

---

# Efficient Differentiable Simulation of Articulated Bodies

---

Yi-Ling Qiao<sup>\*12</sup> Junbang Liang<sup>\*1</sup> Vladlen Koltun<sup>2</sup> Ming C. Lin<sup>1</sup>

## Abstract

We present a method for efficient differentiable simulation of articulated bodies. This enables integration of articulated body dynamics into deep learning frameworks, and gradient-based optimization of neural networks that operate on articulated bodies. We derive the gradients of the contact solver using spatial algebra and the adjoint method. Our approach is an order of magnitude faster than autodiff tools. By only saving the initial states throughout the simulation process, our method reduces memory requirements by two orders of magnitude. We demonstrate the utility of efficient differentiable dynamics for articulated bodies in a variety of applications. We show that reinforcement learning with articulated systems can be accelerated using gradients provided by our method. In applications to control and inverse problems, gradient-based optimization enabled by our work accelerates convergence by more than an order of magnitude.

## 1. Introduction

Differentiable physics enables efficient gradient-based optimization with dynamical systems. It has achieved promising results in both simulated (Hu et al., 2019; Qiao et al., 2020) and real environments (Bern et al., 2019; Song & Boularias, 2020a). Our goal is to make articulated body simulation efficiently differentiable. We aim to maximize efficiency in both computation and memory use, in order to support fast gradient-based optimization of differentiable systems that interact with articulated bodies in physical environments.

Articulated bodies play a central role in robotics, computer graphics, and embodied AI. Many control systems are optimized via experiences collected in simulation (Todorov et al., 2012; Coumans, 2015; Lee et al., 2018; Tedrake et al.,

2019). However, they do not have access to analytic derivatives of the articulated dynamics. Therefore, nearly all gradient-based approaches have to design strategies to compute the gradients indirectly when dealing with articulated bodies.

The straightforward way to differentiate the simulation is to use existing automatic differentiation tools (Griewank & Walther, 2008; Abadi et al., 2016; Paszke et al., 2019). However, autodiff tools consume prohibitive amounts of memory when there are many simulation steps. Autodiff tracks every operation and stores the intermediate results in order to perform backpropagation. In articulated body simulation, the iterative contact solver and the dynamics algorithm (Featherstone, 2007) yield exceedingly long computational graphs. In our experiments, differentiable simulators built with autodiff tools run out of memory after 5,000 simulation steps – just a few seconds of experience. As a result, learning is constrained to short experiences or forced to use large time steps, thus curtailing the scope of behaviors that can be learned or undermining the simulation’s stability.

Furthermore, the overhead of creating and storing the auxiliary variables for autodiff also slows down the forward simulation. Although autodiff tools like DiffTaichi (Hu et al., 2020) and JAX (Bradbury et al., 2018) can accelerate the simulation of fluids and deformable solids by vectorization, it is difficult to achieve the same speedup in articulated body simulation because the articulated dynamics algorithm is highly serialized, unlike inherently parallel computation on grids and particles.

In this paper, we design a differentiable articulated body simulation method that runs 10x faster with 1% of the memory consumption compared to differentiation based on autodiff tools. In order to minimize the overhead of differentiation, we derive the gradients of articulated body simulation using the adjoint method (Giles & Pierce, 2000). The adjoint method has been applied to fluids (McNamara et al., 2004) and multi-body systems (Geilinger et al., 2020), but these applications do not support physically correct differentiable simulation of articulated bodies. Our derivation of the operator adjoints uses spatial algebra (Featherstone, 2008). Our method needs almost no additional computation during the forward simulation and is an order of magnitude faster than autodiff tools.

---

<sup>\*</sup>Equal contribution <sup>1</sup>University of Maryland, College Park <sup>2</sup>Intel Labs. Correspondence to: Yi-Ling Qiao, Junbang Liang, and Ming C. Lin <{yilingq,liangjb,lin}@cs.umd.edu>, Vladlen Koltun <vladlen.koltun@intel.com>.

We further reduce memory requirements by adapting ideas from checkpointing (Griewank & Walther, 2000; Chen et al., 2016) to the differentiable simulation setting. In the forward pass, we store the initial simulation state for each time step. During backpropagation, we recreate intermediate variables by reproducing simulation from the stored state. The overall runtime remains fast, while memory consumption is reduced by two orders of magnitude.

As an application of differentiable dynamics for articulated systems, we show that reinforcement learning (RL) can benefit from the knowledge of gradients in two ways. First, it can make use of the gradients computed by the simulation to generate extra samples using first-order approximation. Second, during the policy learning phase, differentiable physics enables us to perform a one-step rollout of the objective value function so that the policy updates can be more accurate. Both schemes effectively improve the convergence speed and the attained reward. We also demonstrate applications of efficient differentiable articulated dynamics to inverse problems, such as motion control and parameter estimation. Gradient-based optimization enabled by our method accelerates convergence in these settings by more than an order of magnitude. Code is available on our project page: <https://github.com/YilingQiao/diffarticulated>

The contributions of this work are as follows:

- We derive the adjoint formulations for the entire articulated body simulation workflow, enabling a 10x acceleration over autodiff tools.
- We adapt the checkpointing method to the structure of articulated body simulation to reduce memory consumption by 100x, making stable collection of extended experiences feasible.
- We introduce two general schemes for accelerating reinforcement learning using differentiable physics.
- We demonstrate the utility of differentiable simulation of articulated bodies in motion control and parameter estimation, enhancing performance in these scenarios by more than an order of magnitude.

## 2. Related Work

Differentiable programming has been applied to rendering (Li et al., 2018a; Nimier-David et al., 2019; Laine et al., 2020), image processing (Li et al., 2020; 2018b), SLAM (Krishna Murthy et al., 2020), and design (Du et al., 2020). Making complex systems differentiable enables learning and optimization using gradient-based methods. Our literature review focuses on differentiable physics, the adjoint method, and neural approximations of physical systems.

**Differentiable physics.** Differentiable physics provides gradients for learning, control, and inverse problems that in-

volve physical systems. Degraeve et al. (2019) advocate using differentiable physics to solve control problems in robotics. de Avila Belbute-Peres et al. (2018) obtain gradients of 2D rigid body dynamics. Liang et al. (2019) use automatic differentiation tools to obtain gradients for cloth simulation. Qiao et al. (2020) develop a more comprehensive differentiable physics engine for rigid bodies and cloth based on mesh representations. Ingraham et al. (2019) present differentiable simulation of protein dynamics. For volumetric data, ChainQueen (Hu et al., 2019) computes gradients of the MPM simulation pipeline. Bern et al. (2019) use FEM to model soft robots and perform trajectory optimization with analytic gradients. Takahashi et al. (2021) differentiate the simulation of fluids with solid coupling.

Many exciting applications of differentiable physics have been explored (Spielberg et al., 2019; Heiden et al., 2019; 2020; Wang et al., 2020; Krishna Murthy et al., 2021). Huang et al. (2021) propose a soft-body manipulation benchmark for differentiable physics. Toussaint et al. (2018) manipulate tools with the help of differentiable physics. Song & Boularias (2020b) perform system identification by learning from the trajectory, and Song & Boularias (2020a) then use the estimated frictional properties to help robotic pushing.

A related line of work concerns the development of powerful automatic differentiation tools that can be used to differentiate simulation. DiffTaichi (Hu et al., 2020) provides a new programming language and a compiler, enabling the high-performance Taichi simulator to compute the gradients of the simulation. JAX MD (Schoenholz & Cubuk, 2020) makes use of the JAX autodiff library (Bradbury et al., 2018) to differentiate molecular dynamics simulation. These works have specifically made use of vectorization on both CPU and GPU to achieve high performance on grids and particle sets, where the simulation is intrinsically parallel. In contrast, the simulation of articulated bodies is far more serial, as dynamics propagates along kinematic paths rather than acting in parallel on grid cells or particles.

TinyDiffSim (Coumans et al., 2020) provides a templated simulation framework that can leverage existing autodiff tools such as CppAD (Bell et al., 2018), Ceres (Agarwal et al., 2010), and PyTorch (Paszke et al., 2019) to differentiate simulation. However, these methods introduce significant overhead in tracing the computation graph and accumulate substantial computational and memory costs when applied to articulated bodies.

Our method does not need to store the entire computation graph to compute the gradients. We store the initial states of each simulation step and reproduce the intermediate results when needed by the backward pass. This checkpointing strategy is also used in training neural ODEs (Zhuang et al., 2020) and large neural networks (Chen et al., 2016; Gruslys et al., 2016; Kirisame et al., 2021; Shah et al., 2021). We

are the first to adapt this technique to articulated dynamics, achieving dramatic reductions in memory consumption and enabling stable simulation of long experiences.

**Adjoint method.** The adjoint method has been applied to fluid control (McNamara et al., 2004), PDEs (Holl et al., 2020), light transport simulation (Nimier-David et al., 2020), and neural ODEs (Chen et al., 2018; Zhuang et al., 2020). Recently, Geilinger et al. (2020) proposed to use the adjoint method in multi-body dynamics. However, they operate in maximal coordinates and model body attachments using springs. This does not enforce physical validity of articulated bodies. In contrast, we operate in reduced coordinates and derive the adjoints for the articulated body algorithm and spatial algebra operators that properly model the body’s joints. This supports physically correct simulation with joint torques, limits, Coriolis forces, and proper transmission of internal forces between links.

**Neural approximation.** A number of works approximate physics simulation using neural networks (Battaglia et al., 2016; Chang et al., 2017; Mrowca et al., 2018; Schenck & Fox, 2018; Sanchez-Gonzalez et al., 2018; 2020; Li et al., 2019; Belbute-Peres et al., 2020; Ummenhofer et al., 2020; Wandel et al., 2021; Pfaff et al., 2021). Physical principles have also been incorporated in the design of neural networks (Schütt et al., 2017; Anderson et al., 2019; Bogatskiy et al., 2020; Chen et al., 2020; Cranmer et al., 2020). Approximate simulation by neural networks is naturally differentiable, but the networks are not constrained to abide by the underlying physical dynamics and simulation fidelity may degenerate outside the training distribution.

### 3. Preliminaries

**Articulated body dynamics.** For the forward simulation, we choose the recursive Articulated Body Algorithm (ABA) (Featherstone, 2007), which has  $O(n)$  complexity and is widely used in articulated body simulators (Todorov et al., 2012; Coumans, 2015; Lee et al., 2018). In each simulation step  $k$ , the states  $\mathbf{x}^k$  of the system consist of configurations in generalized coordinates and their velocities  $\mathbf{x}^k = [\mathbf{q}^k, \dot{\mathbf{q}}^k] \in \mathbb{R}^{2n_q \times 1}$ , where  $n_q$  is the number of degrees of freedom in the system. Assuming each joint can be actuated by a torque, there is an  $n_q$ -dimensional control signal  $\mathbf{u}^k \in \mathbb{R}^{n_q \times 1}$ . The discretized dynamics at this step can be written as  $\mathbf{f}^{k+1}(\mathbf{u}^k, \mathbf{x}^k, \mathbf{x}^{k+1}) = 0$ .

**Adjoint method.** We can concatenate the states in the entire time sequence into  $\mathbf{x} = [\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^{n_t}] \in \mathbb{R}^{n_t \cdot 2n_q \times 1}$ , with corresponding control input  $\mathbf{u} = [\mathbf{u}^1, \mathbf{u}^2, \dots, \mathbf{u}^{n_t}] \in \mathbb{R}^{n_t \cdot n_u \times 1}$ , where  $n_t$  is the number of simulation steps. The concatenated dynamics equations can be similarly written as

$\mathbf{f} = [\mathbf{f}^1, \mathbf{f}^2, \dots, \mathbf{f}^{n_t}] = 0$ . For a learning or optimization problem, we would usually define a scalar objective function  $\Phi(\mathbf{x}, \mathbf{u})$ . To get the derivative of this function w.r.t. the control input  $\mathbf{u}$ , one needs to compute

$$\frac{d\Phi}{d\mathbf{u}} = \frac{\partial\Phi}{\partial\mathbf{u}} + \frac{\partial\Phi}{\partial\mathbf{x}} \frac{d\mathbf{x}}{d\mathbf{u}}. \quad (1)$$

$\frac{\partial\Phi}{\partial\mathbf{u}}$  is easy to compute for each single simulation step. But it is prohibitively expensive to directly solve  $\frac{\partial\Phi}{\partial\mathbf{x}} \frac{d\mathbf{x}}{d\mathbf{u}}$  because  $\frac{d\mathbf{x}}{d\mathbf{u}}$  will be a  $2n_q n_t \times n_q n_t$  matrix.

Instead, we differentiate the dynamics  $\mathbf{f}(\mathbf{x}, \mathbf{u}) = 0$  and obtain the constraint  $\frac{\partial\mathbf{f}}{\partial\mathbf{x}} \frac{d\mathbf{x}}{d\mathbf{u}} = -\frac{\partial\mathbf{f}}{\partial\mathbf{u}}$ . Applying the adjoint method (Giles & Pierce, 2000),  $\frac{\partial\Phi}{\partial\mathbf{x}} \frac{d\mathbf{x}}{d\mathbf{u}}$  equals to

$$-\mathbf{R}^\top \frac{\partial\mathbf{f}}{\partial\mathbf{u}} \quad \text{such that} \quad \left( \frac{\partial\mathbf{f}}{\partial\mathbf{x}} \right)^\top \mathbf{R} = \left( \frac{\partial\Phi}{\partial\mathbf{x}} \right)^\top \quad (2)$$

Since the dynamics equation for one step only involves a small number of previous steps, the sparsity of  $\frac{\partial\mathbf{f}}{\partial\mathbf{u}}$  and  $\frac{\partial\mathbf{f}}{\partial\mathbf{x}}$  makes it easier to solve for the variable  $\mathbf{R}$  in Eq. 2.

$\frac{\partial\Phi}{\partial\mathbf{u}} - \mathbf{R}^\top \frac{\partial\mathbf{f}}{\partial\mathbf{u}}$  is called the adjoint of  $\mathbf{u}$  and is equivalent to the gradient by our substitution. In the following derivation, we denote the adjoint of a variable  $s$  by  $\bar{s}$ , and the adjoint of a function  $f(\cdot)$  by  $\bar{f}(\cdot)$ .

## 4. Efficient Differentiation

In this section, we introduce our algorithm for the gradient computation in articulated body simulation. For faster differentiation, the gradients of articulated body dynamics are computed by the adjoint method. We derive the adjoint of time integration, contact resolution, and forward dynamics in reverse order. We then adapt the checkpoint method to articulated body simulation to reduce memory consumption.

### 4.1. Adjoint Method for Articulated Dynamics

One forward simulation step can be split into five modules, as shown in Figure 1: perform the forward kinematics from the root link to the leaf links, update the forces from the leaves to the root, update the accelerations from the root to the leaves, detect and resolve collisions, and perform time integration. Backpropagation proceeds through these modules in reverse order.

**Time integration.** Backpropagation starts from the time integration. As an example, for a simulation sequence with  $n_t = 3$  steps and an integrator with a temporal horizon of 2, the constraints  $(\frac{\partial\mathbf{f}}{\partial\mathbf{x}})^\top \mathbf{R} = (\frac{\partial\Phi}{\partial\mathbf{x}})^\top$  in Equation 2 can be expanded as

$$\begin{bmatrix} \left( \frac{\partial\mathbf{f}^1}{\partial\mathbf{x}^1} \right)^\top & \left( \frac{\partial\mathbf{f}^2}{\partial\mathbf{x}^1} \right)^\top & 0 \\ 0 & \left( \frac{\partial\mathbf{f}^2}{\partial\mathbf{x}^2} \right)^\top & \left( \frac{\partial\mathbf{f}^3}{\partial\mathbf{x}^2} \right)^\top \\ 0 & 0 & \left( \frac{\partial\mathbf{f}^3}{\partial\mathbf{x}^3} \right)^\top \end{bmatrix} \begin{bmatrix} \mathbf{R}^1 \\ \mathbf{R}^2 \\ \mathbf{R}^3 \end{bmatrix} = \begin{bmatrix} \left( \frac{\partial\Phi}{\partial\mathbf{x}^1} \right)^\top \\ \left( \frac{\partial\Phi}{\partial\mathbf{x}^2} \right)^\top \\ \left( \frac{\partial\Phi}{\partial\mathbf{x}^3} \right)^\top \end{bmatrix} \quad (3)$$




Figure 1. **The workflow of a simulation step.** Assume there are  $n_q = 6$  DoF. The initial state is an  $n_q \times 3$  matrix containing the position, velocity, and control input of each joint. The forward dynamics will traverse the articulated body sequentially three times. This process is difficult to parallelize and will generate a large number of intermediate results. After the forward dynamics there is a collision resolution step with collision detection and an iterative Gauss-Seidel solver. The size of the initial state (position, velocity, and control input) is much smaller than that of intermediate results but the initial state has all the information to resume all the intermediate variables.

If using the explicit Euler integrator, we have  $\frac{\partial \mathbf{f}^k}{\partial \mathbf{x}^k} = \mathbf{I}$ . Initially  $\mathbf{R}^{n_t} = \frac{\partial \Phi}{\partial \mathbf{x}^{n_t}}$ . The following  $\mathbf{R}^k$  can be computed iteratively by

$$\mathbf{R}^k = \left( \frac{\partial \Phi}{\partial \mathbf{x}^k} \right)^\top - \left( \frac{\partial \mathbf{f}^{k+1}}{\partial \mathbf{x}^k} \right)^\top \mathbf{R}^{k+1}, k = 1, 2, \dots, n_t - 1. \quad (4)$$

When  $\mathbf{R}^k$  backpropagates through time, the gradients  $\overline{\mathbf{u}}^k$  can also be computed by Equation 2. In fact, by the way we calculate  $\mathbf{R}^k$ , it equals the gradients of  $\mathbf{x}^k$ . Other parameters can also be computed in a similar way as  $\overline{\mathbf{u}}^k$ .

**Collision resolution.** The collision resolution step consists of collision detection and a collision solver. Upon receiving the gradients  $\overline{\mathbf{x}}$  from the time integrator, the collision solver needs to pass the gradients to detection, and then to the forward dynamics. In our collision solver, we construct a Mixed Linear Complementarity Problem (MLCP) (Stepien, 2013):

$$\begin{aligned} \mathbf{a} &= \mathbf{A}\mathbf{x} + \mathbf{b} \\ \text{s.t. } \mathbf{0} &\leq \mathbf{a} \perp \mathbf{x} \geq \mathbf{0} \text{ and } \mathbf{c}_- \leq \mathbf{x} \leq \mathbf{c}_+, \end{aligned} \quad (5)$$

where  $\mathbf{x}$  is the new collision-free state,  $\mathbf{A}$  is the inertial matrix,  $\mathbf{b}$  contains the relative velocities, and  $\mathbf{c}_-, \mathbf{c}_+$  are the lower bound and upper bound constraints, respectively. We use the projected Gauss-Seidel (PGS) method to solve this MLCP. This iterative solve trades off accuracy for speed such that the constraints  $\mathbf{a}^\top \mathbf{x} = 0$  might not hold on termination. In this setting, where the solution is not guaranteed to satisfy constraints, implicit differentiation (de Avila Belbute-Peres et al., 2018; Liang et al., 2019; Qiao et al., 2020) no longer works. Instead, we design a reverse version of the PGS solver using the adjoint method to compute the gradients. Further details are in the supplement. In essence, the solver mirrors PGS, substituting the adjoints for the original operators. This step passes the gradients from the collision-free states to the forward dynamics.

**Forward dynamics.** Our articulated body simulator propagates the forward dynamics as shown in the green blocks of Figure 1. Each operation in the forward simulation has its corresponding adjoint operation in the backward pass. To compute the gradients, we derive adjoint rules for the operators in spatial algebra (Featherstone, 2008).

As a simple example, a spatial vector  $\mathbf{p}_i = [\mathbf{w}, \mathbf{v}] \in \mathbb{R}^6$  representing the bias force of the  $i$ -th link is the sum of an external force  $[\mathbf{f}_1, \mathbf{f}_2] \in \mathbb{R}^6$  and a cross product of two spatial motion vectors  $[\mathbf{w}_1, \mathbf{v}_1], [\mathbf{w}_2, \mathbf{v}_2] \in \mathbb{R}^6$ :

$$\begin{bmatrix} \mathbf{w} \\ \mathbf{v} \end{bmatrix} = \begin{bmatrix} \mathbf{f}_1 + \mathbf{w}_1 \times \mathbf{w}_2 + \mathbf{v}_1 \times \mathbf{v}_2 \\ \mathbf{f}_2 + \mathbf{w}_1 \times \mathbf{v}_2 \end{bmatrix}. \quad (6)$$

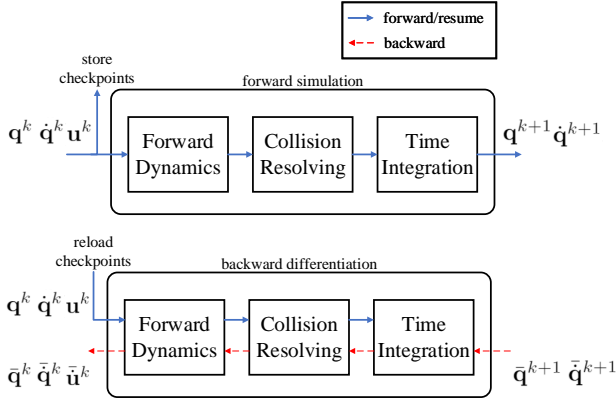
Once we get the adjoint  $[\overline{\mathbf{w}}, \overline{\mathbf{v}}]$  of  $\mathbf{p}_i = [\mathbf{w}, \mathbf{v}]$ , we can propagate it to its inputs:

$$\begin{aligned} \begin{bmatrix} \overline{\mathbf{w}}_1 \\ \overline{\mathbf{v}}_1 \end{bmatrix} &= \begin{bmatrix} -\overline{\mathbf{w}} \times \mathbf{w}_2 - \overline{\mathbf{v}} \times \mathbf{v}_2 \\ -\overline{\mathbf{w}} \times \mathbf{v}_2 \end{bmatrix} \\ \begin{bmatrix} \overline{\mathbf{w}}_2 \\ \overline{\mathbf{v}}_2 \end{bmatrix} &= \begin{bmatrix} \overline{\mathbf{w}} \times \mathbf{w}_1 \\ \overline{\mathbf{w}} \times \mathbf{v}_1 + \overline{\mathbf{v}} \times \mathbf{w}_1 \end{bmatrix}, \quad \begin{bmatrix} \overline{\mathbf{f}}_1 \\ \overline{\mathbf{f}}_2 \end{bmatrix} = \begin{bmatrix} \overline{\mathbf{w}} \\ \overline{\mathbf{v}} \end{bmatrix} \end{aligned} \quad (7)$$

This example shows the adjoint of one forward operation. The time and space complexity of the original operation and its adjoint are the same as the forward simulation. The supplement provides more details on the adjoint operations.

## 4.2. Checkpointing for Articulated Dynamics

The input and output variables of one simulation step have relatively small dimensionalities (positions, velocities, and accelerations), but many more intermediate values are computed during the process. Although this is not a bottleneck for forward simulation because the memory allocation for intermediate results can be reused across time steps, it becomes a problem for differentiable simulation, which needs to store the entire computation graph for the backward pass. This causes memory consumption to explode when simulation proceeds over many time steps, as is the case when



**Figure 2. Differentiation of articulated body dynamics.** *Top:* forward simulation at step  $k$ . The simulator stores the states  $\mathbf{q}^k$ ,  $\dot{\mathbf{q}}^k$ , and control input  $\mathbf{u}^k$ . *Bottom:* During backpropagation, the checkpoints are reloaded and the simulator runs one forward step to reconstruct the intermediate variables. Beginning with the gradients from step  $k+1$ , we use the adjoint method to sequentially compute the gradients of all variables through time integration, the constraint solver, and forward dynamics.

small time steps are needed for accuracy and stability, and when the effects of actions take time to become apparent.

The need for scaling to larger simulation steps motivates our adaptation of the checkpointing method (Griewank & Walther, 2000; Chen et al., 2016). Instead of saving the entire computation graph, we choose to store only the initial states in each time step. We use this granularity because (a) the states of each step have the smallest size among all essential information sufficient to replicate the simulation step, (b) finer checkpoints are not useful because at least one step of intermediate results needs to be stored in order to do the simulation, and (c) sparser checkpoints will use less memory but require multiple steps for reconstructing intermediate variables, costing more memory and computation. We validate our checkpointing policy with an ablation study in the supplement. Figure 2 illustrates the scheme. During forward simulation, we store the simulation state in the beginning of each time step. During backpropagation, we reload the checkpoints (blue arrows) and redo the forward (one-step) simulation to generate the computation graph, and then compute the gradients using the adjoint method in reverse order (red arrows).

In summary, assume the simulation step consists of two parts:

$$\mathbf{Z}^k = G(\mathbf{x}^{k-1}), \quad \mathbf{x}^k = F(\mathbf{Z}^k), \quad (8)$$

where  $\mathbf{Z}^k$  represents all the intermediate results. After each step, we free the space consumed by  $\mathbf{Z}^k$ , only storing  $\mathbf{x}^k$ . During backpropagation, we recompute  $\mathbf{Z}^k$  from  $\mathbf{x}^{k-1}$  and use the adjoint method to compute gradients of the previous

step:

$$\mathbf{Z}^k = G(\mathbf{x}^{k-1}), \quad (9)$$

$$\bar{\mathbf{Z}}^k = \bar{F}(\mathbf{Z}^k, \mathbf{x}^k), \quad (10)$$

$$\bar{\mathbf{x}}_{k-1} = \bar{G}(\mathbf{x}_{k-1}, \bar{\mathbf{Z}}^k). \quad (11)$$

## 5. Reinforcement Learning

Our method can be applied to simulation-based reinforcement learning (RL) tasks to improve policy performance as well as convergence speed. By computing gradients with respect to the actions, differentiable simulation can provide more information about the environment. We suggest two approaches to integrating differentiable physics into RL.

**Sample enhancement.** We can make use of the computed gradients to generate samples in the neighborhood around an existing sample. Specifically, given a sample  $(s, a_0, s'_0, r_0)$  from the history, with observation  $s$ , action  $a_0$ , next-step observation  $s'_0$ , and reward  $r_0$ , we can generate new samples  $(s, a_k, s'_k, r_k)$  using first-order approximation:

$$\begin{aligned} \mathbf{a}_k &= \mathbf{a}_0 + \Delta \mathbf{a}_k, \\ s'_k &= s'_0 + \frac{\partial s'_0}{\partial \mathbf{a}_0} \Delta \mathbf{a}_k, \\ r_k &= r_0 + \frac{\partial r_0}{\partial \mathbf{a}_0} \Delta \mathbf{a}_k, \end{aligned}$$

where  $\Delta \mathbf{a}_k$  is a random perturbation vector. This method, which we call *sample enhancement*, effectively generates as many approximately accurate samples as desired for learning purposes. By providing a sufficient number of generated samples around the neighborhood, the critic can have a better grasp of the function shape (patchwise rather than pointwise), and can thus learn and converge faster.

**Policy enhancement.** Alternatively, the policy update can be adjusted to a format compatible with differentiable physics, usually dependent on the specific RL algorithm. For example, in MBPO (Janner et al., 2019), the policy network is updated using the gradients of the critic:

$$\mathcal{L}'_\mu = -Q(\mathbf{s}, \mu(\mathbf{s})) + Z, \quad (12)$$

where  $\mathcal{L}'_\mu$  is the loss for the policy network  $\mu$ ,  $Q$  is the value function for state-action pairs, and  $Z$  is the regularization term. To make use of the gradients, we can expand the  $Q$  function one step forward,

$$\frac{\partial Q(\mathbf{s}, \mathbf{a})}{\partial \mathbf{a}} = \frac{\partial r}{\partial \mathbf{a}} + \gamma \frac{\partial Q(\mathbf{s}', \mu(\mathbf{s}'))}{\partial \mathbf{s}'} \frac{\partial \mathbf{s}'}{\partial \mathbf{a}}, \quad (13)$$

and substitute it in Eq. 12:

$$\mathcal{L}'_\mu = -\frac{\partial Q(\mathbf{s}, \mathbf{a})}{\partial \mathbf{a}} \mu(\mathbf{s}) + Z. \quad (14)$$

## Efficient Differentiable Simulation of Articulated Bodies

steps	50	100	500	1000	5000
Ceres	100.0	200.0	1100.0	2700.0	23 600.0
CppAD	16.0	26.0	160.0	320.0	3100.0
JAX	200.0	200.0	400.0	700.0	3000.0
PyTorch	1200.0	2400.0	11 700.0	12 400.0	N/A
Ours	<b>0.3</b>	<b>0.3</b>	<b>0.7</b>	<b>1.2</b>	<b>5.0</b>

**Table 1. Peak memory use of different simulation frameworks.** The memory footprint (MB) of our framework is more than two orders of magnitude lower than autodiff methods. PyTorch crashes at 5,000 simulation steps. ADF fails to compute the gradients in a reasonable time (10 min) and is not included here for this reason.

steps	50	100	500	1000	5000
ADF	25.7	25.5	25.1	32.1	58.4
Ceres	27.2	27.5	27.2	34.0	58.2
CppAD	2.4	2.4	2.3	2.3	4.5
JAX	53.3	46.1	43.1	42.7	42.3
PyTorch	195.6	192.2	199.2	192.8	N/A
Ours	<b>0.3</b>	<b>0.3</b>	<b>0.2</b>	<b>0.2</b>	<b>0.2</b>

**Table 2. Simulation time for forward pass.** Ours is at least an order of magnitude faster than autodiff tools (in msec). PyTorch crashes at 5,000 simulation steps. CppAD is the fastest baseline.

Eqs. 12 and 14 yield the same gradients with respect to the action, which provide the gradients of the network parameters. This method, which we call *policy enhancement*, constructs the loss functions with embedded ground-truth gradients so that the policy updates are physics-aware and generally more accurate than merely looking at the Q function itself, thus achieving faster convergence and even potentially higher reward.

## 6. Results

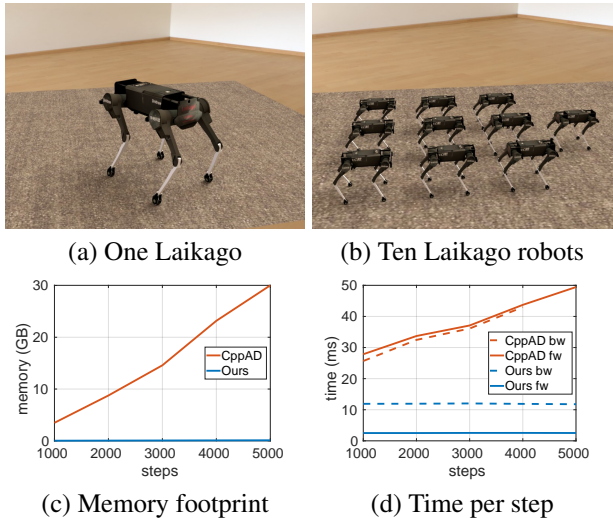
For experiments, we will first scale the complexity of simulated scenes and compare the performance of our method with autodiff tools. Then we will integrate differentiable physics into reinforcement learning and use our method to learn control policies. Lastly, we will apply differentiable articulated dynamics to solve motion control and parameter estimation problems.

### 6.1. Comparison with Autodiff Tools

Using existing autodiff tools is a convenient way to derive simulation gradients. However, these tools are not optimized for articulated dynamics, which adversely affects their computation and memory consumption in this domain. We compare our method with state-of-the-art autodiff tools, including CppAD (Bell et al., 2018), Ceres (Agarwal et al., 2010), PyTorch (Paszke et al., 2019), autodiff (Leal et al., 2018) (referred to as ADF to avoid ambiguity), and JAX (Bradbury et al., 2018). All our experiments are performed on a desktop with an Intel(R) Xeon(R) W-2123 CPU @ 3.60GHz

steps	50	100	500	1000	5000
Ceres	29.5	29.6	29.5	37.6	62.8
CppAD	2.4	2.3	2.3	2.4	4.8
JAX	148.1	125.7	127.6	129.6	126.6
PyTorch	275.7	273.3	280.8	272.1	N/A
Ours	<b>1.2</b>	<b>1.1</b>	<b>1.2</b>	<b>1.1</b>	<b>1.2</b>

**Table 3. Simulation time for the backward pass.** Ours is the fastest (in msec). PyTorch crashes at 5,000 simulation steps.



**Figure 3. Memory and speed of our method vs. CppAD,** the fastest autodiff method. (a,b) Two scenes used in experiments. (c,d) Memory consumption and per-step runtime when simulating ten Laikago robots for different numbers of steps. CppAD crashes when simulating for 5,000 steps due to memory overflow.

with 32GB of memory.

**One robot.** In the first round, we profile all the methods by simulating one Laikago robot standing on the ground. Figure 3(a) provides a visualization. The state vector of the Laikago has 37 dimensions (19-dimensional positions and 18-dimensional velocities). We vary the number of simulation steps from 50 to 5,000. JAX is based on Python and its JIT compiler cannot easily be used in our C++ simulation framework. To test the performance of JAX, we write a simple approximation of our simulator in Python, with similar computational complexity as ours. The simplified test code can be found in the supplement.

The memory consumption of different methods is listed in Table 1. The memory consumed by autodiff tools is orders of magnitude higher than ours. Among all the autodifferentiation tools, JAX and CppAD are the most memory-efficient. ADF fails to backpropagate the gradients in a reasonable time. We show in the supplement that the computation time of backpropagation in ADF is exponential in the depth of the computational graph. Note that articulated body simulation is intrinsically deep. In this experiment, the

depth of one simulation step can reach  $10^3$  due to sequential iterative steps in the forward dynamics and the contact solver. PyTorch crashes at 5,000 simulation steps because of memory overflow.

Table 2 reports the average time for a forward simulation step. The numbers indicate that our method is faster than others by at least an order of magnitude. Our method has a constant time complexity per step. The time per step of ADF, Ceres, and CppAD increases with the number of steps. Note that JAX and Torch are optimized for heavily data-parallel workloads. Articulated body simulation is much more serial in nature and cannot take full advantage of their vectorization. Table 3 reports the time for backward step. Our method is the fastest when computing the gradients.

**Ten robots.** To further test the efficiency of our method, we simulate a scene with 10 Laikago robots. Figure 3(b) provides a visualization. In this experiment, our method is compared with CppAD, the best-performing autodiff framework according to the experiments reported in Tables 1 and 2. The number of simulation steps is varied from 1,000 to 5,000. Figure 3(c) shows that the memory footprint of our method is negligible compared to CppAD. Our method only needs to store 5 KB of data per step, while CppAD needs to save the full data and topology of the computational graph. The forward and backward time per step are plotted in Figure 3(d). Our backward pass takes longer than our forward pass because the method first runs forward simulation to reconstruct the intermediate variables. Nevertheless, our method is much faster than CppAD and the performance gap grows with simulation length.

## 6.2. Integration with Reinforcement Learning

We demonstrate the improvements from the two RL enhancement methods described in Section 5. Note that we cannot use the two methods together because policy enhancement requires the gradients at the sample point, which the extra samples from sample enhancement do not have (unless higher-order gradients are computed). Thus, we test the two techniques separately. We use the model-based MBPO optimizer as the main baseline (Janner et al., 2019). All networks use the PyTorch default initialization scheme.

**N-link pendulum.** We first test our *policy enhancement* method in a simple scenario, where an  $n$ -link pendulum needs to reach a target point by applying torques on each of the joints under gravity (Figure 4(a)). The target point is fixed to be the highest point reachable, and initially the pendulum is positioned horizontally. The reward function is the progress to the target between consecutive steps:

$$r = \|\mathbf{x}_t - \mathbf{x}_g\|^2 - \|\mathbf{x}_{t-1} - \mathbf{x}_g\|^2, \quad (15)$$

where  $\mathbf{x}_t$  is the end-effector location at time  $t$ , and  $\mathbf{x}_g$  is the

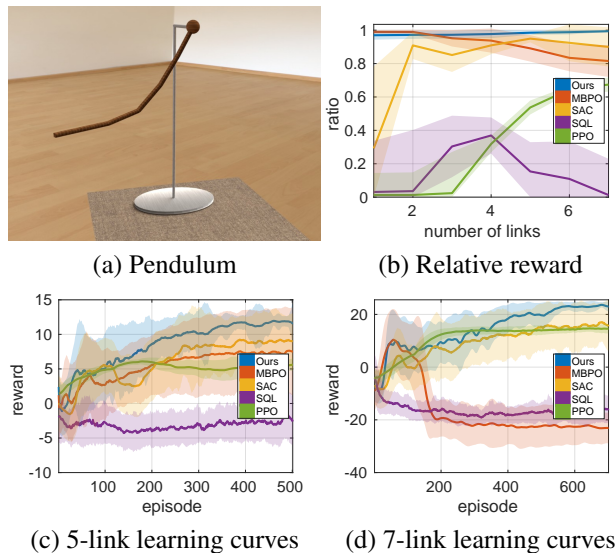


Figure 4. **The  $n$ -link pendulum task.** Our method attains higher reward than MBPO.


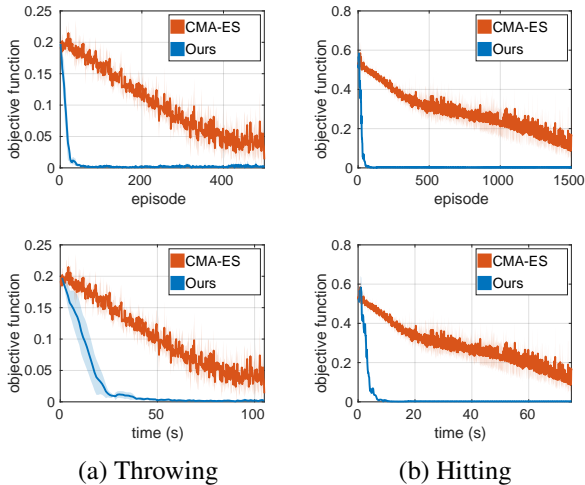


Figure 5. **The MuJoCo Ant task.** Using our differentiable simulator to generate extra samples accelerates learning.

target location.

We trained each model for  $100n$  epochs, where  $n$  is the number of links. The results are shown in Figure 4. In Figure 4(b), we report the relative reward of each task, which is defined as the attained reward divided by the maximal possible reward. MBPO works well in easy scenarios with up to 3 links, but its performance degrades starting at 4 links. In contrast, our model reaches close to the best possible reward for all systems. The learning curves for the 5-link and 7-link systems are shown in Figure 4(c,d). MBPO does not attain satisfactory results for 6- and 7-link systems. We observed that around the 100th epoch, the losses for the Q function of the MBPO method increased a lot. We hypothesize that this is because the complexity of the physical system exceeds the expressive power of the model network in MBPO.

**MuJoCo Ant.** Next, we test our *sample enhancement* method on the MuJoCo Ant. In this scenario, a four-legged robot on the ground needs to learn to walk towards a fixed heading (Figure 5(a)). The scenario is the same as the standard task defined in MuJoCo, except that the simulator is



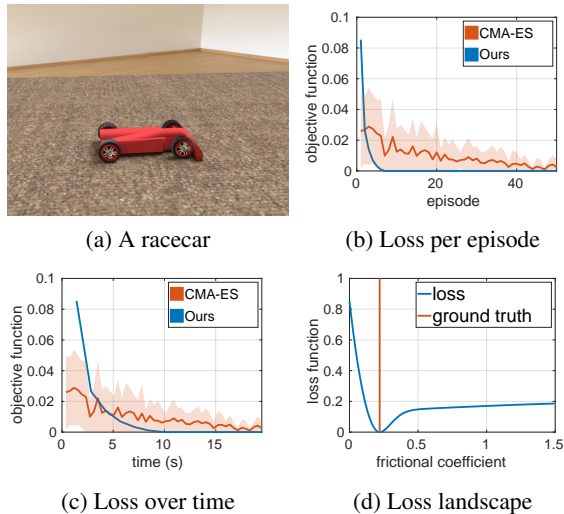
**Figure 6. Motion control.** The task is to optimize the torque vectors to (a) throw a ball with an articulated 9-DoF robot arm and (b) hit a golf ball with an articulated linkage. In both scenarios, the goal is for the ball to settle at a specified target location. SGD with gradients provided by our method converges within 20 and 50 steps, respectively, while gradient-free optimization fails to reach comparable accuracy even with more than an order of magnitude higher number of iterations. The third row indicates that even though CMA-ES does not compute gradients, our method is still faster overall in wall-clock time.

replaced by ours. We also compare with other methods (SAC (Haarnoja et al., 2018), SQL (Haarnoja et al., 2017), and PPO (Schulman et al., 2017) implemented in Ray RLlib (Liang et al., 2018)) for reference. For every true sample we get from the simulator, we generate 9 extra samples around it using our sample enhancement method. Other than that, everything is the same as in MBPO. We repeat the training 4 times for all methods. Figure 5(b) shows the reward over time. Our method exhibits faster convergence.

We can also transfer the policy trained in our differentiable simulator to MuJoCo. This tests the fidelity of our simulator and robustness of the learned policy to simulator details. The results are reported in the supplement.

### 6.3. Applications

**Motion control.** Differentiable physics enables application of gradient-based optimization to learning and control problems that involve physical systems. In Figure 6, we show two physical tasks: (a) throwing a ball to a target location



**Figure 7. Parameter estimation.** The goal is to estimate the sliding friction coefficient that makes the racecar decelerate to a target location. (b,c) Loss curves in episode and wall-clock time, respectively. (d) Loss landscape plot. SGD with gradients computed by our method solves the task within 10 iterations, more accurately and faster than the gradient-free baseline.

using a 9-DoF robot arm and (b) hitting a golf ball to a target location using an articulated linkage. These tasks are also shown in the supplementary video.

For both scenarios, a torque vector is applied to the joints in each time step. Assume there are  $n$  time steps and  $k$  DoFs. The torque variable to be optimized has  $n \times k$  dimensions, and the objective function is the L2 distance from the ball’s actual position to the target. The joint positions and the torques are initialized as  $\mathbf{0}$ . If no gradients are available, a derivative-free optimization algorithm such as CMA-ES (Hansen, 2016) can be used. We plot the loss curves of our method and CMA-ES in Figure 6.

In (a), SGD with gradients from our simulator converges in 20 steps. In contrast, CMA-ES does not reach the same accuracy even after 500 steps. In (b), SGD with gradients from our simulator converges in 50 steps, while CMA-ES does not reach the same accuracy even after 1500 steps.

**Parameter estimation.** Our method can estimate unknown physical parameters from observation. In Figure 7, a racecar starts with horizontal velocity  $1m/s$ . The wheels and the steering system are articulated. We estimate the sliding friction coefficient  $\mu$  between the wheels and the ground such that the car stops at  $x = 0.8m$  at the end of the episode. At the initial guess  $\mu = 0.002$ , the car reaches  $x = 1.0m$ . We use gradient descent to optimize the estimated friction coefficient. After 10 iterations of SGD with gradients provided by our method, the car reaches the goal with  $\mu = 0.21$ . In comparison, the gradient-free baseline takes multiple times longer to reach comparable objective values.



## 7. Conclusion

We have developed a differentiable simulator for articulated body dynamics that runs 10x faster with 100x smaller memory footprint in comparison to existing autodiff tools. To achieve this performance gain, we analyze the workload of each simulation step to better manage computation and memory. The adjoint method is used to compute the gradients of the simulation. We derive the adjoints of spatial algebra and the Gauss-Seidel solver. We then adapt the checkpointing method to deal with the sequential nature of articulated body simulation, reducing memory consumption by two orders of magnitude.

As a supporting contribution, we have explored two applications of differentiable physics to reinforcement learning with physical systems, reporting preliminary results that indicate that differentiable physics can accelerate learning. We have also presented application scenarios that clearly demonstrate the effectiveness of gradient-based optimization in motion control and parameter estimation with articulated physical systems.

**Acknowledgements.** This research is supported in part by Intel Corporation, National Science Foundation, Dr. Barry Mersky and Capital One Endowed Professorship, and Elizabeth Stevinson Iribe Endowed Chair Professorship.

## References

- Abadi, M., Barham, P., Chen, J., Chen, Z., Davis, A., Dean, J., Devin, M., Ghemawat, S., Irving, G., Isard, M., et al. TensorFlow: A system for large-scale machine learning. In *OSDI*, 2016.
- Agarwal, S., Mierle, K., et al. Ceres solver. <http://ceres-solver.org>, 2010.
- Anderson, B. M., Hy, T., and Kondor, R. Cormorant: Covariant molecular neural networks. In *Neural Information Processing Systems*, 2019.
- Battaglia, P. W., Pascanu, R., Lai, M., Rezendes, D., and Kavukcuoglu, K. Interaction networks for learning about objects, relations and physics. In *Neural Information Processing Systems*, 2016.
- Belbute-Peres, F. d. A., Economou, T. D., and Kolter, J. Z. Combining differentiable PDE solvers and graph neural networks for fluid flow prediction. In *ICML*, 2020.
- Bell, B. M. et al. CppAD: C++ algorithmic differentiation. <https://projects.coin-or.org/CppAD>, 2018.
- Bern, J. M., Banzet, P., Poranne, R., and Coros, S. Trajectory optimization for cable-driven soft robot locomotion. In *Robotics: Science and Systems (RSS)*, 2019.
- Bogatskiy, A., Anderson, B. M., Offermann, J. T., Roussi, M., Miller, D. W., and Kondor, R. Lorentz group equivariant neural network for particle physics. In *ICML*, 2020.
- Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. JAX: Composable transformations of Python+NumPy programs. <http://github.com/google/jax>, 2018.
- Chang, M. B., Ullman, T., Torralba, A., and Tenenbaum, J. B. A compositional object-based approach to learning physical dynamics. In *ICLR*, 2017.
- Chen, R. T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K. Neural ordinary differential equations. In *Neural Information Processing Systems*, 2018.
- Chen, T., Xu, B., Zhang, C., and Guestrin, C. Training deep nets with sublinear memory cost. *arXiv:1604.06174*, 2016.
- Chen, Z., Zhang, J., Arjovsky, M., and Bottou, L. Symplectic recurrent neural networks. In *ICLR*, 2020.
- Coumans, E. Bullet physics simulation. In *ACM SIGGRAPH Courses*, 2015.
- Coumans, E. et al. Tiny differentiable simulator. <https://github.com/google-research/tiny-differentiable-simulator>, 2020.
- Cranmer, M., Greydanus, S., Hoyer, S., Battaglia, P., Spergel, D., and Ho, S. Lagrangian neural networks. In *ICLR Workshops*, 2020.
- de Avila Belbute-Peres, F., Smith, K. A., Allen, K. R., Tenenbaum, J., and Kolter, J. Z. End-to-end differentiable physics for learning and control. In *Neural Information Processing Systems*, 2018.
- Degrave, J., Hermans, M., Dambre, J., and wyffels, F. A differentiable physics engine for deep learning in robotics. *Frontiers in Neurorobotics*, 13, 2019.
- Du, T., Wu, K., Spielberg, A., Matusik, W., Zhu, B., and Sifakis, E. Functional optimization of fluidic devices with differentiable stokes flow. *ACM Trans. Graph.*, 2020.
- Featherstone, R. Rigid body dynamics algorithms. 2007.
- Featherstone, R. Spatial vector and dynamics software. <http://royfeatherstone.org/spatial/>, 2008.
- Geilinger, M., Hahn, D., Zehnder, J., Bächer, M., Thomaszewski, B., and Coros, S. ADD: Analytically differentiable dynamics for multi-body systems with frictional contact. *ACM Trans. Graph.*, 39(6), 2020.

- Giles, M. B. and Pierce, N. A. An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion*, 65(3-4):393–415, 2000.
- Griewank, A. and Walther, A. Revolve: An implementation of checkpointing for the reverse or adjoint mode of computational differentiation. *ACM Trans. Math. Softw.*, 26(1):19–45, 2000.
- Griewank, A. and Walther, A. *Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation*. SIAM, 2008.
- Gruslys, A., Munos, R., Danihelka, I., Lanctot, M., and Graves, A. Memory-efficient backpropagation through time. In *Neural Information Processing Systems*, 2016.
- Haarnoja, T., Tang, H., Abbeel, P., and Levine, S. Reinforcement learning with deep energy-based policies. In *ICML*, 2017.
- Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *ICML*, 2018.
- Hansen, N. The CMA evolution strategy: A tutorial. *arXiv:1604.00772*, 2016.
- Heiden, E., Millard, D., Zhang, H., and Sukhatme, G. S. Interactive differentiable simulation. *arXiv:1905.10706*, 2019.
- Heiden, E., Millard, D., Coumans, E., Sheng, Y., and Sukhatme, G. S. NeuralSim: Augmenting differentiable simulators with neural networks. *arXiv:2011.04217*, 2020.
- Holl, P., Koltun, V., and Thuerey, N. Learning to control PDEs with differentiable physics. In *ICLR*, 2020.
- Hu, Y., Liu, J., Spielberg, A., Tenenbaum, J. B., Freeman, W. T., Wu, J., Rus, D., and Matusik, W. Chain-Queen: A real-time differentiable physical simulator for soft robotics. In *ICRA*, 2019.
- Hu, Y., Anderson, L., Li, T.-M., Sun, Q., Carr, N., Ragan-Kelley, J., and Durand, F. DiffTaichi: Differentiable programming for physical simulation. In *ICLR*, 2020.
- Huang, Z., Hu, Y., Du, T., Zhou, S., Su, H., Tenenbaum, J. B., and Gan, C. PlasticineLab: A soft-body manipulation benchmark with differentiable physics. In *ICLR*, 2021.
- Ingraham, J., Riesselman, A., Sander, C., and Marks, D. Learning protein structure with a differentiable simulator. In *ICLR*, 2019.
- Janner, M., Fu, J., Zhang, M., and Levine, S. When to trust your model: Model-based policy optimization. In *Neural Information Processing Systems*, 2019.
- Kirisame, M., Lyubomirsky, S., Haan, A., Brennan, J., He, M., Roesch, J., Chen, T., and Tatlock, Z. Dynamic tensor rematerialization. In *ICLR*, 2021.
- Krishna Murthy, J., Iyer, G., and Paull, L. gradSLAM: Dense SLAM meets automatic differentiation. In *ICRA*, 2020.
- Krishna Murthy, J., Macklin, M., Golemo, F., Voleti, V., Petrini, L., Weiss, M., Consideine, B., Parent-Lévesque, J., Xie, K., Erleben, K., et al. gradSim: Differentiable simulation for system identification and visuomotor control. In *ICLR*, 2021.
- Laine, S., Hellsten, J., Karras, T., Seol, Y., Lehtinen, J., and Aila, T. Modular primitives for high-performance differentiable rendering. *ACM Trans. Graph.*, 39(6), 2020.
- Leal, A. et al. autodiff: automatic differentiation made easier for C++. <https://github.com/autodiff/autodiff>, 2018.
- Lee, J., Grey, M., Ha, S., Kunz, T., Jain, S., Ye, Y., Srinivasa, S., Stilman, M., and Liu, C. DART: Dynamic animation and robotics toolkit. *The Journal of Open Source Software*, 3, 2018.
- Li, T.-M., Aittala, M., Durand, F., and Lehtinen, J. Differentiable monte carlo ray tracing through edge sampling. *ACM Trans. Graph.*, 37(6), 2018a.
- Li, T.-M., Gharbi, M., Adams, A., Durand, F., and Ragan-Kelley, J. Differentiable programming for image processing and deep learning in Halide. *ACM Trans. Graph.*, 37(4), 2018b.
- Li, T.-M., Lukáč, M., Michaël, G., and Ragan-Kelley, J. Differentiable vector graphics rasterization for editing and learning. *ACM Trans. Graph.*, 39(6), 2020.
- Li, Y., Wu, J., Zhu, J.-Y., Tenenbaum, J. B., Torralba, A., and Tedrake, R. Propagation networks for model-based control under partial observation. In *ICRA*, 2019.
- Liang, E., Liaw, R., Nishihara, R., Moritz, P., Fox, R., Goldberg, K., Gonzalez, J. E., Jordan, M. I., and Stoica, I. RLlib: Abstractions for distributed reinforcement learning. In *ICML*, 2018.
- Liang, J., Lin, M. C., and Koltun, V. Differentiable cloth simulation for inverse problems. In *Neural Information Processing Systems*, 2019.

- McNamara, A., Treuille, A., Popovic, Z., and Stam, J. Fluid control using the adjoint method. *ACM Trans. Graph.*, 23(3), 2004.
- Mrowca, D., Zhuang, C., Wang, E., Haber, N., Li, F., Tenenbaum, J., and Yamins, D. L. Flexible neural representation for physics prediction. In *Neural Information Processing Systems*, 2018.
- Nimier-David, M., Vicini, D., Zeltner, T., and Jakob, W. Mitsuba 2: A retargetable forward and inverse renderer. *ACM Trans. Graph.*, 38(6), 2019.
- Nimier-David, M., Speierer, S., Ruiz, B., and Jakob, W. Radiative backpropagation: An adjoint method for lightning-fast differentiable rendering. *ACM Trans. Graph.*, 39(4), 2020.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al. PyTorch: An imperative style, high-performance deep learning library. In *Neural Information Processing Systems*, 2019.
- Pfaff, T., Fortunato, M., Sanchez-Gonzalez, A., and Battaglia, P. Learning mesh-based simulation with graph networks. In *ICLR*, 2021.
- Qiao, Y.-L., Liang, J., Koltun, V., and Lin, M. C. Scalable differentiable physics for learning and control. In *ICML*, 2020.
- Sanchez-Gonzalez, A., Heess, N., Springenberg, J. T., Merel, J., Riedmiller, M., Hadsell, R., and Battaglia, P. Graph networks as learnable physics engines for inference and control. In *ICML*, 2018.
- Sanchez-Gonzalez, A., Godwin, J., Pfaff, T., Ying, R., Leskovec, J., and Battaglia, P. W. Learning to simulate complex physics with graph networks. In *ICML*, 2020.
- Schenck, C. and Fox, D. SPNets: Differentiable fluid dynamics for deep neural networks. In *Conference on Robot Learning (CoRL)*, 2018.
- Schoenholz, S. S. and Cubuk, E. D. JAX M.D.: End-to-end differentiable, hardware accelerated, molecular dynamics in pure Python. In *Neural Information Processing Systems*, 2020.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., and Klimov, O. Proximal policy optimization algorithms. *arXiv:1707.06347*, 2017.
- Schütt, K., Kindermans, P., Felix, H. E. S., Chmiela, S., Tkatchenko, A., and Müller, K. SchNet: A continuous-filter convolutional neural network for modeling quantum interactions. In *Neural Information Processing Systems*, 2017.
- Shah, A., Wu, C.-Y., Mohan, J., Chidambaram, V., and Krähenbühl, P. Memory optimization for deep networks. In *ICLR*, 2021.
- Song, C. and Boularias, A. Learning to slide unknown objects with differentiable physics simulations. In *Robotics: Science and Systems (RSS)*, 2020a.
- Song, C. and Boularias, A. Identifying mechanical models of unknown objects with differentiable physics simulations. In *LADC*, 2020b.
- Spielberg, A., Zhao, A., Hu, Y., Du, T., Matusik, W., and Rus, D. Learning-in-the-loop optimization: End-to-end control and co-design of soft robots through learned deep latent representations. In *Neural Information Processing Systems*, 2019.
- Stepien, J. *Physics-Based Animation of Articulated Rigid Body Systems for Virtual Environments*. PhD thesis, 2013.
- Takahashi, T., Liang, J., Qiao, Y.-L., and Lin, M. C. Differentiable fluids with solid coupling for learning and control. In *AAAI*, 2021.
- Tedrake, R. et al. Drake: Model-based design and verification for robotics. <https://drake.mit.edu>, 2019.
- Todorov, E., Erez, T., and Tassa, Y. MuJoCo: A physics engine for model-based control. In *IROS*, 2012.
- Toussaint, M., Allen, K., Smith, K., and Tenenbaum, J. Differentiable physics and stable modes for tool-use and manipulation planning. In *Robotics: Science and Systems (RSS)*, 2018.
- Ummenhofer, B., Prantl, L., Thürey, N., and Koltun, V. Lagrangian fluid simulation with continuous convolutions. In *ICLR*, 2020.
- Wandel, N., Weinmann, M., and Klein, R. Learning incompressible fluid dynamics from scratch – towards fast, differentiable fluid models that generalize. In *ICLR*, 2021.
- Wang, K., Aanjaneya, M., and Bekris, K. E. A first principles approach for data-efficient system identification of spring-rod systems via differentiable physics engines. In *LADC*, 2020.
- Zhuang, J., Dvornik, N. C., Li, X., Tatikonda, S., Papademetris, X., and Duncan, J. S. Adaptive checkpoint adjoint method for gradient estimation in neural ODE. In *ICML*, 2020.