Generative ODE Modeling with Known Unknowns

Ori Linial

linial04@campus.technion.ac.il
Technion - Israel Institute of Technology
Haifa, Israel

Danny Eytan

danny.eytan@technion.ac.il
Technion - Israel Institute of Technology
Rambam Health Care Campus
Haifa, Israel

Neta Ravid

neta.r@technion.ac.il
Technion - Israel Institute of Technology
Haifa, Israel

Uri Shalit

urishalit@technion.ac.il
Technion - Israel Institute of Technology
Haifa, Israel

ABSTRACT

In several crucial applications, domain knowledge is encoded by a system of ordinary differential equations (ODE), often stemming from underlying physical and biological processes. A motivating example is intensive care unit patients: the dynamics of vital physiological functions, such as the cardiovascular system with its associated variables (heart rate, cardiac contractility and output and vascular resistance) can be approximately described by a known system of ODEs. Typically, some of the ODE variables are directly observed (heart rate and blood pressure for example) while some are unobserved (cardiac contractility, output and vascular resistance), and in addition many other variables are observed but not modeled by the ODE, for example body temperature. Importantly, the unobserved ODE variables are “known-unknowns”: We know they exist and their functional dynamics, but cannot measure them directly, nor do we know the function tying them to all observed measurements. As is often the case in medicine, and specifically the cardiovascular system, estimating these known-unknowns is highly valuable and they serve as targets for therapeutic manipulations. Under this scenario we wish to learn the parameters of the ODE generating each observed time-series, and extrapolate the future of the ODE variables and the observations. We address this task with a variational autoencoder incorporating the known ODE function, called GOKU-net¹ for Generative ODE modeling with Known Unknowns. We first validate our method on videos of single and double pendulums with unknown length or mass; we then apply it to a model of the cardiovascular system. We show that modeling the known-unknowns allows us to successfully discover clinically meaningful unobserved system parameters, leads to much better extrapolation, and enables learning using much smaller training sets.

¹Code available on github.com/orilnial/GOKU

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).
ACM CHIL '21, April 8–10, 2021, Virtual Event, USA
© 2021 Copyright held by the owner/author(s).
ACM ISBN 978-1-4503-8359-2/21/04.
<https://doi.org/10.1145/3450439.3451866>

ACM Reference Format:

Ori Linial, Neta Ravid, Danny Eytan, and Uri Shalit. 2021. Generative ODE Modeling with Known Unknowns. In *ACM Conference on Health, Inference, and Learning (ACM CHIL '21)*, April 8–10, 2021, Virtual Event, USA. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3450439.3451866>

1 INTRODUCTION

Many scientific fields use the language of ordinary differential equations to describe important phenomena. These include microbiology, physiology, ecology, medicine, epidemiology and finance, to name but a few. Typically, an ODE model of the form $\frac{dz(t)}{dt} = f_{\theta_f}(z(t))$ is derived from first principles and mechanistic understanding, where $z(t)$ are time-varying variables and θ_f are static parameters of the ODE model f . Once a model f_{θ_f} is specified, the values $\hat{\theta}_f$ are found that best fit every instance. These estimated values are often of great interest: in ecology these might correspond to the carrying capacity of a species, whereas in medicine they might represent the cardiovascular characteristics of a patient in the course of critical illness. Predictions based on extrapolating the estimated models are also of wide interest, for example predicting how a patient’s state will evolve or respond to specific interventions.

In many scientific applications the assumption is that the dynamic variables $z(t)$ are directly observed, possibly with some independent noise. Alternately, an assumption is often made that the observations, which we denote hereafter as $x(t)$, are a known, fixed mapping of the unobserved $z(t)$. This assumption however is not always realistic: in the case of critically-ill patients for example, while some physiological variables are directly observed such as arterial blood pressure, others that are key determinants of the dynamical system such as cardiac contractility (the heart’s ability to squeeze blood), stroke volume (the volume of blood squeezed during a single heart beat), or systemic vascular resistance are not only unobserved but also have a non-trivial, and possibly unknown, mapping to the observed variables. Estimating these variables is of great clinical importance both diagnostically and in tailoring treatments aimed at their modification.

This work addresses the scenario where, on the one hand we have the mechanistic understanding needed to define the dynamic variables z and a corresponding ODE model f_{θ_f} , but on the other hand we cannot assume that we have a good model for how the variables z tie in to the observations x . In such a scenario, z and

θ_f take on the role of *known-unknowns*: variables with a concrete meaning, which we do not know and wish to infer from data.

Therefore, our goal is to build a learning system that can use the conjunction of mechanistic ODE models together with data-driven methods [3]. Ideally this conjunction will bring out the best of both worlds, allowing us to address problems neither approach can solve on its own, especially focusing on the correct identification of these known-unknowns. Specifically, we propose an autoencoder framework called GOKU-net, standing for Generative ODE Known-Unknown net. This is a VAE architecture with the known differential equation f at its heart, and with an added component that allows us to effectively use standard VAE conditional-Gaussian parameterizations, yet still obtain estimates of the known-unknown quantities which correspond to their natural physical range.

In the next section we frame our task and its relevance to practice with an example of acute care patients. While many approaches exist for learning the parameters of ODEs, and others exist for sequence modeling with latent variables, we believe none of them can jointly address the task we outline in a straightforward way. We therefore give a (necessarily partial) overview of relevant methods and explain why we believe that despite many similarities they are not suited for the learning scenario we describe.

In our experiments we show what gains can be made by successfully harnessing domain knowledge in terms of the ODE form: We can learn the known-unknowns, we can learn with vastly less data, and we can perform the difficult task of sequence extrapolation with much higher accuracy. We show this by comparing our method with several baselines such as LSTM [16] and Latent-ODE [7]² in three domains: a video of a pendulum, a video of double pendulum, and a dynamic model of the cardiovascular system [38].

Our key contributions

- (1) We show how to efficiently integrate mechanistic knowledge into a deep autoencoder framework.
- (2) We develop experimental framework for evaluating models of “known-unknowns”, including one based on a model of the cardiovascular-system derived from first principles.
- (3) We show what benefits could be gained by the integration of mechanistic knowledge over using purely data-driven methods:
 - (a) Accurately inferring latent system parameters. For example, for the haemodynamically unstable patient — estimate the relative contribution of hidden processes such as internal bleeding or reduced vascular resistance as seen in septic shock.
 - (b) Better extrapolation of observed time series.
 - (c) Achieving results which are equivalent to those of data-driven models while using substantially less training samples.

2 TASK DEFINITION

We are given N observed trajectories $X^i = (x_0^i, \dots, x_{T-1}^i)$, $i = 1, \dots, N$, each describing a time evolving phenomena observed at times $t = 0, \dots, T - 1$. We assume each of these time sequences was

generated by a noisy unknown emission process g from underlying latent (unobserved) trajectories $Z^i = (z_0^i, \dots, z_{T-1}^i)$. The dynamics of the latent variables Z^i are governed by an ODE with known functional form f and unknown static parameters θ_f^i . Note that the latent trajectories share the same functional form but have different ODE parameters across the samples $i = 1, \dots, N$:

$$\frac{dz^i(t)}{dt} = f_{\theta_f^i}(z^i(t)) \quad (1)$$

$$x_t^i = g(z^i(t)) + \varepsilon_t^i, \quad \varepsilon_t^i \sim \mathcal{N}(0, \sigma_x I). \quad (2)$$

Given a training set $\{X^i\}_{i=1}^N$, and a new test sequence $X' = (x'_0, \dots, x'_{T-1})$, our task is two-fold:

- (i) Estimate the static parameters θ_f for X' .
- (ii) Extrapolate x'_t for a set of future times $t > T - 1$.

Consider the pixel-pendulum experiment we report in Section 5.2: The latent state parameter Z is the pendulum’s angle and angular velocity; and the ODE system f is the classic pendulum equation, see Eq. (7). We take the parameter θ_f to be a single number, the pendulum’s length. Finally, we assume our observations X are frames in a video of the pendulum, as shown in Fig. 3. That means the emission function g is the function that takes as input the angle of the pendulum and generates a 28×28 pixel image. The image is always scaled so that the length cannot be inferred from a single image. The task here is, given a previously unseen video, to infer the pendulum’s length and to extrapolate the video into the future of the sequence.

In the ICU patient example, X^i would be a time-series of observed vital signs and other measurements such as heart rate and body temperature for patient i . The variables Z^i describe a set of physiological variables such as blood pressure, blood volumes in the heart, cardiac stroke volumes and more. The function f_{θ_f} is an ODE model for these physiological variables [13, 18, 26, 31, 38]. The parameters θ_f^i would be important patient-specific static variables such as arterial and venous compliances. Correct estimation of these variables conveys immediate clinical advantage both by aiding the clinicians in establishing the correct underlying diagnosis (for example hypotension due to reduced cardiac function versus septic shock versus bleeding) and in serving as treatment goals with specific interventions tailored to the identified pathophysiological process.

We note that inferring exactly the true latent parameters might be impossible in some systems due to under-identification, as for example multiple sets of latent z_0 can give rise to the same observed X by way of different emission models g .

3 RELATED WORK

There has been much work recently bringing together ML methods and ODEs. However, we believe we are unique in our setting, as explained below and outlined in Table 1. We divide existing work into several categories: First, work on parameter identification in dynamical systems which assumes both the ODE function f and the emission function g are known. Second, work on latent state sequence modeling. This work does not assume any known dynamics or emission model. Finally, there is recent work tying

²Also called Neural-ODE; we call it Latent-ODE following the usage in Rubanova et al. [29].

Figure 1: GOKU-net model. Left: inference model; an observed signal X_t is taken as input by a bi-directional LSTM to produce estimates of the initial state $\hat{z}_0$ and static parameter $\hat{\theta}_f$. Right: The ODE solver takes these values together with the given ODE function f to produce the latent signal $\hat{Z}$, then reconstruct $\hat{X}$ using an emission network $\hat{g}$. The ODE solver can integrate $\hat{z}_t$ arbitrarily far forward in time, enabling the extrapolation of X for any $t = T + \tau$.

together machine learning models and physical models in a task-specific way.

ODE parameter identification has been the subject of decades of research across many scientific communities. Classic state-space models, including methods such as the Kalman filter [23] and its non-linear extensions [21, 22, 33], can learn the parameters of a dynamic system from observations; however, they are limited to the case where the emission function g is known. Moreover, they usually perform inference on each sequence separately. Many machine learning methods have been proposed for this task, for example using reproducing kernel Hilbert space methods [14] and Gaussian Processes [4, 11, 15], Fast Gaussian Process Based Gradient Matching [36] and recent follow up work [35]. In general these methods assume in practice that the observed signal is simply the state variable $z(t)$ with independent additive noise, whereas we allow a more complicated, and apriori unknown, link from latent states to observations.

Many methods for extrapolation of a given signal assuming there is some unknown arbitrary underlying latent sequence have been proposed. Some prominent examples are LSTM [16], Deep Markov Models [24], Neural ODE (also called Latent-ODE, L-ODE, [7]) and follow up methods, [29, 37], NbedDyn [27], the work by Ayed et al. [2], the Disentangled State Space Model (DSSM, [25]), SVG-LP [9], method using Gaussian Processes [19], and Gaussian Process State Space Models (GPSSM, [12, 39]); motivated by robotics applications, Watter et al. [34] aim to learn latent space such that locally linear dynamics will be useful. These methods do not infer the ODE parameters as they do not learn any intrinsically meaningful latent space. They also do not exploit the prior information embedded in the mechanistic knowledge underlying the derivation of the ODE system f . Of the above methods, DSSM has been shown to learn a latent space which might under the right circumstances correspond to meaningful parameters, but that is not guaranteed, nor is it the goal of the method. In the healthcare regime, Cheng et al. [8] proposed a method for learning a sequence which includes a dynamic

system in the form of a latent force model [1]; this approach builds on learning to fit general basis functions to describe the observed dynamics, and does not take as input an ODE system derived from prior mechanistic understanding.

Closer in spirit to our work is the work by Greydanus et al. [17] on Hamiltonian neural networks. In their model the latent space can be interpreted in the form of learning a conserved physical quantity (Hamiltonian). Although related to our work, we note that not all ODE systems have easily identified conserved quantities. Specifically, the systems that motivate our research do not usually have a Hamiltonian representation. For example in section 5 we experiment with an ODE system of a pendulum with friction as a non-conservative quantity. Another closely related work is SINDy Autoencoders [6]. This work is similar to ours in using an unknown and complex emission function, and learning a physically meaningful latent space. Their model is not given the ODE system f , but instead aims to learn some representation of the underlying ODE as a linear combination of bases functions of the latent state. In the field of learning for healthcare, Soleimani et al. [32] show how a specific ODE model, the linear time-invariant impulse-response model, can be used in conjunction with latent-space models to estimate how a patient’s measurements would react to interventions. It brings together mechanistic modeling in terms of response to impulse treatments, along with data-driven modeling using Gaussian processes. It does not learn emission functions from the latent space, and focuses on the specific ODE model relevant to their task.

4 MODEL AND METHOD

Given N observed trajectories $X^i = (x_0^i, \dots, x_{T-1}^i)$, $i = 1, \dots, N$, our main idea is based on inferring a latent trajectory $\hat{Z}$ while reconstructing the ODE parameters θ_f , in a variational autoencoder approach [10, 28]. As usual, that implies learning both an inference function (encoder) and an emission function (decoder). The inference function takes an observed sequence X^i as input, and has two components: The first infers the ODE parameters $\hat{\theta}_f^i$, and

Method	ODE function	Emission function	θ_f identification	X extrapolation
LSTM [16]	not required	learned	X	✓
L-ODE [7, 29]	not required	learned	X	✓
DMM [24]	not required	learned	X	✓
GPSSM [12]	not required	learned	X	✓
DSSM [25]	not required	learned	X	✓
HNN [17]	can be used	learned	X	✓
NbedDyn [27]	not required	partially given	X	✓
ODIN [35]	required	given	✓	X
UKF [33]	required	given	✓	✓
GOKU-net	required	learned	✓	✓

Table 1: Related Work: for each method we indicate does it require the ODE function f as input; does it assume the emission function g is given or learned; does it allow identification of the “known-unknown” parameters θ_f ; and whether it allows for extrapolating the observed X .

the second infers an initial latent state $\hat{z}_0^i$. We next use the *known* ODE functional form f , the inferred ODE parameters $\hat{\theta}_f^i$ and the inferred initial state $\hat{z}_0^i$ to obtain an estimated trajectory $\hat{Z}^i$ by a numerical ODE solver. We then use $\hat{Z}^i$ as input to a *learned* emission function $\hat{g}$, obtaining a reconstructed signal $\hat{X}^i$. We estimate the log-likelihood of the reconstructed signal, and use stochastic back-propagation [10, 28] through the ODE solver in order to update the parameters of the inference network and emission model – details below. Extrapolating the latent trajectory using the ODE solver lets us make estimates of X^i arbitrarily far forwards or backwards in time. Fig. 1 illustrates the proposed model.

4.1 Generative model and inference

Using the relationships between latent and observed variables given in Eqs. (1) and (2), we define a generative model over the set of ODE parameters θ_f , the latent states Z , and the observations X . While we assume the true ODE function f is given, we estimate the emission by a learned function $\hat{g}$.

A crucial issue we must address is that in standard VAEs the prior distributions of the latent vectors z_0 and θ_f are set to be a zero-mean unit-variance Gaussian. However, we are conducting inference in a latent space where the variables correspond to specific variables with physical meaning and constraints: for example, the variable for blood volume has a limited set of realistic values. We overcome this by defining arbitrary (non-physical) latent vectors $\tilde{z}_0$ and $\tilde{\theta}_f$ with standard Gaussian priors, and then learning deterministic “physically-grounding” transformations h_z and h_{θ_f} such that:

$$\tilde{z}_0 \sim \mathcal{N}(0, I), \quad z_0 = h_z(\tilde{z}_0), \quad (3)$$

$$\tilde{\theta}_f \sim \mathcal{N}(0, I), \quad \theta_f = h_{\theta}(\tilde{\theta}_f). \quad (4)$$

We then have z_t , $t > 0$ and X generated following Eqs. (1) and (2). With the above generative model, we have the following factorized joint distribution over latent and observed variables:

$$p(X, Z, \theta_f, \tilde{z}_0, \tilde{\theta}_f) = p(\tilde{z}_0)p(\tilde{\theta}_f)p(z_0|\tilde{z}_0)p(\theta_f|\tilde{\theta}_f)p(x_0|z_0) \prod_{t=1}^{T-1} p(x_t|z_t)p(z_t|z_{t-1}, \theta_f).$$

This follows due to the conditional independence: for $t' \neq t$: $x_t \perp (z_{t'}, \theta_f, \tilde{\theta}_f, \tilde{z}_0)|z_t$ and for $t' \neq t-1$: $z_t \perp (X, \tilde{\theta}_f, \tilde{z}_0, z_{t'})|z_{t-1}, \theta_f$. The probabilities $p(z_0|\tilde{z}_0)$ and $p(\theta_f|\tilde{\theta}_f)$ are deterministic, meaning they are Dirac functions with the peak defined by Eqs. (3) and (4). The transition distribution $p(z_t|z_{t-1}, \theta_f)$ is also a Dirac function with the peak defined by Eq. (1). Finally, the emission distribution $p(x_t|z_t)$ is defined by Eq. (2).

Inference We define the following joint posterior distribution over the unobserved random variables Z , θ_f , $\tilde{z}_0$ and $\tilde{\theta}_f$, conditioned on a sequence of observations X :

$$q(Z, \theta_f, \tilde{z}_0, \tilde{\theta}_f|X) = q(\tilde{z}_0|X)q(\tilde{\theta}_f|X)q(z_0|\tilde{z}_0)q(\theta_f|\tilde{\theta}_f) \prod_{t=1}^{T-1} q(z_t|z_{t-1}, \theta_f).$$

The inference network conditionals $q(z_t|z_{t-1}, \theta_f)$, $q(z_0|\tilde{z}_0)$ and $q(\theta_f|\tilde{\theta}_f)$ are deterministic and mirror the generative model as defined in Eqs. (1), (3) and (4), respectively. For the posteriors $q(\tilde{z}_0|X)$ and $q(\tilde{\theta}_f|X)$ we use conditional normal distributions where $\phi_{\tilde{z}_0}^{enc}$ and $\phi_{\tilde{\theta}_f}^{enc}$ are learned NNs:

$$q(\tilde{z}_0|X) = \mathcal{N}(\mu_{\tilde{z}_0}, \sigma_{\tilde{z}_0}), \quad [\mu_{\tilde{z}_0}, \sigma_{\tilde{z}_0}] = \phi_{\tilde{z}_0}^{enc}(X),$$

$$q(\tilde{\theta}_f|X) = \mathcal{N}(\mu_{\tilde{\theta}_f}, \sigma_{\tilde{\theta}_f}), \quad [\mu_{\tilde{\theta}_f}, \sigma_{\tilde{\theta}_f}] = \phi_{\tilde{\theta}_f}^{enc}(X).$$

4.2 Objective

We use the evidence lower-bound (ELBO) variational objective [10, 28]:

$$\mathcal{L}(X) = \mathbb{E}_{q(Z, \theta_f, \tilde{z}_0, \tilde{\theta}_f|X)} \left[\log p(X|Z, \theta_f, \tilde{z}_0, \tilde{\theta}_f) \right] - KL[q(Z, \theta_f, \tilde{z}_0, \tilde{\theta}_f|X) || p(Z, \theta_f, \tilde{z}_0, \tilde{\theta}_f)]. \quad (5)$$

Since for all $t' \neq t$ we have $x_t \perp (x_{t'}, z_{t'}, \theta_f, \tilde{z}_0, \tilde{\theta}_f) | z_t$, the first term of Eq. (5) decomposes as:

$$\mathbb{E}_{q(Z, \theta_f, \tilde{\theta}_f, \tilde{z}_0 | X)} \left[\log p(X | Z, \theta_f, \tilde{z}_0, \tilde{\theta}_f) \right] = \sum_{t=0}^{T-1} \mathbb{E}_{q(z_t | X)} [\log p(x_t | z_t)].$$

The KL term decomposes into the following sum of KL terms:

$$KL \left[q(Z, \theta_f, \tilde{z}_0, \tilde{\theta}_f | X) || p(Z, \theta_f, \tilde{z}_0, \tilde{\theta}_f) \right] = KL \left[q(\tilde{\theta}_f | X) || p(\tilde{\theta}_f) \right] + KL \left[q(\tilde{z}_0 | X) || p(\tilde{z}_0) \right].$$

See the appendix for the full derivation.

4.3 Implementation

We model $[\hat{g}, h_z, h_{\theta_f}, \phi_{\tilde{z}_0}^{enc}, \phi_{\tilde{\theta}_f}^{enc}]$ as neural networks: $\hat{g}$, h_z and h_{θ_f} as fully connected neural networks; $\phi_{\tilde{z}_0}^{enc}$ as an RNN which goes over the observed X backwards in time to predict z_0 ; and $\phi_{\tilde{\theta}_f}^{enc}$ as a bi-directional LSTM [20] with fully connected networks from X into $\tilde{\theta}_f$. We use bi-LSTM for θ_f identification since θ_f is time invariant.

In order to perform stochastic backpropagation, we must calculate the gradient through the ODE defined by f and θ_f . To this end, we experimented with two methods: an explicit approach, where the gradients are propagated directly through the operations of the numerical ODE integration, and an implicit approach³ that uses the adjoint method which is more memory efficient [7]. In both methods we used the Runge-Kutta-4 numerical integration. Both methods achieved very similar results, therefore we present in this paper only the latter.

5 EXPERIMENTS

In this section we analyze how GOKU-net can be used for observed signal extrapolation and ODE parameter identification. We first validate the method using two widely known physics models: an OpenAI Gym video simulator of a pendulum [5], introduced as an ODE modeling benchmark in [17], and a similar OpenAI Gym video simulator of a (chaotic) double pendulum. We then apply the method to a model of the cardiovascular system based on the model introduced in [38]. In each case we train the model on a set of sequences with varying ODE parameters (θ_f) and initial conditions (z_0), and test on unseen sequences with parameters and initial conditions sampled from the same distribution as the train set. Each dataset was randomly divided into train, validation and test sets (80%, 10%, 10%). Validation was used for early stopping, and hyper-parameter tuning.

5.1 Baselines

Table 2 summarizes what type of information each of the baselines described below has access to.

Data-driven As baselines with no input from mechanistic models we use (i) LSTM [16], (ii) Latent-ODE [7], originally called

Method	f	Z
LSTM [16]	X	X
Latent-ODE [7]	X	X
DMM [24]	X	X
DI	✓	✓*
GOKU-net	✓	X
HNN [17]	✓	X

Table 2: Each of the methods we compare, and what information does it have access to beyond the observations X : the true ODE function f , and sparse grounding latent states Z . *DI has access to Z at test time as well as train.

Neural-ODE; we denote it L-ODE, (iii) DMM [24]. These methods can only be used to extrapolate X for future time steps, since their latent space has no physical interpretation.

Direct inference (DI) We introduce a baseline which can use both the known model f and the observations X by having access to an additional bit of knowledge: sparse observations from the latent states Z , during both train and test time. Note that none of the other baselines, nor GOKU, has access to such information.

We call this baseline “direct inference” (DI). It has two steps: First inferring the ODE parameters and latent states, and then, separately, learning the connection between the latent states Z and observations X .

Specifically, the ODE parameters and latent states ($\hat{Z}, \hat{\theta}_f$) are estimated directly from a sparse sequence of directly observed latent states Z^{observed} , by minimizing the loss function:

$$\mathcal{L}_{\text{ground}} = \sum_t M(t) \cdot \|\hat{z}_t - z_t^{\text{observed}}\|_2^2, \quad (6)$$

where $M(t) \in \{0, 1\}$ indicates for which time points the latent variables are observed. $\hat{z}_t$ is the latent vector predicted by the model and z_t^{observed} are the observed samples of the latent vectors. We optimize Eq. (6) using gradient descent through an ODE solver with the given ODE model f ; details in the appendix. Then in the second step, DI obtains estimates for the observations X by constructing a training set where the instances are the Z inferred for each training sequence, and the labels are the corresponding observations X . We then learn a function $\hat{g}$ predicting X from Z . Below we denote the method as DI- $q\%$, where $q \in \{1, 5\}$ indicates what percent of the latent Z were available to DI during train and test. We emphasize that this baseline is the only one to have access to observations of Z , and that these are available to it both at train and test time.

5.2 Single Pendulum From Pixels

In this experiment we generate a video of a pendulum, aiming to predict the future frames and to identify its parameter. We use the following non-linear oscillator ODE:

$$\frac{d\theta(t)}{dt} = \omega(t), \quad \frac{d\omega(t)}{dt} = -\frac{g}{l} \sin \theta(t), \quad (7)$$

with gravitational constant $g = 10$. The ODE has a single parameter which is the pendulum’s length l , and the ODE state is $z_t = (\theta(t), \omega(t))$.

³<https://github.com/rtqichen/torchdiffq>

Data Set We followed [17] and used the Pendulum-v0 environment from OpenAI Gym [5]. For training we simulated 500 sequences of 50 time points, with time steps of $\Delta t = 0.05$ and pre-processed the observed data such that each frame is of size 28×28 . We made one important change relative to [17]: the ODE parameter l was uniformly sampled, $l \sim U[1, 2]$ instead of being constant, making the task much harder. The images are scaled so that the length of the pendulum is not identifiable from a single frame. Each test set sequence is 100 time steps long, where the first 50 time steps are given as input, and the following 50 were used only for evaluating the signals extrapolation.

Method	X extrap.	$\ \theta_f - \hat{\theta}_f\ _1$	$\text{corr}(\theta_f, \hat{\theta}_f)$
GOKU-net	7 ± 1	78 ± 8	0.979
Latent-ODE	42 ± 2	N/A	N/A
LSTM	81 ± 4	N/A	N/A
DMM	74 ± 5	N/A	N/A
DI 1%	147 ± 0	511 ± 44	0.046
DI 5%	147 ± 0	77 ± 29	0.751

Table 3: Pixel pendulum mean L_1 error ($\times 10^3$) for extrapolating X with standard error of the mean, and θ_f identification error (both L_1 and correlation coefficient) across test samples. DI 5% and DI 1% indicate the percentage of the Z latent states were observed by DI. The X extrapolation L_1 error of simply predicting all pixels are zero is 147, meaning both DI baselines failed in this task.

Evaluation and results. Table 3 shows the ODE length parameter identification error in terms of L_1 error and the Pearson correlation coefficient r , and mean extrapolation error in terms of L_1 error (L_2 has similar results). Fig. 2 shows how the extrapolation error of the observed signals evolves over time, starting from time $t = 50$. In Fig. 3 we demonstrate how GOKU-net extrapolates on a single, randomly selected, signal when compared to the best performing baseline, the Latent ODE.

In terms of identification, GOKU-net performs much better than the DI 1% baseline, and similar to the DI 5% baseline, even though GOKU-net does not receive access to observations of Z , which are accessible to the DI baselines. When performing extrapolation, we see in Fig. 2 that GOKU-net extrapolates much better than all baselines. Specifically, GOKU-net outperforms HNN [17] which has difficulty with the fact that the ODE parameter is not constant.

We further tested performance as a function of train set size, reasoning that domain knowledge should help reduce the need for many training samples. We evaluated the L_1 extrapolation error of the different methods averaged over time steps $t = 50 \dots 100$, taken on the same test set. As can be seen in Fig. 4, training GOKU-net on as little as 50 train samples provides better X -extrapolation error than Latent ODE achieves with 1000 (20 fold) train samples, showing we can achieve substantial benefits by properly using domain knowledge. Both DI baselines completely failed in this scenario.

5.2.1 Pixel Pendulum with Unknown Unknowns. In this experiment we relax an important assumption we made so far: that we have

Figure 2: Pixel pendulum: mean extrapolation error for observations X over time steps after end of input sequence. Percentages in legend are percent grounding observation in training. HNN by Greydanus et al. [17]. “All black” is predicting all pixels as zero brightness (black).

Figure 3: Predicting the dynamics of pixel pendulum. The first 50 frames are observed, and the next 150 are predicted. The above 3 figures are examples of each method’s predicted frames every 30 time steps.

the full ODE model f for the underlying latent states. Instead, here we experiment with data generated by a true ODE model f that is different from the one given as input to GOKU-net. We do this by adding friction to the pendulum ODE Eq. (7):

$$\frac{d\theta(t)}{dt} = \omega(t), \quad \frac{d\omega(t)}{dt} = -\frac{g}{l} \sin \theta(t) - \frac{b}{m} \omega(t).$$

Importantly, GOKU-net is still given the ODE defined by Eq. (7), without friction. For this task we added a trainable abstract function f_{abs} to the ODE function such that now $\frac{dz_t}{dt} = f_{ODE}(z_t, \theta_f) + f_{abs}(z_t, \theta_f)$, where f_{ODE} is the given friction-less ODE Eq. (7). The

Figure 4: Pixel pendulum X extrapolation error on test set, for different train set sizes. “All black” is simply predicting all pixels as zero brightness (black).

idea is that the abstract function f_{abs} might model the *unknown-unknowns* of the system, in this case the friction. We name this approach GOKU-UU for GOKU with unknown-unknowns. The results show that even with the incorrect ODE model, we were able to extrapolate the signal including the decay of the pendulum’s velocity. We further demonstrate in Fig. 5, that when we zero the inferred f_{abs} at test time the inferred pendulum’s velocity did *not* decay, strengthening the claim that f_{abs} modeled the friction, while f_{ODE} modeled the friction-less pendulum. Results are presented in the appendix, as well as algorithmic details.

Figure 5: Pixel pendulum with friction predicted angle example. Here we demonstrate that zeroing the f_{abs} part of GOKU-UU, results in a friction-less signal.

5.3 Double Pendulum From Pixels

In order to further test the limits of GOKU, we created a dataset comprised of videos of a double pendulum. The dynamics of the double pendulum are much more complex than those of the single pendulum, and are known to have chaotic nature [30], making this task a much harder one. The dataset creation was similar to the single pendulum experiment, using the mass of the second pendulum as the unknown parameter θ_f we wish to identify. More details about the ODE system, and the dataset creation process in the appendix.

Table 4 shows the ODE parameters identification error in terms of the (Pearson) correlation coefficient and L_1 error, and mean extrapolation error in terms of L_1 error. Although the double pendulum has a chaotic nature, GOKU-net was still able to extrapolate the observed video signals well into the future, and better than the compared baselines. In terms of identification of the mass parameter, GOKU-net managed to correlate very well with the ground truth, but was off in its scale, predicting masses which were too small.

Method	X extrapol.	$\ \theta_f - \hat{\theta}_f\ _1$	$corr(\theta_f, \hat{\theta}_f)$
GOKU-net	12 ± 0	1.378 ± 0.030	0.972
Latent-ODE	16 ± 1	N/A	N/A
LSTM	32 ± 1	N/A	N/A
DMM	26 ± 1	N/A	N/A
DI 1%	37 ± 0	0.261 ± 0.041	0.536
DI 5%	37 ± 0	0.057 ± 0.005	0.998

Table 4: Double pendulum mean L_1 error ($\times 10^3$) for extrapolating X with standard error of the mean, and θ_f identification error (both L_1 and correlation coefficient) across test samples. DI 5% and DI 1% indicate the percentage of the Z latent states were observed by DI. The X extrapolation L_1 error of simply predicting all pixels are zero is 37, meaning both DI baselines failed in this task.

5.4 Cardiovascular System

As mentioned in the introduction, a major clinical challenge in the context of critical care lies in elucidating the causes for haemodynamic instability and deriving an appropriate treatment plan. We approached this challenge using a simplified mechanistic model of the cardiovascular system (CVS), which is a modified version of a model first suggested in Zenker et al. [38]. It is a multi-compartment model comprising the heart, the venous and arterial subsystems together with a reflex loop component representing the nervous

system control of blood pressure:

$$\begin{aligned}\frac{dSV(t)}{dt} &= I_{\text{external}}, \\ \frac{dP_a(t)}{dt} &= \frac{1}{C_a} \left(\frac{P_a(t) - P_v(t)}{R_{TPR}(S)} - SV \cdot f_{HR}(S) \right) \\ \frac{dP_v(t)}{dt} &= \frac{1}{C_v} \left(-C_a \frac{dP_a(t)}{dt} + I_{\text{external}} \right) \\ \frac{dS(t)}{dt} &= \frac{1}{\tau_{Baro}} \left(1 - \frac{1}{1 + e^{-k_{width}(P_a(t) - P_{aset})}} - S \right),\end{aligned}$$

where

$$\begin{aligned}R_{TPR}(S) &= S(t)(R_{TPR_{Max}} - R_{TPR_{Min}}) + R_{TPR_{Min}} + R_{TPR_{Mod}} \\ f_{HR}(S) &= S(t)(f_{HR_{Max}} - f_{HR_{Min}}) + f_{HR_{Min}}.\end{aligned}$$

In this model the variables have a directly interpretable mechanistic meaning: SV , P_a , P_v , S which are respectively cardiac stroke volume (the amount of blood ejected by the heart), arterial blood pressure, venous blood pressure and autonomic baroreflex tone (the reflex responsible for adapting to perturbations in blood pressure, keeping homeostasis). Table 6 in the appendix provides a full glossary of the above terms. The ODE state is $z_t = (SV(t), P_a(t), P_v(t), S(t))$. The observed state is the patient's vital signs and defined to be: $x_t = (P_a(t), P_v(t), f_{HR}(t))$, where $f_{HR}(t)$ is the patient's heart-rate at time t . Note that two of the three observed variables are the same as two of the latent variables, albeit with added noise as described in the dataset subsection. The CVS ODE model has several other parameters which we treat as known, setting them to the values stated in [38]. The model is a 4-variable ODE system with 3-dimensional observations, which has much more involved dynamics than the pixel-pendulum above. Although far from comprehensive, this model can capture the prototypical behaviour of the cardiovascular system and its responses to pathological insults such as internal bleeding or septic shock that manifests as a reduction in peripheral vascular resistance.

In this system we wish to find two parameters: $\theta_{CVS} = (I_{\text{external}}, R_{\text{Mod}})$ which are known-unknowns describing recognized clinical conditions: $I_{\text{external}} < 0$ implies a patient is losing blood, and $R_{\text{Mod}} \geq 0$ implies that their total peripheral resistance is getting lower, which is a condition of distributive shock. Both conditions can lead to an observed drop in blood-pressure, and discerning the relative contribution of each to such a drop is important clinically: Often the underlying causes are not immediately clear and the choice of treatment relies on their accurate estimation.

Data Set We simulated 1000 sequences of length 400, with time steps of $\Delta_t = 1$. The parameter I_{external} was randomly sampled to be either -2 or 0 , and the parameter $R_{TPR_{Mod}}$ was randomly sampled to be either 0.5 or 0 . Initial ODE states uniformly sampled from $SV(0) \sim U[90, 100]$, $P_a(0) \sim U[75, 85]$, $P_v \sim [3, 7]$ and $S \sim [0.15, 0.25]$, where the intervals were set to the values given in Zenker et al. [38]. The observations were additionally corrupted with white Gaussian noise with standard deviation of $\sigma_x = 5$ for P_a , $\sigma_x = 0.5$ for P_v and $\sigma_x = 0.05$ for f_{HR} (standard deviation matches scale of the observed signal).

Evaluation and results In addition to identification and extrapolation, we attempt to classify each X^i series according to the sign

Method	I_{external}	R_{Mod}	Class.	X extrap.
GOKU	24 ± 2	$3 \pm .0$	0%	29 ± 1
K-Means	N/A	N/A	13%	N/A
LSTM	N/A	N/A	N/A	90 ± 2
DMM	N/A	N/A	N/A	85 ± 3
DI-5%	1 ± 1	$0 \pm .0$	0%	11 ± 10

Table 5: CVS parameter identification and extrapolation error ($\times 10^3$), and classification error of clinical conditions. See text for K-means method. L-ODE and DI-1% failed and are not shown.

of the inferred I_{external} and R_{Mod} . These correspond to one of four possible clinical conditions: (1) Healthy (both non-negative), (2) Hemorrhagic shock ($I_{\text{external}} < 0$, $R_{\text{Mod}} \geq 0$), (3) Distributive shock ($I_{\text{external}} \geq 0$, $R_{\text{Mod}} < 0$) and (4) Combined shock ($I_{\text{external}} < 0$, $R_{\text{Mod}} < 0$). In addition to the previously mentioned baselines, we also compare with the following baseline: cluster the observations using K-means with $K = 4$, and assign each cluster to the most common true clinical condition.

Table 5 shows results on all the above tasks. L-ODE, and DI-1% failed in this scenario and are not presented. We see that without any access to the latent space, GOKU-net successfully classifies which of the four clinical conditions the signal corresponds to, and extrapolates much better than the LSTM and DMM baselines. DI with mask rate lower than 5% failed completely in reconstructing X .

6 CONCLUSION

We explore the advantages of creating a model which is a hybrid of mechanistic and data-driven approaches [3]. Compared to purely data-driven models such as Latent ODE [7] and LSTM, our model has an important advantage: it has access to part of the true mechanism underlying the data generating process, allowing us to reason about “known-unknown” variables which are often of crucial interest in clinical and other applications. We show the extra knowledge encoded in the ODE structure allows us to correctly identify meaningful ODE parameters θ_f , which is impossible with the methods above. We also show this extra knowledge translates to much better time-series extrapolation and significantly lower sample complexity. Comparing with direct inference of the ODE parameters, we show that adding the data-driven neural-net component allows GOKU-net to correctly learn the emission model, and yields much better estimates of the model parameters.

Scientists have been developing mechanistic understanding of biological and physiological systems for centuries. We believe finding new ways of combining this understanding into data-driven modeling is a valuable avenue for future research, especially in domains where data is scarce or where causal reasoning is crucial.

7 ACKNOWLEDGMENTS

We wish to thank Guy Tennenholtz, Hagai Rossman and Rahul Krishnan for their useful comments on the manuscript. This research was partially supported by the Israel Science Foundation (grant No. 1950/19).

REFERENCES

- [1] Mauricio Alvarez, David Luengo, and Neil D Lawrence. 2009. Latent force models. (2009), 9–16.
- [2] Ibrahim Ayed, Emmanuel de Bézenac, Arthur Pajot, Julien Brajard, and Patrick Gallinari. 2019. Learning dynamical systems from partial observations. *arXiv preprint arXiv:1902.11136* (2019).
- [3] Ruth E Baker, Jose-Maria Pena, Jayaratnam Jayamohan, and Antoine Jérusalem. 2018. Mechanistic models versus machine learning, a fight worth fighting for the biological community? *Biology letters* 14, 5 (2018), 20170660.
- [4] David Barber and Yali Wang. 2014. Gaussian processes for Bayesian estimation in ordinary differential equations. (2014), 1485–1493.
- [5] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. 2016. OpenAI gym. *arXiv preprint arXiv:1606.01540* (2016).
- [6] Kathleen Champion, Bethany Lusch, J Nathan Kutz, and Steven L Brunton. 2019. Data-driven discovery of coordinates and governing equations. *Proceedings of the National Academy of Sciences* 116, 45 (2019), 22445–22451.
- [7] Tian Qi Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. 2018. Neural ordinary differential equations. (2018), 6571–6583.
- [8] Li-Fang Cheng, Bianca Dumitrescu, Michael Zhang, Corey Chivers, Michael Draugelis, Kai Li, and Barbara E Engelhardt. 2019. Patient-Specific Effects of Medication Using Latent Force Models with Gaussian Processes. *arXiv preprint arXiv:1906.00226* (2019).
- [9] Emily Denton and Rob Fergus. 2018. Stochastic video generation with a learned prior. *arXiv preprint arXiv:1802.07687* (2018).
- [10] P Kingma Diederik, Max Welling, et al. 2014. Auto-encoding variational bayes. 1 (2014).
- [11] Frank Dondelinger, Dirk Husmeier, Simon Rogers, and Maurizio Filippone. 2013. ODE parameter inference using adaptive gradient matching with Gaussian processes. (2013), 216–228.
- [12] Stefanos Eleftheriadis, Tom Nicholson, Marc Deisenroth, and James Hensman. 2017. Identification of Gaussian process state space models. (2017), 5309–5319.
- [13] LM Ellwein, SR Pope, A Xie, JJ Batzel, CT Kelley, and MS Olufsen. 2013. Patient-specific modeling of cardiovascular and respiratory dynamics during hypercapnia. *Mathematical biosciences* 241, 1 (2013), 56–74.
- [14] Javier González, Ivan Vujčić, and Ernst Wit. 2014. Reproducing kernel Hilbert space based estimation of systems of ordinary differential equations. *Pattern Recognition Letters* 45 (2014), 26–32.
- [15] Nico S Gorbach, Stefan Bauer, and Joachim M Buhmann. 2017. Scalable variational inference for dynamical systems. (2017), 4806–4815.
- [16] Alex Graves. 2013. Generating sequences with recurrent neural networks. *arXiv preprint arXiv:1308.0850* (2013).
- [17] Samuel Greydanus, Misko Dzamba, and Jason Yosinski. 2019. Hamiltonian neural networks. (2019), 15353–15363.
- [18] Arthur C Guyton, Thomas G Coleman, and Harris J Granger. 1972. Circulation: overall regulation. *Annual review of physiology* 34, 1 (1972), 13–44.
- [19] Markus Heinonen, Cagatay Yildiz, Henrik Mannerström, Jukka Intosalmi, and Harri Lähdesmäki. 2018. Learning unknown ODE models with Gaussian processes. (2018), 1959–1968.
- [20] Zhiheng Huang, Wei Xu, and Kai Yu. 2015. Bidirectional LSTM-CRF models for sequence tagging. *arXiv preprint arXiv:1508.01991* (2015).
- [21] Andrew H Jazwinski. 2007. *Stochastic processes and filtering theory*. Courier Corporation.
- [22] Simon J Julier and Jeffrey K Uhlmann. 1997. New extension of the Kalman filter to nonlinear systems. 3068 (1997), 182–193.
- [23] Rudolf Emil Kalman et al. 1960. Contributions to the theory of optimal control. *Bol. soc. mat. mexicana* 5, 2 (1960), 102–119.
- [24] Rahul G Krishnan, Uri Shalit, and David Sontag. 2017. Structured inference networks for nonlinear state space models. (2017).
- [25] Đorđe Miladinović, Muhammad Waleed Gondal, Bernhard Schölkopf, Joachim M Buhmann, and Stefan Bauer. 2019. Disentangled state space representations. *arXiv preprint arXiv:1906.03255* (2019).
- [26] Mette S Olufsen and Johnny T Ottesen. 2013. A practical approach to parameter estimation applied to model predicting heart rate regulation. *Journal of mathematical biology* 67, 1 (2013), 39–68.
- [27] Said Ouala, Duong Nguyen, Lucas Drumetz, Bertrand Chapron, Ananda Pascual, Fabrice Collard, Lucile Gaultier, and Ronan Fablet. 2019. Learning latent dynamics for partially-observed chaotic systems. *arXiv preprint arXiv:1907.02452* (2019).
- [28] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. (2014), 1278–1286.
- [29] Yulia Rubanova, Tian Qi Chen, and David K Duvenaud. 2019. Latent Ordinary Differential Equations for Irregularly-Sampled Time Series. (2019), 5321–5331.
- [30] Troy Shinbrot, Celso Grebogi, Jack Wisdom, and James A Yorke. 1992. Chaos in a double pendulum. *American Journal of Physics* 60, 6 (1992), 491–499.
- [31] Bram W Smith, J Geoffrey Chase, Roger I Nokes, Geoffrey M Shaw, and Graeme Wake. 2004. Minimal haemodynamic system model including ventricular interaction and valve dynamics. *Medical engineering & physics* 26, 2 (2004), 131–139.
- [32] Hossein Soleimani, Adarsh Subbaswamy, and Suchi Saria. 2017. Treatment-response models for counterfactual reasoning with continuous-time, continuous-valued interventions. (2017).
- [33] Eric A Wan and Rudolph Van Der Merwe. 2000. The unscented Kalman filter for nonlinear estimation. (2000), 153–158.
- [34] Manuel Watter, Jost Springenberg, Joshka Boedecker, and Martin Riedmiller. 2015. Embed to control: A locally linear latent dynamics model for control from raw images. *Advances in neural information processing systems* 28 (2015), 2746–2754.
- [35] Philippe Wenk, Gabriele Abbati, Stefan Bauer, Michael A Osborne, Andreas Krause, and Bernhard Schölkopf. 2020. ODIN: ODE-Informed Regression for Parameter and State Inference in Time-Continuous Dynamical Systems. (February 2020).
- [36] Philippe Wenk, Alkis Gotovos, Stefan Bauer, Nico S. Gorbach, Andreas Krause, and Joachim M. Buhmann. 2019. Fast Gaussian process based gradient matching for parameter identification in systems of nonlinear ODEs. 89 (2019), 1351–1360.
- [37] Cagatay Yildiz, Markus Heinonen, and Harri Lähdesmäki. 2019. ODE2VAE: Deep generative second order ODEs with Bayesian neural networks. (2019), 13412–13421.
- [38] Sven Zenker, Jonathan Rubin, and Gilles Clermont. 2007. From inverse problems in mathematical physiology to quantitative differential diagnoses. *PLoS computational biology* 3, 11 (2007).
- [39] Zheng Zhao, Filip Tronarp, Roland Hostettler, and Simo Särkkä. 2020. State-Space Gaussian Process for Drift Estimation in Stochastic Differential Equations. (2020), 5295–5299.

A OBJECTIVE FUNCTION

Derivation of the likelihood term in the objective (left term in Eq. (5)):

$$\begin{aligned} \mathbb{E}_{q(Z, \theta_f, \tilde{\theta}_f, \tilde{z}_0 | X)} \left[\log p(X | Z, \theta_f, \tilde{z}_0, \tilde{\theta}_f) \right] &= \mathbb{E}_{q(Z, \theta_f, \tilde{\theta}_f, \tilde{z}_0 | X)} \left[\log \prod_{t=0}^{T-1} p(x_t | z_t) \right] \\ &= \sum_{t=0}^{T-1} \mathbb{E}_{q(Z, \theta_f, \tilde{\theta}_f, \tilde{z}_0 | X)} [\log p(x_t | z_t)] \\ &= \sum_{t=0}^{T-1} \mathbb{E}_{q(z_t | X)} [\log p(x_t | z_t)]. \end{aligned}$$

This follows since for $t' \neq t$: $x_t \perp (z_{t'}, \theta_f, \tilde{\theta}_f, \tilde{z}_0) | z_t$.

Before decomposing the KL term, we note that the conditionals $p(z_0 | \tilde{z}_0)$ and $p(\theta_f | \tilde{\theta}_f)$ are deterministic, meaning they are Dirac functions with the peak defined by Eqs. (3) and (4). The transition distribution $p(z_t | z_{t-1}, \theta_f)$ is also a Dirac function with the peak defined by Eq. 1 as stated in Section 4.

Therefore, The KL term from Eq. (5) can be written as:

$$\begin{aligned} KL \left[q(Z, \theta_f, \tilde{z}_0, \tilde{\theta}_f | X) || p(Z, \theta_f, \tilde{z}_0, \tilde{\theta}_f) \right] &= \\ \int_Z \int_{\theta_f} \int_{\tilde{z}_0} \int_{\tilde{\theta}_f} q(\tilde{z}_0 | X) q(\tilde{\theta}_f | X) q(z_0 | \tilde{z}_0) q(\theta_f | \tilde{\theta}_f) \prod_{t=1}^{T-1} q(z_t | z_{t-1}, \theta_f) \cdot \\ \cdot \log \left[\frac{p(\tilde{z}_0) p(\tilde{\theta}_f) p(z_0 | \tilde{z}_0) p(\theta_f | \tilde{\theta}_f) \prod_{t=1}^{T-1} p(z_t | z_{t-1}, \theta_f)}{q(\tilde{z}_0 | X) q(\tilde{\theta}_f | X) q(z_0 | \tilde{z}_0) q(\theta_f | \tilde{\theta}_f) \prod_{t=1}^{T-1} q(z_t | z_{t-1}, \theta_f)} \right]. \end{aligned}$$

The KL term we got, decomposes into the sum of 3 terms:

(i) The first term:

$$\begin{aligned} \int_Z \int_{\theta_f} \int_{\tilde{z}_0} \int_{\tilde{\theta}_f} q(\tilde{z}_0 | X) q(\tilde{\theta}_f | X) q(z_0 | \tilde{z}_0) q(\theta_f | \tilde{\theta}_f) \prod_{t=1}^{T-1} q(z_t | z_{t-1}, \theta_f) \log \left[\frac{p(\tilde{z}_0) p(z_0 | \tilde{z}_0)}{q(\tilde{z}_0 | X) q(z_0 | \tilde{z}_0)} \right] &= \\ \int_{\tilde{z}_0} q(\tilde{z}_0 | X) \log \left[\frac{p(\tilde{z}_0)}{q(\tilde{z}_0 | X)} \right] \int_Z \int_{\theta_f} \int_{\tilde{\theta}_f} q(\tilde{\theta}_f | X) q(z_0 | \tilde{z}_0) q(\theta_f | \tilde{\theta}_f) \prod_{t=1}^{T-1} q(z_t | z_{t-1}, \theta_f) &= \\ KL(q(\tilde{z}_0 | X) || p(\tilde{z}_0)), \end{aligned}$$

where $p(z_0 | \tilde{z}_0) = q(z_0 | \tilde{z}_0)$ by construction since both are determined exactly by Eq. (3).

(ii) In the same way, we get:

$$\begin{aligned} \int_Z \int_{\theta_f} \int_{\tilde{z}_0} \int_{\tilde{\theta}_f} q(\tilde{z}_0 | X) q(\tilde{\theta}_f | X) q(z_0 | \tilde{z}_0) q(\theta_f | \tilde{\theta}_f) \prod_{t=1}^{T-1} q(z_t | z_{t-1}, \theta_f) \log \left[\frac{p(\tilde{\theta}_f) p(\theta_f | \tilde{\theta}_f)}{q(\tilde{\theta}_f | X) q(\theta_f | \tilde{\theta}_f)} \right] &= \\ KL(q(\tilde{\theta}_f | X) || p(\tilde{\theta}_f)), \end{aligned}$$

where $p(\theta_f | \tilde{\theta}_f) = q(\theta_f | \tilde{\theta}_f)$ by construction since both are determined exactly by Eq. (4).

(iii) The last term:

$$\int_Z \int_{\theta_f} \int_{\tilde{z}_0} \int_{\tilde{\theta}_f} q(\tilde{z}_0 | X) q(\tilde{\theta}_f | X) q(z_0 | \tilde{z}_0) q(\theta_f | \tilde{\theta}_f) \prod_{t=1}^{T-1} q(z_t | z_{t-1}, \theta_f) \log \left[\frac{\prod_{t=1}^{T-1} p(z_t | z_{t-1}, \theta_f)}{\prod_{t=1}^{T-1} q(z_t | z_{t-1}, \theta_f)} \right] = 0,$$

where $p(z_t | z_{t-1}, \theta_f) = q(z_t | z_{t-1}, \theta_f)$ by construction since both are determined exactly by the ODE system f . Thus the logarithmic term equals 0.

B ALGORITHMS

The algorithms below give the training procedure for a single iteration and a batch of size 1. The extension to larger batch sizes is straightforward. X^i denotes a time sequence of length T for an observed signal.

At inference time, we are given an observed signal X^i of length T , and extrapolate the signal to time $T + \tau$. Meaning, the ODE-solver for-loop is from time $t = 1$ to $t = T + \tau - 1$, and the output sequence $\hat{X}^i$ is of length $T + \tau$.

Algorithm 1 gives the training procedure for GOKU-net.

Algorithm 2 gives the training procedure for the Direct Inference (DI) baseline described in Section 5. An important note is that this baseline is also given a latent space trajectory denoted as Z^i , and a grounding mask indicator denoted as M^i (Eq. (6)). These additional signals are not accessed by other methods. In this baseline, we first learn the parameters $\hat{\theta}_f^i$ and the initial state of the ODE $\hat{z}_0$ of every signal in the train and test sets. We then evaluate $\hat{Z}^i$ for all signals in the train set, and use these predictions to learn the emission function $\hat{g}$, using the given train set signals X^i . In some cases, learning $\hat{z}_0^i$ was too difficult for the baseline so we tried a different approach. We used the learned parameters $\hat{\theta}_f^i$, the given ODE function f and the first observed latent vector $z_{t'}^i$, (meaning the mask $M^i(t') = 1$ and $M^i(t < t') = 0$), and used the ODE solver to calculate $\hat{z}_0^i$ backwards in time.

Algorithm 3 addresses the unknown-unknowns task. This algorithm is very similar to Algorithm 1, with the changes highlighted in blue. The main changes are that this method also includes an abstract function f_{abs} which models the *Unknown Unknowns* part of the ODE. I.e., the ODE is changed to be: $\frac{dz_t}{dt} = f_{ODE}(z_t, \theta_f) + f_{abs}(z_t)$.

Algorithm 1 GOKU-net

Input:

1. sequence $X^i = (x_0, \dots, x_{T-1})$
2. ODE function f
3. ODE solver
4. hyper-parameter λ_1

initialize the neural nets $\phi_{z_0}^{enc}$, $\phi_{\hat{\theta}_f}^{enc}$, h_z , h_θ and $\hat{g}$.

$$[\mu_{\hat{z}_0}, \sigma_{\hat{z}_0}] = \phi_{z_0}^{enc}(X^i), \quad \hat{z}_0 \sim \mathcal{N}(\mu_{\hat{z}_0}, \sigma_{\hat{z}_0}), \quad \hat{z}_0 = h_z(\hat{z}_0)$$

$$[\mu_{\hat{\theta}_f}, \sigma_{\hat{\theta}_f}] = \phi_{\hat{\theta}_f}^{enc}(X^i), \quad \hat{\theta}_f \sim \mathcal{N}(\mu_{\hat{\theta}_f}, \sigma_{\hat{\theta}_f}), \quad \hat{\theta}_f = h_\theta(\hat{\theta}_f)$$

for $t = 1, \dots, T - 1$ **do**

$$\hat{z}_t = \text{ODEsolver}(f, \hat{\theta}_f, \hat{z}_{t-1})$$

end for

$$\hat{X}^i = \hat{g}(\hat{Z}^i); \quad \{\hat{Z}^i = (\hat{z}_0, \dots, \hat{z}_{T-1}) \text{ for sample } i\}$$

$$ll_loss = \mathcal{L}_{likelihood}(X^i, \hat{X}^i); \quad \{\text{see first term in Eq. (5)}\}$$

$$kl_loss = \mathcal{L}_{kl}(\mu_{\hat{z}_0}, \sigma_{\hat{z}_0}, \mu_{\hat{\theta}_f}, \sigma_{\hat{\theta}_f}); \quad \{\text{see second term in Eq. (5)}\}$$

$$loss = ll_loss + \lambda_1 kl_loss$$

$$\text{backpropagate}(loss)$$

Algorithm 2 Direct Inference (DI)

Input:

1. ODE function f
2. ODE solver
3. train and test sets of observed signals X^i , Z^i and M^i .

for train and test sets **do**

Initialize $\hat{\theta}_f, \hat{z}_0$

for $t = 1, \dots, T - 1$ **do**

$$\hat{z}_t = \text{ODEsolver}(f, \hat{\theta}_f, \hat{z}_{t-1})$$

end for

$$loss = ||Z^i - \hat{Z}^i||_2$$

$$\hat{\theta}_f := \hat{\theta}_f + \lambda \frac{\partial loss}{\partial \hat{\theta}_f} \quad \{\text{backpropagate loss through the ODE solver}\}$$

$$\hat{z}_0 := \hat{z}_0 + \lambda \frac{\partial loss}{\partial \hat{z}_0}$$

end for

for train set **do**

$$\hat{X}^i = \hat{g}(\hat{Z}^i)$$

$$generative_loss = ||X - \hat{X}^i||_2$$

$$\text{backpropagate}(generative_loss)$$

end for

C EXPERIMENTS

We provide here more information about the experiments described in Section 5. We ran all of the experiments on a desktop CPU. Full code implementation for creating the datasets, implementing GOKU-net, and implementing baselines is available on github.com/orilnial/GOKU.

Algorithm 3 GOKU with Unknown Unknowns (GOKU-UU)**Input:**

1. sequence $X^i = (x_0, \dots, x_{T-1})$
2. ODE function f
3. ODE solver
4. hyper-parameter λ_1

initialize the neural nets f_{abs} , $\phi_{z_0}^{enc}$, $\phi_{\theta_f}^{enc}$, h_z , h_θ and $\hat{g}$.

$[\mu_{z_0}, \sigma_{z_0}] = \phi_{z_0}^{enc}(X^i)$, $\tilde{z}_0 \sim \mathcal{N}(\mu_{z_0}, \sigma_{z_0})$, $z_0 = h_z^{ODE}(\tilde{z}_0)$
 $[\mu_{\theta_f}, \sigma_{\theta_f}] = \phi_{\theta_f}^{enc}(X^i)$, $\tilde{\theta}_f \sim \mathcal{N}(\mu_{\theta_f}, \sigma_{\theta_f})$, $\hat{\theta}_f = h_\theta^{ODE}(\tilde{\theta}_f)$

for $t = 1, \dots, T - 1$ **do**

$z_t^{ODE} = \text{ODEsolver}(f + f_{abs}, \hat{\theta}_f, z_{t-1})$

end for

$\hat{X}^i = \hat{g}(\hat{Z}^i)$

$\text{loss} = \mathcal{L}_{\text{likelihood}}(X^i, \hat{X}^i) + \lambda_1 \mathcal{L}_{kl}(\mu_{z_0}, \sigma_{z_0}, \mu_{\theta_f}, \sigma_{\theta_f})$ {see Eq. (5)}

backpropagate(loss)

C.1 Hyper parameter selection

All hyper parameters used for dataset creation, and for all methods, are explicitly declared in the code. In this section we describe the methods for selecting the hyper parameters.

KL distance. In GOKU-net and L-ODE, we set the KL hyper parameter initial value to 10^{-5} , with a KL annealing scheme so that it would gradually increase to the value of 1. To choose the KL hyper parameter initial value, we first set the KL hyper parameter to initial value of 1, and divided by 10 until finding the initial value that provided the best results on the validation set.

Optimization. In all methods, we experimented with batch sizes from the set $\{16, 32, 64, 128\}$. We found that batch size of 64 provided the best results on the validation set for all methods in all experiments. In all methods we used Adam optimizer with learning rate 10^{-3} . We also experimented with other learning rates 10^{-2} and 10^{-4} , which did not provide better results.

C.2 Single Pendulum From Pixels

Algorithm implementation details. For all algorithms we used an input-to-RNN network and emission function exactly as suggested in Greydanus et al. [17], composed of four fully-connected layers with ReLU activations and residual connections. The output of the input-to-RNN net dimension is 32.

In GOKU-net, for $\phi_{z_0}^{enc}$ we used an RNN with hidden dimension of 16 followed by a linear transformation to μ_{z_0} and another linear transformation to σ_{z_0} , both with dimension of 16. $\phi_{\theta_f}^{enc}$ is very similar to $\phi_{z_0}^{enc}$ except for using a bi-directional LSTM instead of an RNN. The h functions are implemented with an MLP with one hidden layer with 200 neurons. The output of h_z is linear, and the output of h_θ is forced to be positive using the softplus activation, so that θ_{pendulum} would be physically feasible. In Latent ODE, we used input dimension of 16 for the RNN, which then transforms linearly to μ_{z_0} and σ_{z_0} with dimension of 16 as well. The ODE function f_{abs} is modeled as a neural network of sizes $4 \rightarrow 200 \rightarrow 200 \rightarrow 4$ with ReLU activation. In the LSTM baseline, we used an LSTM with 4 layers and a hidden size of 16, followed by the same emission function as GOKU. In DMM, the inference model is built by first using the same input-to-RNN network as GOKU, followed by a combiner function that averages the previous z_{t-1} and the RNN output h_t to produce z_t . We set z dim to be 32. The generative model is built by first using a transition function which is a $32 \rightarrow 200 \rightarrow 200 \rightarrow 32$ NN with ReLU activation function, and then using the same emission function as GOKU. In HNN, We used the code provided by [17]. The only change we made is in the dataset creation process, l is uniformly sampled instead of being constant.

C.3 Pixel Pendulum with Unknown Unknowns

In this experiment we aimed to show how GOKU can be modified to handle unknown unknowns in the ODE: We are given an ODE system that only partially describes the system that created the data. Specifically in this scenario, the pixel-pendulum data is created with a friction model:

$$\frac{d\theta(t)}{dt} = \omega(t), \quad \frac{d\omega(t)}{dt} = -\frac{g}{l} \sin \theta(t) - \frac{b}{m} \omega(t),$$

and we are only given with the friction-less ODE system in Eq. (7). Our method (Algorithm 3) models the time derivatives of the unknown part, making the ODE functional form as:

$$\frac{dz_t}{dt} = f_{ODE}(z_t, \theta_f) + f_{abs}(z_t),$$

where f_{abs} is modeled as a neural network.

Data set. We created this data set in the same way as in the friction-less pixel-pendulum experiment. Here we set $l \sim U[1, 2]$ as in the non-friction experiment, and we set in addition $m = 1$, $b = 0.7$.

Algorithm implementation details. The only difference between this experiment and the pixel pendulum experiment, is that in GOKU-UU, we added a neural network that models f_{abs} , which is implemented as a fully connected network with $2 \rightarrow 200 \rightarrow 200 \rightarrow 2$ layers and ReLU activations.

Results. In Fig. 6 we compare the X extrapolation error between GOKU-UU and the baselines. The observed signal is of length $T = 50$, and we extrapolated the signals until $T = 200$. In Fig. 7 we demonstrate the extrapolation of X , by randomly selecting one test sample and showing the pendulum’s predicted angle for future times. In both we observe that GOKU-UU achieved much better results than the compared baselines. In Fig. 5 we demonstrate that GOKU-UU’s added function f_{abs} , learned only the friction part. This is done by first training using GOKU-UU, and then zeroing f_{abs} during test time. Fig. 5 shows that the signal with f_{abs} zeroed extrapolates as if there was no friction at all, suggesting that we successfully separated the friction model from the pendulum model. We also tested if GOKU-UU could perform θ_f identification, and obtained the following results: correlation coefficient = 0.967, and L_1 error of 0.109 ± 0.013 .

These results show that using GOKU-net with the unknown-unknowns modification can successfully identify the ODE parameters and extrapolate the observed signal, although it does not observe the full ODE functional form. Moreover, it demonstrates capability to separate the Known-Unknowns (the given ODE’s state and parameters) from the Unknown-Unknowns (the friction).

Figure 6: Pixel pendulum with friction - mean extrapolation error for observations X over time steps after end of input sequence.

C.4 Double Pendulum From Pixels

In this experiment we used a double-pendulum ODE system, which has significantly more complicated dynamics than the single pendulum one. Indeed, double pendulums are known to be chaotic systems. The ODE we used is exactly the one described in Acrobot-v1 environment from OpenAI Gym [5]. The ODE state in time t is defined by the angles and angular velocities of the two pendulums: $z_t = (\theta_1, \omega_1, \theta_2, \omega_2)$.

Figure 7: Pixel pendulum with friction predicted angle example. Comparing GOKU-UU to the baselines.

The ODE system is therefore:

$$\begin{aligned}
 \frac{d\theta_1(t)}{dt} &= \omega_1 \\
 \frac{d\theta_2(t)}{dt} &= \omega_2 \\
 \frac{d\omega_1(t)}{dt} &= -\left(d_2 \frac{d\omega_2(t)}{dt} + \phi_1\right) \frac{1}{d_1} \\
 \frac{d\omega_2(t)}{dt} &= \frac{\frac{d_2}{d_1} \phi_1 - \phi_2}{m_2 c_2^2 + I_2 - \frac{d_2^2}{d_1}},
 \end{aligned}$$

where the auxiliary functions d_1, d_2, ϕ_1, ϕ_2 are declared as:

$$\begin{aligned}
 d_1 &= m_1 c_1^2 + m_2 (l_1^2 + (c_2)^2 + 2l_1 c_2 \cos(\theta_2)) + I_1 + I_2 \\
 d_2 &= m_2 (c_2^2 + l_1 c_2 \cos(\theta_2)) + I_2 \\
 \phi_2 &= m_2 c_2 g \cos(\theta_1 + \theta_2 - 0.5\pi) \\
 \phi_1 &= -m_2 l_1 c_2 \omega_2^2 \sin(\theta_2) - 2m_2 l_1 c_2 \omega_2 \omega_1 \sin(\theta_2) + (m_1 l_1 + m_2 l_1) g \cos(\theta_1 - 0.5\pi) + \phi_2.
 \end{aligned}$$

The parameters m_1, m_2 are the masses of the pendulums, l_1, l_2 are the lengths, I_1, I_2 are the moments of inertia, and c_1, c_2 are the position of the center of mass for each pendulum.

Dataset. We created this data set in a similar way to the pixel-pendulum experiment. For training we simulated 500 sequences of 50 time points, with time steps of $\Delta_t = 0.05$ and pre-processed the observed signals in exactly the same way as in the single pendulum experiment, but with frame size of 32×32 . The ODE parameter we aimed to infer is m_2 and was uniformly sampled $m_2 \sim U[1, 2]$. The other parameters were assumed to be constant and known as described in Acrobot-V1 (all parameters with values 1.0, except $c_1 = c_2 = 0.5$ and $g = 9.8$). The ODE initial state z_0 was uniformly sampled, $z_0 \sim U[\frac{\pi}{10}, \frac{\pi}{6}]^4$. Each test set sequence is 100 time steps long, where the first 50 time steps are given as input, and the following 50 were used only for evaluating the signals extrapolation.

Algorithm implementation details and Results. The implementation of this task is exactly as in the pixel pendulum task, except the obvious difference of incorporating the double pendulum ODE instead of the single pendulum one for GOKU and DI. Results are described fully in Section 5.

C.5 CVS

We first give the complete ODE we used in this experiment. This is simplified form of the model given in Zenker et al. [38]:

$$\begin{aligned}\frac{dSV(t)}{dt} &= I_{\text{external}}, \\ \frac{dP_a(t)}{dt} &= \frac{1}{C_a} \left(\frac{P_a(t) - P_v(t)}{R_{TPR}(S)} - SV \cdot f_{HR}(S) \right) \\ \frac{dP_v(t)}{dt} &= \frac{1}{C_v} \left(-C_a \frac{dP_a(t)}{dt} + I_{\text{external}} \right) \\ \frac{dS(t)}{dt} &= \frac{1}{\tau_{Baro}} \left(1 - \frac{1}{1 + e^{-k_{width}(P_a(t) - P_{aset})}} - S \right),\end{aligned}$$

where

$$\begin{aligned}R_{TPR}(S) &= S(t)(R_{TPR_{Max}} - R_{TPR_{Min}}) + R_{TPR_{Min}} + R_{TPR_{Mod}} \\ f_{HR}(S) &= S(t)(f_{HR_{Max}} - f_{HR_{Min}}) + f_{HR_{Min}}.\end{aligned}$$

In this model the variables have a directly interpretable mechanistic meaning: SV , P_a , P_v , S are respectively cardiac stroke volume (the amount of blood ejected by the heart), arterial blood pressure, venous blood pressure and autonomic baroreflex tone (the reflex responsible for adapting to perturbations in blood pressure, keeping homeostasis). In Table 6 we give information about the meaning of the ODE state variables and ODE parameters.

Beyond I_{external} and $R_{TPR_{Mod}}$ (denoted R_{Mod} for brevity in the main text), the CVS ODE model has several other parameters which we treat as known, setting them to the values stated in Zenker et al. [38]. The ODE state is $z_t = (SV(t), P_a(t), P_v(t), S(t))$. The observed state is the patient's vital signs and defined to be: $x_t = (P_a(t), P_v(t), f_{HR}(t))$, where $f_{HR}(t)$ is the patient's heart-rate at time t . Note that two of the three observed variables are the same as two of the latent variables, albeit with added noise as we explain now.

Dataset. We simulated 1000 sequences of length 400. The parameter I_{external} was randomly sampled to be either -2 or 0 , and the parameter $R_{TPR_{Mod}}$ was randomly sampled to be either 0.5 or 0 . Initial ODE states uniformly sampled from $SV(0) \sim U[90, 100]$, $P_a(0) \sim U[75, 85]$, $P_v \sim [3, 7]$ and $S \sim [0.15, 0.25]$, where the intervals were set to the values given in [38]. Observations were additionally corrupted with white Gaussian noise with standard deviation $\sigma_x = 5$ for P_a , $\sigma_x = 0.5$ for P_v and $\sigma_x = 0.05$ for f_{HR} (standard deviation matches scale of the observed signal).

Algorithm implementation details. For all algorithms we used an input-to-rnn network of 2 fully connected layers with ReLU activation and 64 hidden units, and output with dimension of 64.

In GOKU, The RNN and LSTM are implemented as in the LV experiment with output of dimension 64, followed by a linear transformation to μ_{z_0} and σ_{z_0} of dimension 64 as well. The h functions are implemented as in the LV, except that their output has a sigmoid activation layer, to bound them to a physically feasible solution. The emission function is a takes P_a and P_v from the latent trajectories, and a fully connected $4 \rightarrow 200 \rightarrow 1$ network with ReLU activation layer to compute f_{HR} . In Latent ODE the emission function is a fully connected $4 \rightarrow 200 \rightarrow 3$ network with ReLU activation layer. In LSTM, we used the same network as in the LV experiment. In DMM, the inference model is built by first using the same input-to-RNN network as GOKU, followed by a combiner function that averages the previous z_{t-1} and the RNN output h_t to produce z_t . We set z dim to be 8. The generative model is build by first using a transition function which is a $8 \rightarrow 200 \rightarrow 200 \rightarrow 8$ NN with ReLU activation function, and then using the same emission function as GOKU.

Symbol	Description	Unit
SV	Stroke volume, the volume of blood ejected during 1 cardiac cycle/ejection period	ml
P_a	Pressure in arterial compartment	mm Hg
P_v	Pressure in venous compartment	mm Hg
S	Autonomic baroreflex tone, i.e., the reflex responsible for adapting to perturbations in blood pressure, keeping homeostasis	-
f_{HR}	Heart rate, i.e., the number of complete cardiac cycles per unit time	Hz
R_{TPR}	Total peripheral/systemic vascular hydraulic resistance, i.e., the hydraulic resistance opposing the flow through the capillary streambed that is driven by the arterio-venous pressure difference	mm HG s/ml
C_a, C_v	Compliance of arterial, venous compartment	ml/mm Hg
τ_{Baro}	Time constant of the baroreflex response, i.e., of the linear low pass characteristic of the physiological negative feedback loop controlling arterial pressure	s
k_{width}	Constant determining the shape and maximal slope of the logistic baroreflex nonlinearity	mm Hg ⁻¹
P_{aset}	Set point of the baroreflex feedback loop	mm Hg
$I_{external}$	Possible external blood withdrawal or fluid infusion to or from the venous compartment	ml/s
$R_{TPR_{Mod}}$	Possible modification in R_{TPR}	mm Hg s/ml

Table 6: Glossary of Variables and Parameters of the Cardiovascular Model, as shown in Zenker et al. [38]