations can be run at the same time), interpretability (less parameters to examine a-posteriori), and identifiability (i.e., identifying the optimal parameters from the measured data). They also suffer less from the curse of dimensionality, in particular in situations where the number of data samples is limited, which is a common issue in medical image analysis.

We illustrate this in Fig. 7.1 with two representative examples, one from biophysical modelling and another one from machine learning. They depict two situations in which a simple model is already very informative, either to understand a physiological behaviour with modelling (relation between cardiac volume and deformation), or to perform a simple diagnosis with machine learning from well-chosen and very informative input variables.

Historically-speaking, machine learning and modelling are not disconnected at all, as shown by early works not necessarily in the medical domain (Crutchfield and McNamara, 1987; Nelles, 2001). In its premises, Artificial Intelligence (AI) referred to knowledge-based decisions, in algorithmic models that aimed to model specific aspects of human thinking for problem solving. Unfortunately, this meaning strongly deviated from its original focus by abuse of language, using the term "AI" in place of its sub-branches "machine learning" and even "deep learning", which both tend to privilege knowledge discovery from large amounts of data. More recently, machine learning and biophysical modelling also both benefited the drastic evolution of technology, which provided access to powerful computational resources and therefore the ability to develop new types of computational models, therefore enhancing their relevance to solve complex problems. In machine learning, the last decade witnessed the advent of deep neural networks, which consist of several layers of neurons (in practice, more than three to be considered "deep"). Each of them adds non-linearity and a higher level of complexity in modelling the relationship between the inputs and outputs, associated with an increase of performance with sufficient amounts of data. Similarly, the understanding of living organs was boosted by the use of multiscale biophysical models, which can represent an organ's function from the microscale or cellular level up to the macroscale, challenges being to relate these different levels of observation together (Hunter et al., 2013). Naturally, as for the model complexity and the required database size, the need for highly-performing computing resources depends on the dimensional complexity of the question: for example, lumped parameters models of the cardiovascular circulation are very fast as they represent zero-dimensional data, compared to biomechanical models of the cardiac contraction that operate on 3D volumetric meshes.

Biophysical models and machine learning therefore share many similarities, and open the way to many interactions. This is precisely what we will develop in this chapter, which consists of a (subjective) attempt to provide a didactic and up-to-date overview of some major research tracks where these two fields can collaborate and benefit each other, as briefly summarised in Table 7.1. We refer the reader to other reviews specific to cardiac modelling (Mansi et al., 2020) or with a broader view on biological and biomedical modelling (Alber et al., 2019) for complementary overviews of this topic.

Here, we specifically articulate contents around two complementary points-of-view of the question set in the title ("Machine learning and biophysical models: how to benefit each other?"), which represent the two main sections of this chapter. In each of them, we detail specific questions of relevance with examples majoritarily in the context of computational cardiology, which is our field of interest. In this sense, the details given here are highly complementary from the previous chapters of this book, in particular Chapter 2 ("Linear and nonlinear dimensionality reduction of biomechanical models"), Chapter 4 ("Data-driven modelling and artificial intelligence") and Chapter 5 ("Deep Physics: combining physics-based modelling with deep learning for real-time soft tissue simulation), although prior reading of these chapters is not mandatory to apprehend our chapter, and conversely.

¹ Note that in the machine learning community, the terms used are slightly different, as "state variables" and "model parameters" are referred to as "parameters" and "hyperparameters", respectively.

Table 7.1: Main advantages and drawbacks from machine learning and biophysical modelling, which suggest areas of research where these two modelling strategies can benefit each other, as we develop in this chapter.

	Strengths	Weaknesses	Actions
Machine learning (ML)	<ul style="list-style-type: none"> - Efficient analysis of large amounts of data - Identification of data correlations and dynamics inference - Wide availability of Open Source tools 	<ul style="list-style-type: none"> - Black-boxness of some ML methods - Lack of interpretability and causality - Disregard laws of physics - Data hungriness 	<ul style="list-style-type: none"> - Using BM to whitening ML and improve interpretability - Physics-based regularisation of ML solutions - Generation of synthetic data with BM
Biophysical models (BM)	<ul style="list-style-type: none"> - Well-understood, controlled and interpretable causality - Based on laws of physics - Capture multi-scale nature of physical phenomena - Reliable predictions on unseen data 	<ul style="list-style-type: none"> - Long computational costs of detailed BM - Difficult assimilation of large amounts of data - Non-observable unknown model parameters set to generic values - Potentially complex 3D mesh generation 	<ul style="list-style-type: none"> - Surrogates of BM results with ML - Constrain space of BM solutions with ML - BM parameter optimisation with ML - Uncertainty quantification of BM parameters with ML - ML-based mesh generation

1. Machine learning as a support to biophysical models

The basic principle of supervised machine learning is to approximate complex relationships between input data x and output labels y . Complexity first lies in the dimensionality of these inputs and outputs, which can be high-dimensional. For example, an input image of 100×100 pixels has a dimensionality of 10^4 . Regarding outputs, automatic diagnosis generally means predicting a single output value (of dimensionality 1), while segmenting an image means predicting a categorical label at each pixel of this image (again, of dimensionality 10^4 if the image is composed 100×100 pixels). Complexity also lies in the actual link between inputs and outputs. A linear relationship can be represented through matrix multiplication (for example, $y = Wx + \epsilon$, where W is the matrix encoding the actual relationship, and ϵ corresponds to residual noise), while a non-linear relationship means $y = f(x) + \epsilon$, where f is a nonlinear function.

In this section, we will develop how the powerful approximation capacities of machine learning can be used to support the two main phases of optimisation involved in biophysical models: (i) to reduce the computational burden when solving the equations of a given biophysical model; and (ii) to optimally adapt the parameters of a biophysical model to a given patient's data, a process often referred to as parameter personalization. In addition, we also develop how machine learning can provide convenient data representations to better examine the simulation results. Specifically, the machine learning algorithms of interest will mostly be supervised², and for regression purposes³. Figure 7.2 shows how a standard fully-connected neural network can be adapted to perform regression at different levels of complexity.

Finally, we will briefly comment on how other types of machine learning models (mostly unsupervised, meaning that no labels are used for learning) can offer convenient and simplified representations of the data (either real data, or the output of biophysical models) to quantitatively assess the realism of simulations.

1.1 Accelerating computations in biophysical models

There are different ways in which machine learning techniques can help to reduce the large computational burden usually associated with complex biophysical models. The literature groups them depending on the considered type of bias (observational, inductive, or learning bias, following Karniadakis et al. (2021)), at which the integration of machine learning

² This means that outputs (also referred to as labels) are known for a subset of the data, which serves for the training phase of machine learning.

³ This means that the output to predict is a continuous variable (or several variables), as opposed to classification where the output is binary or categorical.

and domain knowledge is performed (supervised, loss-terms, or interleaved, as in Thuerey et al.⁴) or their application context (reduced-order models, super-resolution, upsampling, style transfer, or compression, as in Bai et al. (2021)). Here we give some examples that encompass the most relevant approaches (the reader is referred to Willard et al. (2021) for a more detailed overview):

- Estimating surrogates of biophysical model outputs with machine learning
- Solving biophysical model equations with differential operators based on machine learning
- Reduced-order models based on machine learning

1.1.1 Estimating surrogates of biophysical model outputs with machine learning

The “learning” process of machine learning refers to its training phase, which means that from representative sample data the machine learning model parameters are learnt, which may take a while depending on the complexity of the model and the data involved. During the testing phase, the optimal parameters are applied to estimate an output prediction from new input data, which is much faster compared to the training phase. For this reason, machine learning regression models naturally appear as convenient tools to approximate the solutions of complex equations involved in biophysical models, and therefore reach very fast computations.

The simplest and direct approach to accelerate computations in biophysical models is to learn the relation between the input of the biophysical models and the simulation results with advanced machine learning techniques. Indeed, in theory, a sufficiently complex machine learning model might be able to approximate complex non-linear relationships; it needs to find a convenient representation of the data that makes this approximation easier to estimate and have enough representative data samples for training. Usually, learning is performed in a supervised manner, thus requiring the generation of a lot of simulation results capturing the space of solutions, which can be computationally expensive in the training phase. On the other hand, once trained, Machine learning-based surrogates of complex system dynamics can be evaluated with significantly less computational resources and several orders of magnitude faster than conventional finite element solvers. For instance, deep learning was used to estimate haemodynamics parameters derived from blood velocity vector fields of fluid simulations in the aorta (Liang et al., 2018), including coronary arteries (Li et al., 2021), reducing calculation times 600-fold.

However, most widespread neural networks are optimised to work with regular Cartesian-like domains such as images, while biophysical model outputs are usually represented in non-Euclidean domains such as graphs and meshes. Geometrical deep learning approaches (Bronstein et al., 2017), initially developed to operate with graph data, have therefore been applied to biomedical meshes. For example, recent work proposed to enhance electrical activation patterns from sparse data, after training a graph convolutional neural network with synthetically generated electrophysiological simulations (Meister et al., 2021). The relation between the morphological complexity of the left atrial appendage and in-silico haemodynamics indices of thrombogenic risk, derived from fluid simulations (Morales Ferez et al., 2021), was also better learnt with geometric deep learning than with other n approaches working in the Euclidean domain (therefore insufficiently exploiting graph data) such as combinations of principal component analysis and fully convolutional layers or 3D flattening followed by U-net models. More advanced end-to-end (i.e., learning from known solutions but without a solver-in-the-loop) geometric deep learning methods (Pfaff et al., 2021) are being proposed for fluid simulations, where rich physical states are represented by graphs of interacting particles and complex dynamics are approximated by learned message-passing among nodes. These fully learned simulators often only work well with conditions similar to the training distribution, but generalisation can be improved with training noise tuning, temporal downsampling, added loss constraints and convolution architectures (Stachenfeld et al., 2022).

1.1.2 Solving biophysical model equations with differential operators based on machine learning

There are several ways to use machine learning techniques to facilitate the solving of biophysical model equations. The pioneering attempts (see Karniadakis et al. (2021) for a recent review from the original authors, and Lu et al. (2021) for related code libraries), often referred to as physics-informed neural networks (PINNs), were based on the incorporation of physical dynamics in the loss functions of neural networks, using automatic differentiation to represent all the differential operators of the physical system to be solved (i.e., differentiating neural network with respect to its input coordinates and model parameters). These approaches are based on simultaneously optimising the parameters of a neural network from loss-terms fitting to observed data and to yield predictions that approximately satisfy a given set of physical soft constraints such as mass or momentum conservation, penalising deviations from the target values (e.g., zero residuals for the conservation laws). The main advantages of PINNs are two-fold: 1) they can easily handle irregular and complex domains since they do

⁴ <https://physicsbaseddeeplearning.org>

not explicitly require mesh generation; and 2) they do not need a large training dataset since they use physical equations to guide the training process, together with discrepancies between the predicted and actual initial and boundary conditions. In consequence, PINNs have been used in multiple complex biophysical problems (Karniadakis et al., 2021), notably for fluid mechanics (Cai et al., 2022).

Using a deep learning-based method with physically based soft constraints makes the approach dependent on the learned representation of model derivatives that, furthermore, can be very expensive since each derivative requires backpropagation through the full network. Differentiable physics approaches⁴ (Holl et al., 2020) were recently proposed to solve some of the drawbacks associated to PINNs, tightly coupling the physical model in the training of the deep neural network (i.e., solver-in-the-loop), which requires a fully differentiable simulator and allows efficient evaluation of the simulation and model derivatives. The main principles of differential physics methods is to rely on the stability and robustness of existing numerical solvers through the problem discretization, integrating ways to compute gradients with respect to their inputs; later, backpropagation schemes are used to communicate gradient information from the simulator to the neural network and vice versa. On the other hand, differential physics implementations are not straightforward and ask for a good understanding of the problem to solve to appropriately choose a correct discretization.

1.1.3 Reduced-order models based on machine learning

Reduced order models (ROMs), based on a projection process (e.g., principal component analysis, proper orthogonal decomposition) can estimate a basis to reduce the state-space dimensionality of the solution manifold. Such models have been proposed as more computationally tractable versions of the full order models (FOMs). However, traditional ROM approaches are intrinsically linear, and are therefore limited given the usual non-linearity of complex biophysical problems. Deep learning techniques can be used to learn non-linear representations of these problems, providing high accuracy with fewer degrees of freedom compared to projection-based techniques. Typically, autoencoders are appropriate for such tasks since they build a reduced latent space and in the meantime compress (encoder) and reconstruct (decoder) the input data. For instance, a deep learning-based ROM methodology was recently proposed (Fresca et al., 2020) for cardiac electrophysiology, where a convolutional autoencoder is combined with a deep feedforward neural network to exploiting examples of FOM solutions obtained for different parameter values and to learn both, the nonlinear manifold where the ROM solution is sought, and the nonlinear reduced dynamics. The authors reported deep learning-based ROMs can be 275 times faster compared to classical projection-based approaches in testing times. A similar approach (Kashtanova et al., 2021) using deep neural networks (i.e., ResNets) showed good generalisation ability with unseen configurations with different cardiac tissue properties (e.g., scars, multiple onsets, conduction velocities).

1.2 Estimating unobservable parameters with biophysical models

Apart from helping to understand a physiological phenomenon, biophysical models are extremely interesting as they involve parameters that are not directly quantifiable from real patient data, in particular from imaging. Thus, if a model can be “personalised” (see detailed explanations below) to existing data from a given patient, the fitted model will provide access to such physiological parameters and shed new light into this patient’s assessment.

For example, in cardiac imaging, myocardial strain can reflect subtle abnormalities related to diseases (e.g. locally reduced deformation in case of myocardial infarction) but only remains a limited surrogate for myocardial contractility. Several works have started to personalise a biomechanical model to patient data to estimate hidden parameters such as myocardial contractility (Rumindo et al., 2020), material properties of the aorta (Zhang, 2021), or thrombus (Yin et al., 2021), or to add consistency in estimation of more accessible parameters that can be missing on an heterogeneous patient database (Mollero et al., 2019).

Personalization can be complex. A simple sequential approach would be to slightly adjust the biophysical model’s parameters, run the model, observe the fit to the real data, and repeat this process until a satisfactory fit is reached. Machine learning can be of great help to improve this process.

Gradient-free methods are variational optimization methods that are close to a machine learning process, and very interesting when an analytical formulation of the models is complex or impossible. Among these, the evolutionary algorithm CMA-ES (Covariance Matrix Adaptation – Evolution Strategy) (Hansen and Ostermeier, 1996) is considered state-of-the-art in a large variety of difficult non-linear non-convex optimization problems. This algorithm runs several simulations at each iteration (generation), and adapts across iterations the covariance matrix that encodes the distribution of these synthetic data. Its performance was demonstrated for the individual personalization of complex biomechanical models

of the heart (Mollero et al., 2018), which can be reinforced with population-based priors (Mollero et al., 2019), and even for population-based personalization as illustrated on simple geometrical models of myocardial infarct (Mom et al., 2021).

Data-driven methods were also proposed for personalization. Some of them can estimate the relationship between the model parameters and the outputs, and eventually simplify the optimization involved during personalization. An interesting approach lies in probabilistic methods that consider uncertainties both on the data and on the model parameters. They model interactions between observations, parameters and state variables, the personalization process consisting in estimating the posterior distribution for the parameters given the observations. Efficient Bayesian inference strategies using polynomial chaos (a way to represent a random variable as a polynomial function of other random variables) were specifically proposed for the personalization of cardiac electrophysiology (Konukoglu et al., 2011) and cardiac electromechanical (Neumann et al., 2014) models. Another interesting strategy was adapted from the field of reinforcement learning in Neumann et al. (2016). It considers an “intelligent” agent that takes actions (adapting the model parameters) with the aim of learning a policy that maximises a cumulative reward (defining personalization success). All these methods are very interesting as they are faster than sequential methods, with the limitation that they require building a relevant database by previously running a large amount of well-selected simulations.

1.3 Examining the outputs of simulations

The output of biophysical models generally consists of complex data in the form of high-dimensional descriptors. For example, in the case of cardiac models, such data can be temporal signals (pressure, volume, etc.) for 0D models of the cardiac circulation (Morris et al., 2016), temporal sequences of 3D shapes for cardiac mechanical models (Suinesiaputra et al., 2016), or temporal sequences of 3D patterns embedded into a 3D shape for more advanced models of the cardiac circulation (Morris et al., 2016) or models of the cardiac conduction system (Camara et al., 2011).

A subfield of machine learning called representation learning offers a convenient way to examine a set of simulations obtained by varying the model parameters, either to generate a realistic population or to examine the model sensitivity to its parameters. Representation learning mostly consists in estimating a simplified/latent (low-dimensional) space where the (high-dimensional) data from a population is organised according to specific statistical properties. This space can then be exploited further for statistical comparisons, interpretations, or even serve as input for other diagnosis or prognosis algorithms. Methods for this are mostly unsupervised, and can be linear or nonlinear. The most popular one is Principal Component Analysis (PCA), which organises the data samples according to the main directions of variance in the data. Many examples of non-linear algorithms can be related to the generalised framework presented in Yan et al. (2007), for which samples with similar values are grouped together in the latent space, while dissimilar samples are kept more distant from each other.

The latent space can still consist of several dimensions, which may be challenging for visualisation purposes but relevant to perform statistical comparisons between samples. A subset of popular methods can explicitly estimate a two-dimensional representation of a population, tailored for visualisation purposes, which explains their popularity: UMAP (McInnes et al., 2018) and t-SNE (vanderMaaten and Hinton, 2008). However with these two methods, one needs to keep in mind that distances in the latent space are not necessarily statistically meaningful (Wang et al., 2021), and therefore their use to compare real vs. synthetic populations has to be taken with care.

Figure 7.3 illustrates this with two representative examples from two cardiovascular biophysical models. First, a sensitivity analysis was run on a 0D version of 3D cardiac electromechanical model⁵. We changed two hyperparameters among 14 active hyperparameters, which resulted in a set of 130 different pressure and volume curves. Using either pressure or volume curves as input, standard non-linear representation learning was able to estimate a low-dimensional representation of the 130 samples, where these were consistently arranged according to the values of the two modified parameters (see Fig. 7.3.A). In a second example, we designed a synthetic dataset of infarct images that was optimised to best resemble a real dataset of similar images. Standard non-linear representation learning provided a simplified low-dimensional representation of these two datasets, which allowed visualising the matching of their distributions and therefore confirming the realism of the synthetic dataset that was designed (see Fig. 7.3.B).

We can even go further by exploiting the latent spaces encoding the output from two different biophysical models. There may be known correspondences between the two latent spaces if each model was personalised to the same patients. These correspondences can be used to wrap one latent space onto the other one, and therefore match the output of the first model to the second one, so that they can be compared. We used this concept inspired from the machine learning field of domain

⁵ CellML model: <https://models.cellml.org/workspace/44c> ; we used a script version provided by the authors.

adaptation to quantify differences between simulations from two different cardiac biomechanical models, assessed through myocardial deformation patterns (Duchateau et al., 2019).

2. Biophysical models as a support to machine learning

In theory, modelling is the ultimate tool for learning about something, since being able to model it in a realistic way means that it is (almost) fully understood. To a lesser extent, modelling can be seen as a highly valuable source of knowledge for a purely data-driven approach, which by definition depends on the input data (amount and representativeness of samples) and the way they are handled during learning.

An ideal machine learning algorithm should be able to perfectly solve a given problem with infinite amounts of data. Given that the machine learning remains a model (namely an approximation) and that data are not infinite, biophysical modelling offers very interesting support for machine learning either on the machine learning model itself by introducing physiologically-relevant knowledge, or on the input data by enabling to generate new synthetic realistic data. These consist of the two practical situations we will develop in the remainder of this section.

2.1 Constraining machine learning algorithms with physiologically-relevant knowledge

Modelling means following physical principles. Adding prior knowledge and constraints to the machine learning can be seen as adding a portion of human learning into the data-based learning. The benefits are two-fold. First, prior knowledge brings robustness in the case of limited population sizes or with missing data. Then, in the meantime, it can reinforce the physiological relevance of the outputs, which is not guaranteed depending on the machine learning model and the representativeness of the data samples.

In cardiac imaging, the value of considering physical principles of conservation was demonstrated on two major quantification tasks: segmentation and tracking. The myocardium is supposed to be incompressible, which can be formulated as conserving the myocardial volume across a sequence. For a tracking or segmentation algorithm, this could consist of an additional penalty in the energy to minimise (also referred to as the loss function), expressing for example how far the estimated myocardial volume deviates from the initial volume. For registration/tracking, some algorithms express the transformation that wraps one image to another one by means of dense velocity fields, as in the popular Demons algorithm (which belongs to variational methods and not machine learning). This allows formulating incompressibility in a finer manner, using the divergence of such velocities (Mansi et al., 2011).

Another frequent issue arises with the recent segmentation algorithms based on deep neural networks. These are highly performing, up to the level of experienced annotators, but mostly process 2D images. This comes from the drastic increase in computational complexity when handling 3D images with neural networks, and the access to much more data samples when considering all slices in a 3D image as several independent 2D images. However, these suffer from spatial and temporal consistency issues: there is no guarantee of consistency across the segmentations of the slices of an MRI stack, or across the temporal instants of the cardiac cycle. Recent works demonstrated that adapting the learning architecture to include simple prior knowledge can solve this issue. For example, segmentation results can be propagated across slices to reach better spatial consistency (Zheng et al., 2018), or constrained to be anatomically-consistent (Oktay et al., 2018). They can also be represented in a simpler low-dimensional latent space with additional constraints to order them according to physiological measurements such as the volume of the cardiac chambers, which allows smoothing temporal trajectories in a physiologically-relevant manner (Painchaud et al., 2022). The recent improvements in the realism and portability of biophysical models, and in constraining machine learning algorithms with prior knowledge will offer the possibility to go much further in the next future, where such prior knowledge could consist of shape, motion, or other data extracted from the simulations.

2.2 Synthetic data generation for training and validation of machine learning algorithms

Another substantial asset of biophysical modelling comes from the generative nature of the models: one can easily generate many variants of a simulation by changing a subset of hyperparameters.

In the machine learning context, the first application that comes to mind is validation with ground truth data. Indeed, simulations represent a controlled environment conditioned by few variables, with known output data (of better quality compared to real data) that can stand as ground truth data to evaluate supervised machine learning algorithms. For biomedical applications, this purpose is therefore close to using experimental models before testing methods on clinical

data, except that simulated data allows in-silico evaluation and with a much larger population size, which suits the needs of machine learning algorithms.

Figure 7.4 illustrates this on a challenging regression application, namely predicting the location of a myocardial infarct from the local deformation (strain) of the myocardium. On real data, deformation can be extracted from echocardiography, with the objective of assessing its potential for infarct localization against magnetic resonance, which is more expensive, may require the injection of a contrast agent, and is not possible in all situations. Using a state-of-the-art electromechanical model of the heart composed of the left and right ventricles, we altered the contraction parameters of the model at specific locations that mimic an infarct to generate abnormal deformation patterns. We built a large population of 500 synthetic cases with infarcts of random shape and location, which served to evaluate our machine learning algorithm for infarct localization, formulated as a non-linear regression (Duchateau et al., 2016). Although several aspects of the simulations may be improved, such as the variety of infarct and ventricular shapes, or the heterogeneity of deformation patterns, this model served to understand and validate the behaviour of the machine learning algorithm on a controlled environment before applying it to real data (both training and testing).

Apart from validation, another major challenge for machine learning lies in the amount and quality of data, and more specifically annotated data. The generative capability of biophysical simulations means that they can be employed to obtain synthetic data to augment the input database with realistic samples, and in the meantime have ground truth data from the model's theoretical values, to use as labels for supervised methods. However, the main condition for this is being able to generate realistic-enough data that can be mixed with real data.

Without going into detailed simulations, a partially similar (but much simpler) strategy referred to as “data augmentation” is already part of many standard machine learning pipelines. For image processing applications (not necessarily medical applications), it consists in generating many variants of the existing input images by changing their rotation, scaling/zoom, brightness, etc. Similarly, for text processing applications, many variants of the input text data can be generated by changing words by their synonyms or applying positional swaps in a sentence. The purpose is to help the learning algorithm to be invariant to these irrelevant confounding factors.

For medical applications, this process can be moved further by taking advantage of realistic simulations to artificially augment a given database. A pioneering example of this demonstrated how realistic synthetic images from catheters allowed training a catheter localization algorithm with ground truth data, without the need for manually annotating the catheter location in representative images (Heimann et al., 2014). Similarly, a popular work from the field of computer vision reported how synthetic images from video games can be used to train self-driving cars (Richter et al., 2016). Recent works demonstrated how electrophysiological simulations of the heart can boost the performance of anomaly localization algorithms (Doste et al., 2019), or ECG-delineation algorithms based on neural networks (Jimenez-Perez et al., 2021). The last decade also witnessed a substantial evolution in the realism of synthetic image sequences generated from an electromechanical model of the heart, either by warping existing image sequences according to the model's motion (Prakosa et al., 2013; Duchateau et al., 2018) or by mimicking the imaging process to generate realistic textures (Alessandrini et al., 2015; Alessandrini et al., 2018; Zhou et al., 2018). Figure 7.5 illustrates how we used this framework to generate a large database of 465 synthetic cine MR sequences (15 healthy and 450 with abnormal tissue viability, mimicking lesions of random location, extent and grade), ready to serve as training data for cardiac segmentation or motion estimation algorithms (Duchateau et al., 2018).

Naturally, a key issue with these approaches remains how far simulations are from real data, which can be quantified using some of the techniques described in Section 1.3. Bridging the gap between synthetic and real data is a whole field of research in itself, referred to as “domain adaptation”. We refer the reader to the very interesting reviews on this topic (Cheplygina et al., 2019; Kouw and Loog, 2021; Guan and Liu, 2022). In any case, a perfect match between simulations and real data is impossible, as there remains uncertainty both in the biophysical models (which are an approximation of reality) and in the acquired data (which consist of a snapshot of a real organ's behaviour).

Conclusions

We have provided an overview of the main areas where machine learning and biophysical modelling, two areas of research that share many more common points that one tends to think, can benefit each other. Both rely on models, which are the cornerstone of challenges still to address in both fields: simplicity against realism, computational efficiency, interpretability,

and medical relevance, among the main ones. Both also rely on data, and not only Big data but fully curated and unbiased datasets, to train the machine learning models and accurately personalise the biophysical models.

In the current context of research dependence on local data and tuned models, one should keep in mind that the explosion of machine learning has been possible thanks to the open mind and widespread policies of the community to release as Open Access tons of labelled data for training and evaluation, as well as the state-of-the-art machine learning algorithms as Open Source codes. This has happened to some extent for biophysical modelling (see for example popular initiatives such as the Cardiac Atlas Project (Fonseca et al., 2011), cellML (Garny et al., 2008), or recent challenges to foster reproducibility⁶), but to a lower extent compared to the machine learning community, and efforts towards this should be widely encouraged.

In the end, both machine learning and biophysical modelling pursue the same application objective: refining our understanding of disease and improving patients' care. In this context, being able to translate biophysical and machine learning models from moderate or even large controlled datasets to real-life environments is crucial, which raises additional issues on the computational complexity and generalisation ability of the models, not directly addressed in this overview. Historically-speaking, some years ago one bottleneck for applying biophysical models to real data was meshing; the recent introduction of geometrical aspects into deep learning methods (Bronstein et al., 2017) is now opening new perspectives not only to construct the meshes but also for the analysis of the data descriptors attached to them. Huge steps forward are therefore expected, similar to those observed when convolutional neural networks were widely deployed for the analysis of medical images. Still in this historical perspective, the application scenarios for biophysical models have substantially evolved from constructing a Virtual Physiological Human (Hunter et al., 2013), integrating many scales of modelling not necessarily specific to a patient. The scientific community now refers to much more personalised versions referred to as the Digital Twin of a given patient (Corral-Acero et al., 2020), which enables running many possible virtual scenarios such as tens of different treatments (for a more concrete example, many different choices of pacing sites and delays between the electrodes to treat asynchrony or rhythm disorders). Again, such simulations are not possible with computationally costly models, and the advent of machine learning has certainly a role to play, as briefly commented in this overview.

However, in the end, the impact of such models is and will be strongly conditioned by the final user and interfaces set to interact with their outputs. One may actually wonder how the combination of machine learning and biophysical models will be shown to the clinician, whose role is not to run machine learning and biophysical modelling tools. The growing fields of visual analytics and interactive tools (e.g., web-based, augmented/virtual reality headsets) represent promising solutions to handle this, although they need to be developed by multi-disciplinary teams including clinicians, data scientists and modellers. Alternatively, if this is set up by (device, drug, imaging, etc.) companies, these industrial actors should be able to settle and run in-silico trials consisting of hundreds of simulations to obtain FDA or CE certifications of the computational outputs (Pappalardo et al., 2019).

Definitely, the combination of machine learning and biophysical modelling is going to have a major impact and will force researchers to think differently since many applications that were impossible five-to-ten years ago are now possible, paving the way for more interactive relations between the final users and these computational tools, ultimately leading to better clinical decisions and optimised patient management.

Acknowledgements

The research of Prof. O. Camara on in-silico modelling and machine learning is funded by the Spanish Ministry of Economy and Competitiveness under the Programme Retos Investigación project (RTI2018-101193-B-I00) and the H2020 EU SimCardioTest project (Digital transformation in Health and Care SC1-DTH-06-2020; grant agreement No. 101016496). The research of Dr. N. Duchateau on these topics is partially funded by the French ANR (LABEX PRIMES of Univ. Lyon [ANR-11-LABX-0063] within the program "Investissements d'Avenir" [ANR-11-IDEX-0007], the JCJC project "MIC-MAC" [ANR-19-CE45-0005]).

⁶ <http://stacom.cardiacatlas.org/>

References

1. Alber M, Buganza Tepole A, Cannon WR, De S, Dura-Bernal S, Garikipati K, Karniadakis G, Lytton WW, Perdikaris P, Petzold L, Kuhl E. Integrating machine learning and multiscale modeling-perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. *NPJ Digit Med*. 2019;2:115.
2. Alessandrini M, De Craene M, Bernard O, Giffard-Roisin S, Allain P, Waechter-Stehle I, Weese J, Saloux E, Delingette H, Sermesant M, D'hooge J. A pipeline for the generation of realistic 3D synthetic echocardiographic sequences: methodology and open-access database. *IEEE Trans Med Imaging*. 2015;34:1436-51.
3. Alessandrini M, Chakraborty B, Heyde B, Bernard O, De Craene M, Sermesant M, D'hooge J. Realistic vendor-specific synthetic ultrasound data for quality assurance of 2-D speckle tracking echocardiography: simulation pipeline and open access database. *IEEE Trans Ultrason Ferroelectr Freq Control*. 2018;65:411-22.
4. Bai K, Wang C, Desbrun M, Liu X. Predicting high-resolution turbulence details in space and time. *ACM Trans Graph*. 2021;40:200.
5. Bijmens B, Cikes M, Butakoff C, Sitges M, Crispi F. Myocardial motion and deformation: What does it tell us and how does it relate to function? *Fetal Diagn Ther*. 2012;32:5-16.
6. Bronstein MM, Bruna J, LeCun Y, Szlam A, Vandergheynst P. Geometric deep learning: going beyond Euclidean data. *IEEE Signal Process Mag*. 2017;34:18-42.
7. Cai S, Mao Z, Wang Z, Yin M, Karniadakis GE. Physics-informed neural networks (PINNs) for fluid mechanics: a review. *Acta Mech Sin*. 2022.
8. Camara O, Sermesant M, Lamata P, Wang L, Pop M, Relan J, De Craene M, Delingette H, Liu H, Niederer S, Pashaei A, Plank G, Romero D, Sebastian R, Wong KCL, Zhang H, Ayache N, Frangi AF, Shi P, Smith NP, Wright GA. Inter-model consistency and complementarity: learning from ex-vivo imaging and electrophysiological data towards an integrated understanding of cardiac physiology. *Prog Biophys Mol Biol*. 2011;107:122-33.
9. Cheplygina V, de Bruijne M, Pluim JPW. Not-so-supervised: A survey of semi-supervised, multi-instance, and transfer learning in medical image analysis. *Med Image Anal*. 2019;54:280-96.
10. Corral-Acero J, Margara F, Marciniak M, Rodero C, Loncaric F, Feng Y, Gilbert A, Fernandes JF, Bukhari HA, Wajdan A, Villegas Martinez M, Sousa Santos M, Shamohammdi M, Luo H, Westphal P, Leeson P, DiAchille P, Gurev V, Mayr M, Geris L, Pathmanathan P, Morrison T, Cornelussen R, Prinzen F, Delhaas T, Doltra A, Sitges M, Vigmond EJ, Zacur E, Grau V, Rodriguez B, Remme EW, Niederer S, Mortier P, McLeod K, Potse M, Pueyo E, Bueno-Orovio A, Lamata P. The 'Digital Twin' to enable the vision of precision cardiology. *Eur Heart J*. 2020;41:4556-64.
11. Crutchfield JP, McNamara B. Equations of motion from a data series. *Complex Systems*. 1987;1:417-52.
12. Di Folco, Mocerri P, Clarysse P, Duchateau N. Characterizing interactions between cardiac shape and deformation by non-linear manifold learning. *Med Image Anal*. 2022;75:102278.
13. Doste R, Lozano M, Gomez JF, Alcaine A, Mont L, Berruezo A, Camara O, Sebastian R. Predicting the origin of outflow tract ventricular arrhythmias using machine learning techniques trained with patient-specific electrophysiological simulations. *Proc. Computing in Cardiology (CinC)*. 2019;1-4.
14. Duchateau N, De Craene M, Allain P, Saloux E, Sermesant M. Infarct localization from myocardial deformation: Prediction and uncertainty quantification by regression from a low-dimensional space. *IEEE Trans Med Imaging*. 2016;35:2340-52.
15. Duchateau N, Sermesant M, Delingette H, Ayache N. Model-based generation of large databases of cardiac images: synthesis of pathological cine MR sequences from real healthy cases. *IEEE Trans Med Imaging*. 2018;37:755-66.

16. Duchateau N, Rumindo K, Clarysse P. Domain adaptation via dimensionality reduction for the comparison of cardiac simulation models. *Proc. Functional Imaging and Modeling of the Heart (FIMH), LNCS* 2019;11504:276-84.
17. Fonseca CG, Backhaus M, Bluemke DA, Britten RD, Chung JD, Cowan BR, Dinov ID, Finn JP, Hunter PJ, Kadish AH, Lee DC, Lima JAC, Medrano-Gracia P, Shivkumar K, Suinesiaputra A, Tao W, Young AA. The Cardiac Atlas Project--an imaging database for computational modeling and statistical atlases of the heart. *Bioinformatics*. 2011;27:2288-95.
18. Fresca S, Manzoni A, Dedè L, Quarteroni A. Deep learning-based reduced order models in cardiac electrophysiology. *Plos One*. 2020;15:e0239416.
19. Garny A, Nickerson DP, Cooper J, dos Santos RW, Miller AK, McKeever S, Nielsen PMF, Hunter PJ. CellML and associated tools and techniques. *Philos Trans A Math Phys Eng Sci*. 2008;366:3017-43.
20. Guan H, Liu M. Domain adaptation for medical image analysis: a survey. *IEEE Trans Biomed Eng*. 2022;69:1173-85.
21. Hansen A, Ostermeier A. Adapting arbitrary normal mutation distributions in evolution strategies: The covariance matrix adaptation. *Proc. IEEE Conference on Evolutionary Computation (ICEC)*, 1996;312-7.
22. Heimann T, Mountney P, John M, Ionasec R. Real-time ultrasound transducer localization in fluoroscopy images by transfer learning from synthetic training data. *Med Image Anal*. 2014;18:1320-8.
23. Holl P, Koltun V, Thuerey N. Learning to control PDEs with differentiable physics. *Proc. International Conference on Learning Representations (ICLR)*. 2020.
24. Hunter P, Chapman T, Coveney PV, de Bono B, Diaz V, Fenner J, Frangi AF, Harris P, Hose R, Kohl P, Lawford P, McCormack K, Mendes M, Omholt S, Quarteroni A, Shublaq N, Skår J, Stroetmann K, Tegner J, Thomas SR, Tollis I, Tsamardinos I, van Beek JHGM, Viceconti M. A vision and strategy for the virtual physiological human: 2012 update. *Interface Focus*. 2013;3:20130004.
25. Jimenez-Perez G, Acosta J, Alcaine A, Camara O. Generalizing electrocardiogram delineation: training convolutional neural networks with synthetic data augmentation. *arXiv*. 2021. <https://arxiv.org/abs/2111.12996>
26. Karniadakis GE, Kevrekidis IG, Lu L, Perdikaris P, Wang S, Yang L. Physics-informed machine learning. *Nat Rev Phys*. 2021;3:422-40.
27. Kashtanova V, Ayed I, Cedilnik N, Gallinari P, Sermesant M. EP-Net 2.0: out-of-domain generalisation for deep learning models of cardiac electrophysiology. *Proc. Functional Imaging and Modeling of the Heart (FIMH), LNCS*. 2021;12738:482-92.
28. Konukoglu E, Relan J, Cilingir U, Menze BH, Chinchapatnam P, Jadidi A, Cochet H, Hocini M, Delingette H, Jaïs P, Haïssaguerre M, Ayache N, Sermesant M. Efficient probabilistic model personalization integrating uncertainty on data and parameters: Application to eikonal-diffusion models in cardiac electrophysiology. *Prog Biophys Mol Biol*. 2011;107:134-46.
29. Kouw WM, Loog M. A review of domain adaptation without target labels. *IEEE Trans Pattern Anal Mach Intell*. 2021;43:766-785.
30. Li G, Wang H, Zhang M, Tupin S, Qiao A, Liu Y, Ohta M, Anzai H. Prediction of 3D Cardiovascular hemodynamics before and after coronary artery bypass surgery via deep learning. *Commun Biol*. 2021;4:99.
31. Liang L, Liu M, Martin C, Sun W. A deep learning approach to estimate stress distribution: a fast and accurate surrogate of finite-element analysis. *J R Soc Interface*. 2018;15:20170844.

32. Lu L, Meng X, Mao Z, Karniadakis GE. DeepXDE: a deep learning library for solving differential equations. *SIAM Rev.* 2021;63:208:28.
33. Mansi T, Pennec X, Sermesant M, Delingette H, Ayache N. iLogDemons: a demons-Based Registration Algorithm for Tracking Incompressible Elastic Biological Tissues. *Int J Comput Vis.* 2011;92:92-111.
34. Mansi T, Passerini T, Comaniciu D. Artificial intelligence for computational modeling of the heart. *Elsevier.* 2020.
35. McInnes L, Healy J, Melville J. Uniform manifold approximation and projection for dimension reduction. *arXiv.* 2018. <https://arxiv.org/abs/1802.03426>
36. Meister F, Passerini T, Audigier C, Lluch È, Mihalef V, Ashikaga H, Maier A, Halperin H, Mansi T. Extrapolation of ventricular activation times from sparse electroanatomical data using graph convolutional neural networks. *Front Physiol.* 2021;12:694869.
37. Molléro R, Pennec X, Delingette H, Garny A, Ayache N, Sermesant M. Multifidelity-CMA: a multifidelity approach for efficient personalisation of 3D cardiac electromechanical models. *Biomech Model Mechanobiol.* 2018;17:285-300.
38. Molléro R, Pennec X, Delingette H, Ayache N, Sermesant M. Population-based priors in cardiac model personalisation for consistent parameter estimation in heterogeneous databases. *Int J Numer Method Biomed Eng.* 2019;35:e3158.
39. Mom K, Clarysse P, Duchateau N. Population-based personalization of geometric models of myocardial infarction. *Proc. Functional Imaging and Modeling of the Heart (FIMH), LNCS* 2021;12738:3-11.
40. Morales Ferez X, Mill J, Juhl KA, Acebes C, Iriart X, Legghe B, Cochet H, De Backer O, Paulsen RR, Camara O. Deep learning framework for real-time estimation of in-silico thrombotic risk indices in the left atrial appendage. *Front Physiol.* 2021;12:694945.
41. Morris PD, Narracott A, von Tengg-Kobligk H, Silva Soto DA, Hsiao S, Lungu A, Evans P, Bressloff NW, Lawford PV, Hose DR, Gunn JP. Computational fluid dynamics modelling in cardiovascular medicine. *Heart.* 2016;102:18-28.
42. Nelles O. Nonlinear system identification. *Springer Berlin Heidelberg.* 2001.
43. Neumann D, Mansi T, Georgescu B, Kamen A, Kayvanpour E, Amr A, Sedaghat-Hamedani F, Haas J, Katus H, Meder B, Hornegger J, Comaniciu D. Robust image-based estimation of cardiac tissue parameters and their uncertainty from noisy data. *Proc. Med Image Comput Comput Assist Interv (MICCAI), LNCS.* 2014;17:9-16.
44. Neumann D, Mansi T, Itu L, Georgescu B, Kayvanpour E, Sedaghat-Hamedani F, Amr A, Haas J, Katus H, Meder B, Steidl S, Hornegger J, Comaniciu D. A self-taught artificial agent for multi-physics computational model personalization. *Med Image Anal.* 2016;34:52-64.
45. Oktay O, Ferrante E, Kamnitsas K, Heinrich M, Bai W, Caballero J, Cook SA, De Marvao A, Dawes T, O'Regan DP, Kainz B, Glocker B, Rueckert D. Anatomically constrained neural networks (ACNNs): application to cardiac image enhancement and segmentation. *IEEE Trans Med Imaging.* 2018;37:384-95.
46. Painchaud N, Duchateau N, Bernard O, Jodoin PM. Echocardiography segmentation with enforced temporal consistency. *arXiv.* 2021. <https://arxiv.org/abs/2112.02102>
47. Pappalardo F, Russo G, Tshinanu FM, Viceconti M. In silico clinical trials: concepts and early adoptions. *Brief Bioinform.* 2019;20:1699-1708.
48. Pfaff T, Fortunato M, Sanchez-Gonzalez A, Battaglia P. Learning mesh-based simulation with graph networks. *Proc. International Conference on Learning Representations (ICLR).* 2021.

49. Prakosa A, Sermesant M, Delingette H, Marchesseau S, Saloux E, Allain P, Villain N, Ayache N. Generation of synthetic but visually realistic time series of cardiac images combining a biophysical model and clinical images. *IEEE Trans Med Imaging*. 2013;32:99-109.
50. Richter RR, Vineet V, Roth S, Koltun V. Playing for Data: Ground Truth from Computer Games. *Proc. European Conference on Computer Vision (ECCV), LNCS*. 2016;9906:102-18.
51. Rumindo GK, Ohayon J, Croisille P, Clarysse P. In vivo estimation of normal left ventricular stiffness and contractility based on routine cine MR acquisition. *Med Eng Phys*. 2020;85:16-26.
52. Stachenfeld K, Fielding DB, Kochkov D, Cranmer M, Pfaff T, Godwin J, Cui C, Ho S, Battaglia P, Sanchez-Gonzalez A. Learned coarse models for efficient turbulence simulation. *arXiv*. 2022. <https://arxiv.org/abs/2112.15275>
53. Suinesiaputra A, McCulloch AD, Nash MP, Pontre B, Young AA. Cardiac image modelling: Breadth and depth in heart disease. *Med Image Anal*. 2016;33:38-43.
54. van der Maaten LJP, Hinton GE. Visualizing high-dimensional data using t-SNE. *J Mach Learn Res*. 2008;9:2579-605.
55. Wang Y, Huang H, Rudin C, Shaposhnik Y. Understanding how dimension reduction tools work: an empirical approach to deciphering t-SNE, UMAP, TriMap, and PaCMAP for data visualization. *J Mach Learn Res*. 2021;22:1-73.
56. Willard J, Jia X, Xu S, Steinbach M, Kumar V. Integrating scientific knowledge with machine learning for engineering and environmental systems. *arXiv*. 2020. <https://arxiv.org/abs/2003.04919>
57. Yan S, Xu D, Zhang B, Zhang HJ, Yang Q, Lin S. Graph embedding and extensions: a general framework for dimensionality reduction. *IEEE Trans Pattern Anal Mach Intell*. 2007;29:40-51.
58. Yin M, Zheng X, Humphrey JD, Karniadakis GE. Non-invasive inference of thrombus material properties with physics-informed neural networks. *Comput Methods Appl Mech Eng*. 2021;375:113603.
59. Zhang R. Patient specific imaging-based modeling of vessels. Application to the assessment of pathologies in the thoracic aorta. *PhD thesis, CREATIS lab*, 2021.
60. Zheng Q, Delingette H, Duchateau N, Ayache N. 3D consistent & robust segmentation of cardiac images by deep learning with spatial propagation. *IEEE Trans Med Imaging* 2018;37:2137-48.
61. Zhou Y, Giffard-Roisin S, De Craene M, Camarasu-Pop S, D'Hooge J, Alessandrini M, Friboulet D, Sermesant M, Bernard O. A framework for the generation of realistic synthetic cardiac ultrasound and magnetic resonance imaging sequences from the same virtual patients. *IEEE Trans Med Imaging*. 2018;37:741-54.

Figure legends

Figure 7.1: (A) A simple elliptical model of the left ventricle is already very informative on the relationship between cardiac volume and myocardial deformation to maintain the same amount of ejected blood (the stroke volume), larger hearts need less contraction (myocardial strain). Adapted from Bijmens et al. (2012). (B) A simple classification model based on two well-chosen input variables is already very accurate to classify normal against diseased myocardial shapes⁷ (e.g., patients with hypertrophic cardiomyopathy): few or no improvements are observed with more complex classification models or more input variables, with the additional risk of overfitting the training data and the reduction of interpretability. Adapted from code available online⁸.

Figure 7.2: Predicting a continuous (1-dimensional) output variable y from a (1-dimensional) input variable x , using machine learning for regression (here, the multi-layer perceptron regressor from the scikit-learn⁹). Top row: schematic view of the multi-layer perceptron used to solve this problem. Its input and output layers consist of 2 neurons and 1 neuron, while the (hidden) layers consist of 100 neurons each. Bottom row: estimated regression for noisy sinusoidal data, using multi-layer perceptrons of increasing complexities with 0, 1, or 2 layers (0 layer corresponds to a single linear relationship between the input and the output).

Figure 7.3: (A) Sensitivity analysis for a 0D cardiac model: representation learning allows examining the distribution of the generated samples, and confirms their consistency with the changed hyperparameters. Adapted from Di Folco et al. (2021). (B) Fit between real and synthetic datasets of infarct images: representation learning allows visualising the distribution of both populations and their overlap, which confirms that the simulations have been correctly personalised to the real population. Adapted from Mom et al. (2021).

Figure 7.4: Examples of synthetic cases generated from a cardiac electromechanical model to validate a regression-based machine learning algorithm that localises myocardial infarct from myocardial deformation. The data consisted of the deformation patterns in the main anatomical directions (here, only radial strain displayed for the sake of simplicity), and the ground truth infarct locations, which were the labels used for regression. The columns on the right indicate the predicted infarct location and its uncertainty. Adapted from Duchateau et al. (2016).

Figure 7.5: (A) Overlapped view of the 3D mesh used in a cardiac electromechanical model and the corresponding synthetic magnetic resonance image that was generated from it. (B) Healthy and diseased variants that were generated from the same real image sequence. Adapted from Duchateau et al. (2018).

⁷ Beware that this result is relevant for this disease and data under study, as more complex models may work better in other situations.

⁸ <https://deepimaging2019.sciencesconf.org/resource/page/id/7>

⁹ https://scikit-learn.org/stable/modules/generated/sklearn.neural_network.MLPRegressor.html

Figure 7.1

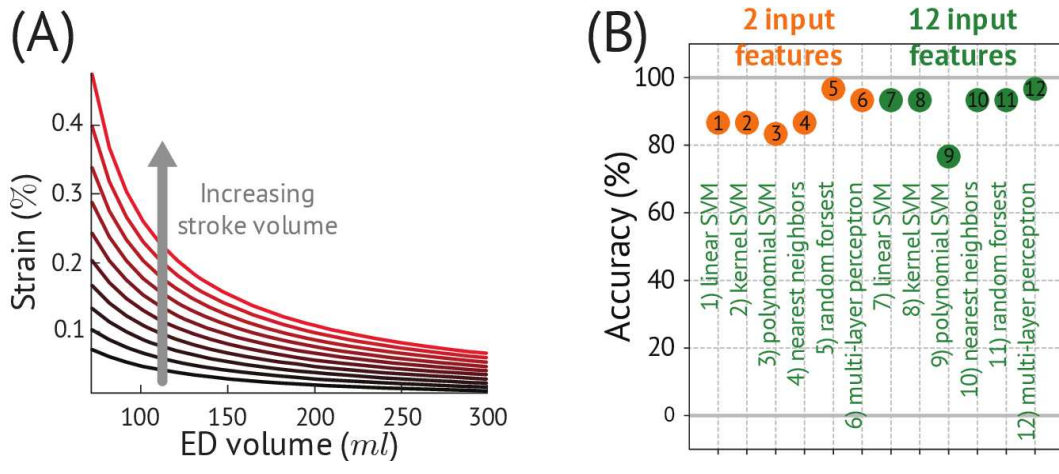


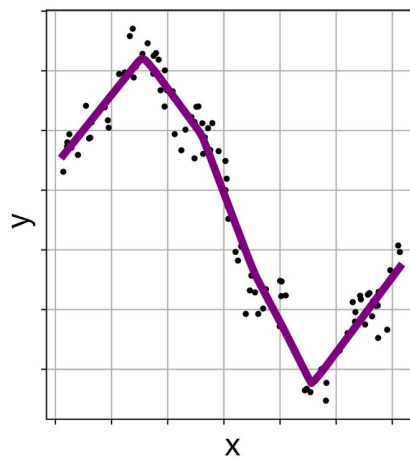



Figure 7.2



MLP - 0 layer



MLP - 1 layer



MLP - 2 layers

Figure 7.3





Figure 7.4




Figure 7.5

(A)



(B)

