

---

# Spring-Rod System Identification via Differentiable Physics Engine

---

Kun Wang, Mridul Aanjaneya and Kostas Bekris

Rutgers University

{kun.wang2012, mridul.aanjaneya, kostas.bekris}@rutgers.edu

## Abstract

We propose a novel differentiable physics engine for system identification of complex spring-rod assemblies. Unlike black-box data-driven methods for learning the evolution of a dynamical system *and* its parameters, we modularize the design of our engine using a discrete form of the governing equations of motion, similar to a traditional physics engine. We further reduce the dimension from 3D to 1D for each module, which allows efficient learning of system parameters using linear regression. The regression parameters correspond to physical quantities, such as spring stiffness or the mass of the rod, making the pipeline explainable. The approach significantly reduces the amount of training data required, and also avoids iterative identification of data sampling and model training. We compare the performance of the proposed engine with previous solutions, and demonstrate its efficacy on tensegrity systems, such as NASA’s icosahedron.

## 1 Introduction

Cable-driven robots are gaining increasing attention due to their adaptiveness and safety. Tensegrity structures have many applications: from manipulation [13], locomotion [21], morphing airfoil [5] to spacecraft lander [4]. While useful and versatile, they are difficult to accurately model and control. Identifying system parameters is necessary, either to learn controllers in simulation (as real-world experiments are time-consuming, expensive and dangerous), or for traditional model-based control. In all these cases, the spring-rod representation considered in this work is the basic modeling element.

However, the spring-rod system has high degrees of freedom for system identification. Physics-based methods for simulation require accurate models that capture non-linear material behavior, which are difficult to construct. In contrast, data-driven methods can simulate any system from observed data, with sufficient training data. But the large number of variables and non-linear material properties necessitate copious amounts of training data.

Motivated by these issues, we propose a data-driven differentiable physics engine that combines the benefits of data-driven and physics-based models, while alleviating most of their drawbacks, and is designed from first principles. Previous data-driven models have required large amounts of data, because they learn the parameters *and* the physics of the system. Furthermore, the hidden variables and black box nature of these models are not explainable, and difficult to transfer to new environments. Our approach is based on the observation that the equations that govern the motion of such systems are well-understood, and can be directly baked into the data-driven model. Such a design can reduce demands on training data and can also generalize to new environments, as the governing principles remain the same. We further simplify the differentiable engine by modularization, which compartmentalizes the problem of learning the dynamics of the whole system to smaller well-contained problems. For each module, we also reduce the dimension from 3D to 1D, by taking advantage the properties of spring-rod systems, which allows for efficient parameter inference using linear regression. As a side benefit, the regression parameters correspond to physical quantities, such as the spring stiffness or the mass of the rod, making the framework explainable.

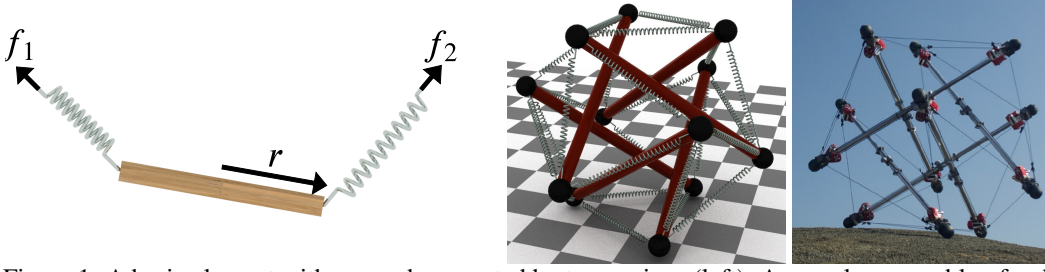


Figure 1: A basic element with one rod connected by two springs (left). A complex assembly of rods and springs forming a tensegrity robot in simulation (middle), and the real world (right).

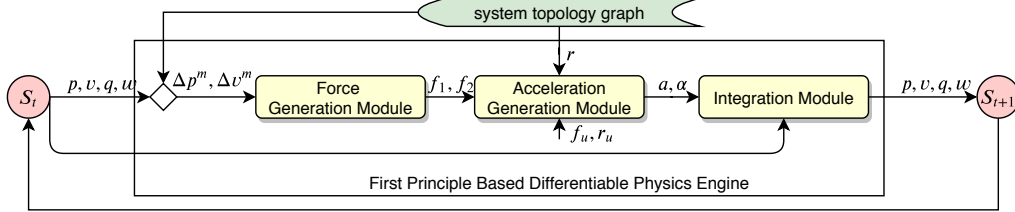


Figure 2: Flow chart showing the data flow when simulating one time step with our physics engine.

## 2 Related Work

Traditional methods for system identification build a dynamics model by minimizing the prediction error [27] [8]. These methods require parameter refinement and data sampling in an iterative fashion, to decrease the prediction error. This iterative process can be avoided using data-driven techniques that directly fit a physics model to data [19, 20, 2]. However, these techniques treat the dynamics as a black box, are data hungry, and require retraining in a new environment.

Instead of treating the environment as a black box, *interaction network* [3] took the first step to modularize objects and their interactions. Later, researchers extended this idea by a hierarchical relation network for graph-based object representation [17] and a propagation network [14] for the multi-step propagation network. While these methods are an improvement over previous approaches, they still treat the interactions between different objects as black boxes and try to learn them from data, even though the governing equations of motion are well-understood.

Quite a few authors have recently introduced differentiable physics engines that focus on many aspects not central to our work. For example, forward dynamics [9], Material Point Method (MPM) [11], linear complementarity problems (LCP) [6], augmented Lagrangian [12], differentiable programming [10], augmented traditional physics simulators [1], and LSTM dynamics model without system topology [7]. Researchers have also proposed differentiable engines specific to certain kinds of objects, such as molecules [23], fluids [22], and cloth [15]. Recent works on tensegrity robots [26, 16, 25] make major improvement on locomotion in simulation and have great challenges on policy transfer to real world system. All these works motivate us to mitigate the reality gap for cable driven robots between simulation and reality via system identification.

## 3 Methods

Our system views a spring-rod system as a composition of basic *elements* (see Fig. 1(left)), where springs generate forces that influence rod dynamics. We subdivide each time step of the simulation into three modules: force generation, acceleration computation, and state update/integration (see Fig. 2). The physics engine takes as input the current rod state  $S_t = \{p, v, q, \omega\}$ , where  $p$  is position,  $v$  is linear velocity,  $q$  is orientation (expressed as a quaternion), and  $\omega$  is the angular velocity. Based on  $S_t$ , the position and linear velocity  $p^m, v^m$  of the two rod endpoints is computed, and is used to compute the relative compression (or expansion) and velocity  $\Delta p^m, \Delta v^m$  of the two attached springs. Then, the first module predicts the spring forces  $f$ , the second module computes the linear and angular accelerations  $a, \alpha$ , and the third module integrates the new state  $S_{t+1}$ .

### 3.1 System Topology Graph

We use a topology graph to represent interconnections between different components of the spring-rod system. Each rod and spring has a corresponding vertex, and directed edges represent relations between them. Figure 3 shows an example topology graph for the basic spring-rod element shown in Figure 1(left).

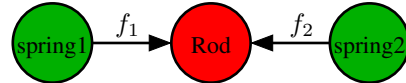


Figure 3: Element topology graph.

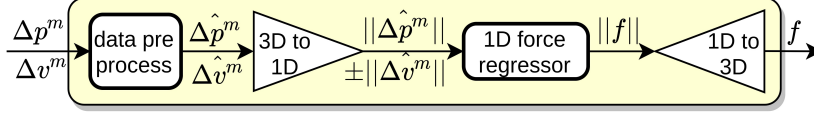


Figure 4: Force generation module, which uses dimensionality reduction to compute spring forces.

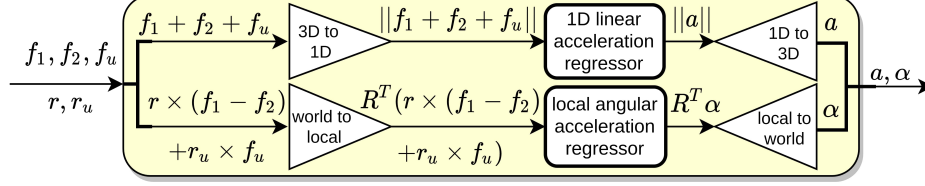


Figure 5: Acceleration generation module

### 3.2 Force Generation Module

The relative compression (or expansion)  $\Delta p^m$  and velocity  $\Delta v^m$  of each spring is given as input to the force generation module, which outputs the spring forces  $f$  by Hooke's law. As shown in Fig.4, two unknown parameters, stiffness  $K$  and damping  $c$ , can be easily learned using linear regression.

### 3.3 Acceleration Generation Module

The acceleration generation module takes the spring forces  $f$  and control force  $f_u$  as input and outputs each rod's linear and angular accelerations  $a, \alpha$  as shown in Fig 5.  $f_1$  and  $f_2$  are spring forces on the two rod ends,  $f_u$  is control force,  $r$  is the half-length rod vector,  $r_u$  is control force arm,  $R$  is the rod local/world frame rotation matrix. The rod mass  $M$  and inertia  $I$  are unknown parameters to identify.

### 3.4 Integration Module and Method Implementation

The integration module computes forward dynamics of each rod using the current accelerations  $a, \alpha$ . We apply the semi-implicit Euler method [24] to compute the updated state  $S_{t+1} = \{p_{t+1}, v_{t+1}, q_{t+1}, \omega_{t+1}\}$  at the end of the current time step.

The learning module receives the current state  $S_t$  and returns a prediction  $\hat{S}_{t+1}$ . The loss function is the MSE between the predicted  $\hat{S}_{t+1}$  and ground truth state  $S_{t+1}$ . The proposed decomposition, first-principles approach and the cable's linear nature allow the application of linear regression, which helps with data efficiency. This linear regression step has been implemented as a single layer neural network without activation function on pyTorch [18].

## 4 Experiments

The task is to identify parameters including spring stiffness  $K$ , damping  $c$  and rod mass  $M$ , inertia  $I$ . We evaluate these parameters by predict trajectories of the rods and compare with the ground truth.

### 4.1 Simple Spring-Rod Model Identification

First, we identify these parameters in a simple spring-rod system as shown in Figure 1 (a). **Interaction** is an improved version of the Interaction Network [3] as shown in Fig. 6. It has two Multilayer Perceptrons (MLPs), one to generate spring forces  $f$  and the other to generate rod state  $S_{t+1}$ . Unlike [3], which takes raw state  $S_t$  as input, we generate  $\Delta \hat{p}_t^m, \Delta \hat{v}_t^m$  as input. **Interaction+Int** appends the integration module to the Interaction Network, and replaces input  $S_t$  of  $MLP_2$  by  $r$ . **Koopman+Int** is to use the Koopman operator to predict accelerations and apply the Integration Module to map them to  $S_{t+1}$ . **Interaction** only predicts a  $S_{t+1}$  in training data that is close to  $S_t$ . **Interaction+Int** experiences increasing error from accumulated prediction errors. The Koopman operator **Koopman+Int** designed from first principles gives accurate predictions similar to **Ours** in this simple system. Comparison of errors is shown in Figure 8 a).

### 4.2 Complex Tensegrity Model Identification

We consider an icosahedron tensegrity system as shown in Fig. 1 (c). It is composed of 6 rods and 24 springs. Each rod is connected to 8 springs and has a length of 1.04m. Each spring's rest length is 0.637m. We set the gravity constant to  $g = -9.81$  in Mujoco. We collect 1000 trajectories with



Figure 6: Interaction Network

Figure 7: Koopman with Integration Module

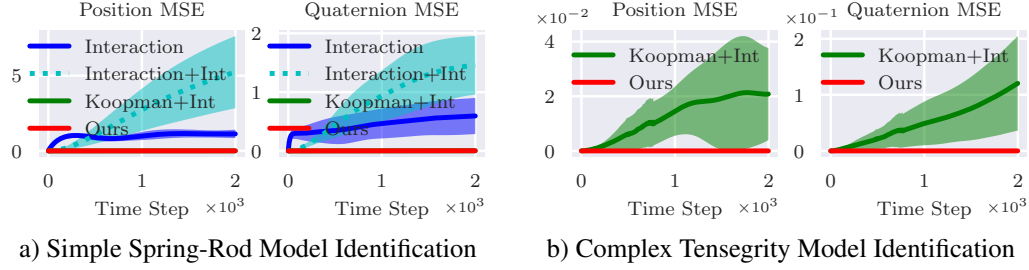


Figure 8: Physical Parameters Estimation

different initial conditions for training, 200 for validation and 100 for testing. The result is shown in Figure 8 b). **Our** approach outperforms **Koopman+Int** because designing basis functions for The Koopman operator has an increased data requirement relative to our approach.

#### 4.3 Data Efficiency Experiment

The proposed method has relatively small data requirements as shown in Fig. 9 a). Instead of training on 1000 trajectories, which have 736,167 time steps in total, we train our model with less data and evaluate performance. We randomly select 10%, 1%, 0.1%, 0.01% of the 736,167 time steps for training. The model achieves good performance even with 73 time steps for training. All trajectories are from the complex tensegrity setup with different parameters.

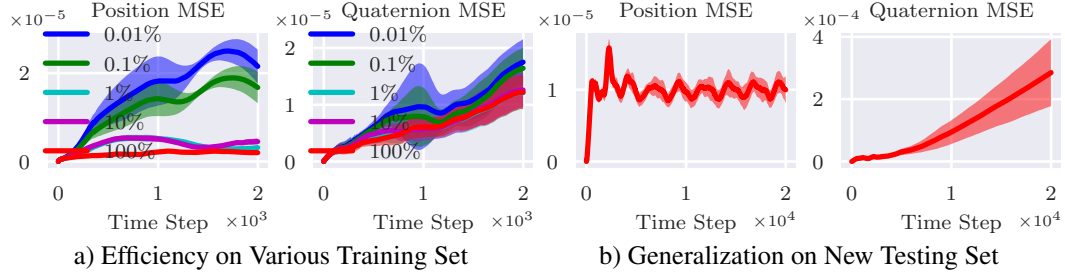


Figure 9: Data Efficiency and Model Generalization Experiment.

The proposed solution achieves very low error at a magnitude of  $10^{-5}$ , since it: 1) introduces a first-principles approach in learning physical parameters; 2) removes redundant data from regression; 3) operates -for now- over relatively clean data from simulation before moving to real-world data.

#### 4.4 Model Generalization Experiment

This section generalizes the physics engine trained with a dataset without external forces to a dataset with such forces. We are interested in evaluating: 1) how the physics engine performs for longer time horizons (e.g., after 2000 time steps); 2) if it can adapt to new scenarios. We generate a new dataset with 20,000 time steps trajectories with random directed perturbation forces  $f_u$ . The external force  $f_u$  does not have the same scale as the internal spring forces, so we add a new scalar module with only one parameter  $h$ , as in Fig. 10. We also apply dimensionality reduction to improve data efficiency. The tuning process is to freeze all other modules' weights and train this by the new dataset. The error graphs are shown in Fig. 9 b).

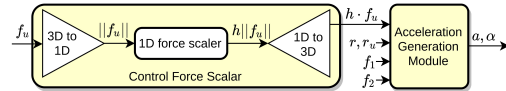


Figure 10: Control Force Scalar

## 5 Conclusion and Future Work

This paper proposes a differentiable physics engine for system identification of spring-rod systems based on first principles. The engine has three modules: force generation, acceleration generation and integration, which express the corresponding physical processes of spring-rod systems. This results in reduced data requirements and improved parameter accuracy. It also provides an explainable, accurate and fast physics engine. In the future, we plan to address contacts and friction. This will involve replacing the linear regressor with nonlinear models in the existing modules. To overcome noise in real data, we plan the addition of a residual network along with the nonlinear model. These changes may also help with temporal scalability. The ultimate mission, we generate policy from our identified engine and evaluate on the real platform to finally mitigate the reality gap.

## References

- [1] Anurag Ajay, Maria Bauza, Jiajun Wu, Nima Fazeli, Joshua B Tenenbaum, Alberto Rodriguez, and Leslie P Kaelbling. Combining Physical Simulators and Object-Based Networks for Control. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2019.
- [2] Martin Asenov, Michael Burke, Daniel Angelov, Todor Davchev, Kartic Subr, and Subramanian Ramamoorthy. Vid2param: Online system identification from video for robotics applications. *arXiv preprint arXiv:1907.06422*, 2019.
- [3] Peter Battaglia, Razvan Pascanu, Matthew Lai, Danilo Jimenez Rezende, et al. Interaction networks for learning about objects, relations and physics. In *Advances in neural information processing systems*, pages 4502–4510, 2016.
- [4] Jonathan Bruce, Andrew P Sabelhaus, Yangxin Chen, Dizhou Lu, Kyle Morse, Sophie Milam, Ken Caluwaerts, Alice M Agogino, and Vytas SunSpiral. Superball: Exploring tensegrities for planetary probes. *12th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS)*, 2014.
- [5] Muhao Chen, Jiacheng Liu, and Robert E Skelton. Design and control of tensegrity morphing airfoils. *Mechanics Research Communications*, 103:103480, 2020.
- [6] Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, and J Zico Kolter. End-to-end differentiable physics for learning and control. In *Advances in Neural Information Processing Systems*, pages 7178–7189, 2018.
- [7] Florian Golemo, Adrien Ali Taiga, Aaron Courville, and Pierre-Yves Oudeyer. Sim-to-real transfer with neural-augmented robot simulation. In *Conference on Robot Learning*, pages 817–828, 2018.
- [8] Nikolaus Hansen and Andreas Ostermeier. Completely derandomized self-adaptation in evolution strategies. *Evolutionary computation*, 9(2):159–195, 2001.
- [9] Eric Heiden, David Millard, Hejia Zhang, and Gaurav S Sukhatme. Interactive differentiable simulation. *arXiv preprint arXiv:1905.10706*, 2019.
- [10] Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Frédo Durand. DiffTaichi: Differentiable programming for physical simulation. *arXiv preprint arXiv:1910.00935*, 2019.
- [11] Yuanming Hu, Jiancheng Liu, Andrew Spielberg, Joshua B Tenenbaum, William T Freeman, Jiajun Wu, Daniela Rus, and Wojciech Matusik. Chainqueen: A real-time differentiable physical simulator for soft robotics. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 6265–6271. IEEE, 2019.
- [12] Benoit Landry, Zachary Manchester, and Marco Pavone. A differentiable augmented lagrangian method for bilevel nonlinear optimization. *arXiv preprint arXiv:1902.03319*, 2019.
- [13] Steven Lessard, Dennis Castro, William Asper, Shaurya Deep Chopra, Leya Breanna Baltaxe-Admony, Mircea Teodorescu, Vytas SunSpiral, and Adrian Agogino. A bio-inspired tensegrity manipulator with multi-dof, structurally compliant joints. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5515–5520. IEEE, 2016.
- [14] Yunzhu Li, Jiajun Wu, Jun-Yan Zhu, Joshua B Tenenbaum, Antonio Torralba, and Russ Tedrake. Propagation networks for model-based control under partial observation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 1205–1211. IEEE, 2019.
- [15] Junbang Liang, Ming C. Lin, and Vladlen Koltun. Differentiable cloth simulation for inverse problems. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [16] Z. Littlefield, D. Surovik, M. Vespignani, J. Bruce, W. Wang, and K. E. Bekris. Kinodynamic planning for spherical tensegrity locomotion with effective gait primitives. *International Journal of Robotics Research (IJRR)*, accepted 2019.
- [17] Damian Mrowca, Chengxu Zhuang, Elias Wang, Nick Haber, Li F Fei-Fei, Josh Tenenbaum, and Daniel L Yamins. Flexible neural representation for physics prediction. In *Advances in Neural Information Processing Systems*, pages 8799–8810, 2018.
- [18] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimeshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035, 2019.
- [19] Frank Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958.
- [20] David E Rumelhart, Geoffrey E Hinton, and Ronald J Williams. Learning representations by back-propagating errors. *nature*, 323(6088):533–536, 1986.

- [21] Andrew P Sabelhaus, Lara Janse van Vuuren, Ankita Joshi, Edward Zhu, Hunter J Garnier, Kimberly A Sover, Jesus Navarro, Adrian K Agogino, and Alice M Agogino. Design, simulation, and testing of a flexible actuated spine for quadruped robots. *arXiv preprint arXiv:1804.06527*, 2018.
- [22] C. Schenck and D. Fox. Spnets: Differentiable fluid dynamics for deep neural networks. In *Proceedings of the Second Conference on Robot Learning (CoRL)*, Zurich, Switzerland, 2018.
- [23] Samuel S. Schoenholz and Ekin D. Cubuk. Jax m.d.: End-to-end differentiable, hardware accelerated, molecular dynamics in pure python. <https://github.com/google/jax-md>, <https://arxiv.org/abs/1912.04232>, 2019.
- [24] David Stewart and Jeffrey C Trinkle. An implicit time-stepping scheme for rigid body dynamics with coulomb friction. In *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, volume 1, pages 162–169. IEEE, 2000.
- [25] D. Surovik, J. Bruce, K. Wang, M. Vespignani, and K. E. Bekris. Any-axis tensegrity rolling via bootstrapped learning and symmetry reduction. In *International Symposium on Experimental Robotics (ISER)*, Buenos Aires, Argentina, 11/2018 2018.
- [26] D Surovik, K Wang, M Vespignani, J Bruce, and K E Bekris. Adaptive Tensegrity Locomotion: Controlling a Compliant Icosahedron with Symmetry-Reduced Reinforcement Learning. *International Journal of Robotics Research (IJRR)*, 2019.
- [27] Jan Swevers, Chris Ganseman, D Bilgin Tukul, Joris De Schutter, and Hendrik Van Brussel. Optimal robot excitation and identification. *IEEE transactions on robotics and automation*, 13(5):730–740, 1997.