

Forced Variational Integrator Networks for Prediction and Control of Mechanical Systems

Aaron Havens

AHAVENS2@ILLINOIS.EDU

Girish Chowdhary

GIRISHC@ILLINOIS.EDU

Department of ECE & Coordinated Science Laboratory

University of Illinois at Urbana-Champaign, Champaign, IL, USA, 61820

Abstract

As deep learning becomes more prevalent for prediction and control of real physical systems, it is important that these overparameterized models are consistent with physically plausible dynamics. This elicits a problem with how much inductive bias to impose on the model through known physical parameters and principles to reduce complexity of the learning problem to give us more reliable predictions. Recent work employs discrete variational integrators parameterized as a neural network architecture to learn conservative Lagrangian systems. The learned model captures and enforces global energy preserving properties of the system from very few trajectories. However, most real systems are inherently non-conservative and, in practice, we would also like to apply actuation. In this paper we extend this paradigm to account for general forcing (e.g. control input and damping) via discrete d'Alembert's principle which may ultimately be used for control applications. We show that this forced variational integrator networks (FVIN) architecture allows us to accurately account for energy dissipation and external forcing while still capturing the true underlying energy-based passive dynamics. We show that in application this can result in highly-data efficient model-based control and can predict on real non-conservative systems.

Keywords: Variational Integrators, Physics-constrained learning, Model-based Control

1. Introduction

Deep neural networks (DNN) have seen increasing success in scaling and generalizing to large data sets for classification, prediction and even control (Krizhevsky et al. (2012); Silver et al. (2018)). However when it comes to physic systems like autonomous robots, generalization is especially important so that interpolations/extrapolations of training data still results in physically plausible and safe behavior (Guiochet et al. (2017)). In these physical applications we often choose to engineer carefully parameterized physics-based models over deep overparametrized models to ensure this level of generalization. These expertly engineered models impose a great amount of inductive bias (fundamental assumptions about the output function), but as we know incorrect bias can be detrimental to decision-making agents and fully modeling every contributing process can be costly and laborious (Dullerud and Paganini (2013); Janner et al. (2019)). Clearly these systems all share similar physical principles such as conservation laws and symmetries. How can we use highly flexible deep architectures and impose inductive bias in a way that generally applies to a large number of physical systems? Better yet, can they be interpretable through existing physical theory?

This is a challenging problem since there are real conserved and invariant quantities which make the difference between a system being stable or unstable and have major consequences in both long-term prediction and control. Greydanus et al. (2019) highlights these issues and provides a method

of directly parameterizing a system Hamiltonian with a DNN and then reconstructing the dynamics with classical mechanics techniques.

On the other hand [Chen et al. \(2018\)](#) interprets deep residual networks as Euler discretized continuous time ordinary differential equations. [Saemundsson et al. \(2020\)](#) provides a natural link from this perspective to a class of discretization techniques from discrete mechanics called Variational Integrators (VI) ([Marsden and West \(2001\)](#)). Variational integrators formally discretize the action of a Lagrangian system, rather than the derivative, which approximately preserves energy and momentum (otherwise known as a symplectic integrator). [Saemundsson et al. \(2020\)](#) uses an explicit variational integrator structure to propose VI networks (VIN) that can learn dynamics of conservative Lagrangian systems with very accurate energy behavior. The resulting inductive bias from this VI discretization yields data-efficient generalization for a large class of systems. However, none of the aforementioned methods take into account external forcing and non-conservative systems which are abundant in the physical world. [Lutter et al. \(2019\)](#) proposes an architecture that incorporates forces by learning the inertial matrix and the potential function by minimizing violations of Lagrangian dynamics. However the proposed loss function does not account for passive energy dissipation and doesn't explicitly enforce any kind of learned conservation of preserved quantities. Perhaps most related to our work, [Hochlehnert et al. \(2021\)](#) concurrently proposed methods for learning conservative VIN's with discontinuous contact dynamics which could be complementary to this work.

Contributions In this paper we propose a direct extension of [Saemundsson et al. \(2020\)](#) by applying the forced discrete variational principle *Discrete Lagrange-d'Alembert Principle* ([Junge et al. \(2005\)](#); [Marsden and West \(2001\)](#)). With this extension we are able to introduce general non-conservative forcing into an explicit VI that yields a Forced Variational Integrator Network (FVIN) learning architecture. FVIN can handle control inputs and dissipation while still enjoying the long term energy tracking properties of the original VIN, enabling a much wider scope of applications in control. We also highlight the efficient generalization properties of FVIN in low data regimes as few as 5 trajectories on two nonlinear nonconservative systems. We further demonstrate the importance of this extension by applying FVIN to predict a real non-conservative cart-pole system from data.

2. Preliminaries

We are motivated to introduce a class of integrators which define a discrete-time dynamical system which preserves global energy and momentum properties of the corresponding continuous system. A brief background of Lagrangian mechanics with forcing will be presented along with their lesser-known discrete analogs in the form of variational integrators. Specifically, we look to control mechanical systems whose Lagrangian $L : TQ \rightarrow \mathbb{R}$ is given as the difference of kinetic and potential energies.

$$L(q, \dot{q}) = T(\dot{q}) - V(q) = \frac{1}{2} \dot{q}^T M \dot{q} - V(q) \quad (1)$$

Here we have a configuration state q belonging to Q which, in general, could be an n -dimensional configuration manifold. However in this paper we restrict our focus to integrators on the euclidean space \mathbb{R}^n . The state space is then given by the tangent bundle TQ where at a point q we have position and velocity (q, \dot{q}) . Using this interpretation the state space TQ is the euclidean space \mathbb{R}^{2n} . M is a symmetric positive definite mass matrix and $V(q)$ is the energy potential function. We also consider external forcing to account for actuators and passive dissipation with the Lagrangian force

$f_L : TQ \rightarrow T^*Q$, where T^*Q is the cotangent bundle. We can obtain the governing dynamics by finding extremizing curves in $\mathcal{C}(Q)$ that minimize an action functional, where $\mathcal{C}(Q) = \{q : [0, T] \rightarrow Q \mid q \in \mathcal{C}^2\}$ is the path space (essentially twice differentiable curves). These extreme curves can be characterized by *Lagrange-d'Alembert principle* (Vujanovic (1978)) and satisfy

$$\delta \int_0^T L(q(t), \dot{q}(t)) dt + \int_0^T f_L(q(t), \dot{q}(t), u(t)) \cdot \delta q(t) dt = 0 \quad (2)$$

where δ are variations of $\mathcal{C}(Q)$ such that $\delta q(0) = \delta q(T) = 0$. This necessary condition gives us a description of the configuration dynamics resembling a forced variant of the standard Euler-Lagrange equations.

$$\frac{\partial L}{\partial q}(q, \dot{q}) - \frac{d}{dt} \left(\frac{\partial L(q, \dot{q})}{\partial \dot{q}} \right) + f_L(q, \dot{q}, u) = 0 \quad (3)$$

At this point, one could discretize Eq 3 directly to obtain an explicit forward Euler scheme with time step h . This could be interpreted as a residual dynamic network update rule. But then what can we say about aforementioned energy conservation properties? Since discretization was performed directly on the resulting ODE, in some sense we lose information given by the Lagrangian. We can instead apply discretization to the integral Eq. 2 directly and derive a new integration rule which will more accurately conserve total energy under the action (sometimes called symplectic integrators). This is the purpose of *Variational Integrators* (Wendlandt and Marsden (1997); Marsden and West (2001)) which define implicit and explicit discrete update rules using a discrete analog to the Lagrange-d'Alembert principle.

2.1. Variational Integrators

Variational integrators are aimed to provide a discrete analog to the continuous time Euler-Lagrange principle, which involves taking an approximation of the action integral. This idea is introduced by the *Discrete Lagrangian* defined by a quadrature-based approximation of the action integral found in Eq. 2.

$$L_d(q_k, q_{k+1}) \approx \int_{t_k}^{t_{k+1}} L(q(t), \dot{q}(t)) dt$$

We approximate the Lagrangian forcing by defining left $f_d^-(q_k, q_{k+1}, u_k)$ and right $f_d^+(q_k, q_{k+1}, u_k)$ *discrete Lagrangian forces* which are then used to approximate the variation of the continuous external force in Eq. 2 over a small interval:

$$f_d^-(q_k, q_{k+1}, u_k) \cdot \delta q_k + f_d^+(q_k, q_{k+1}, u_k) \cdot \delta q_{k+1} \approx \int_{t_k}^{t_{k+1}} f(q(t), \dot{q}(t), u(t)) \cdot \delta q dt$$

Using these approximations with a similar variational argument yields a discrete analog to our forced Euler-Lagrange equations. For each k the following must hold, defining an integrator implicitly. Here D_i denotes the slot derivative differentiating with respect to the i -th argument

$$D_2 L_d(q_{k-1}, q_k) + D_1 L_d(q_k, q_{k+1}) + f_d^+(q_{k-1}, q_k, u_k) + f_d^-(q_k, q_{k+1}, u_k) = 0 \quad (4)$$

In view of the discrete Legendre Transform, described in great detail in Marsden and West (2001), we can obtain an equivalent position-momentum form corresponding to the discrete Hamiltonian where the momentum for the Newtonian system is given by $p = \frac{\partial L}{\partial \dot{q}} = M\dot{q}$.

$$p_k = -D_1 L_d(q_k, q_{k+1}) - f_d^-(q_k, q_{k+1}, u_k), \quad p_{k+1} = D_2 L_d(q_k, q_{k+1}) + f_d^+(q_k, q_{k+1}, u_k) \quad (5)$$

Given the Lagrangian in the form of Eq. 1, we can actually write both integrators (4, 5) explicitly as a function of (q_k, q_{k-1}, u_k) and (q_k, \dot{q}_k, u_k) respectively. These explicit integrators will be introduced in the following section.

2.2. Explicit Integrators with Forcing

The integrators that will be ultimately useful in designing a feed-forward network architecture will be those which involve no numerical scheme and give the next state explicitly in terms of previous states and controls. To obtain this explicit integrator we take advantage of our chosen approximation to the discrete Lagrangian. We can use a symmetric quadrature rule to approximate the discrete Lagrangian to derive both the position and position-momentum integrators explicitly:

$$L_d(q_k, q_{k+1}) = \frac{h}{2} \left(L \left(q_k, \frac{q_{k+1} - q_k}{h} \right) + L \left(q_{k+1}, \frac{q_{k+1} - q_k}{h} \right) \right) \quad (6)$$

Here we refer to the Lagrangian given in Eq. 1 and a uniform time-step h although in principle it can vary. For the position integrator, we choose a simplified form of the discrete forces under a *zero-order hold* assumption distributed only over f_d^+ , where $F(q_k, u_k)$ maps control inputs approximately to discrete Lagrangian forces

$$\begin{aligned} f_d^-(q_k, q_{k+1}, u_k) &= hF(q_k, u_k) \\ f_d^+(q_k, q_{k+1}, u_k) &= 0 \end{aligned}$$

This choice in particular will give us an explicit position integrator akin to the **Stormer-Verlet** (SV) integrator Verlet (1967)

$$q_{k+1} = 2q_k - q_{k-1} + h^2 M^{-1} (F(q_k, u_k) - \nabla V(q_k)) \quad (7)$$

For our position-velocity integrator we choose a symmetric representation of the discrete forces which are distributed evenly over both f_d^+ and f_d^-

$$f_d^\pm(q_k, q_{k+1}, u_k) = \frac{h}{2} F(q_k, u_k)$$

This yields an explicit position-velocity integrator otherwise known as **Velocity-Verlet** (VV) integrator. Note again for our Newtonian system that momentum and velocity relate by $p = M\dot{q}$

$$\begin{aligned} q_{k+1} &= q_k + h\dot{q} + \frac{h^2}{2} M^{-1} (F(q_k, u_k) - \nabla V(q_k)) \\ \dot{q}_{k+1} &= \dot{q}_k + hM^{-1} \left(F(q_k, u_k) - \frac{\nabla V(q_k) + \nabla V(q_{k+1})}{2} \right) \end{aligned} \quad (8)$$

Now that we have an explicit expression for the state evolution, we can inspire a feed-forward architecture where functions V and F be learned directly from data from possibly non-conservative systems. Focusing on the VV integrator, we show that this structure can effectively capture these terms in physically meaningful ways, being able to precisely account for multiple types of forcing even in very low data regimes.

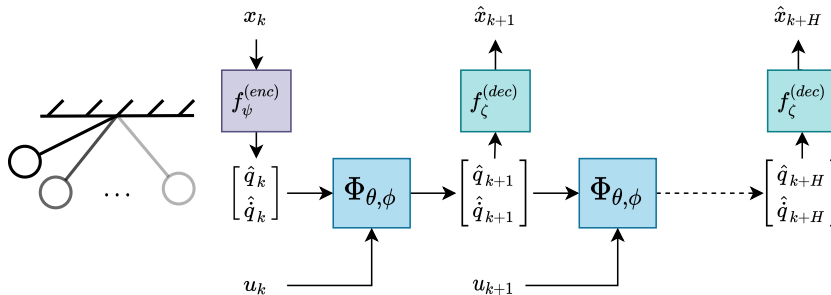


Figure 1: The FVIN predicts ahead multiple steps from an initial condition and a sequence of controls in an open-loop fashion. The loss is then computed against these roll-outs to encourage long-term consistency.

3. Forced Variational Integrator Networks

With an explicit integrator that supports forcing in hand, we may now discuss how this structure can be used to design a general purpose neural network predictor for non-conservative Newtonian systems. Enforcing that a feed-forward neural network takes on the form of a variational integrator lets us arrive at predictors with correct global energy and interpretable forcing (e.g. damping, actuation) with much fewer samples. As a result, such architectures (like the integrators they are inspired by) have better long term prediction and can be used to implement higher performing model-based controllers for a broad class of systems. We begin by considering the case where configuration states are measured directly. The SV (7) integrator takes the argument (q_k, q_{k-1}) where VV (8) takes in (q_k, \dot{q}_k) . The potential function and forcing function are parameterized by V_θ and F_ϕ as simple feed-forward neural networks. We choose to absorb the mass matrix M into these learned objects, although, in principle, it may be learned as well. Although one can differentiate V we found that estimating $M^{-1}\nabla V_\theta$ directly makes no difference in practice and is more efficient to evaluate in the forward pass.

$$V_\theta(q_k) \approx M^{-1}V(q_k), \quad F_\phi(x_k, u_k) \approx M^{-1}F(q_k, u_k) \tag{9}$$

3.1. Encoding Observations into Configurations

Oftentimes, a proper configuration state cannot be observed directly and the state must be estimated from observations. We may do this using a type of dynamic autoencoder scheme (related to Kingma and Welling (2013)), where the VIN dynamics are applied to encoded states before being decoded. Note that one could use a variational autoencoder scheme as done by Saemundsson et al. (2020), however we focus on the deterministic setting for exposition. We can learn all of these functions simultaneously simply by minimizing the multi-step “open-loop” loss in Eq. 10 where we decode a sequence of configuration states rolled out from the initial encoded state. The encoder and decoder networks are denoted by $f_\psi^{(enc)}$ and $f_\zeta^{(dec)}$ respectively and we aggregate parameters under $\Theta = (\theta, \phi, \psi, \zeta)$. The loss is computed under the squared l_2 loss up to time T . We denote both SV and

VV integrators using $\Phi_{\theta,\phi}$.

$$\begin{aligned}
 & \underset{\Theta}{\text{minimize}} && \sum_{k=0}^{T-1} \|\hat{y}_{k+1} - y_{k+1}\|_2^2 \\
 & \text{subject to} && \hat{x}_0 = f_{\psi}^{(enc)}(y_0) \\
 & && \hat{y}_k = f_{\zeta}^{(dec)}(\hat{x}_k) \\
 & && \hat{x}_{k+1} = \Phi_{\theta,\phi}(\hat{x}_k, u_k) \quad \forall k \in \{0, \dots, T-1\}
 \end{aligned} \tag{10}$$

where here our labels y_k are simply the true next configuration states, $y_k = (q_k)$ for SV and $y_k = (q_k, \dot{q}_k)$ for VV integrators. In the case that configuration come directly as observations, the encoding and decoding networks are simply identity operators.

3.2. Inductive Bias in Forcing

Forcing can be introduced in several ways, but most commonly in mechanical systems forcing is comes in the form of actuation, damping and heat. In some cases, we have additional information such that damping is only a function of velocity or that forcing is independent of state entirely. In either of these cases, we can decouple the parameters ϕ into ϕ_1 and ϕ_2 for control and damping respectively. For the general case going forward we will use the following structure for the VV-FVIN forces

$$F_{\phi_{1,2}}(q_k, \dot{q}_k, u_k) = F_{\phi_1}^{(control)}(q_k, u_k) + F_{\phi_2}^{(damping)}(q_k, \dot{q}_k) \tag{11}$$

For the SV-FVIN we chose to capture velocity information in damping using the current and previous states where control remains consistent with the VV-FVIN.

$$F_{\phi_{1,2}}(q_{k-1}, q_k, u_k) = F_{\phi_1}^{(control)}(q_k, u_k) + F_{\phi_2}^{(damping)}(q_{k-1}, q_k) \tag{12}$$

4. Experimental Results

In this section we will investigate several aspects of performance for the proposed FVIN, specifically the VV-FVIN variant (performance results for SV-FVIN can be found in supplementary material). The prediction and control performance will be evaluated under two forms of forcing mentioned previously, actuation and damping. We will also consider generalization characteristics such as the FVIN’s ability to capture the true energy conserving behavior in the absence of forcing. This would show that the model is able to truly capture the properties that we claim to describe in the FVIN, rather than simply being a black-box residual model. We study the damped pendulum (1D rotational configuration state), the damped cartpole (2D configuration state with rotation and translation) and a real rotational cartpole system. Although these systems may seem simplistic, they are highly-nonlinear and canonical examples of real mechanical systems, yet are easy to interpret.

4.1. Baseline Residual Dynamic Network

To evaluate both prediction and control of the proposed FVIN framework, we compare against a similar architecture where the FVIN dynamics are replaced by a general residual dynamic network (ResNN) as in [Saemundsson et al. \(2020\)](#), trained over the H -step loss function in Eq. 10. Additionally we give the residual network the same additive forcing assumption as the FVIN so that, with

encoding and decoding network structure being held constant, we can evaluate whether or not the inductive bias we introduce through FVIN improves prediction accuracy and control performance. ResNN is a competitive baseline, even on lower data regimes (5-10 trajectories), however it lacks any structure. This makes ResNN likely to over-fit and not capturing long term dynamics.

4.2. Prediction with Forcing: Generalization and Interpretability

Although variational integrators are no longer exactly symplectic in the presence of heavy forcing, they still often account for forcing significantly better than other discretizations Ober-Blöbaum et al. (2010). Ideally the FVIN components $F^{(control)}$, $F^{(damping)}$ and V should truly represent the corresponding physical properties of the mechanical system. To study these benefits in the VV-FVIN we consider a damped-pendulum undergoing random control input with the observation states of the form $z = [\cos \theta, \sin \theta, \dot{\theta}]$. Only 5 trajectories, each 50 time steps and random control inputs, are taken from different initial conditions. After which we test the learned models in two settings: (a) A test trajectory with random control forcing and (b) A test trajectory with no torque inputs, dissipating energy passively. In both cases models produce open-loop trajectories given the same control sequences of the test trajectory. Results for FVIN and ResNN are shown in Fig. 2 and show that both models are able to track accurately up until the trained horizon length where ResNN starts to diverge. Although the ResNN tracks for a significant period of time, the FVIN almost exactly accounts for forcing due to its variational integrator structure. More interestingly in Fig. 2, we find

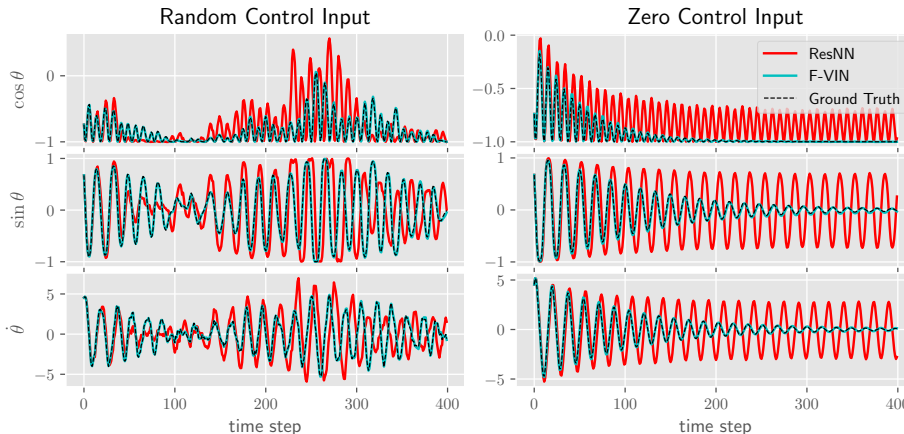


Figure 2: Damped Pendulum: Both VV-FVIN and ResNN models are trained on 5 trajectories with uniform random control inputs. We compare the prediction performance on a new initial condition with control inputs and without any control. We see that FVIN has better long term prediction error in both settings and captures the true passive dynamics despite having never observed a zero control trajectory. This is an important generalization property of FVIN

that even though the models have never been trained on a trajectory without forcing, that under no control inputs the FVIN still tracks very closely to the true passive dynamics of the pendulum where the ResNN captures a very different dissipation behavior. This tells us that the FVIN is actually accounting for forcing due to control allowing it to extract the true passive dynamics from very few trajectories. In addition to these two cases, we can further investigate the FVIN by removing the additive term $F^{(control)}$ and comparing against the same pendulum simulation without damping.




Figure 3: Taking the same learned VV-FVIN from Fig. 2, we scale the damping terms in both the FVIN and the pendulum simulation by $\alpha \in \{-0.3, 0, 1.0, 1.5\}$. Dotted lines denote the ground truth generated from simulation. This shows that the model has identified the true damping dissipation and the inductive biases imposed on forcing and passive dynamics are meaningful/interpretable.

We can even scale the respective damping terms with the parameter α which demonstrates that our model is interpretable with respect to our true mechanical system. Note that we simply use the same FVIN model learned from the same 5 forced pendulum trajectories with damping. In Fig. 3 we test both positive and negative values of α where $\alpha = 0.0$ corresponds to zero damping dissipation.

4.3. Model-based Control with Forced-VIN

In the previous section we have shown that FVIN can very accurately predict non-conservative systems with control even with only a few trajectories. In principle, this prediction performance combined with a model-based planning method should translate to superior control performance with respect to a cost function. We aim to emphasize the benefit of accurately accounting for passive dynamics and forcing in a model-predictive control (MPC) setting. Namely, we utilize a general sample-based zeroth-order optimization called the cross-entropy method (CEM) [Rubinstein \(1997\)](#) to compute the optimal control at each time step. CEM uses the learned model to sample and evaluate trajectories and then fits a multivariate Gaussian to the predicted optimal control sequence. In the MPC setting the CEM optimization is performed at every time step and only the first optimal control in the sequence is applied. We consider two control tasks: (a) A 1-degree of freedom damped pendulum swing up and stabilize task, (b) A 2-degree of freedom damped cartpole swing up and stabilize task. Additionally we compare three increasing data regimes for trajectories of length 50. We expect that both models will learn similar control policies in the highest data regime, however the FVIN should benefit from its inductive biases in the lower data settings. The pendulum trajectories are collected with random sampling and cartpole trajectories use additional policy refinement trajectories for exploration. In Fig. 4 we plot the cost difference between FVIN and ResNN MPC on the pendulum system over a grid of initial conditions (higher is better). In Table 1 we see that FVIN is able to successfully solve the damped pendulum at all initial conditions for all data regimes, where ResNN struggles to stabilize at all at in 5 trajectories and about 74% of the initial conditions at 10. In the 20 trajectory setting both models perform well at almost all initial conditions with similar resulting cost. However, according to Fig. 4 the control effort for achieving the same total cost is still significantly less using FVIN. We also compare on the damped cartpole task with similar results, but with 2-degrees of freedom it requires a greater amount of data. The

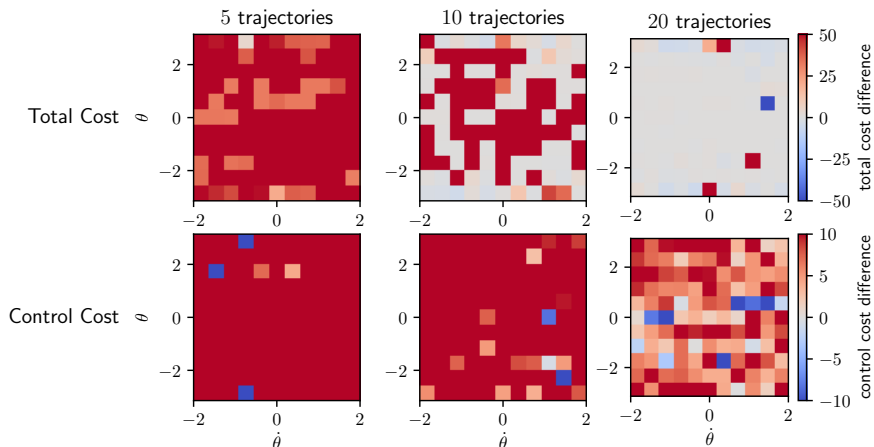


Figure 4: Damped pendulum: We compare MPC total cost (top row) and control effort (bottom row) for VV-FVIN and ResNN models on three data regimes. Each square is an initial condition and color indicates the cost difference (higher red values are in favor of FVIN).

	Damped Pendulum Swing-up			Damped Cart-Pole Swing-up		
	5 traj.	10 traj.	20 traj.	20 traj.	30 traj.	50 traj.
FVIN success %	1.0	1.0	1.0	0.88	0.96	1.0
ResNN success %	0.12	0.74	0.99	0.02	0.43	0.88

Table 1: We compare the percentage of successful attempts of the MPC tasks out of 100 sampled initial conditions using a planning horizon of $H = 15$. Success is indicated by achieving a state within an $\epsilon = 0.1$ ball of zero position and velocity of the pole.

FVIN structure enables the MPC controller to take advantage of the underlying passive dynamics even in low-data regimes, making it more likely to produce a high-performing control policy.

4.4. Evaluation of Real Quanser Pendulum Platform

To further provide justification of why the FVIN is a viable choice when modeling mechanical system, we test the prediction performance on measurements from a real system. We use the Quanser-Qube Servo 2 (QQS2) as seen in Fig. 5 which consist of 2 degrees of freedom and 1 actuation at the base. The QQS2 is very similar to a “cart-pole” system, but with rotational motion at the base rather than translation. The pendulum angle is given by θ_1 and the base angle θ_2 . The QQS2 system provides both of these measurements and their velocities via optical encoders. The observation is given as $(\cos \theta_1, \sin \theta_1, \theta_2, \dot{\theta}_1, \dot{\theta}_2)$, so that the autoencoder is necessary to infer the configuration state. We run similar prediction experiments where each model is provided 10 trajectories of 50 steps each using random control inputs. We then test on a new trajectory of length 200 where controls are turned off after step 100. The idea is that the model should predict the dissipation behavior. In Fig. 6 we show that although the ResNN accurately predicts accurately for the first 50 steps, it does not capture any dissipation and has over-fit the training data. On the contrary, FVIN has lower accuracy in the short terms, but is able to predict well past the 50 step horizon and predict dissipation due to damping in the QQS2.

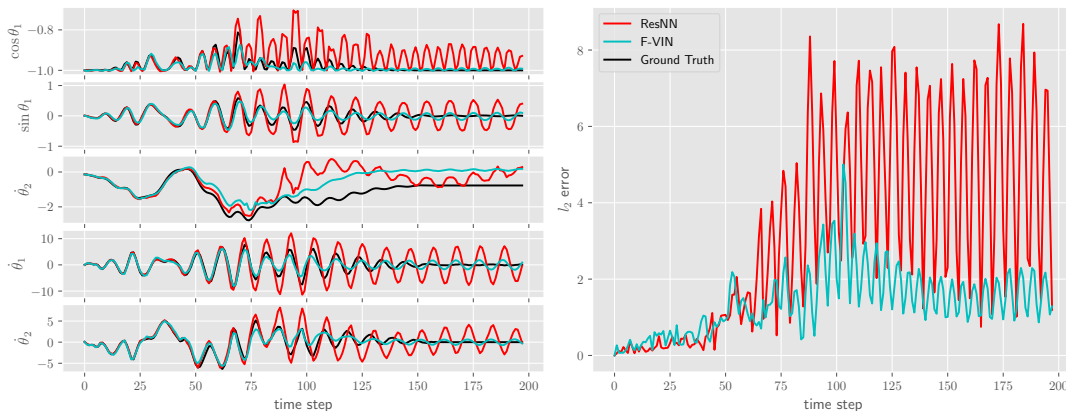


Figure 6: Prediction results on the real QQS2 (depicted in Fig. 5) from 10 50 step trajectories with random control. In this test trajectory control is turned off after 100 steps and the dissipation behavior is predicted successfully by the VV-FVIN where ResNN over-fits the first 50 steps. Note that θ_2 the rotational base does not necessarily go to 0.

5. Conclusion

In this paper we introduce a necessary extension of Variational Integrators Networks (Saemundsson et al. (2020)) to enable applications to general non-conservative and controlled systems using the modified variational *Discrete Lagrange-d'Alembert* principle. We then parameterized these new forcing terms in a way where the the non-conservative passive dynamics and forcing could be accurately learned simultaneously from random trajectories. We have shown that these forcing terms can be interpretable by generalizing the learned terms to other levels of damping forcing using the same model. The FVIN is highly data-efficient compared to the baseline ResNN which makes it ideal for model-based control, consistently obtaining more efficient control even in high-data regimes. The FVIN is also capable of predicting on a real non-conservative cart-pole and proves to capture energy behavior better than the baseline. Going forward we would like to scale this model to larger and more complicated systems and apply it to real-time control applications. Additionally we think it would be important to combine FVIN with the rich field of discrete mechanics and optimal control (DMOC) (Ober-Blöbaum et al. (2010)) and implicit definitions for better end-to-end learning control. FVIN may serve as a building block for future efforts in this area.

Acknowledgments

This work is supported by the NSF/USDA National AI Institute AIFARMS, AFRI grant no. 2020-67021-32799/project accession no.1024178 9:48 and by USDA/NSF NRI 2.0 project USDA#2019-67021-28989/NSF#1830343.




Figure 5: The Quanser-Qube Servo 2 Quanser (2020) uses a reinforcement learning-like interface developed by Polzounov et al. (2020).

References

- Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems 31*, pages 6571–6583. Curran Associates, Inc., 2018.
- Geir E Dullerud and Fernando Paganini. *A course in robust control theory: a convex approach*, volume 36. Springer Science & Business Media, 2013.
- Samuel Greydanus, Misko Dzamba, and Jason Yosinski. Hamiltonian neural networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 15379–15389. Curran Associates, Inc., 2019.
- Jérémie Guiochet, Mathilde Machin, and H el ene Waeselynck. Safety-critical advanced robots: A survey. *Robotics and Autonomous Systems*, 94:43–52, 2017.
- Andreas Hochlehnert, Alexander Terenin, Steind or S amundsson, and Marc Deisenroth. Learning contact dynamics using physically structured neural networks. In *International Conference on Artificial Intelligence and Statistics*, pages 2152–2160. PMLR, 2021.
- Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d Alch e-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 12519–12530. Curran Associates, Inc., 2019.
- Oliver Junge, Jerrold E Marsden, and Sina Ober-Bl obaum. Discrete mechanics and optimal control. *IFAC Proceedings Volumes*, 38(1):538–543, 2005.
- Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.
- Michael Lutter, Christian Ritter, and Jan Peters. Deep lagrangian networks: Using physics as model prior for deep learning. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019.
- J. E. Marsden and M. West. Discrete mechanics and variational integrators. *Acta Numerica*, 10: 357–514, 2001. doi: 10.1017/S096249290100006X.
- Sina Ober-Bl obaum, Oliver Junge, and Jerrold E. Marsden. Discrete mechanics and optimal control: An analysis. *ESAIM: Control, Optimisation and Calculus of Variations*, 17(2):322–352, Mar 2010. ISSN 1262-3377. doi: 10.1051/cocv/2010012.
- Kirill Polzounov, Ramitha Sundar, and Lee Redden. Blue river controls: A toolkit for reinforcement learning control systems on hardware, 2020.

- Quanser. Quanser-cube servo 2, 2020. [Online; accessed July 25, 2020].
- Reuven Y Rubinstein. Optimization of computer simulation models with rare events. *European Journal of Operational Research*, 99(1):89–112, 1997.
- Steindor Saemundsson, Alexander Terenin, Katja Hofmann, and Marc Deisenroth. Variational integrator networks for physically structured embeddings. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 3078–3087, Online, 26–28 Aug 2020. PMLR.
- David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dhharshan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. 362(6419):1140–1144, 2018. ISSN 0036-8075. doi: 10.1126/science.aar6404.
- Loup Verlet. Computer ”experiments” on classical fluids. i. thermodynamical properties of lennard-jones molecules. *Phys. Rev.*, 159:98–103, Jul 1967. doi: 10.1103/PhysRev.159.98.
- B Vujanovic. Conservation laws of dynamical systems via d’alembert’s principle. *International Journal of Non-Linear Mechanics*, 13(3):185–197, 1978.
- Jeffrey M Wendlandt and Jerrold E Marsden. Mechanical integrators derived from a discrete variational principle. *Physica D: Nonlinear Phenomena*, 106(3-4):223–246, 1997.