

# Data-Driven Simulator for Mechanical Circulatory Support with Domain Adversarial Neural Process

**Sophia Sun**

*University of California, San Diego*

SHS066@UCSD.EDU

**Wenyuan Chen**

*Stanford University*

WENYUANC@STANFORD.EDU

**Zihao Zhou**

*University of California, San Diego*

ZIZ244@UCSD.EDU

**Sonia Fereidooni**

*University of California, San Diego*

SFEREIDOONI@UCSD.EDU

**Elise Jortberg**

*AbioMed*

EJORTBERG@ABIOMED.COM

**Rose Yu**

*University of California, San Diego*

ROSEYU@UCSD.EDU

## Abstract

We propose a data-driven simulator for Mechanical Circulatory Support (MCS) devices, implemented as a probabilistic deep sequence model. Existing mechanical simulators for MCS rely on oversimplifying assumptions and are insensitive to patient-specific behavior, limiting their applicability to real-world treatment scenarios. To address these shortcomings, our model Domain Adversarial Neural Process (DANP) employs a neural process architecture, allowing it to capture the probabilistic relationship between MCS pump levels and aortic pressure measurements with uncertainty. We use domain adversarial training to combine simulation data with real-world observations, resulting in a more realistic and diverse representation of potential outcomes. Empirical results with an improvement of 19% in non-stationary trend prediction establish DANP as an effective tool for clinicians to understand and make informed decisions regarding MCS patient treatment. The code for DANP is open-sourced at <https://github.com/Rose-STL-Lab/DANP>.

**Keywords:** neural process, domain adaptation, medical time series forecasting, physical simulator

## 1. Introduction

Machine learning has shown great promise to revolutionize patient care by improving outcome prediction and supporting clinical decisions (Davenport and Kalakota, 2019). However, it still faces significant challenges in learning accurate dynamics models from medical time series data, due to their non-stationary, uncertain, and confounding nature (Gottesman et al., 2019). There is also a distribution shift issue when sequentially suggesting optimal treatment plans, as the predicted future scenarios may not be in the training data.

In this paper, we aim to simulate the impact of Mechanical Circulatory Support (MCS) support levels on patient blood pressure. Forward flow MCS devices are mechanical devices designed to assist the heart in pumping blood from the left ventricle into the ascending aorta to deliver oxygenated blood to the body. They ensure that blood continues to circulate properly throughout the body when the heart fails to perform this function adequately. MCS devices play a pivotal role in maintaining blood pressure, maintaining organ perfusion and aiding heart muscle recovery.

Accurate simulation of MCS support is highly valuable for clinical decision support in treatment plan optimization. The simulator would allow physicians to understand how different MCS support levels would change blood pressure in a hypothetical setting. Traditionally, while clinicians can control the support level of MCS, there are limited recommendations and no universal protocols for optimally operating or weaning

the device (B. Geller, 2022). In practice, clinicians often rely on empirical experiences, which are difficult to transfer or standardize. A simulator would allow clinicians to ask “what-if” questions, providing critical information of plausible scenarios to consider before implementing the treatment change.

While there exist mechanical simulators for MCS support levels (Schampaert et al., 2011), they are based on over-simplified assumptions of the human arterial system (Bonfanti et al., 2017). Since the simulation only relies on mechanistic parameters as input, they are insensitive to patient-specific behavior (van de Meerakker et al., 2019), making them less applicable in real-life treatment practices. Furthermore, biomechanical simulators are deterministic and cannot capture the uncertainty common in the treatment of MCS patients. Finally, many of the simulation parameters are theoretical and not readily observable in regular practice. Directly training a simulation model on real-world data would suffer from the distribution shift issue as many of the risky treatments would not appear in the training set.

Our paper presents a data-driven simulator for Mechanical Circulatory Support (MCS) using data from the *Impella CP* device (Impella, 2014). We developed a novel model, Domain Adversarial Neural Process (DANP), to address these challenges. DANP employs an encoder-decoder neural process model to learn the probabilistic relationship between the MCS pump support level (P-level) and key aortic pressure measurements. The recurrent neural network architecture of the encoder allows it to take varying context lengths as input. It can be conditioned on future P-level, allowing the user to ask “what if” questions. The neural process architecture can capture the uncertainty in the forecast. To mitigate the gap between simulation data and real world observations, we use domain adversarial training (Ganin et al., 2016) to incorporate diverse simulation data and enable sim-to-real transfer.

Empirical results from DANP validate its efficacy, showcasing substantial improvements in forecasting accuracy and trend prediction compared to existing methods. The outcome forecasts generated by DANP prove to be more realistic and diverse, addressing the limitations of previous simulators. Our model acts effectively as a data-driven simulator and provides clinicians with an effective tool with confidence boundaries for optimal MCS treatment planning.

In summary, our contributions are:

- We propose a data-driven simulator for MCS devices, implemented as a probabilistic deep sequence model, domain-adversarial neural process (DANP).
- We present a novel sim-to-real transfer approach to integrate both simulated and real-world data from MCS devices. This allows our DANP model to answer “what-if” questions and generate diverse scenarios that are rarely seen in the data. The probabilistic nature of our model also helps clinicians to better address uncertainty in the treatment of MCS patients.
- We provide qualitative and quantitative experiment results to show the efficacy of our model. MAP forecasts by DANP are 5% more accurate overall and 19% more accurate for non-stationary data compared to the best-performing baseline. DANP’s more expressive and accurate forecasts can provide more insights compared to existing models and simulators.

## 2. Related Works

**Existing simulators used in MCS.** Traditional biomechanical simulators are based on over-simplified assumptions of the human arterial system (Bonfanti et al., 2017). A significant gap exists between what happens in a simulator and what works in the real world. Furthermore, biomechanical simulators are deterministic. They cannot capture the uncertainty common in the treatment of MCS patients. Directly training a simulation model on real-world data would suffer from the distribution shift issue, as many of the risky treatments would not appear in the training set.

Designing a hypothesis-driven healthcare simulator that encompasses all possible information is challenging (Gottesman et al., 2019), and even leading synthetic data generation methods may miss some deviations,

leading to invalid data (Chen et al., 2019). For instance, Prakosa et al. (2012) uses the traditional motion simulator to generate visually realistic cardiac images, while Mahmood et al. (2018) improves the simulator trained on synthetic data using rendered real-world data. However, less attention has been paid to improving the traditional cardiac motion model’s ability to capture accurately. While intricate simulators have been developed for hemodynamic modeling of patient phenotypes and cardiovascular treatment (D. Doshi, 2015), they rely on strong modeling assumptions and require clinical validation, necessitating the development of improved approaches that bridge the gap between mechanistic (hypothesis-driven) and phenomenological (real-world data-driven) models within the healthcare domain (Frangi et al., 2018).

**Time series forecasting in the medical applications.** Most works that address medical time-series forecasting focus their efforts on large-scale electronic health records (EHRs); we refer the readers to (Morid et al., 2023) for a detailed survey. Many works use Gaussian processes to impute and model multivariate medical time series data (Schulam and Saria, 2015; Cheng et al., 2017; Prasad et al., 2017) because they naturally capture irregular time series observations and can estimate prediction uncertainties in a probabilistic framework (Roberts et al., 2013). In this work, we use a Neural Process model (Garnelo et al., 2018) that combines Gaussian processes with deep neural networks. It inherits many desirable properties of GP but is more expressive and efficient to optimize.

**Sim-to-real adaptation.** Developing data-driven simulators has been mostly prominent in robotic applications (Zhao et al., 2020). Recent works have shown that by transferring knowledge learned in simulation towards their deployment in the real world, domain adaptation can improve the model’s convergence and robustness for a diverse set of tasks (Cheng et al., 2019). Common approaches include meta learning (Arndt et al., 2020), continual learning (Traoré et al., 2019), and using domain-adversarial training (Ganin et al., 2016) to learn domain-agnostic feature representations. The studies of sim-to-real transfer to build better simulators for medical applications remain scarce.

**Background of Impella.** The Impella CP is a left ventricular assist device (LVAD) that supports the heart’s function via a percutaneously inserted catheter. It pumps blood from the left ventricle (LV) to the ascending aorta. The device has positive effects on cardiovascular parameters, including cardiac output, coronary blood flow, LV end-diastolic pressure, and myocardial workload. The Automated Impella Controller (AIC) provides continuous high-resolution data at 25 Hz to monitor aortic pressure, motor current, motor speed, and pump flow. One of the primary indications for the use of Impella is during high-risk percutaneous coronary interventions (HR-PCI) (Impella, 2014). A PCI is a cardiac catheterization lab-based procedure to address coronary artery disease (CAD). PCIs considered high risk are a function of cardiovascular risk factors, including the complexity of CAD, hemodynamic compromise, and patient comorbidities. The Impella CP is used in HR-PCI for better clinical outcomes and to avert hemodynamic instability.

### 3. Problem Setting and Data

Our work models the effect of different levels of MCS support on patients’ hemodynamic status. We provide a visualization of the system of interest in Figure 1 (left). The main variable to be predicted is future Mean Aortic Pressure (MAP), a key hemodynamic parameter that serves as a critical endpoint for assessing organ perfusion and determining the readiness for de-escalation of support. The desiderata of our data-driven simulator are: (1) to accurately forecast future MAP based on past data, and (2) to be able to reason about alternative input of pump speed and produce sensible hypothetical outcomes.

Mathematically, the simulator is a function that maps the historical MAP along with other hemodynamic features and the proposed future P-levels to future MAP values. We can define the problem as a multi-step time series forecasting task. Let  $\mathbf{x} = x^{1:k} = (x^1, \dots, x^k) \in \mathbb{R}^{k \times d_x}$  denote the  $k$  time steps of  $d_x$  input features,  $\mathbf{pl} = pl^{1:h} = (pl^1, \dots, pl^h) \in \{1, \dots, 9\}^h$  the hypothetical pump support level (P-level) of the

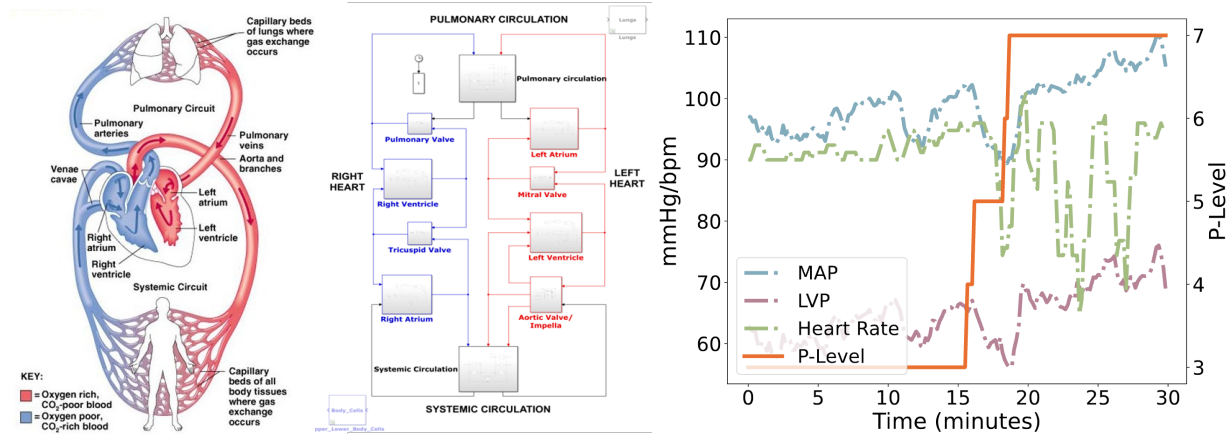


Figure 1: **Left:** Cardiovascular System overview (adapted from Pearson’s “Essentials of Human Anatomy and Physiology” Ch. 11 “The Cardiovascular System”). **Middle:** Translation to Lumped Parameter circuitry model, the source of our simulated data. **Right:** Visualizations of a patient data sample. Note that as the clinician increases the **pump level**, the **mean aortic pressure (MAP)** increases as a result.

future  $h$  steps, and  $\mathbf{y} = \mathbf{y}^{1:h} = (y^1, \dots, y^h) \in \mathbb{R}^h$  as the MAP values of the future  $h$  time steps. The superscript indicates the data time step. Our model learns the function:

$$f : x^{1:k}, pl^{1:h} \rightarrow y^{1:h} \tag{1}$$

The model takes an input sequence of 15 minutes ( $k = 90$ ) and outputs a sequence of the next 15 minutes ( $h = 90$ ). The horizon lengths and are chosen based on suggestions from physicians. For training and testing our models, the “future” support level aligns with the actions taken by physicians in real-life scenarios, enabling us to calculate the errors in MAP prediction retrospectively. We selected the following  $d_x = 7$  features as simulator input, based on medical relevance and their availability for potential real-time decision-making. Figure 1 (right) visualizes some features of a real data sample.

1. **Aortic Pressure (MAP):** The Aortic Pressure (mmHg) is the blood pressure at the aorta’s root and is directly measured by an optical sensor on the Impella CP’s outlet shown in Figure ?? in the Appendix.
2. **P-level (Motor Speed):** Impella operates at 9 different speed levels, from P1-P9, each with a constant motor speed (rpm). The P-level proportionally determines the blood flow provided to the patient by the motor’s speed and current. Clinicians can control the P-level while the patient is on support.
3. **Pump Flow:** (L/min) is the calculated forward blood flow through the Impella cannula.
4. **Left Ventricular Pressure (LVP):** (mmHg), the pressure inside the heart’s left ventricle during each cardiac cycle. (AbioMed, 2020).
5. **Heart Rate:** (BPM) Heart rate is estimated from Aortic Pressure over the cardiac cycle via a smooth filter and Fourier transform.
6. **Tau LV:** (ms) Time constant of isovolumic relaxation ( $\tau$ ) representing the time LVP takes to reduce by  $1/e$  from aortic valve closure. Estimated from an exponential fit on the downstroke of the LVP curve.
7. **Contractility:** A measure of load-dependent ventricular contractility (mmHg/s), estimated using a linear transformation of the maximum rate of LV pressure rise.

### 3.1. Simulated Data

We utilize data generated from an electrical circuit-inspired simulator. The Windkessel Lumped Parameter circuitry simulator (T. Heldt, 2001) models cardiac function through a parameterized set of differential equations. In the electrical circuit analogy, potentials represent the differential blood pressure, and currents represent blood flow, see the chart in the center of Figure 1.

Among the input parameters in the cardiac simulator, four were determined to be derivable from sensor data within an Impella CP: P-level, heart rate, contractility, and the time constant of relaxation (Tau). A grid search was performed across the physiological ranges for these four parameters and passed into the simulator. The simulator outputs included MAP, pump flow rate, and the Left Ventricular Pressure (LVP). An artificial time series was created from the simulator point estimates by sampling from the simulator with random noise. Data augmentation techniques of shifting and rescaling were performed on the time-series samples to imitate the effect of unused input parameters from the simulator.

### 3.2. Real world data

The simulator alone does not adequately address clinical decision support during patient treatment, as many of the input parameters are theoretical and not readily observable in regular practice. In training our data-driven simulator, we perform a sim-to-real transfer using both simulated data and real data.

Real-world data was sourced from 211 High-Risk Percutaneous Coronary Intervention (HR-PCI) patients. Among the features, heart rate, end-systolic elastance, and the time constant of relaxation  $\tau$  were derived features from the 25Hz signals. The raw waveform data from the Impella pump are sampled at 25 Hz, and we process the raw waveform data into a rolling average of 0.1 Hz time series. We filter out samples with outlier features and those with medication or intervention alarms to avoid confounding. In total, the HR-PCI data set consists of 43,814 30-minute-long time series.

## 4. Domain Adversarial Neural Process (DANP)

We present Domain Adversarial Neural Process (DANP), a probabilistic deep sequence model for MCS simulation. We tailor the approach with a custom-designed Neural Process encoder and a sequence decoder for time series forecasting. Our method effectively bridges the domain differences between simulation and real-world data, acting as an adept data-driven simulator.

**Domain-Adversarial Training.** The high-level idea behind domain-adversarial training Ganin et al. (2016) is to learn domain-invariant features by simultaneously optimizing two tasks: the primary prediction task (e.g., classification/regression) and a domain classification task. This is achieved through a shared feature extractor followed by a Gradient Reversal Layer (GRL) that encourages the network to learn features invariant to domain differences. In the context of our application, the simulation data helps the model learn from a diverse set of scenarios not present in the real dataset, without overly dominating the training data. The objective function for domain adversarial training is as follows:

$$L_{\text{DANP}} = L_y + \lambda \cdot L_d \quad (2)$$

where  $L_y$  is the task loss in the target domain (i.e. the task of forecasting MAP),  $L_d$  is the domain classification loss (between simulation and real data), and  $\lambda$  is a hyperparameter controlling the trade-off between domain adaptation and task performance.

**Neural Process.** Neural processes (NPs) (Garnelo et al., 2018) are a class of models that parameterize a stochastic process with neural networks. They are structured similarly to a Variational Autoencoder (VAE) (Kingma and Welling, 2013) where a latent variable model is employed to estimate the conditional distribution

over the outputs of the unlabeled (target) points given the set of labeled (context) points. We show the model architecture of DANP in Figure 2, for which we will discuss in detail in the following sections.

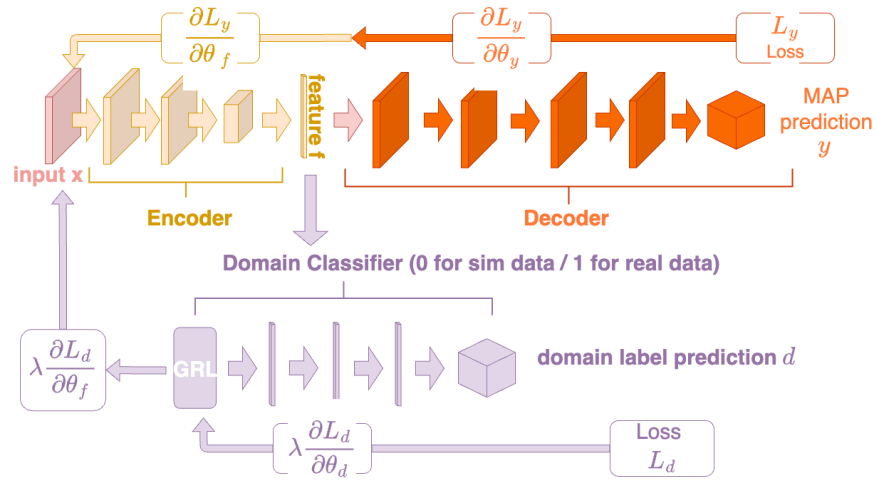


Figure 2: Architecture of DANP with the neural process **encoder**, **decoder**, and a **domain classifier**. Input  $x$  (along with P-level  $pl$ , combined for conciseness) is passed into the **encoder**. The output (feature  $z$ ) is then passed separately into the **decoder** and **domain classifier** to obtain regression  $y$  prediction (MAP) and domain prediction  $d$  (0 or 1), respectively. The Gradient Reversal Layer (GRL) for the domain classifier acts as an identity transform. The loss of  $y$  prediction ( $L_y$ ) and domain prediction ( $L_d$ ) will be calculated separately but backpropagated together.

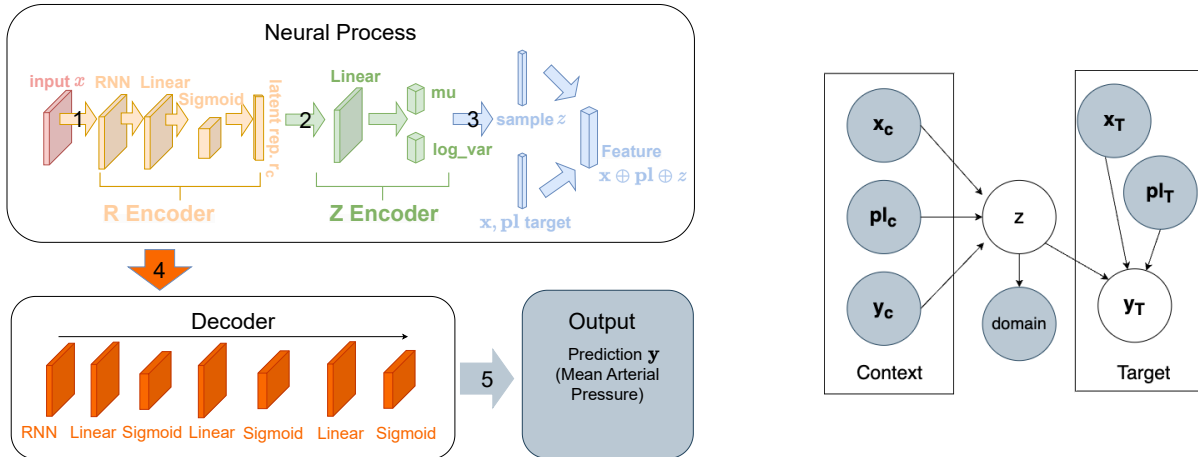


Figure 3: **Left:** Neural Process Architecture Details and **Decoder** Details. The input data  $x$  is initially processed by an **R Encoder** to obtain its latent representation. This latent representation is then further processed by a **Z Encoder**, which parametrizes it into the **latent space**  $z$ . When concatenated with the target  $x$ , the sampled values from this latent space result in a combined feature  $x_T \oplus pl_T \oplus z$ . This is passed through a sequential **Decoder**, where it undergoes a series of transformations to ultimately generate predictions for MAP. **Right:** Graphical model of DANP. A grey background indicates that the variable is observed.

**Neural Process Encoder.** DANP’s neural process encoder converts the input sequence into latent representations as the **Decoder** in Figure 2. The neural process model captures the intrinsic stochasticity of the data in

a distribution over latent variables. This unique composition permits the neural process to derive a predictive distribution over functions, reflecting the uncertainty in patient data.

Specifically, the neural process encoder model learns a variational posterior distribution  $q_\phi(z|\mathbf{x}_c, \mathbf{pl}_c, \mathbf{y}_c)$  over a latent variable  $z$ , based on a set of context (denoted as  $\mathcal{C}$ ):  $(\mathbf{x}_c, \mathbf{pl}_c, \mathbf{y}_c) = \{(\mathbf{x}_i, \mathbf{pl}_i, \mathbf{y}_i)\}_{i \in \mathcal{C}}$ , and parameterised by a neural network  $\phi$ . In training, each batch of data is randomly split 0.9/0.1 into context and target data. We use subscript here to index samples, and  $i \in \mathcal{C}$  indicates that the sample is used as a context point. At the testing stage, all training data are used as context samples, and the test data as target samples. We adopt a two two-part encoder design consisting of a sequential model and a distribution embedding model, as illustrated detailedly in Figure 3. The **R-Encoder** is a recurrent neural network (RNN) sequential model that transforms the sequence data into a representation vector  $r$ . Other than dimension reduction, the recurrent architecture has the advantage of enabling us to potentially encode longer or shorter history for prediction. The **R-Encoder** takes in the context data  $(\mathbf{x}_c, \mathbf{pl}_c, \mathbf{y}_c)$ , target point inputs  $(\mathbf{x}_T, \mathbf{pl}_T)$ , and yields the latent function representation  $r$  as depicted in Figure 3.

This latent representation  $r$  then serves as input to the **Z-Encoder** implemented as a multi-layer perceptron. We assume that  $z$  follows a multivariate standard normal distribution that captures the uncertainty in the function space. **Z-Encoder** then outputs the mean  $\mu_z$  and variance  $\sigma_z^2$  of the variational distribution:

$$q(z|\mathbf{x}_c, \mathbf{pl}_c, \mathbf{y}_c) = \mathcal{N}(\mu_z(r), \sigma_z^2(r))$$

The concatenated feature of  $\mathbf{x}_T \oplus \mathbf{pl}_T \oplus z$  plays a dual role: it is used by the Domain Classifier to distinguish between simulation and real data, and also by the decoder to make predictions on  $\mathbf{y}$ . The domain classifier utilizes a gradient reversal function (GRL) on feature  $\mathbf{x}_T \oplus \mathbf{pl}_T \oplus z$  before directing the output to a domain classifier, as shown in Figure 2.

**Sequential Decoder.** We captured the complexity of the prediction with the **Sequential Decoder** in Figure 2 (also shown in more detail as the **Decoder** of Figure 3), a neural network model  $\theta$  that parameterized

$$p_\theta(y_{n+1}|\mathbf{x}_T, \mathbf{pl}_T, z) = \mathcal{N}(\mu_\theta(\mathbf{x}_T \oplus \mathbf{pl}_T \oplus z), \sigma_\theta^2(\mathbf{x}_T \oplus \mathbf{pl}_T \oplus z))$$

for a new observation  $y_{n+1}$ . The **Sequential Decoder** is implemented as two RNN layers, followed by four linear layers with a sigmoid activation function.

In the training phase, the neural process model maximizes the probability over the evidence lower bound (ELBO) loss, defined as

$$\text{ELBO}(\theta, q) = \mathbb{E}_{z \sim q(z; \mathbf{x}_c, \mathbf{pl}_c, \mathbf{y}_c)} [\log p(\mathbf{x}_T, \mathbf{pl}_T | z; \theta)] - \text{KL}[q(z; \mathbf{x}_c, \mathbf{pl}_c, \mathbf{y}_c) || p(z | \mathbf{x}_T, \mathbf{pl}_T; \theta)]$$

where  $\theta$  are the model parameters,  $z$  indicates latent variables, and  $\mathbf{x}, \mathbf{pl}, \mathbf{y}$  represents observed data in the context and target sets. This is the  $L_y$  loss in equation 2. During inference, a vector  $z$  is sampled from encoded  $p(z)$ . The architecture resonates with the depiction in Figure 3, wherein the  $\mathbf{x}$  target within the neural process amalgamates features of the input sequence and the future sequence’s feature - the P-level, to generate the feature  $\mathbf{x}_T \oplus \mathbf{pl}_T \oplus z$ . This is then channeled to the decoder to predict  $\mathbf{y}_T$ .

**Domain Classifier.** As demonstrated in the purple part in Figure 2, the real data is labeled as 1, and the sim data is labeled as 0. Similar to the sequential decoder, the feature  $\mathbf{x}_T \oplus \mathbf{pl}_T \oplus z$ , extracted by the Neural Process, undergoes a series of transformations. It begins with a gradient reversal layer(GRL) that multiplies the gradient by a certain negative constant during back-propagation. It effectively makes the network learn features that are insensitive to domain-specific variations. The output from GRL is passed to a series of transformations in the domain classifier, which is responsible for classifying the domain of the input data. The domain classifier learns to distinguish between features from different domains.

To train the model, we employ the loss function in equation 2. First, the ELBO loss, represented as  $L_y$ , is employed to reduce the gap between the model’s predictions and the actual MAP values. It is the sum of the

reconstruction loss, Mean Absolute Error (MAE), and the regularization term, the KL divergence between the approximate posterior and the prior of  $\mathbf{z}$ . The domain classification loss  $L_d$  is a Negative Log Likelihood Loss (NLLLoss) applied to the predicted domain labels. We aggregate these losses and back-propagate them jointly, as illustrated in Figure 2. This training approach equips the sim-to-real model with the capability to generate outputs that closely mirror real-world scenarios, exhibiting minimal error with uncertainty.

## 5. Experiments

### 5.1. Baselines

We compare our model to the following baselines:

- MLP: Multi-layer perceptron neural network that takes all the input time series data and outputs the prediction both as a flattened vector.
- CLMU: Li et al. (2022) Conditional Legendre Memory Unit (LMU) that employs an LMU encoder (Voelker et al., 2019) with a neural process decoder trained on both simulation and real data. The conditional LMU has a fixed initial hidden state based on the MAP standard deviations from the specific domain and encodes the initial time steps of a sample. During training, real data was incorporated, and the label “simulated” or “real” was passed as a feature to the CLMU encoder as a condition.
- NP direct transfer: DANP’s neural process model, trained only on *simulation* data, without domain adaptation.
- NP no sim: DANP’s neural process model, trained only on *real* data, without domain adaptation.

### 5.2. Experiment Results

**Prediction Accuracy.** We present the result of our experiments in Table 1. To better evaluate the behavior and expressively of the models, we split the test data into three trend categories: increasing (inc), decreasing (dec), or stationary (stat), and report the MAE calculated on each category as well. In both of our training and testing sets, stationary streams make up for around 65% of the data, increasing trends 20%, and decreasing trends 15%. We refer the readers to the appendix <sup>1</sup> for details on how the trends are calculated.

Overall, DANP showed superior accuracy in MAP prediction in all our metrics, improving 20% over the best baselines for increase-trending data and 13% in terms of trend prediction accuracy. Such improvements are significant because the non-stationary scenarios are clinically actionable and for which an accurate simulator is most helpful. We also show in the appendix a data distribution analysis showing that DANP is able to better match the distribution of real data compared to the baseline models.

Method	MAE (mmHg) ↓	MAE (inc) ↓	MAE (dec) ↓	MAE (stat) ↓	Trend Acc ↑
MLP	7.97 ± .26	9.04 ± .68	10.96 ± .61	6.78 ± .43	0.57 ± .03
CLMU	6.93 ± .11	8.65 ± .56	8.47 ± .24	5.51 ± .04	0.65 ± .01
NP direct transfer	7.36 ± .91	9.72 ± 1.23	8.79 ± 1.06	6.25 ± .95	0.64 ± .00
NP no sim	8.68 ± .06	<b>6.90</b> ± .01	15.34 ± .02	7.63 ± .01	0.52 ± .00
DANP (ours)	<b>6.65</b> ± .13	<b>6.94</b> ± .10	<b>8.46</b> ± .17	<b>5.36</b> ± .09	<b>0.70</b> ± .01

Table 1: Empirical results in terms of Mean Average Error (MAE) for data with increasing (inc), decreasing (dec), stationary (stat) trends, and trend prediction accuracy. DANP achieves significantly lower performs significantly better on trending data compared to baselines.

1. Please refer to <https://github.com/Rose-STL-Lab/DANP> for full paper with appendix.



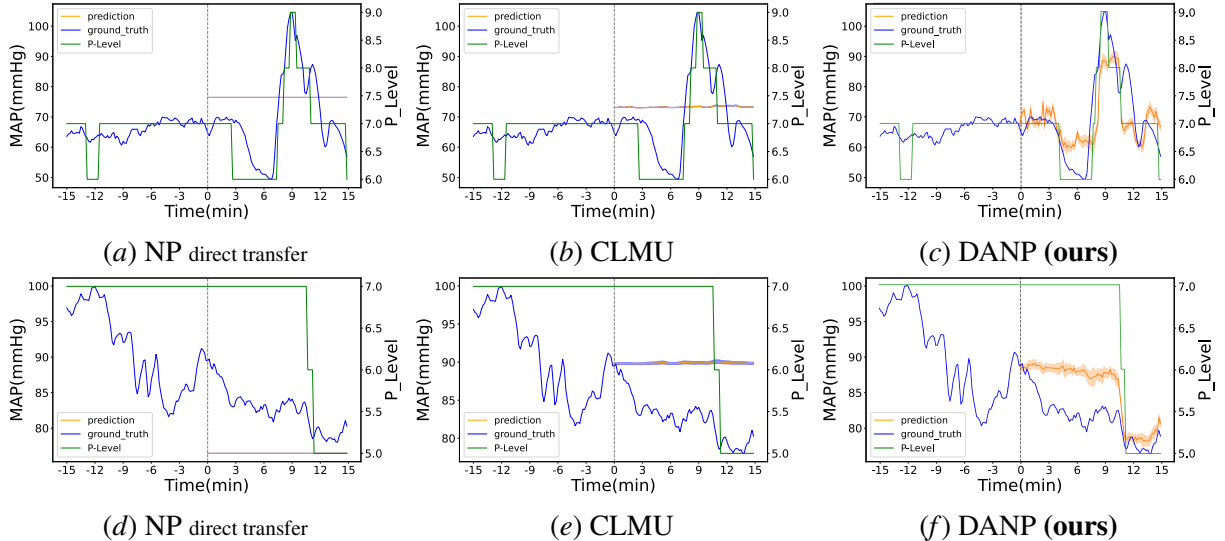


Figure 4: Sample Prediction using Model Trained and Tested on HR-PCI Cohort. In scenarios where both direct transfer and CLMU can only produce flat forecasts, DANP is able to produce sensible MAP that corresponds to p-level shifts.

**Scenario Creation.** We show two qualitative examples in Figure 4 to illustrate DANP’s ability for “what-if” scenario creation. The first row (a)-(c) shows an increase of support level varying from P6-P9 and a simultaneous change in MAP. Only the DANP model is able to predict an increase in MAP in response to the increase in support. The second row (d)-(f) has a decreasing trend and support reduction from P7 to P5. Again, DT and CLMU predict a straight line and do not capture the trend over the forecast window, while the DANP captures the true trend in response to the p-level drop. These two samples illustrate the much-improved expressiveness of our DANP compared to existing baselines, and that DANP is able to produce sensible MAP in response to P-level shifts.

### 5.3. Ablation Studies

We compare DANP with other alternative approaches to address the distribution shift problem of learning from real datasets. For the neural process model (NP no sim), we experimented with the following modifications to address the challenges presented by our data.

- `seq model`: The R-Encoder in our architecture encodes varying-length context into a representation vector, in addition to the neural-process-specific Z-Encoder.
- `sampling`: subsampling training data to balance the trend distribution. (Branco et al., 2016)
- `DA`: domain adversarial training with simulation data, see section 4.1 for details.
- `meta-regression`: We explore the meta-regression settings that are common to neural process baselines (Garnelo et al., 2018; Nguyen and Grover, 2022), where each time series is considered a unique function with its own context and predictions, instead of learning the problem as one single function as in Equation 1 in our setting. In the meta-regression setting,  $x$  of the neural process is the time stamp and P-level, and  $y$  is the hemodynamic features. The input time steps are used as context points to obtain the neural process that produces forecasts for future time steps.

Our ablation experiment result shows that domain adversarial training (DA) allows the model to significantly improve in expressiveness, evident in the high trend accuracy and lower MAE values for increasing and

Method	MAE (mmHg) ↓	MAE (inc) ↓	MAE (dec) ↓	MAE (stat) ↓	Trend Acc ↑
NP no sim	8.68 ± 0.06	6.90 ± .01	15.34 ± .02	7.63 ± .01	0.52 ± .00
NP sim + real	7.55 ± .87	9.23 ± .93	8.91 ± .77	6.45 ± .89	0.61 ± .02
– seq model	7.19 ± .10	8.73 ± 1.26	10.28 ± 1.42	5.92 ± .02	0.04 ± .05
+ sampling	8.84 ± .72	9.49 ± .91	11.86 ± 2.05	7.85 ± .64	0.45 ± .18
+ DA + sampling	8.90 ± .83	8.07 ± .58	14.72 ± 1.13	7.91 ± .71	0.46 ± .12
+ DA (final)	6.65 ± .13	<b>6.94</b> ± .10	<b>8.46</b> ± .17	5.36 ± .09	<b>0.70</b> ± .01
meta-regression	<b>5.57</b> ± 0.20	8.19 ± 1.23	9.74 ± 1.06	<b>4.23</b> ± .95	0.70 ± .08

Table 2: Ablation study on the current architecture. The + and – signs indicate addition and deletion of the modifications. DANP (2nd to last row) was selected as our final model for its expressivity and prediction accuracy on non-stationary data.

decreasing trend data. Subsampling, an otherwise popular method to handle imbalanced data in machine learning, proved to severely harm prediction accuracy. This is potentially due to over-sampling noisy data in our datasets, which are prevalent among the non-stationary time series. The meta-regression model exhibits the behavior of capturing the dominant stationarity in the dataset and under-performing on the more clinically important non-stationary data.

## 6. Conclusion and Discussion

In this paper, we developed a data-driven simulator to predict MAP time series given different levels of pump speed in mechanical circulatory support (MCS) devices. Our innovation lies in the introduction of a Domain-Adversarial Neural Process (DANP), a generative probabilistic deep sequence model, which integrates both simulation data from the Lumped Parameter cardiac simulator with real-world sensor data from Impella CP. This integrative approach not only facilitates simulating various scenarios, including those rarely observed in the data, but also allows the DANP model to be able to answer “what-if” questions regarding pump level changes and subsequent patient MAP response. This enhances DANP model’s applicability in clinical settings for MCS patients. Furthermore, incorporating both the diversity of simulation data and the expressiveness of the real-world data, the DANP model demonstrates superior accuracy, with a 5% overall improvement and a 19% increase in accuracy in non-stationary data compared to the best-performing baselines. Its performance offers more insightful and expressive forecasts than existing models and simulators. These attributes make the DANP model an invaluable tool in scenarios where detecting dynamic and non-stationary change is crucial.

**Future Works.** Additional medical metrics are required to evaluate DANP’s performance appropriately and assess its generalizability and robustness before applying them in clinical settings.

One direction to improve our model is to incorporate other modes of data to create a fuller picture of patient status to decide how much MCS support is needed. This could include doctor notes (text), lab values (tabular style data), or other important measures such as cardiac output, oxygen saturation, and lactate (measured less consistently at lower resolution). In our work, we rely on aortic pressure signals and heart rate as they are available as high-resolution signals.

Another interesting future direction is to explore the optimal decision-making problem using the DANP simulator to automatically recommend the MCS support levels. We plan to couple the data-driven simulator developed in this work with a reinforcement learning (RL) algorithm. The RL algorithm will interact with the simulator as the environment and optimize the policy of changing the MCS support levels for improved patient outcomes.

## Acknowledgments

This work was supported in part by a research grant from Abiomed, Inc. (Danvers, MA), the U.S. Army Research Office under Army-ECASE award W911NF-23-1-0231, the U.S. Department Of Energy, Office of Science, IARPA HAYSTAC Program, CDC-RFA-FT-23-0069, DARPA AIE FoundSci, NSF Grants #2205093, #2146343, and #2134274.

## References

- AbioMed. Automated impella controller™ (aic), 2020. URL <https://www.heartrecovery.com/education/education-library/abiomed-automated-impella-controller-aic>.
- Karol Arndt, Murtaza Hazara, Ali Ghadirzadeh, and Ville Kyrki. Meta reinforcement learning for sim-to-real domain adaptation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2725–2731. IEEE, 2020.
- et. al. B. Geller. Escalating and de-escalation temporary mechanical circulatory support in cardiogenic shock: A scientific statement from the american heart association. *Circulation*, 146:e50–68, August 2022.
- Mirko Bonfanti, Stavroula Balabani, John P Greenwood, Sapna Puppala, Shervanthi Homer-Vanniasinkam, and Vanessa Díaz-Zuccarini. Computational tools for clinical support: a multi-scale compliant model for haemodynamic simulations in an aortic dissection based on multi-modal imaging data. *Journal of The Royal Society Interface*, 14(136):20170632, 2017.
- Paula Branco, Luís Torgo, and Rita P Ribeiro. A survey of predictive modeling on imbalanced domains. *ACM computing surveys (CSUR)*, 49(2):1–50, 2016.
- Junqiao Chen, David Chun, Milesh Patel, Epsion Chiang, and Jesse James. The validity of synthetic clinical data: a validation study of a leading synthetic data generator (synthea) using clinical quality measures. *BMC medical informatics and decision making*, 19(1):1–9, 2019.
- Li-Fang Cheng, Gregory Darnell, Bianca Dumitrascu, Corey Chivers, Michael E Draugelis, Kai Li, and Barbara E Engelhardt. Sparse multi-output gaussian processes for medical time series prediction. *arXiv preprint arXiv:1703.09112*, 2017.
- Richard Cheng, Gábor Orosz, Richard M Murray, and Joel W Burdick. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 3387–3395, 2019.
- D. Burkhoff D. Doshi. Cardiovascular simulation of heart failure pathophysiology and therapeutics. *Journal of Cardiac Failure*, 22:303–311, December 2015.
- Thomas Davenport and Ravi Kalakota. The potential for artificial intelligence in healthcare. *Future healthcare journal*, 6(2):94, 2019.
- Alejandro F Frangi, Sotirios A Tsaftaris, and Jerry L Prince. Simulation and synthesis in medical imaging. *IEEE transactions on medical imaging*, 37(3):673–679, 2018.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.

- Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J Rezende, SM Eslami, and Yee Whye Teh. Neural processes. *arXiv preprint arXiv:1807.01622*, 2018.
- Omer Gottesman, Fredrik Johansson, Matthieu Komorowski, Aldo Faisal, David Sontag, Finale Doshi-Velez, and Leo Anthony Celi. Guidelines for reinforcement learning in healthcare. *Nature medicine*, 25(1):16–18, 2019.
- CP Impella. Instructions for use & clinical reference manual (united states only), abimed. *Inc.*, Jul, 2014.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Alan Li, Zihao Zhou, Elise Jortberg, and Rose Yu. Forecasting aortic pressure cross-cohort with deep sequence models. In *2022 Computing in Cardiology (CinC)*, volume 498, pages 1–4. IEEE, 2022.
- Faisal Mahmood, Richard Chen, Sandra Sudarsky, Daphne Yu, and Nicholas J Durr. Deep learning with cinematic rendering: fine-tuning deep neural networks using photorealistic medical images. *Physics in Medicine & Biology*, 63(18):185012, 2018.
- Mohammad Amin Morid, Olivia R Liu Sheng, and Joseph Dunbar. Time series prediction using deep learning methods in healthcare. *ACM Transactions on Management Information Systems*, 14(1):1–29, 2023.
- Tung Nguyen and Aditya Grover. Transformer neural processes: Uncertainty-aware meta learning via sequence modeling. *arXiv preprint arXiv:2207.04179*, 2022.
- Adityo Prakosa, Maxime Sermesant, Hervé Delingette, Stéphanie Marchesseau, Eric Saloux, Pascal Allain, Nicolas Villain, and Nicholas Ayache. Generation of synthetic but visually realistic time series of cardiac images combining a biophysical model and clinical images. *IEEE transactions on medical imaging*, 32(1): 99–109, 2012.
- Niranjani Prasad, Li-Fang Cheng, Corey Chivers, Michael Draugelis, and Barbara E Engelhardt. A reinforcement learning approach to weaning of mechanical ventilation in intensive care units. *arXiv preprint arXiv:1704.06300*, 2017.
- Stephen Roberts, Michael Osborne, Mark Ebdon, Steven Reece, Neale Gibson, and Suzanne Aigrain. Gaussian processes for time-series modelling. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 371(1984):20110550, 2013.
- Stéphanie Schampaert, Marcel van’t Veer, Frans N van de Vosse, Nico HJ Pijls, Bas A de Mol, and Marcel CM Rutten. In vitro comparison of support capabilities of intra-aortic balloon pump and impella 2.5 left percutaneous. *Artificial Organs*, 35(9):893–901, 2011.
- Peter Schulam and Suchi Saria. A framework for individualizing predictions of disease trajectories by exploiting multi-resolution structure. *Advances in neural information processing systems*, 28, 2015.
- et. al. T. Heldt. Computational modeling of cardiovascular response to orthostatic stress. *J Appl Physiol*, 92: 1239–1254, 2001.
- René Traoré, Hugo Caselles-Dupré, Timothée Lesort, Te Sun, Natalia Díaz-Rodríguez, and David Filliat. Continual reinforcement learning deployed in real-life using policy distillation and sim2real transfer. *arXiv preprint arXiv:1906.04452*, 2019.
- T van de Meerakker, F van de Vosse, M Rutten, D van Dort, J Foolen, and K Peij. Comparing heart assist devices, using a lumped parameter model and mock-loop. 2019.

Aaron Voelker, Ivana Kajić, and Chris Eliasmith. Legendre memory units: Continuous-time representation in recurrent neural networks. *Advances in neural information processing systems*, 32, 2019.

Wenshuai Zhao, Jorge Peña Queraltá, and Tomi Westerlund. Sim-to-real transfer in deep reinforcement learning for robotics: a survey. In *2020 IEEE symposium series on computational intelligence (SSCI)*, pages 737–744. IEEE, 2020.