PINNs and Inverse Problems using Non-Invasive Medical Imaging for Digital Twins

Presented by: Franny Dean*

Joint work with: Maggie Kuang*, Jack Benarroch Jedlicki*, David Ouyang, Anthony Philippakis, David Sontag, and Ahmed Alaa

March 1, 2024

Project Goal

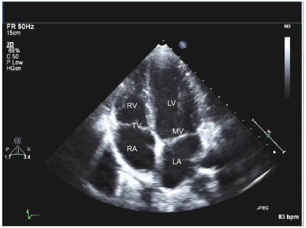

↓

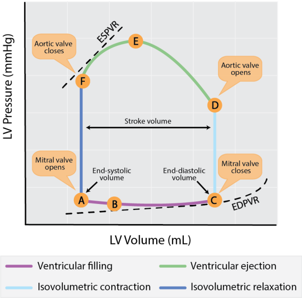

Motivations

- Non-invasive cardiovascular measurements
- Digital twins

Non-invasive cardiovascular measurements

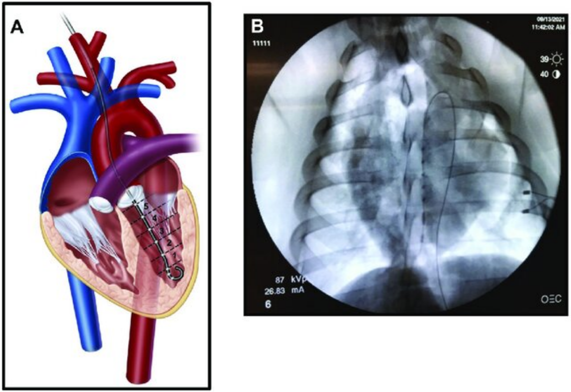

Figure: Placement of PV catheter for generating PV loops [SEA+22]

Non-invasive cardiovascular measurements

Figure: Ultrasound is a non-invasive technique to obtain patient data.

Non-invasive cardiovascular measurements

Figure: Ultrasound is a non-invasive technique to obtain patient data.

Digital Twins

Definition: Digital Twin

A **digital twin** is a virtual representation of a system that can be personalized with data and then modified to observe how the true system would change in response to interventions or over time.

- Concept being increasingly applied and developed for healthcare.

  [BSdSvdH18, BBE$^+$20, FDV$^+$20, OCNY23, WZL$^+$24, HJM$^+$19, CGMR20, CDLH21, Sub20, VID$^+$21, CAMM$^+$20, EMD20, SHL23]

Digital Twins

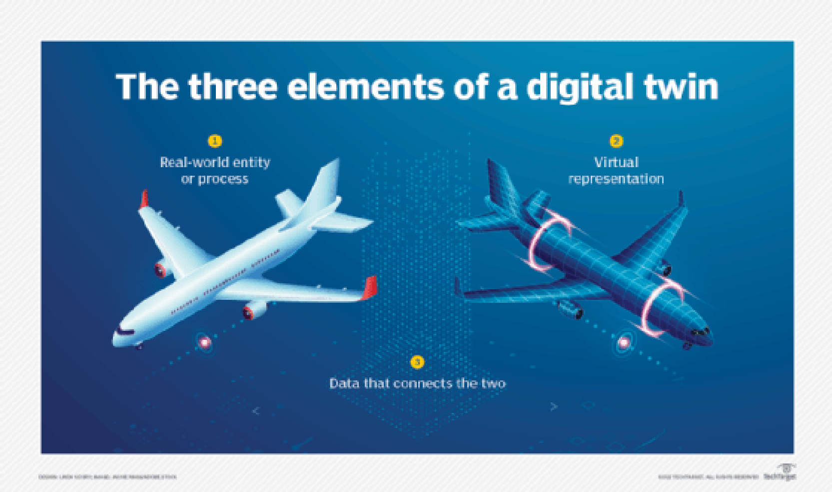

Digital Twins

A digital twin built from non-invasive medical images can reducing the cost and health burden of:

1. evaluating patient health status and prognosis

Digital Twins

A digital twin built from non-invasive medical images can reducing the cost and health burden of:

1. evaluating patient health status and prognosis
2. providing timely and accurate diagnoses

Digital Twins

A digital twin built from non-invasive medical images can reducing the cost and health burden of:

1. evaluating patient health status and prognosis
2. providing timely and accurate diagnoses
3. performing *in-silico* trials for determining correct treatments and guiding intervention

Physics-Informed Approach

With a lot of data, we can train a model to *implicitly learn* biophysical constraints.

In smaller data settings, biophysical models can be used to aide in model training:

- **Physics constrained learning** requires solutions to satisfy physical models (hard constraint).
- **Physics-informed neural networks (PINNs)** incorporate physical models in the loss function (soft constraint).

Physics-Informed Neural Networks

Data driven approach to

1. finding solutions **u** to differential equations (**forward model**), and
2. finding parameters θ that *discover* the differential equations (**inverse model**).

$$\frac{du}{dt} + f[\theta, u(x,t)] = 0$$

Physics-Informed Neural Networks

Data driven approach to

1. finding solutions **u** to differential equations (**forward model**), and
2. finding parameters θ that *discover* the differential equations (**inverse model**).

$$\frac{du}{dt} + f[\theta, u(x,t)] = 0$$

$$\mathcal{F} := \frac{du}{dt} + f[\theta, u(x,t)] \approx 0$$

Physics-Informed Neural Networks

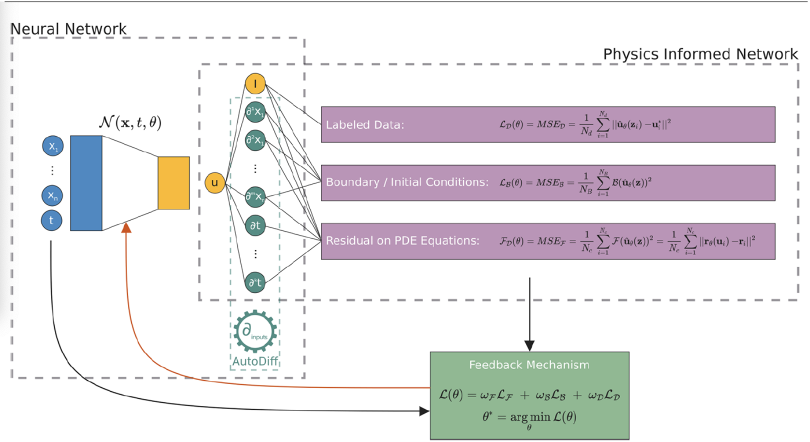

Figure: Physics-informed neural network [CDCG+22]

Physics-Informed Neural Networks

People usually cite Raissi et 2019 as pioneering the approach [RPK19].

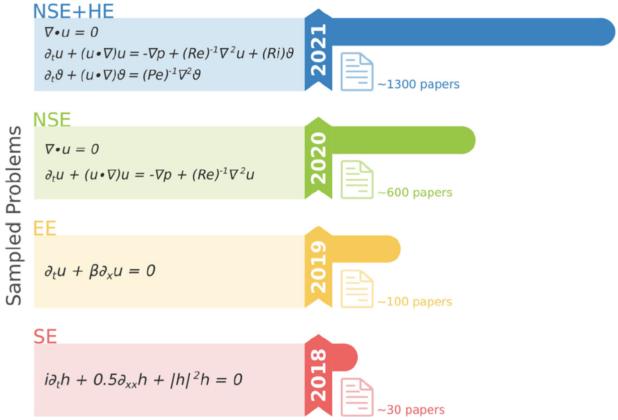

Figure: PINN literature is prolific [CDCG+22]

Physics-Informed Approach

We assume that cardiac pressures and volumes are governed by a 5 variable ***electric circuit model*** of the heart.

Patient specific parameters for the resulting ODE system give us a patient specific PV-loop (unique ODE solution).

Parameters $\longrightarrow$ 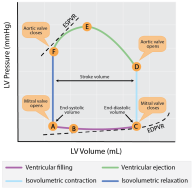

Electric Circuit Model of the Heart

Describe pressures as voltages:

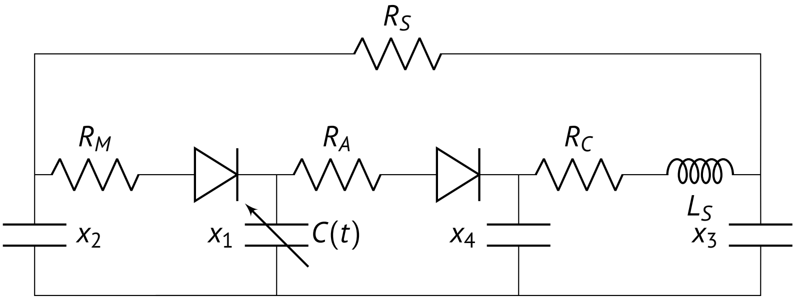

Let $\mathbf{x} = [x_1, x_2, x_3, x_4, x_5] = [P_{LV}(t), P_{LA}(t), P_A(t), P_{Ao}(t), Q(t)]$ describe the voltages x_1, x_2, x_3, x_4 and total flow x_5.

Volume-Pressure Relationship in Model

Elastance relates pressure and volume:

$$E(t) = \frac{P_{LV}(t)}{V_{LV}(t) - V_d}$$

Modeled as

$$E(t) = (E_{MAX} - E_{MIN}) \cdot 1.55 \cdot \left[\frac{\left(\frac{t_n}{0.7}\right)^{1.9}}{1 + \left(\frac{t_n}{0.7}\right)^{1.9}} \right] \cdot \left[\frac{1}{1 + \left(\frac{t_n}{1.17}\right)^{21.9}} \right] + E_{MIN},$$

with $t_n = t/T_{max}$ for $T_{max} = 0.2 + 0.15T_c$, and T_c the duration of a cardiac cycle.

Visualizations of this relationship with toy volume:
`https://www.desmos.com/calculator/dgfbaot4zf`

`https://www.desmos.com/calculator/4dfyy8ao9o`

ODE System Representing the Model

Let $\mathbf{x} = [V_{LV}(t) - V_d, x_2, x_3, x_4, x_5]$.

$$\mathbf{x}' = A_1\mathbf{x} + D_1 p(\mathbf{x}) \tag{1}$$

Where A_1 is a 5×5 time-independent matrix, and D_1 is a 5×2 time-independent matrix representing the activity of the diodes:

$$A_1 = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & \frac{-1}{R_S C_R} & \frac{1}{R_S C_R} & 0 & 0 \\ 0 & \frac{1}{R_S C_S} & \frac{-1}{R_S C_S} & 0 & \frac{1}{C_S} \\ 0 & 0 & 0 & 0 & \frac{-1}{C_A} \\ 0 & 0 & \frac{-1}{L_S} & \frac{1}{L_S} & \frac{-R_C}{L_S} \end{bmatrix} ; \; D_1 = \begin{bmatrix} 1 & -1 \\ \frac{-1}{C_R} & 0 \\ 0 & 0 \\ 0 & \frac{1}{C_A} \\ 0 & 0 \end{bmatrix} \tag{2}$$

With the vector $p(\mathbf{x})$ given by:

$$p(\mathbf{x}) = \begin{bmatrix} \frac{\max\{x_2 - x_1 \cdot E(t), 0\}}{R_M} \\ \frac{\max\{x_1 \cdot E(t) - x_4, 0\}}{R_A} \end{bmatrix} \tag{3}$$

Project Goal

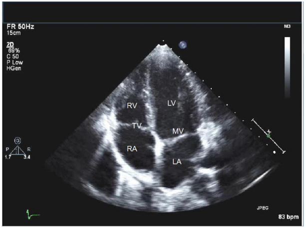

↓

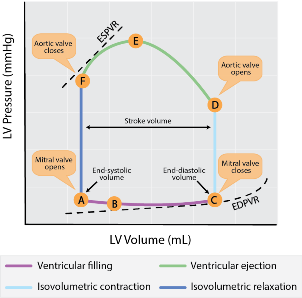

Two Inverse Problems

1. Obtaining **patient specific parameters** θ is an **inverse problem**:

$$\mathcal{M}(\theta) = \mathbf{x}$$

for ODE solution $\mathbf{x}$ and *forward model* $\mathcal{M}$ (obtaining the ODE solution).

Two Inverse Problems

1. Obtaining **patient specific parameters** θ is an **inverse problem**:

$$\mathcal{M}(\theta) = \mathbf{x}$$

for ODE solution $\mathbf{x}$ and *forward model* $\mathcal{M}$ (obtaining the ODE solution).

This inverse problem is **identifiable**:

- Existence and uniqueness of $\mathbf{x}$ are guarenteed by Picard-Lindelöf.
- Uniqueness of θ: We can show that in our case, *complete knowledge* of the system's state variables (pressures, volumes, and blood flow) allows us to **identify model parameters uniquely**.

It makes sense to talk about $\mathcal{M}^{-1}$ and $\mathcal{M}$.

Two Inverse Problems

2. We have a second inverse problem:

$$\mathcal{K}(\mathbf{x}) = \mathbf{y}$$

where unknown ***forward model*** $\mathcal{K}$ maps the cardiovascular system state $\mathbf{x}$ to the observed echo videos $\mathbf{y}$.

State Variables (Pressure, Volume, Flow) $\rightarrow$ 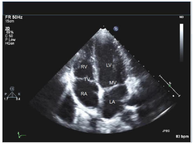

Two Inverse Problems

2. We have a second inverse problem:

$$\mathcal{K}(\mathbf{x}) = \mathbf{y}$$

where unknown ***forward model*** $\mathcal{K}$ maps the cardiovascular system state $\mathbf{x}$ to the observed echo videos $\mathbf{y}$.

State Variables (Pressure, Volume, Flow) $\rightarrow$ 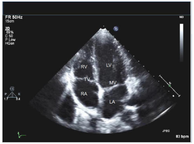

This problem is **ill-posed**. We may have many of the same $\mathbf{y}$ originating from different $\mathbf{x}$'s (not **injective**).

Two Inverse Problems

We combine the two inverse problems to

$$\mathcal{A}(\theta) = \mathbf{y}$$

for ultrasound images $\mathbf{y}$ and ODE parameters θ.

The unknown forward model $\mathcal{A}$ is the composition of $\mathcal{M}$ and $\mathcal{K}$.

Model Parameters $\rightarrow$ State Variables $\rightarrow$ 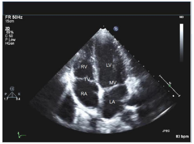

$$\mathcal{A}(\theta) = \mathcal{K}(\mathcal{M}(\theta)) = \mathbf{y}$$

If we can train a model to invert this task, we also obtain patient specific PV loops.

Physics-Informed Transfer Architecture for $\mathcal{A}^{-1}$

We have an inverse problem with *unknown forward model* $\mathcal{A}$ and partially labeled data $(\mathbf{y}_i, \tilde{\mathbf{x}}_i)$ where $\tilde{\mathbf{x}}_i$ are partial knowledge of $\mathbf{x}_i$ at some time points.

We train a **3D-CNN** to obtain a **generative** solution to the **inverse problem** $\mathcal{A}^{-1}$.

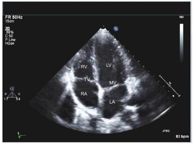

$\rightarrow$ State Variables $\rightarrow$ Model Parameters

$$\mathcal{M}^{-1}\mathcal{K}^{-1}(\mathbf{y}) = \theta$$

$$\mathcal{A}^{-1}(\mathbf{y}) = \theta$$

Physics-Informed Transfer Architecture for $\mathcal{A}^{-1}$

We have partial information about $\mathbf{x}_i$ associated to each video $\mathbf{y}_i$. This means we have information about the intermediate step of $\mathcal{A}$:

$$\mathbf{x} = \mathcal{K}^{-1}(\mathbf{y}) = \mathcal{M}(\theta).$$

Physics-Informed Transfer Architecture for $\mathcal{A}^{-1}$

We have partial information about $\mathbf{x}_i$ associated to each video $\mathbf{y}_i$. This means we have information about the intermediate step of $\mathcal{A}$:

$$\mathbf{x} = \mathcal{K}^{-1}(\mathbf{y}) = \mathcal{M}(\theta).$$

We use **fixed NN interpolator** approximation of $\mathcal{M}$ to model the dynamics of the electric circuit system.

Physics-Informed Transfer Architecture for $\mathcal{A}^{-1}$

We have partial information about $\mathbf{x}_i$ associated to each video $\mathbf{y}_i$. This means we have information about the intermediate step of $\mathcal{A}$:

$$\mathbf{x} = \mathcal{K}^{-1}(\mathbf{y}) = \mathcal{M}(\theta).$$

We use **fixed NN interpolator** approximation of $\mathcal{M}$ to model the dynamics of the electric circuit system.

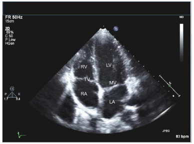

$\rightarrow$ Parameters $\theta \rightarrow \mathcal{M}(\theta) = \mathbf{x} \rightarrow$ Training Loss

Fixed NN interpolator

Our approach for incorporating physics:

- Use numerical ODE solvers to generate synthetic data: $(\theta_i, \tilde{\mathbf{x}}_i)$.
- Train a NN to learn the map $\mathcal{M} : \theta \to \mathbf{x}$ by minimizing loss between true and predicted $\tilde{\mathbf{x}}_i$.
- **Fix** the weights of this model.

Fixed NN interpolator

Our approach for incorporating physics:

- Use numerical ODE solvers to generate synthetic data: $(\theta_i, \tilde{\mathbf{x}}_i)$.
- Train a NN to learn the map $\mathcal{M} : \theta \to \mathbf{x}$ by minimizing loss between true and predicted $\tilde{\mathbf{x}}_i$.
- **Fix** the weights of this model.

PINN approach to inverse problems:

- Define the differential equation f as $\mathcal{F}(t) := \mathbf{x}' - A_1\mathbf{x} - D_1 p(\mathbf{x})$ for solution $\mathbf{x}(t)$.
- Train a NN to **approximate x** and learn parameters θ by minimizing $MSE_{\mathbf{x}} + MSE_{\mathcal{F}}$ (data loss + soft physics constraint).

Motivation for fixed NN interpolator approach

1. Resembles supervised learning used in input-output (IO) control theory systems to identify the effects of parameters on system dynamics.

Motivation for fixed NN interpolator approach

1. Resembles supervised learning used in input-output (IO) control theory systems to identify the effects of parameters on system dynamics.

2. Scales well with increasingly complex physical models (once interpolator is trained no increase in cost).

Details on Architecture and Loss

Step 1: Train a neural network interpolator to automate solving the 5x5 ODE system using synthetic data and output only $V_{LV}(t_{ES}), V_{LV}(t_{ED})$. This is the **NN Interpolator**.

$$[T_c, start_v, E_{\max}, E_{\min}, R_M, R_A, V_d] \rightarrow \text{solution: } \mathbf{x}(\mathbf{t}) \rightarrow V_{LV}(t_{ES}), V_{LV}(t_{ED})$$

Details on Architecture and Loss

Step 1: Train a neural network interpolator to automate solving the 5x5 ODE system using synthetic data and output only $V_{LV}(t_{ES}), V_{LV}(t_{ED})$. This is the **NN Interpolator**.

$$[T_c, start_v, E_{\max}, E_{\min}, R_M, R_A, V_d] \rightarrow \text{solution: } \mathbf{x}(\mathbf{t}) \rightarrow V_{LV}(t_{ES}), V_{LV}(t_{ED})$$

- 3,840 synthetic data points linearly sampled from realistic parameter ranges mapping parameters

Details on Architecture and Loss

Step 2:

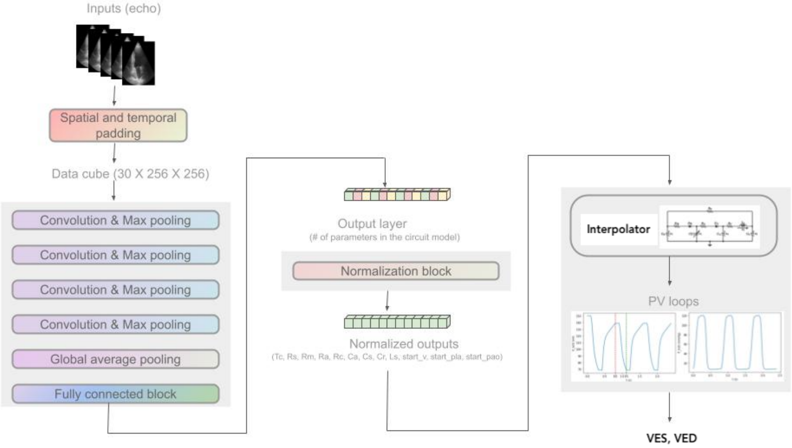

Experiments: Data and MAE

Data:

- **EchoNet:** 10,030 apical-4-chamber ultrasound videos from routine clinical care at Stanford University Hospital [OHG+20]
- **CAMUS:** 500 fully annotated cardiac ultrasound videos in 2-chamber view [Lec19]
- each with left-ventricle volume for end systole and diastole

Experiments: Data and MAE

Data:

- **EchoNet:** 10,030 apical-4-chamber ultrasound videos from routine clinical care at Stanford University Hospital [OHG+20]
- **CAMUS:** 500 fully annotated cardiac ultrasound videos in 2-chamber view [Lec19]
- each with left-ventricle volume for end systole and diastole

dataset	MAE (%)	MAE without PINN (%)
CAMUS	7.50	7.16
EchoNet	5.59	5.53

Table: Mean Absolute Error (MAE) Achieved with and without PINN: True vs. Simulated EF labels

Experiments: Personalized PV loops

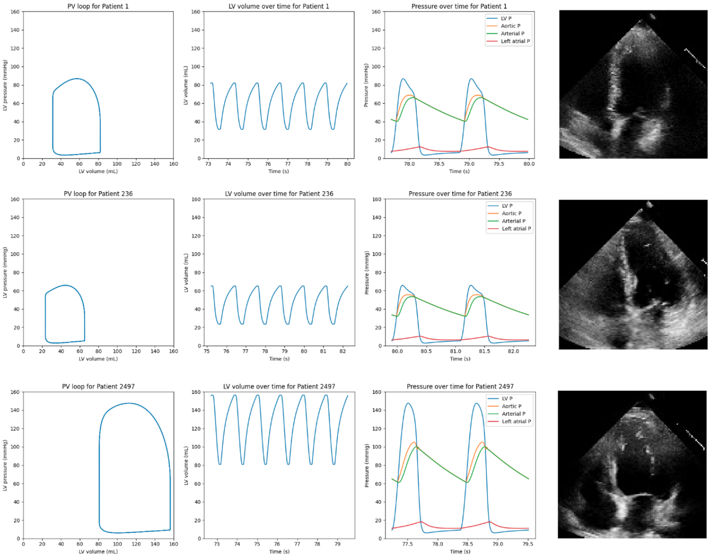

Experiments: Parameters predicted (EchoNet)

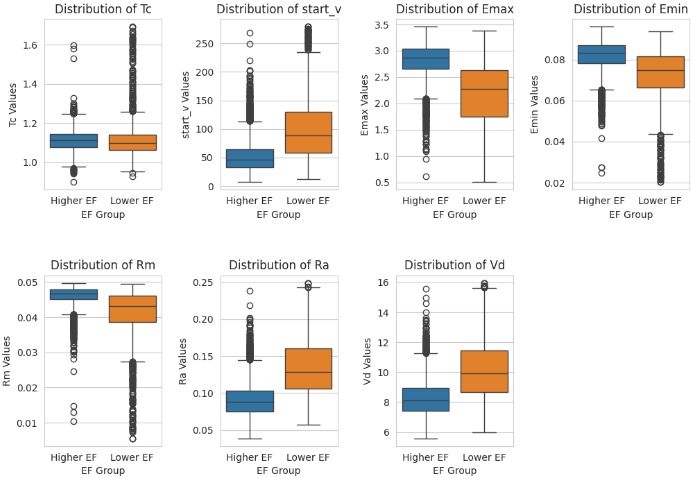

Experiments: Parameters predicted (CAMUS)

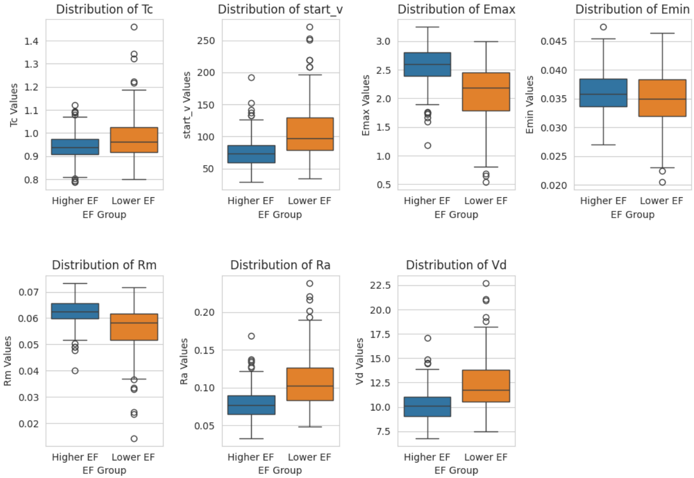

Experiments: Learning pressure differentials

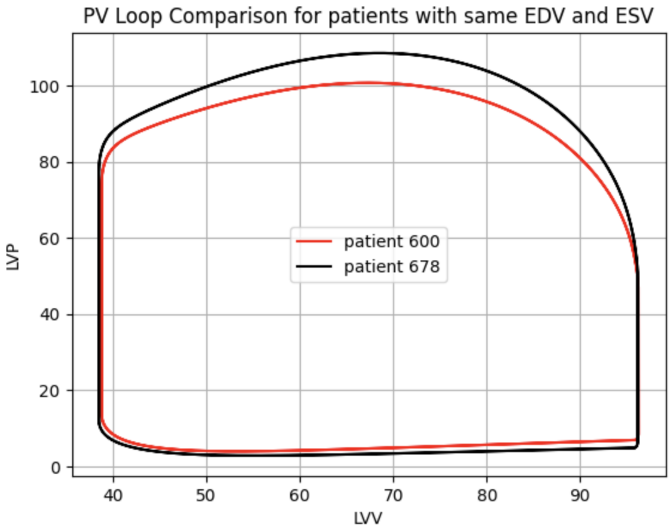

Experiments: Predicting disease labels

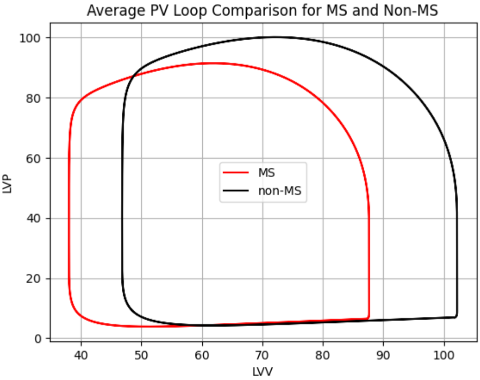

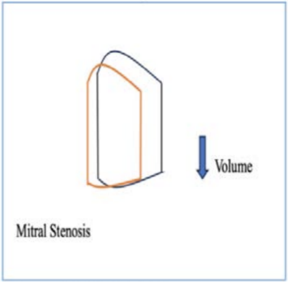

Experiments: AUC for predicting MS labels

Settings:

- Supervised learning task to predict MS labels directly from videos
- Predicting MS labels from learned parameters (trained with EDV and ESV loss)
- Predicting MS labels from learned parameters (trained with EF loss)

Experiments: AUC for predicting MS labels

Settings:

- Supervised learning task to predict MS labels directly from videos
- Predicting MS labels from learned parameters (trained with EDV and ESV loss)
- Predicting MS labels from learned parameters (trained with EF loss)

AUC	videos to MS	parameters to MS (V loss)	parameters to MS (EF loss)
EchoNet	0.8	0.59	0.5

Experiments: In-silico trial for left-ventricular assistance device (LVAD)

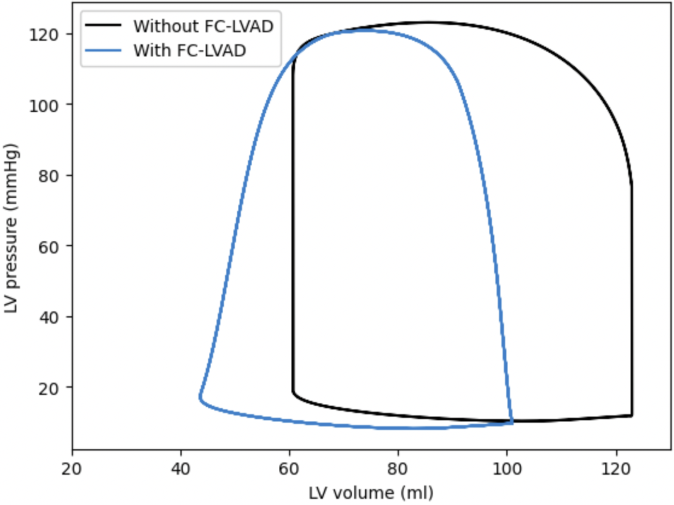

Experiments: In-silico trial for left-ventricular assistance device (LVAD) with different pumping speeds

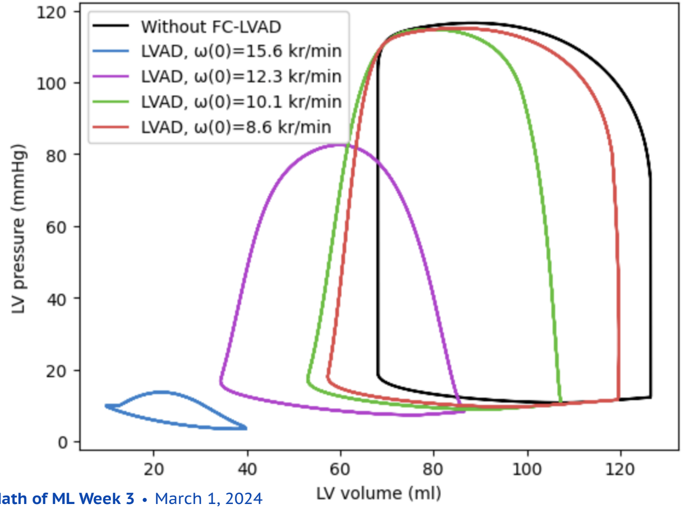

Summary of the Novelty of Our Approach

1. We demonstrate the ability to learn **patient specific, latent** ODE parameters.

Summary of the Novelty of Our Approach

1. We demonstrate the ability to learn **patient specific, latent** ODE parameters.

2. We show the utility of a **novel PINN-like approach** with *real medical image data directly*!

Summary of the Novelty of Our Approach

1. We demonstrate the ability to learn **patient specific, latent** ODE parameters.

2. We show the utility of a **novel PINN-like approach** with *real medical image data directly*!

3. Our fixed NN interpolator increases speed and convergence over using ODE solvers directly in model training.

References

[BBE+20] Bergthor Björnsson, Carl Borrebaeck, Nils Elander, Thomas Gasslander, Danuta R Gawel, Mika Gustafsson, Rebecka Jörnsten, Eun Jung Lee, Xinxiu Li, Sandra Lilja, et al.
Digital twins to personalize medicine.
Genome medicine, 12:1–4, 2020.

[BSdSvdH18] Koen Bruynseels, Filippo Santoni de Sio, and Jeroen van den Hoven.
Digital twins in health care: Ethical implications of an emerging engineering paradigm.
Frontiers in Genetics, 9, 2018.

[CAMM+20] Jorge Corral-Acero, Francesca Margara, Maciej Marciniak, Cristobal Rodero, Filip Loncaric, Yingjing Feng, Andrew Gilbert, Joao F Fernandes, Hassaan A Bukhari, Ali Wajdan, et al.
The ‘digital twin’to enable the vision of precision cardiology.
European heart journal, 41(48):4556–4564, 2020.

[CDCG+22] Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli.
Scientific machine learning through physics–informed neural networks: Where we are and what’s next.
Journal of Scientific Computing, 92(3):88, 2022.

[CDLH21] Michelangelo Canzoneri, Alessandro De Luca, and Jakob Harttung.
Digital twins: A general overview of the biopharma industry.
Digital Twins: Applications to the Design and Optimization of Bioprocesses, pages 167–184, 2021.

[CGMR20] Angelo Croatti, Matteo Gabellini, Sara Montagna, and Alessandro Ricci.
On the integration of agents and digital twins in healthcare.
Journal of Medical Systems, 44:1–8, 2020.

[EMD20] Tolga Erol, Arif Furkan Mendi, and Dilara Doğan.
The digital twin revolution in healthcare.
In *2020 4th International Symposium on Multidisciplinary Studies and Innovative Technologies (ISMSIT)*, pages 1–7, 2020.

[FDV+20] Marzia Di Filippo, Chiara Damiani, Marco Vanoni, Davide Maspero, Giancarlo Mauri, Lilia Alberghina, and Dario Pescini.
Single-cell Digital Twins for Cancer Preclinical Investigation, pages 331–343