ADD: Analytically Differentiable Dynamics for Multi-Body Systems with Frictional Contact

MORITZ GEILINGER* and DAVID HAHN*, ETH Zürich

JONAS ZEHNDER, Université de Montréal

MORITZ BÄCHER, Disney Research

BERNHARD THOMASZEWSKI, ETH Zürich and Université de Montréal

STELIAN COROS, ETH Zürich

Fig. 1. Applications of our differentiable simulation framework (left-to-right): estimation of stiffness, damping and friction properties from real-world experiments, manipulation of multi-body systems, self-supervised learning of control policies for a throwing task, and physics-based motion planning for robotic creatures with compliant motors and soft feet.

We present a differentiable dynamics solver that is able to handle frictional contact for rigid and deformable objects within a unified framework. Through a principled mollification of normal and tangential contact forces, our method circumvents the main difficulties inherent to the non-smooth nature of frictional contact. We combine this new contact model with fully-implicit time integration to obtain a robust and efficient dynamics solver that is analytically differentiable. In conjunction with adjoint sensitivity analysis, our formulation enables gradient-based optimization with adaptive trade-offs between simulation accuracy and smoothness of objective function landscapes. We thoroughly analyse our approach on a set of simulation examples involving rigid bodies, visco-elastic materials, and coupled multi-body systems. We furthermore showcase applications of our differentiable simulator to parameter estimation for deformable objects, motion planning for robotic manipulation, trajectory optimization for compliant walking robots, as well as efficient self-supervised learning of control policies.

CCS Concepts: • **Computing methodologies** → **Physical simulation**; **Motion processing**.

Additional Key Words and Phrases: Differentiable simulation, Optimization, Deformable Models, Rigid-body dynamics, Contact mechanics, Robotics

*Joint first authors.

Authors' addresses: Moritz Geilinger, moritz.geilinger@inf.ethz.ch; David Hahn, david.hahn@inf.ethz.ch, CRL, ETH Zürich; Jonas Zehnder, jonas.zehnder@umontreal.ca, LIGUM, Université de Montréal; Moritz Bächer, moritz.baecher@disneyresearch.com, Disney Research, Zürich; Bernhard Thomaszewski, bthomasz@inf.ethz.ch, CRL, ETH Zürich, LIGUM, Université de Montréal; Stelian Coros, scoros@inf.ethz.ch, CRL, ETH Zürich.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0730-0301/2020/12-ART190 \$15.00

<https://doi.org/10.1145/3414685.3417766>

ACM Reference Format:

Moritz Geilinger, David Hahn, Jonas Zehnder, Moritz Bächer, Bernhard Thomaszewski, and Stelian Coros. 2020. ADD: Analytically Differentiable Dynamics for Multi-Body Systems with Frictional Contact. *ACM Trans. Graph.* 39, 6, Article 190 (December 2020), 15 pages. <https://doi.org/10.1145/3414685.3417766>

1 INTRODUCTION

Simulation tools are crucial to a variety of applications in engineering and robotics, where they can be used to test the performance of a design long before the first prototype is ever built. Based on forward simulation, however, these virtual proving grounds are typically limited to trial-and-error approaches, where the onus is on the user to painstakingly find appropriate control or design parameters. Inverse simulation tools promise a more direct and powerful approach, as they can anticipate and exploit the way in which a change in parameters affects the performance of the design. Unlike forward simulation, this inverse approach hinges on the ability to compute derivatives of simulation runs.

Differentiable simulations have recently seen increasing attention in the vision, graphics, and robotics communities, e.g., in the context of fluid animation [Holl et al. 2020], soft-body dynamics [Hu et al. 2019b], and light transport [Loubet et al. 2019]. As a characteristic trait of real-world interactions, however, the strong non-linearities and quasi-discontinuities induced by collisions and frictional contact pose substantial challenges for differentiable simulation.

The equations of motion for mechanical systems with frictional contact can be cast as a non-linear complementarity problem that couples the tangential forces to the magnitude of the normal forces. While the exact form depends on the friction model being used, this formulation typically postulates different regimes—approaching or separating and stick or slip—with different bounds on the admissible forces. These dichotomies induce discontinuities in forces and their

derivatives, which, already for the forward problem, necessitate complex numerical solvers that are computationally demanding. For inverse simulations that rely on gradient-based optimization, however, this lack of smoothness is a major stumbling block.

To overcome the difficulties of the non-smooth problem setting, we propose a differentiable simulator that combines fully implicit time-stepping with a principled mollification of normal and tangential contact forces. These contact forces, as well as the coupled system mechanics of rigid and soft objects, are handled through a soft constraint formulation that is simple, numerically robust and very effective. This formulation also allows us to analytically compute derivatives of simulation outcomes through adjoint sensitivity analysis. Our simulation model enables an easy-to-tune trade-off between accuracy and smoothness of objective function landscapes, and we showcase its potential by applying it to a variety of application domains, which can be summarized as follows:

Real2Sim: we perform data-driven parameter estimation to find material constants that characterize the stiffness, viscous damping, and friction properties of deformable objects. This enables the creation of simulation models that accurately reproduce and predict the behaviour of physical specimens.

Control: our differentiable simulator is ideally suited for control problems that involve coupled dynamical systems and frictional contact. We explore this important problem domain in the context of robotic manipulation, focusing, through simulation and physical experiments, on behaviours that include controlled multi-bounce tossing, dragging, and sliding of different types of objects. We also use our simulator to generate optimal motion trajectories for a simulated legged robot with compliant actuators and soft feet.

Self-supervised learning: by embedding our simulator as a specialized node in a neural network, we show that control policies that map target outcomes to appropriate control inputs can be easily learned in a self-supervised manner; this is achieved by defining the loss directly as a function of the results generated with our differentiable simulation model. We explore this concept through a simple game where a robot arm learns how to toss a ball such that it lands in an interactively placed cup after exactly one bounce.

Our main contributions can be summarized as follows:

- A stable, differentiable, two-way coupled simulation framework for both rigid and soft bodies.
- A smooth frictional contact model that can effectively balance differentiability requirements for inverse problems with accurate modelling of real-world behaviour.
- An evaluation of our friction formulations on a diverse set of applications that highlight the benefits of differentiable simulators.

2 RELATED WORK

Contact modelling is a well-studied problem in mechanics (see e.g. [Brogliato 1999]), and has received a significant amount of attention in graphics due to its importance in physics-based animation. Here, we focus our review on frictional contact dynamics, its differentiability, and its use in solving inverse problems.

2.1 Frictional Contact Dynamics

For rigid bodies, early work focused on acceleration-level [Baraff 1994] and velocity-level [Anitescu and Potra 1997; Stewart and Trinkle 1996] linear complementarity problem (LCP) formulations and isotropic Coulomb friction. Later, iterative LCP approaches were explored [Dureiz et al. 2006; Erleben 2007], and LCP formulations extended to quasi-rigid objects [Pauly et al. 2004], deformable solids [Dureiz et al. 2006; Otaduy et al. 2009], and fast simulation of large sets of rigid bodies [Kaufman et al. 2005]. Harmon et al. [2009] propose an asynchronous treatment of contact mechanics for deformable models, discretizing the contact barrier potential. Unifying fast collision detection and contact force computations, Allard et al. [2010] target deformable contact dynamics where deeper penetrations can arise. Recent work [Ding and Schroeder 2020] has integrated penalty-based friction models with optimization-based time integration. Exact Coulomb friction paired with adaptive node refinement at contact locations is used for accurate cloth simulation [Li et al. 2018]. Anisotropic Coulomb friction in the deformable [Pabst et al. 2009] and rigid-body setting [Erleben et al. 2019] has also been considered.

Like Kaufman et al. [2008], our frictional contact formulation applies to coupled rigid and deformable bodies. For this setting, Macklin et al. [2019] solve a *non-linear* complementarity problem for an exact Coulomb friction formulation with a non-smooth Newton method. A similar technique was applied to fibre assemblies [Bertails-Descoubes et al. 2011]. However, in contrast to the above body of work, we target inverse contact applications [Chen et al. 2017a; Coros et al. 2012; Popović et al. 2000; Tan et al. 2012; Twigg and James 2007], rather than animation. Interestingly, our hybrid approach can be interpreted as the opposite of staggered projections [Kaufman et al. 2008], as we soften contact constraints with penalty forces, but can enforce static friction with hard constraints. Softening contact has also proven useful for the control of physics-based characters [Jain and Liu 2011]. In the context of identifying frictional parameters, we share goals with Pai et al. [2001]. However, in contrast to measuring frictional textures with a robotic system, we aim at characterizing material properties and initial conditions from motion data. For parameter estimation, we draw inspiration from physics-guided reconstruction [Monszpart et al. 2016], who consider frictionless contact between rigid bodies. Our formulation extends to frictional contact and deformable bodies.

2.2 Inverse Contact

Ly et al. [2018] solve for the rest configuration and friction forces of a shell model such that the corresponding static equilibrium state best approximates a user-specified target shape. In a first pass, contact is handled using frictionless bilateral constraints whereas frictional forces are found in a second pass. While Ly et al. focus on quasi-static inverse problems, our focus is on inverse *dynamics*. To enable design optimization, Chen et al. [2017b] introduce an implicitly integrated elasto-dynamics solver that can handle complex contact scenarios. We share the goal of making our frictional contact model smooth and invertible with previous work in optimal control [Erez and Todorov 2012; Todorov 2011; Yunt and Glocker 2006]. However, instead of relying on finite-differencing

and explicit time integration [Todorov et al. 2012], our approach builds on analytically-computed derivatives and fully implicit time integration, allowing us to handle both soft and stiff systems.

2.3 Differentiable Simulation

Emerging differentiable physics-based simulators are increasingly used for applications in machine learning, physics-based control, and inverse design. While common physics engines such as NVIDIA’s PhysX or Bullet Physics enjoy widespread use, they prioritize speed over accuracy, especially in the context of dynamics of deformable objects.

The “MuJoCo” framework [Todorov et al. 2012] has been designed as a differentiable simulator for articulated rigid-body dynamics. However, derivatives are computed through finite differences, and the use of an explicit time integration method in conjunction with a penalty-based contact model restricts the size of the time steps that can be used for reasonably accurate results. Targeting end-to-end learning, de Avila Belbute-Peres et al. [2018] compute analytical gradients of LCPs for simple rigid-body systems. In contrast, our approach is able to handle multi-body systems that couple rigid and deformable objects and is very robust with respect to the coefficients used for the contact model even under large time steps. Learning-based approaches have also been applied to fluids [Schenck and Fox 2018], cloth [Liang et al. 2019], and robotic manipulation [Toussaint et al. 2018].

Hu et al. [2019b] base their differentiable simulation on the material point method and target specific applications in soft robotics for which explicit time integration is sufficient. Degraeve et al. [2019] rely on auto-differentiation to compute derivatives of rigid-body dynamics. Similarly, for end-to-end learning of controllers, recording and reversing simulations for back-propagation, DiffTaichi [2019a] also relies on auto-differentiation. While auto-differentiation can work well in some settings, the stiff problems we are considering demand fully implicit integration schemes that must run to convergence. Consequently, our non-linear solver requires a variable number of iterations in the outer loop, accurate and efficient sparse linear solvers in the inner loop, and line-search mechanisms to enforce monotonicity. In this setting, it is very unclear how to apply auto-differentiation, as it would need to act on the precise sequence of numerical operations performed during the solve. Our approach, on the other hand, analytically computes derivatives via sensitivity analysis, similar to recent work in graphics [Hahn et al. 2019; Hoshiyari et al. 2019], which describes differentiable simulators for visco-elastic materials and constrained flexible multi-body dynamics. However, as a significant departure from prior work, frictional contact is at the core of our paper.

3 DIFFERENTIABLE MULTI-BODY DYNAMICS: PRELIMINARIES

We begin by deriving the mathematical formulation underlying our analytically differentiable simulation model. While this derivation follows general concepts recently described in the literature [Hahn et al. 2019; Zimmermann et al. 2019b], we note that our quest for a unified simulation framework for multi-body systems composed of

arbitrary arrangements of rigid and soft objects demands a somewhat different methodology, as detailed below.

3.1 Implicit time-stepping for multi-body systems

We start from a time-continuous dynamics problem described in terms of generalized coordinates $\mathbf{q}(t)$, see also [Kaufman et al. 2008]. These generalized coordinates can directly represent the nodal positions of a finite element mesh or the degrees of freedom that describe the position and orientation of a rigid body in space. Without loss of generality, here we consider dynamical systems that are characterized by a set of *time independent* input parameters $\mathbf{p}$, which could represent, for instance, material constants or a fixed sequence of control actions. The dynamics of such a system can be succinctly represented through a differential equation of the form:

$$\mathbf{r}(\mathbf{q}, \dot{\mathbf{q}}, \ddot{\mathbf{q}}, \mathbf{p}) = \mathbf{0} \quad \forall t, \quad \text{where } \mathbf{r} := \hat{\mathbf{M}}(\mathbf{q}, \mathbf{p})\ddot{\mathbf{q}} - \hat{\mathbf{f}}(\mathbf{q}, \dot{\mathbf{q}}, \mathbf{p}). \quad (1)$$

For soft bodies, the generalized mass $\hat{\mathbf{M}}$ is identical to the constant FEM mass matrix, whereas for each rigid body, the corresponding 6×6 block is defined as

$$\hat{\mathbf{M}}_{\text{RB}} := \begin{pmatrix} m\mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_q \end{pmatrix},$$

with m and $\mathbf{I}_q$ representing its mass and configuration-dependent moment of inertia tensor in generalized coordinates [Liu and Jain 2012] (see also our supplements).

The generalized force vector $\hat{\mathbf{f}}$ aggregates all internal and external forces applied to the simulation’s degrees of freedom. Sections 4 and 5 detail the exact mathematical formulation for these forces, but here we note that for each rigid body, forces must be mapped from world to generalized space, and an additional term $\mathbf{C}(\mathbf{q}, \dot{\mathbf{q}})$ that captures the effect of fictitious centrifugal and Coriolis forces must also be added.

Turning to a time-discrete setting, we opt to discretize the residual $\mathbf{r}$ using implicit numerical integration schemes. This decision warrants a brief discussion. While implicit integration schemes are standard when it comes to simulation of elastic objects, rigid-body dynamics is typically handled with explicit methods such as symplectic Euler. This misalignment in the choice of integrators often necessitates multi-stage time-stepping schemes for two-way coupling of rigid bodies and elastic objects [Shinar et al. 2008]. One notable exception is “RedMax” [Wang et al. 2019], which directly couples rigid and deformable objects. However, their system is only efficient for a small number of deformable DOFs, and they employ staggered projections for ground contacts, which is not well suited for differentiable simulation. Although effective, such a specialized time-stepping scheme complicates matters when differentiating simulation results. Furthermore, as discussed in §4, the smooth contact model we propose relies on numerically stiff forces arising from a unilateral potential. Robustly resolving contacts in this setting demands the use of implicit integration for both rigid and elastic objects. As an added bonus, two-way coupling between these two classes of objects becomes straightforward and numerically robust even under large simulation time steps.

We begin by approximating the first and second time derivatives of the system’s generalized coordinates, $\dot{\mathbf{q}}$ and $\ddot{\mathbf{q}}$, according to an implicit time-discretization scheme of our choice. For instance with

BDF1 (implicit Euler), they are $\dot{\mathbf{q}}(t_i) \approx \dot{\mathbf{q}}^i = (\mathbf{q}^i - \mathbf{q}^{i-1})/\Delta_t$ and $\ddot{\mathbf{q}}(t_i) \approx \ddot{\mathbf{q}}^i = (\mathbf{q}^i - 2\mathbf{q}^{i-1} + \mathbf{q}^{i-2})/\Delta_t^2$, where Δ_t is a constant finite time step and superscript i refers to the time-step index for time t_i .

For each forward simulation step i we then employ Newton's method to find the end-of-time-step configuration $\mathbf{q}^i$ such that $\mathbf{r}^i := \mathbf{r}(\mathbf{q}^i, \dot{\mathbf{q}}^i, \ddot{\mathbf{q}}^i, \mathbf{p}) = \hat{\mathbf{M}}(\mathbf{q}^i, \mathbf{p})\dot{\mathbf{q}}^i - \hat{\mathbf{f}}(\mathbf{q}^i, \dot{\mathbf{q}}^i, \mathbf{p}) = 0$. This process requires the Jacobian $d\mathbf{r}^i/d\mathbf{q}^i$, which combines derivatives of the generalized forces, the discretized generalized velocities and accelerations, as well as the generalized mass matrix w.r.t. $\mathbf{q}^i$. Recall that for rigid bodies the mass matrix is configuration-dependent, see also our supplements for details. We note that in contrast to previous work [Hahn et al. 2019; Zimmermann et al. 2019b], we cannot treat time-stepping as an energy minimization problem due to this state dependent nature of the mass matrix. A recent alternative formulation for rigid bodies [Pan and Manocha 2018] could overcome this limitation. Nevertheless, in conjunction with a line-search routine that monitors the magnitude of the residual, directly finding the root of $\mathbf{r}^i$ works very well in practice.

3.2 Simulation Derivatives

With the basic procedure for implicit time-stepping in place, we now turn our attention to the task of computing derivatives of simulation outcomes. To keep the exposition brief, we concatenate the generalized coordinates for an entire sequence of time steps computed through forward simulation into a vector $\tilde{\mathbf{q}} := (\mathbf{q}^{1\top}, \dots, \mathbf{q}^{n_t\top})^\top$, where n_t is the number of time steps. Similarly, we let $\tilde{\mathbf{r}}$ be the vector that concatenates the residuals for all time steps of a simulation run. Differentiating $\tilde{\mathbf{r}}$ with respect to input parameters $\mathbf{p}$ gives:

$$\frac{d\tilde{\mathbf{r}}}{d\mathbf{p}} = \frac{\partial \tilde{\mathbf{r}}}{\partial \mathbf{p}} + \frac{\partial \tilde{\mathbf{r}}}{\partial \tilde{\mathbf{q}}} \frac{d\tilde{\mathbf{q}}}{d\mathbf{p}}. \quad (2)$$

Since for *any* $\mathbf{p}$ we compute a motion trajectory such that $\tilde{\mathbf{r}} = \mathbf{0}$ during the forward simulation stage, the total derivative (2) is also $\mathbf{0}$. This simple observation lets us compute the Jacobian $d\tilde{\mathbf{q}}/d\mathbf{p}$, which is also called the sensitivity matrix, as:

$$\mathbf{S} := \frac{d\tilde{\mathbf{q}}}{d\mathbf{p}} = - \left(\frac{\partial \tilde{\mathbf{r}}}{\partial \tilde{\mathbf{q}}} \right)^{-1} \frac{\partial \tilde{\mathbf{r}}}{\partial \mathbf{p}}. \quad (3)$$

The sensitivity matrix $\mathbf{S}$ captures the way in which an entire trajectory computed through forward simulation changes as the input parameters $\mathbf{p}$ change. Note that $\mathbf{S}$ can be computed efficiently by exploiting the sparsity structure of $\partial \tilde{\mathbf{r}}/\partial \tilde{\mathbf{q}}$, a matrix that can be easily evaluated by re-using most of the same ingredients already computed for forward simulation. For example, in the case of BDF1, each sub-block $\mathbf{s}_j^i := d\mathbf{q}^i/d\mathbf{p}_j$ of $\mathbf{S}$, which represents the change in the configuration of the dynamical system at time t_i with respect to the j -th input parameter $\mathbf{p}_j$, reduces to:

$$\mathbf{s}_j^i = - \left(\frac{\partial \mathbf{r}^i}{\partial \mathbf{q}^i} \right)^{-1} \left(\frac{\partial \mathbf{r}^i}{\partial \mathbf{p}_j} + \frac{\partial \mathbf{r}^i}{\partial \mathbf{q}^{i-1}} \mathbf{s}_j^{i-1} + \frac{\partial \mathbf{r}^i}{\partial \mathbf{q}^{i-2}} \mathbf{s}_j^{i-2} \right). \quad (4)$$

Note that the dependency on past sensitivities follows the same time-stepping scheme as the forward simulation. We refer the interested reader to [Zimmermann et al. 2019a] for more details on this topic, but for completeness we briefly describe a variant of this formulation, the adjoint method, as it can lead to additional gains in computational efficiency.

3.3 The adjoint method

For most applications we do not need to compute the sensitivity matrix explicitly. Instead, what we are interested in, is finding input parameters $\mathbf{p}^*$ that minimize an objective defined as a function of the simulation result:

$$\mathbf{p}^* = \arg \min_{\mathbf{p}} \Phi(\mathbf{p}, \tilde{\mathbf{q}}(\mathbf{p})), \quad (5)$$

where $\tilde{\mathbf{q}}(\mathbf{p})$ is the entire trajectory of the dynamical system.

In order to solve this optimization problem using gradient-based methods such as ADAM [Kingma and Ba 2014; O'Hara 2019] or L-BFGS [Nocedal 1980; Qiu 2019], we need to compute the derivative

$$\frac{d\Phi}{d\mathbf{p}} = \frac{\partial \Phi}{\partial \mathbf{p}} + \frac{\partial \Phi}{\partial \tilde{\mathbf{q}}} \mathbf{S}. \quad (6)$$

The adjoint method enables the computation of this objective function gradient without computing $\mathbf{S}$ directly. While Bradley [2013] derives the adjoint method for the time-continuous case, here we describe it directly in the time-discretized setting.

We first express the sensitivity matrix as $\mathbf{S} = -\tilde{\mathbf{B}}^{-1}\tilde{\mathbf{A}}$, with $\tilde{\mathbf{A}} := \partial \tilde{\mathbf{r}}/\partial \mathbf{p}$ and $\tilde{\mathbf{B}} := \partial \tilde{\mathbf{r}}/\partial \tilde{\mathbf{q}}$, respectively. Introducing the notation $\tilde{\mathbf{y}} := \partial \Phi/\partial \tilde{\mathbf{q}}$, Eq. (6) becomes

$$\frac{d\Phi}{d\mathbf{p}} = \frac{\partial \Phi}{\partial \mathbf{p}} - \tilde{\mathbf{y}} \tilde{\mathbf{B}}^{-1} \tilde{\mathbf{A}}, \quad (7)$$

Instead of evaluating this expression directly, we first solve for the *adjoint state* $\tilde{\boldsymbol{\lambda}}$:

$$\tilde{\mathbf{B}}^\top \tilde{\boldsymbol{\lambda}} = \tilde{\mathbf{y}}^\top, \quad (8)$$

and then compute the objective function gradient as:

$$\frac{d\Phi}{d\mathbf{p}} = \frac{\partial \Phi}{\partial \mathbf{p}} - \tilde{\boldsymbol{\lambda}}^\top \tilde{\mathbf{A}}. \quad (9)$$

Note that both $\tilde{\mathbf{y}}$ and $\tilde{\boldsymbol{\lambda}}$ are vectors of length $|\mathbf{q}| n_t$, whereas the size of $\tilde{\mathbf{A}}$ is $|\mathbf{q}| n_t \times |\mathbf{p}|$, where n_t is the number of time steps, $|\mathbf{q}|$ is the number of degrees of freedom, and $|\mathbf{p}|$ is the number of parameters. Furthermore, the sparsity of $\tilde{\mathbf{A}}$ (or lack thereof) depends on the time interval and spatial region where each parameter influences the simulation. For instance, homogeneous material parameters of deformable objects generally affect the entire simulation and lead to dense (parts of) $\tilde{\mathbf{A}}$, whereas control inputs for specific time steps influence only a small subset of the simulation and result in a much sparser structure. As a final note, it is worth mentioning that the block-triangular structure of $\tilde{\mathbf{B}}$ can also be leveraged to speed up the computation of the adjoint state. This block-matrix form emphasizes when (and where) data must be stored during the forward simulation, regardless of how the system is solved. While computing the adjoint state, information propagates backwards in time (a common feature of time-dependent adjoint formulations).

4 DIFFERENTIABLE FRICTIONAL CONTACT MODEL

With the differentiable simulation framework in place, our challenge now is to formulate a smooth frictional contact model that approaches, in the limit, the discontinuous nature of physical contacts. In this context, we always refer to the smoothness of the resulting trajectory, rather than the contact forces. Specifically, an impulsive force would result in a non-smooth trajectory, which we must avoid. As we show in this section, a soft-constraint approach

to modelling frictional contact fits seamlessly into the theory presented in §3, and enables an easy-to-tune trade-off between accuracy (e.g. solutions satisfying contact complementarity constraints and Coulomb's law of friction) and smoother, easier to optimize objective function landscapes.

Our work investigates different friction models from multiple points of view: accuracy of forward simulation results, smoothness of optimization landscapes, and suitability for gradient-based optimization. It is this mission statement that sets our work apart from previous efforts. Our extensive virtual and real-world experiments show that our simulation model can be used for various applications in graphics and robotics.

We begin by formulating the contact conditions and response forces in terms of a single contact point $\mathbf{x}(\mathbf{q})$. For deformable objects, we handle contacts on the nodes of the FEM mesh. For rigid bodies, we implement spherical or point-based collision proxies and then map the resulting world-space contact forces $\mathbf{f}$ into generalized coordinate space $\hat{\mathbf{f}}$, as described in §5.

For the normal component of the contact we must find a state of the dynamical system such that

$$g(\mathbf{x}) \geq 0, \quad (10)$$

where we refer to g as the *gap* function, measuring the distance from $\mathbf{x}$ to the closest obstacle. Assuming that g is a signed distance function (and negative if $\mathbf{x}$ lies inside of an obstacle), we define the outward unit normal $\mathbf{n} := (\partial g / \partial \mathbf{x})^\top$. Note that we use the convention that the derivative of a scalar function w.r.t. a vector argument is a row-vector throughout this paper, consequently $\mathbf{n}$ is a column-vector. Furthermore, we assume that g is available in closed form; we mostly use planar obstacles in our examples. We will also use the notation $\mathbf{N} := \mathbf{n}\mathbf{n}^\top$ for the matrix projecting to this normal. The force required to maintain non-negative gaps (non-penetration), denoted $\mathbf{f}_n$, must be oriented along $\mathbf{n}$, i.e. $\mathbf{f}_n = f_n \mathbf{n}$. The *normal force magnitude* f_n must be non-negative and can only be non-zero if there is a contact, consequently (Hertz-Signorini-Moreau condition, see also Eq. (2.10) in [Wriggers 2006]):

$$f_n \geq 0, \quad f_n g = 0 \quad (11)$$

In the tangential direction, we must determine a friction force $\mathbf{f}_t$ whose magnitude is constrained by the Coulomb limit

$$\|\mathbf{f}_t\| \leq c_f f_n, \quad (12)$$

where c_f is the *coefficient of friction*. This inequality distinguishes two regimes: *static* and *dynamic* friction (also referred to as *stick* and *slip* respectively). In the former case, the friction force magnitude is below the Coulomb limit and the tangential velocity vanishes (stick). In the latter case, following the principle of maximum dissipation [Stewart 2000], the friction force is oriented opposite the tangential velocity and Eq. (12) is satisfied as an equality (slip).

More formally we have either

$$\mathbf{T}\dot{\mathbf{x}} = 0 \text{ (stick), or} \quad (13)$$

$$\mathbf{f}_t = -c_f f_n \mathbf{t} \text{ (slip),} \quad (14)$$

where $\mathbf{t} := (\mathbf{T}\dot{\mathbf{x}}) / \|\mathbf{T}\dot{\mathbf{x}}\|$ and $\mathbf{T} := \mathbf{I} - \mathbf{N}$ projects to the tangent plane. Note that in either case $\mathbf{f}_t \cdot \mathbf{n} = 0$ must hold.

4.1 Sequential Quadratic Programming

Here, we briefly outline a hard constrained formulation of frictional contacts, which we employ for comparison. Introducing Lagrange multipliers (representing contact forces) for the contact constraints (10) and (13), along with the dynamic friction forces, into the residual (1) leads to a system of equations representing the Karush-Kuhn-Tucker conditions of an optimization problem:

$$\begin{aligned} \mathbf{r} + \mathbf{n}\lambda + \bar{\mathbf{T}}^\top \boldsymbol{\mu} &= 0, \\ g(\mathbf{x}) &\geq 0, \text{ and} \\ \bar{\mathbf{T}}\dot{\mathbf{x}} &= 0, \end{aligned} \quad (15)$$

where λ and $\boldsymbol{\mu}$ are the Lagrange multipliers for the normal and static friction constraints respectively. For a single contact point in static friction, $\boldsymbol{\mu}$ has two components and $\bar{\mathbf{T}}$ contains two orthogonal tangent vectors (corresponding to the two non-zero eigenvalues of $\mathbf{T}$). In the case of dynamic friction (sliding) we remove the constraint on the tangential velocity and replace the friction force $\bar{\mathbf{T}}^\top \boldsymbol{\mu}$ with $\mathbf{f}_t$ according to Eq. (14).

In principle, the KKT system (15) can be linearized and solved as a sequence of quadratic programs (SQP). In this way, each quadratic program takes the role of the linear solve in Newton's method. However, there are some limitations to this approach. In particular, linearizing the dynamic friction force resulting from Eq. (14) is not straightforward, as this force depends not only on the simulation state and velocity, but also on the normal force, in this case given by the Lagrange multiplier λ . Furthermore, the direction of the friction force should be opposing the current sliding velocity, which is ill-defined if this velocity approaches zero.

One way to address these issues is to assume that $\mathbf{f}_t$ is constant rather than linear, and compute its value based on the previous iteration. Doing so however reduces the convergence of the SQP iteration. In some cases the ill-conditioned behaviour of the unit vector $\mathbf{t}$ can even cause this approach to not converge at all. To help mitigate this problem, we first approximately solve the problem assuming sticking contacts, and then only update the friction force magnitude according to the Coulomb limit, but maintain its direction. Similarly, the approach of Tan et al. [2012] formulates the frictional contact problem as a quadratic program with linear complementarity constraints by linearizing the friction cone. Their solver also iteratively updates a set of active constraints, which effectively selects the direction of the sliding force. Finally, the SQP iteration is not as straightforward to stabilize by a line-search procedure as a Newton iteration would be. Macklin et al. [2019] report similar issues and present approaches to improve upon them.

Note that derivatives of the KKT conditions (15), once the solver has converged, could still be used to compute simulation state derivatives by direct differentiation. However, it is currently not clear whether this approach would immediately integrate with our *adjoint* formulation. We refer to [Amos and Kolter 2017] for further details on derivatives of QP solutions.

In the light of these limitations, we have implemented an SQP forward simulation approach only for the sake of comparisons on simple deformable examples, where stability and convergence is not an issue; see also Fig. 2. We show that our hybrid method, §4.3, yields visually indistinguishable results.

4.2 Penalty methods

Another approach to solve dynamics with contacts is to convert the constraints (10) and (13) to *soft* constraints, and introduce penalty forces if the constraints are violated. For the contact constraint, this means we need to add a force that is oriented along the outward normal and vanishes when $g > 0$. There are various types of functions, such as soft-max or truncated log-barriers, that we can use for this purpose. We find that even the simplest choice of a piece-wise *linear penalty function* is fast and sufficiently accurate for our applications:

$$\mathbf{f}_n = \mathbf{n} k_n \max(-g, 0), \quad (16)$$

where k_n is a penalty factor that must be chosen large enough to sufficiently enforce the normal constraint.

Note that k_n effectively controls the steepness of the resulting force and therefore the smoothness of the collision response. Of course the derivative of this force contains a discontinuity due to the max operator. While this could easily be replaced by a soft-max, our experiments indicate that doing so is not necessary. To find itself exactly at the kink in the contact force, a point on the multi-body system would have to be on the $g = 0$ manifold at the end of the time step due only to gravity, inertia, and internal force, which is extremely unlikely. The force Jacobian, which is needed for both forward simulation and to compute derivatives is well-defined everywhere else.

Similarly, we can formulate Coulomb friction as a clamped linear *tangential* penalty force with corresponding penalty factor k_t :

$$\mathbf{f}_t = -\mathbf{t} \min(k_t \|\mathbf{T}\dot{\mathbf{x}}\|, c_f f_n). \quad (17)$$

This expression reduces to (14) for dynamic friction, but results in the linear penalty force $k_t \mathbf{T}\dot{\mathbf{x}}$ for static friction.

Instead of solving the KKT conditions (15), we now only need to solve a formally unconstrained system that has the same form as Eq. (1), with additional penalty forces according to (16) and (17): $\mathbf{r} + \mathbf{f}_t + \mathbf{f}_n = 0$. From now on, we assume that these penalty forces (mapped to generalized coordinates in the case of rigid bodies) are part of the residual itself. Consequently, we can still use a standard Newton method to solve for the dynamic behaviour (after applying an adequate time discretization scheme). Stabilizing the solver with a line-search method that enforces decreasing residuals is sufficient to handle the non-smooth points of the penalty forces introduced by the max and min operators.

We can also smooth out the transition between stick penalty and dynamic friction force and replace $\mathbf{f}_t$ of Eq. (17) with

$$\mathbf{f}_t = -\mathbf{t} c_f f_n \tanh\left(k_t \|\mathbf{T}\dot{\mathbf{x}}\| / (c_f f_n)\right). \quad (18)$$

Using a hyperbolic tangent function introduces slightly softer friction constraints and can help improve the performance of optimization methods. Note that for tangential velocities close to zero $\tanh(v) \approx v$, which resolves any numerical issues with computing $\mathbf{t}$ for small velocities.

In our examples we always choose $k_t = k_n \Delta t$, because k_n penalizes a positional error, whereas k_t penalizes a velocity error. Consequently, scaling by a time step admits a similar positional error per time step for both components.

4.3 Hybrid method

Solving frictional contacts with hard constraints is both computationally expensive in terms of the forward simulation, as well as more challenging to differentiate, especially in an adjoint formulation. Typically, the normal contact constraint is satisfied exactly, but the friction model is simplified to enable more efficient simulation. Even then, the resulting complementarities are technically a combinatorial problem and heuristics are employed to choose the active constraints.

Penalty formulations, on the other hand, are formally unconstrained and therefore relatively straightforward to simulate and differentiate. While the choice of penalty stiffness does affect numerical conditioning, using implicit integration enables stable simulation for a wide range of stiffness. The main drawback of penalties is that one must allow some small constraint violations.

While a small violation of the normal constraint (10) is often visually imperceptible, softening the static friction constraint can introduce unacceptable artefacts in some situations. In particular, if a (heavy) object is supposed to *rest* on an inclined surface *under static friction* and the simulation runs for a sufficiently long time, the tangential slipping introduced by softening the stick constraint will inevitably become visually noticeable.

We address this problem by formulating a hybrid method using linear penalty forces, Eq. (16) and (17), combined with *equality* constraints for the static friction case. In order to take one time step in the forward simulation, we proceed as follows: first, we approximately solve the linear penalty problem to a residual of $\|\mathbf{r}\| < \epsilon^{1/2}$. Then, we apply hard constraints of the form (13) if the friction force according to (17) is below the Coulomb limit. Note that these *equality* constraints are analogous to standard Dirichlet boundary conditions. Consequently, we continue with the Newton iteration, including these constraints. However, if enforcing a hard constraint leads to a tangential force exceeding the Coulomb limit, we revert that contact point back to the clamped penalty force (17). We only allow this change if the intermediate state satisfies $\|\mathbf{r}\| < \epsilon^{1/2}$. In rare cases, we eventually fall back to the penalty formulation. These cases typically arise right at the transition point when a sliding object comes to rest; once at rest, the hard constraints keep it in place reliably. Even if we revert to penalties, we know that hard constraints would have violated the Coulomb limit in that time step, so allowing a small sliding velocity is acceptable at this point. Finally, we continue the solver iterations until convergence, i.e. $\|\mathbf{r}\| < \epsilon$. In this way, we enforce sticking with hard constraints whenever possible, and guarantee that the Coulomb limit is never violated.

Note that employing a penalty formulation for the *normal* forces means that both, the normal force and the Coulomb limit, depend *directly* on the simulation state (as opposed to representing normal forces as Lagrange multipliers). Consequently, we can easily compute the derivatives of these forces, maintain the quadratic convergence of Newton's method, and stabilize the solver with a line search.

Fig. 2. Dropping a cylinder onto an inclined plane. The resulting motion is generated using different contact and friction methods: linear penalty (green), tanh penalty (blue), hybrid (yellow), SQP (red). The final state of the simulation (a) after 500 time steps ($t = 2.5$ s) is shown on the top left, and the mean tangential velocity over time (b) is shown in the top right sub-figure. Solver convergence for time steps 60 (sliding phase, solid), and 150 (sticking phase, dashed) are shown in sub-figure (c).

Similarly, when computing sensitivities or adjoint objective function gradients, we can incorporate the additional equality constraints in the same way as Dirichlet boundary conditions. That is, the sensitivities or the adjoint state must fulfil analogous boundary conditions to the forward simulation. Specifically, for BDF1, the static friction constraint at time t_i becomes $\bar{\mathbf{T}}^i \mathbf{x}^i = \bar{\mathbf{T}}^i \mathbf{x}^{i-1}$ and consequently, we have $\bar{\mathbf{T}}^i \mathbf{s}^i = \bar{\mathbf{T}}^i \mathbf{s}^{i-1}$ for sensitivities of $\mathbf{x}$, or $\bar{\mathbf{T}}^i \lambda^{i-1} = \bar{\mathbf{T}}^i \lambda^i$ for the adjoint state (note that the earlier state is unknown in the latter case).

4.4 Summary and evaluation

In this section we have described various ways of handling contacts with Coulomb friction, focusing our discussion on a single contact point. Now we briefly summarize how these considerations integrate with our differentiable simulation framework. Recall that for deformable objects the nodes of the FEM mesh directly define the degrees of freedom. In this case, the contact handling extends naturally from a single point to each mesh node independently. For rigid bodies, on the other hand, we handle spheres or point sets defining the collision proxy. The resulting contact forces are then mapped from each contact point to the generalized force vector as described in §5.

In Fig. 2, we compare the different contact models described above through an experiment where a soft cylindrical puck is dropped onto an inclined plane. We choose a relatively low penalty factor of $k_n = 10^2$ for this example to highlight the differences between soft and hard constrained methods, especially in the static friction phase. The differences in the normal direction are imperceptible, even with a fairly low penalty factor. With soft friction constraints the cylinder slides slightly further (Fig. 2a) and does not come to a

complete stop (Fig. 2b) (though the tangential speed is on the order of 1 mm/s). Also of note is the fact that the SQP solver converges much slower during the sliding phase (Fig. 2c), while our hybrid method performs similarly to pure penalty methods in most cases, but delivers equivalent results to the SQP version under static friction. Conversely, tanh penalty forces are highly non-convex around the Coulomb limit, which results in slightly slower convergence during the sticking phase as compared to the linear friction forces. Overall, the linear penalty approach to friction forces converges faster than tanh penalties or hybrid contacts, while the SQP method takes about twice as long in terms of total CPU time; see also Table 1.

In summary, our validation tests show that penalty-based models of frictional contact approach the ground truth defined by complementary contact constraints and Coulomb's friction law, and they lead to better convergence rates for forward simulation than alternatives based exclusively on hard constraints. In this context, employing an implicit integration method in conjunction with a line-search routine maintains simulation stability at every time step even for very stiff penalties. Furthermore, our treatment of normal and friction forces makes sensitivity analysis (both the direct and adjoint variation) easy to apply to the simulated motion trajectories. We note that this would be much more challenging to achieve if we had to take derivatives of the general KKT conditions of the underlying linear complementarity problem.

To understand how the different contact models affect the types of inverse problems we aim to solve with our differentiable simulator, we perform another experiment. Here we exhaustively sample the objective function on the task of tossing a ball to a specific target location (see also Fig. 10a). In particular, we evaluate the objective function value and its gradients on a regular grid in input parameters v_x and v_z (i.e. the initial linear velocity in the forward and upward direction, respectively). Note that the ball is initially spinning with a non-zero angular velocity that is kept fixed for all tosses. Figure 3 illustrates these results. The two local minima correspond to one-bounce and two-bounce solutions to this control problem. As can be seen, contacts introduce noise into the objective function (visible as wiggly isolines and somewhat incoherent gradients); this is inevitable as contacts are inherently discontinuous events. Slight changes in the object's initial velocity can lead to a different order in which the nodal degrees of freedom impact the ground. When both the object and the ground are stiff, the noise in the objective function landscape caused by these discretization artefacts can be significant and lead to reduced performance of gradient-based optimization methods. Nevertheless, our experiment shows that the smoothness of the objective function landscape can be effectively controlled through the parameters used by the contact model. This is because smoother contact models enlarge the window of time over which contacts are resolved, and they avoid the use of large impulsive forces. Sensitivities with respect to the exact timing and order of collisions are therefore reduced. This observation, which is supported by the objective function landscapes visualized in Fig. 3, can be exploited to improve convergence rates for the inverse problems that leverage our differentiable simulator.

For the example in Fig. 4, we evaluate a simple continuation approach. This time around, the task is to throw a geometrically complex object (a bunny) such that it lands upright in a particular

Fig. 3. Objective function (shading and isolines) and gradients (arrows) corresponding to the control problem presented in Fig. 10a, using soft (a; $k_n = 100$) and stiff (b; $k_n = 10^3$) linear penalty contacts, compared to hybrid contacts (c; $k_n = 10^3$). Note that the overall nature of the objective landscape (i.e. the local minima it exhibits) remains the same, although these minima do shift slightly when softening the contacts. This behaviour promotes the use of continuation methods, whereby solutions obtained with a soft, smoother contact model are used as initial guess for simulations with stiffer, more realistic, parameters.

Fig. 4. Optimization convergence for throwing a bunny, similar to Fig. 11b but without the wall, such that it lands upright, close to a target pose. A continuation strategy applied to the stiffness of the contact model can drastically improve optimization results for such challenging control problems. Direct optimization uses $k_n = 10^3$; continuation uses k_n : 100 (yellow), 200 (purple), 400 (green), 800 (blue), and 10^3 (red).

location. When the optimization problem is solved using a stiff contact model, an unfavourable local minimum is quickly reached. In contrast, if the optimization problem starts out with a soft contact model which gets progressively stiffer over time, a much better solution is found.

Based on the experimental results described above, we conclude that the linear friction force model offers a favourable trade-off between simplicity, accuracy and practical performance, and as such it is our default choice for the results we present in this paper.

5 INTERNAL AND EXTERNAL FORCES IN GENERALIZED COORDINATES

In this section, we describe the models used to generate the forces acting on the multi-body systems simulated within our framework. We also present basic validation tests of our forward simulation.

5.1 Soft bodies

For deformable elastic objects, we employ a standard Neo-Hookean material model, given by homogeneous Lamé parameters (μ , λ) and constant mass density ρ . As is standard, this material model describes the energy density as a function of the deformation gradient

Fig. 5. Dropping an object composed of three tori. Selected frames from the animation (a), and overlay of the wall contact configuration (b) simulated with different contact methods: linear penalty (green), tanh penalty (blue), hybrid (yellow). Solver convergence (c) for the shown time steps with ground contact (solid), and wall contact (dashed).

F. Internal shape-restoring forces that arise in response to induced deformations are then computed as the negative gradient of the energy density integrated over each element with respect to the nodal degrees of freedom.

To model the behaviour of real-world objects, the elastic forces described above must be complemented by internal damping forces. Most viscosity models, such as the ones described in [Hahn et al. 2019], define the viscous stress (and consequently the damping force) based on the linear strain rate ($\dot{\mathbf{F}} + \dot{\mathbf{F}}^T$). One major drawback of these models is that they are not invariant to rotational motion, and therefore damp out the angular velocity of a deformable object during free flight. Brown et al. [2019], on the other hand, describe a family of rotation invariant viscosity models. Here, we employ a quadratic model, similar to their power-law damping, defining the viscous stress as a function of the Green strain rate $\mathbf{D}$. In particular, we define the viscous stress as

$$\sigma_v := \frac{\nu}{2} \frac{d(\text{tr}(\mathbf{D}^T \mathbf{D}))}{d\dot{\mathbf{F}}}, \quad \mathbf{D} := \frac{1}{2}(\dot{\mathbf{F}}^T \dot{\mathbf{F}} + \dot{\mathbf{F}} \dot{\mathbf{F}}^T), \quad (19)$$

where $\mathbf{F}$ is the deformation gradient, $\dot{\mathbf{F}}$ is the velocity gradient, and ν is the material's viscosity. We compute this derivative, as well as the corresponding damping matrix entries, using symbolic

Fig. 6. Total energy over time for bouncing rigid cubes (BDF2) with various ground contact damping coefficients (increasing left-to-right in the inset image).

differentiation. As this viscosity model is based on a quadratic strain rate, it behaves like a power-law model with flow index $h = 2$.

Figure 5 shows a deformable object composed of three tori in a drop test. The ground is inclined by 20° , while the friction coefficients of the ground and the wall are 0.4 and 0.8 respectively. In this example, we use BDF2 time integration. The rotation invariant viscosity model allows the object to rotate freely in the absence of contacts, but damps out elastic oscillations. Our three contact methods converge reliably to very low residuals; Fig. 5c shows convergence for two representative time steps during ground and wall contact respectively. Again, the linear penalty method is fastest, while tanh and hybrid contacts are closely matched. These tests further confirm our conclusion that penalty methods are sufficiently accurate when using implicit integration, which allows a high penalty stiffness.

5.2 Rigid bodies

Cartesian-space forces $\mathbf{f}$ and torques $\boldsymbol{\tau}$ applied to a rigid body project to generalized coordinates via the standard transformation

$$\hat{\mathbf{f}} = \begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{J}_\omega^\top \end{pmatrix} \begin{pmatrix} \mathbf{f} \\ [\dot{\mathbf{x}}]\mathbf{f} + \boldsymbol{\tau} \end{pmatrix}, \quad (20)$$

where the Jacobian $\mathbf{J}_\omega$ maps the rate of change of a rigid body's rotational degrees of freedom to changes in its world-space angular velocity. We use this expression, for example, to apply the contact and friction forces computed in §4 to rigid bodies that are in contact with the environment. Note that this operation demands the computation of the world coordinates of a contact point, $\mathbf{x}(\mathbf{q})$, as well as its time derivative $\dot{\mathbf{x}}(\mathbf{q}) = (\partial \mathbf{x} / \partial \mathbf{q}) \dot{\mathbf{q}}$. We parameterize rotations with exponential coordinates $\boldsymbol{\theta}$ and compute derivatives as in Gallego and Yezzi [2015]. Force Jacobians, which are needed for both forward simulation and sensitivity analysis, can be easily computed analytically by using the chain rule in conjunction with the derivatives presented in our supplements.

One important concept that must still be modelled is the restitution behaviour of rigid-body contacts. While post-impact velocities for deformable objects are governed by the material's elastic parameters and internal viscosity, for rigid bodies we must explicitly include a damping force in the normal direction in the event of a contact:

$$\mathbf{f}_d = -k_d \mathbf{N} \dot{\mathbf{x}} \quad \text{if } g(\mathbf{x}) \leq 0, \quad \mathbf{f}_d = 0 \text{ otherwise}, \quad (21)$$

where k_d is the damping coefficient. For our implicit soft contacts, this contact damping model replaces the common Moreau impact law used in explicit rigid-body engines to model restitution behaviour. In the absence of external forces, the contact phase for a single one-dimensional point mass x against a wall at $x = 0$ can be described as a damped harmonic oscillator $m\ddot{x} + k_d\dot{x} + k_n x = 0$. Analysing the exact solution for this oscillator, with initial conditions $x_0 = 0$ and $\dot{x}_0 = -v_{\text{in}}$, we find the following relation between the damping coefficient and the restitution ratio:

$$\frac{v_{\text{out}}}{v_{\text{in}}} = \exp\left(\frac{-\pi k_d}{\sqrt{4k_n m - k_d^2}}\right), \quad (22)$$

where v_{out} is the outgoing velocity measured after the first half-period of oscillation. Restitution occurs only below the critical damping factor, $k_d^2 < 4k_n m$. This relation can be useful to either determine the restitution coefficient having estimated k_d using our system, or to set k_d manually for a desired restitution behaviour.

Figure 6 shows a basic test case for our fully implicit rigid-body system using BDF2 time integration. Without additional damping, numerical damping is barely noticeable when time-stepping at $\Delta t = 1/60$ s. This corresponds to almost perfectly elastic collisions. We can effectively control the restitution via our linear contact damping model. Note that the symmetry of the contact is maintained over many bounces. In our video, we also show that rotating the cube slightly to the left quickly breaks this symmetry for comparison.

5.3 Multi-body systems

As we use implicit integration schemes for time-stepping, we employ (stiff) generalized springs to couple the individual constituents of a multi-body system to each other. This is a simple, general, and drift-free technique that can be shown to be closely related to Baumgarte-stabilized velocity-level constraints for rigid-body dynamics [Catto 2011]. In general, for non-dissipative coupling elements, we define a potential energy as a function of points or vectors anchored on the multi-body system. Taking the gradient of this potential energy with respect to the system's generalized coordinates directly outputs the resulting generalized forces. More formally, constraints $\mathbf{c}(\mathbf{q})$ are enforced through potentials of the form $E(\mathbf{q}) = (k_c/2) \mathbf{c}(\mathbf{q})^\top \mathbf{c}(\mathbf{q})$.

For instance, the constraint

$$\mathbf{c}_s := \|\mathbf{x}_1(\mathbf{q}) - \mathbf{x}_2(\mathbf{q})\| - l_0 = 0 \quad (23),$$

asks that a specific distance is maintained between two points on the multi-body system. Its resulting potential models a stiff linear spring of rest length l_0 . We use zero-length springs to formulate ball-and-socket joints. Furthermore, unilateral springs (of non-zero length), which do not produce a force under compression, model cables and elastic strings (similar to Bern et al. [2019]).

Hinge joints (i.e. 1-DOF revolute joints) connecting two rigid bodies are defined through attachment points ($\hat{\mathbf{x}}_1, \hat{\mathbf{x}}_2$) and rotation axes ($\hat{\mathbf{a}}_1, \hat{\mathbf{a}}_2$), specified in the local coordinate frame of each rigid body respectively. We model hinge joints with two constraints: a zero-length spring connecting the attachment points ($\hat{\mathbf{x}}_1, \hat{\mathbf{x}}_2$), and a second constraint that aligns the rotation axes ($\hat{\mathbf{a}}_1, \hat{\mathbf{a}}_2$):

$$\mathbf{c}_h := \mathbf{w}(\hat{\mathbf{a}}_2) - \mathbf{w}(\hat{\mathbf{a}}_1) = 0, \quad (24)$$

where $\mathbf{w}(\hat{\mathbf{a}})$ denotes the mapping from local to world coordinates.

We model active motors by extending hinge joints with a second set of local axes ($\hat{\mathbf{b}}_1, \hat{\mathbf{b}}_2$) that are orthogonal to the rotation axes $\hat{\mathbf{a}}_1$ and $\hat{\mathbf{a}}_2$ respectively. A motor constraint, $\mathbf{c}_m$, enforces a specific relative angle α between $\mathbf{w}(\hat{\mathbf{b}}_1)$ and $\mathbf{w}(\hat{\mathbf{b}}_2)$:

$$\mathbf{c}_m := \mathbf{w}(\hat{\mathbf{b}}_1) - \mathbf{R}(\alpha)\mathbf{w}(\hat{\mathbf{b}}_2). \quad (25)$$

To model position-controlled motors that are driven by typical Proportional-Derivative controllers, it is also important to add a damping component to the torques these motors generate. To this end we directly define a world-space torque as a function of end-of-time-step angular velocities:

$$\boldsymbol{\tau}_{md} := k_{md}(\boldsymbol{\omega}_1(\mathbf{q}, \dot{\mathbf{q}}) - \boldsymbol{\omega}_2(\mathbf{q}, \dot{\mathbf{q}})), \quad (26)$$

which we then project into generalized coordinates using Eq. 20.

6 RESULTS

We now turn our attention to *optimization problems* based on our dynamics system. In this section, we show the effectiveness of our differentiable simulation for gradient-based optimization compared to gradient-free alternatives on various examples. We present results on various applications: material parameter estimation including contacts, optimizing initial conditions, machine learning with our differentiable simulator directly built into the loss function, and trajectory optimization for robotics. Unless stated otherwise we use a penalty factor of $k_n = 10^3$ and convergence tolerance $\varepsilon = 10^{-10}$.

6.1 Material parameter estimation

Our system allows us to estimate material parameters such as stiffness and damping of deformable objects. We capture the real-world behaviour of our specimens using either an optical motion-capture system, or a *Kinect* v2 depth camera.

In the former case, we track up to six labelled optical markers on the specimen at a frame rate of 120 Hz using an array of 10 *OptiTrack Prime 13* cameras. The system calibration ensures that the world-space coordinates align with the ground and wall planes in order to include these rigid obstacles in the simulation. For motion capture data, the objective function measures the sum of squared distances between the tracked marker position and the corresponding location on the simulated mesh for all time steps. In each time step, we only consider markers that are currently visible to the tracking system.

In the latter case, we read 3D point clouds from the *Kinect* at 30 Hz and then identify the ground and wall planes in a manual post-processing step. We apply box filters in both world and colour space to identify which points correspond to the surface of the specimen. As we do not have a direct correspondence between tracked points and mesh locations in this case, the objective function instead measures the distances of all filtered points to their respective closest point on the surface of the simulated mesh.

For parameter estimation, we employ a continuation strategy in time, first optimizing the initial conditions of the simulation (in terms of position, orientation, velocity, and spin) to match the recorded motion during the ballistic phase (before the first contact). In the second phase, we keep the initial conditions fixed, and optimize the material parameters such as to best approximate the first bounce of the recorded motion. Finally, we add a third phase

Fig. 7. Parameter estimation for throwing a sphere. The 3D image shows motion-capture trajectories for all six markers (two along each coordinate axis) and snapshots of the best fitting simulation. The inset graph shows captured and simulated trajectories for the front facing marker.

Fig. 8. Parameter estimation for throwing a cube. The 3D image shows motion-capture trajectories for all six markers (one in the centre of each face) and snapshots of the best fitting simulation. The inset graph shows captured and simulated trajectories for the front facing marker.

Fig. 9. Parameter estimation with *Kinect* data. Image shows input point cloud (time colour-coded purple to yellow) and representative time steps of the simulation result (green). Graph shows average coordinates; axes represent the camera orientation (z forward, y up).

where all parameters (material and initial conditions) are optimized simultaneously for the entire motion. In our experiments, we find that L-BFGS is well suited for these optimization tasks.

Here, we show three results for material parameter estimation of real-world specimens. We prepare two custom-made elastic foam

specimens, a sphere and a cube, for motion capture with six slightly inset motion capture markers each, see Fig. 7 and 8. The motion capture data provides direct correspondences between the tracked markers and their simulated counterparts, allowing us to find the initial orientation and angular velocity in the first phase. In the first example, Fig. 7, we then optimize material parameters for the duration of the first bounce, and finally all parameters over the entire recorded trajectory. The second example, Fig. 8, uses a more automated approach, optimizing all parameters over increasing time horizons. Apart from the material parameter optimization in the first example, which includes a short ADAM phase, we use L-BFGS for all these optimizations. In our accompanying video we also show verification tests for both of these results, where we use the material parameters obtained via these optimizations, and then only fit the initial conditions to the ballistic phase of a different recorded motion. Please also refer to Table 1 for details on material parameters and runtime.

Finally, we show an example for a foam ball without additional markers, where we record the real-world motion with a *Kinect* depth camera, Fig. 9. In this case, we minimize the distance from the simulated surface to the recorded point cloud, which means that we do not have any rotational information about the real-world specimen. Nevertheless, by allowing the optimization to change initial conditions during later stages where contacts are taken into account, we find a good match between real and simulated motion. We employ the same approach to estimate parameters of a common tennis ball, which we subsequently throw with a robot as discussed in the next section.

6.2 Throwing

We can parameterize, and optimize for, the initial conditions of our simulation, such as in the examples shown in Fig. 10. In these cases, we must account for the contribution of the initial conditions to the objective function gradient in Eq. (9). While previous work provides an adjoint formulation for general, implicitly defined, initial conditions [Bradley 2013], when directly parameterizing initial conditions we find it more convenient to calculate the corresponding derivatives explicitly.

Parameters that define initial conditions, $\mathbf{p}_0$, are parameter variables that affect only the initialization of the time-integration scheme, but do *not* directly affect any of the *unknown* states $\tilde{\mathbf{q}}$. We can therefore compute the derivatives of the residuals w.r.t. these parameters analytically: $\partial \mathbf{r}^i / \partial \mathbf{p}_0 = (\partial \mathbf{r}^i / \partial \chi)(\partial \chi / \partial \mathbf{p}_0)$, where χ refers to the initial state of the time integrator. The first term follows directly from the choice of time-integration method, while the second term follows from the parameterization of initial conditions. Finally, these derivatives are added to the matrix $\tilde{\mathbf{A}}$ of the sensitivity system, where each block now becomes $\mathbf{A}^i := \partial \mathbf{r}^i / \partial \mathbf{p} + \partial \mathbf{r}^i / \partial \mathbf{p}_0$. Note that only the first few time steps receive a non-zero update, depending on the chosen time-discretization scheme.

In our tests, we optimize initial linear and angular velocities for throwing a deformable ball, Fig. 10. The objective function measures the distance from the ball's centre of mass to a specific target point at the end of the simulation (Fig. 10a), or to a target line over a specified time range (Fig. 10b), respectively.

In order to compare our results to a gradient-free sampling method, we run CMA-ES on the optimization problem in Fig. 10a.

Fig. 10. Throwing a deformable ball: a point target for the ball's centre of mass (a) admits multiple exact solutions with either zero, one, or two bounces off the floor. Asking the second half of the c.o.m. trajectory (red) to be as close to a vertical line as possible (b) requires a trade-off between forward motion and back-spin (black arrow) such that friction slows the ball down when it bounces off the ground.

Fig. 11. Throwing with multiple contacts: (a) multiple paths for the ball's c.o.m. to reach the target point, labelled by contact sequence ('w' wall, 'f' floor); (b) throwing the bunny to a specific target pose (wireframe) after bouncing off the floor and the wall (time colour-coded from dark to bright).

Qualitatively, gradient-based approaches are less likely to traverse a saddle point, whereas sampling methods explore the parameter space more randomly in the early stages before settling into a local minimum. In this particular case, using L-BFGS and linear penalty contacts requires 28 simulation runs to find an exact solution for the two-bounce motion (relative objective function value $\Phi/\Phi_0 < 10^{-25}$) in a matter of minutes (Table 1), while CMA-ES requires 1042 simulations and 2h 15m to find an approximate solution with $\Phi/\Phi_0 \approx 10^{-6}$.

On a more complex example (Fig. 11b), L-BFGS finds a better solution ($\Phi/\Phi_0 \approx 4 \cdot 10^{-4}$) after 319 simulations, whereas CMA-ES returns a noticeably worse result ($\Phi/\Phi_0 \approx 10^{-2}$) even after running over 8000 simulations. We observe the same behaviour for a trajectory optimization test, similar to Fig. 14, where CMA-ES fails to produce an acceptable solution after multiple hours, whereas our system yields a good result in a few minutes using direct sensitivity analysis and Gauss-Newton optimization; see our supplements for details.

We then show results for artistic control of animations. In particular, we find optimal throwing velocities such that an elastic object hits a specified target after multiple bounces. We first extend the example of Fig. 10a by including a wall and increasing the distance to the target, Fig. 11a. Depending on the initial conditions, we can now find multiple paths to the target with various bounce patterns, as labelled in the image. In these examples, using smoother tanh friction forces (blue) yields slightly better results than linear ones (green).

We can also throw the Stanford bunny (again including a contact with a wall) such that it lands at a specific target location (Fig. 11b). In this example, the objective function measures the squared distance

Fig. 12. Integrating differentiable simulation with a neural network.

to the target pose for each mesh node, and also includes a regularizer that additionally penalizes solutions where the bunny falls over.

After performing parameter estimation for a tennis ball, as described in the previous section, we also optimize initial conditions for a new throw such that the tennis ball hits a specific location on the wall after bouncing off a table once. For the resulting initial position and velocity, we then generate a throwing motion (shown in our video) for a UR5 robot using a standard inverse kinematics model.

6.3 Self-supervised learning of control policies

Learning-based methods that leverage neural networks and simulation data to train control policies have achieved impressive results for various control applications. Forward simulation is commonly used as an infinite data source that is sampled over initial conditions and control parameters to generate training data. While this data-driven approach effectively decouples learning from simulation and thus simplifies implementation, it critically relies on the parameter space sampling reflected in the training data to yield an appropriate coverage in performance space. We pursue an alternative strategy that, rather than using a fixed training set, integrates simulation directly into the loss function, as illustrated in Fig. 12, thus enabling the learning algorithm to exploit the map between parameter and performance space provided by our differentiable simulator. Note that this is not our core contribution as similar approaches have been explored in related work, nevertheless, we demonstrate that our framework is well-suited for this important application. As in the other applications, the objective (or *loss*) function measures the simulation result $\tilde{\mathbf{q}}$ against a desired target behaviour $\mathbf{q}^*$, i.e. $\Phi(\mathbf{q}^*, \tilde{\mathbf{q}})$. We then train the neural network $\phi(\mathbf{q}^*, \mathbf{w})$ to return simulation parameters $\mathbf{p}$ that achieve the given target. The result is a weight vector $\mathbf{w}$ for the network that minimizes the training loss:

$$\min_{\mathbf{w}} \frac{1}{n} \sum_{l=1}^n \Phi(\mathbf{q}_l^*, \tilde{\mathbf{q}}(\phi(\mathbf{q}_l^*, \mathbf{w}))). \quad (27)$$

Note that the simulated trajectory is a function of the parameters returned by the neural network, i.e. $\tilde{\mathbf{q}}(\phi(\mathbf{q}^*, \mathbf{w}))$. Consequently, a *differentiable* simulation is key to computing gradients during training while avoiding costly finite differencing. We demonstrate this approach on a simple game, where we train a neural network to find the throwing velocity for a rigid ball such that, after a single impact with the ground, it hits a given target position. The objective (or *loss*) function measures proximity to the target location using a soft minimum over the descending part of the trajectory, and includes a penalty term that discourages solutions without contact. Rather than measuring distance at a specific time, this approach provides

Fig. 13. Convergence of the learning process for the throw controller (left) and the corresponding network architecture (right). The neural network outputs the initial velocity of the simulation such as to hit the given target after one bounce. (ELU: exponential linear unit.)

more flexibility in terms of timing, allowing the learning algorithm to find better solutions.

We select $n = 1000$ target positions for training and 100 for testing, uniformly sampled in a rectangular region. We train using ADAM [Kingma and Ba 2014] with $\beta_1 = 0.95$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, and a mini-batch size of 5. We start with a learning rate of 10^{-2} , which is reduced by a factor of 0.5 after each epoch. On average, each epoch takes about 450 s of CPU time to compute. The architecture of the network is shown in Fig. 13. It is worth noting that, even in this comparatively simple example, accounting for friction in the simulation is crucial for accurate control; a controller trained in a friction-less environment will systematically fail to hit the target.

6.4 Trajectory optimization

We present various applications to robot control using trajectory optimization. We optimize for per-time-step control parameters $\mathbf{p}^i$ representing either the position and orientation of a robotic end-effector, or target motor angles. Note that each subset of parameters affects only one time step, while the entire parameter vector $\mathbf{p}$ remains formally time-independent.

We manually define target trajectories for specific feature points. The objective function Φ again measures squared distances between the simulated and target trajectories of these features. To ensure temporal smoothness of parameters, we add the regularization term $\beta \sum_i \|\mathbf{p}^{i+1} - \mathbf{p}^i\|^2$ to Φ .

We first optimize the trajectory for a 6-DOF robotic end-effector over a time interval of 1 s, and test our results on a real-world *Universal Robotics UR5* robot running at a controller frequency of 125 Hz, Fig. 14. This test shows that we can effectively optimize for a large number of parameters, as each time step ($\Delta_t = 1/60$ s) has its own set of end-effector coordinates. We also demonstrate that our simulated results carry over to the real world by having the robot perform the resulting motion repeatedly for comparison (see our video).

Another example using our differentiable simulator in a trajectory optimization setting is the manipulation of a sheet modelled as a mass-spring system, Fig. 1. The control parameters specify the positions of two handles, each attached to one corner of the sheet. Initially, the sheet lies on the floor facing upwards. The objective function asks for the following motion of this sheet: after 1 s, it

Fig. 14. Optimizing end-effector trajectories for robotic control of coupled dynamic systems. Our robot drags a rigid cube attached with nylon strings over a distance of 5 cm (a) or 11 cm (b) respectively, and actuates a coupled system composed of two rigid cubes and four elastic rubber bands (c) such that the lower cube is tipped over.

should be flipped, facing downwards, and later, at $t = 2$ s it should be flipped back to its original orientation, as well as translated to the right. Adding a smoothness regularizer for the control parameters, the trajectory optimization finds a solution after a few Gauss-Newton iterations, as can be seen in the accompanying video.

In our most complex example, we optimize for control inputs of a legged robot actuated by compliant motors. For this experiment, we model the robot's actuators as PD controllers with relatively low gains, a reasonable model of position-controlled motors, implemented as soft, damped, angular constraints between the coordinate frames of adjacent rigid links. We use the motion synthesis tool of Geilinger et al. [2018] to create a nominal motion trajectory for this robot. However, their trajectory optimization model is based on a relatively coarse approximation of the robot's dynamics and assumes precise and strong actuators. Unsurprisingly, due to these modelling simplifications, when our compliant robot attempts to track the nominal trajectory, it fails to locomote effectively.

We note that this use-case is motivated by real-world challenges. Compliance, whether parasitic (e.g. motors that are not infinitely strong) or purposely built in (e.g. rubber feet designed to soften impacts), is a defining characteristic of physical robots. Most existing motion-planning and trajectory-optimization algorithms, however, are not able to account for it. Our differentiable simulator on the other hand allows us to optimize the robot's motion by directly considering its full body dynamics, including the compliant nature of its actuators and feet. The result, as shown in Fig. 15 and the accompanying video, is a successful locomotion gait for this robotic creature with compliant actuators and soft feet.

Again, the nominal motion trajectory synthesized with an idealized robot model does not carry over to the compliant robot with soft feet, as shown in the accompanying video. Using our differentiable simulator, we apply trajectory optimization to the full coupled multi-body dynamics and attain optimized controls such that the robot reaches the target distance travelled. We refer to Table 2 for details on simulation parameters and runtime.

7 DISCUSSION

We present an analytically differentiable dynamics solver that handles frictional contact for soft and rigid bodies. A key aspect of

Fig. 15. Coupling soft and rigid bodies allows us to equip this compliant robot with soft feet. We can perform trajectory optimization on the entire robot to find control inputs that account for the compliant motors, as well as the deformable end-effectors.

our approach is the use of a soft constraint formulation of frictional contact, which makes our simulation model straightforward to differentiate. Our results show that penalty-based contact models, especially in the normal component, are sufficiently accurate when combined with implicit time integration, and also enable tuneable, sufficiently smooth contact treatment for gradient-based optimization. We also analyse the effects of penalties against hard constraints with respect to static and dynamic friction. For dynamic motion, where static friction persists only for a short contact duration, penalty-based methods perform adequately and improve the performance of optimization methods built upon these simulations. When persistent static friction is necessary, our hybrid method adds the corresponding equality constraints, but maintains soft contacts (rather than adding inequality constraints) in the normal component, and therefore still fits into our differentiable simulation framework.

Our optimization examples show that using this framework, gradient-based optimization methods greatly outperform sampling-based methods such as CMA-ES. We demonstrate the effectiveness of our approach on a wide range of applications, including parameter estimation, robotic locomotion and manipulation tasks, as well as learning-based control.

All of our examples assume that contact forces obey the isotropic Coulomb model. While this assumption is valid for a large class of surfaces, anisotropy is an important characteristic of various physical systems such as textiles [Pabst et al. 2009]. Furthermore, in our modelling, we assume a functional representation of the distance metric between different bodies to be readily available, and to be sufficiently smooth. For complex contact scenarios where the geometric representations of the bodies involved in collisions are highly detailed or non-convex, more elaborate collision handling methods are required [Allard et al. 2010].

Our experiment on learning-based control is an initial investigation into combining our differentiable simulator with machine-learning techniques. In the future, we see great promise in leveraging this concept in the context of deep reinforcement learning. By eliminating the need for random exploration, for example, the analytic gradient information that our framework provides is likely to improve sample efficiency.

Similarly, our experiments on performing throwing motions on a real-world robot demonstrate that our simulation results translate to the physical specimens. However, there are still numerous sources of error such as aligning the robot in its environment,

Table 1. Overview of material parameters, simulation and optimization performance on our FEM examples. Timings obtained on a 4×3.5 GHz CPU with 16 GB RAM. Columns: Young's modulus E , viscosity ν (* indicates BDF1 time integration), mass density ρ , coefficient of friction c_f (values in braces refer to walls), number of elements in the mesh m , runtime of a single simulation t_{sim} , runtime of the entire optimization t_{opt} (in sec. or hh:mm:ss), number of simulation runs during optimization i_{opt} . For parameter estimation examples (Figs. 7 and 8, the first row gives the initial guess material parameters (in italics) with timings for optimization of initial conditions (braces) and material parameters. The second row gives final material parameters with timings for optimization of initial conditions on a different recorded trajectory.

Figure	E [Pa]	ν [Pa s]	ρ [kg/m ³]	c_f [1]	m	t_{sim} [s]	t_{opt}	i_{opt}
2 lin.	2.1E+03	0 *	150	0.4	879	26.24		
2 tanh	2.1E+03	0 *	150	0.4	879	30.86		
2 hyb.	2.1E+03	0 *	150	0.4	879	31.2		
2 SQP	2.1E+03	0 *	150	0.4	879	66.78		
5 lin.	2.1E+04	0.1	150	0.4 (0.8)	1930	119.19		
5 tanh	2.1E+04	0.1	150	0.4 (0.8)	1930	178.5		
5 hyb.	2.1E+04	0.1	150	0.4 (0.8)	1930	180.75		
10a (0)	2.0E+04	0.0125	150	0.4	451	6.92	27.33	4
10a (1)	2.0E+04	0.0125	150	0.4	451	8.65	222.54	24
10a (2)	2.0E+04	0.0125	150	0.4	451	8.62	265.21	28
10b	2.0E+04	0.0125	150	0.4	451	8.21	01:01:11	453
11b	2.4E+04	0 *	90	0.2 (0.4)	2729	71.28	05:59:32	322
7 (initial mat. param.)	2.1E+05	0.025	86	0.4	1351	(4.13) 38.35	05:08:42	(104) 586
Second motion (video)	3.3E+04	6.225	86	0.47	1351	13.38	158.04	105
8 (initial mat. param.)	2.0E+04	0.525	77	0.4	1501	(3.08) 421.70	06:58:22.30	(150) 350
Second motion (video)	1.5E+04	0.7183	77	1.16	1501	1.89	430.35	112

Table 2. Overview of simulation parameters and performance on our trajectory optimization examples. Columns: number of degrees of freedom $|\mathbf{q}|$, number of parameters per time step $|\mathbf{p}|/n_t$, number of time steps n_t , time step size Δt , contact stiffness k_n , contact damping k_d , stiffness of constraints k with stiffness of compliant motors (braces), motor damping k_{md} , coefficient of friction c_f , runtime for a single simulation t_{sim} , runtime for the entire optimization t_{opt} , number of optimization iterations i_{opt} .

Example	Figure	$ \mathbf{q} $	$ \mathbf{p} /n_t$	n_t	Δt [s]	k_n	k_d	k	k_{md}	c_f [1]	β	t_{sim} [s]	t_{opt}	i_{opt}
Cube dragging	video	12	6	60	1/60	200	1e-3	5	n.a.	0.5	1e-2	0.044	00:03.3	20
Mass-spring flipping	1	75	6	120	1/60	1e3	1e-3	1000	n.a.	0.5	1	0.109	00:54	25
Compliant robot	1	78	12	192	1/60	1e5	0.1	5e5 (1e5)	10	0.8	1e-3	0.550	19:21	60
Compliant robot, soft feet	15	414	12	144	1/60	1e5	1e-3	5e5 (1e5)	10	0.8	1e-3	1.552	27:54	17

unwanted movements of the robot's base, and latency of the hardware controller, which require further investigation.

In summary, our experiments demonstrate that our system enables efficient inverse problem solving for various applications in graphics and robotics. For many of these applications, soft constraints with linear penalty forces, combined with implicit integration, lead to physically meaningful and analytically differentiable simulations. Furthermore, we explore options for smoother friction forces, which help reduce noise in the objective function, as well as equality constraints for static friction in cases where physical accuracy is key. In either situation, a soft contact in the normal component enables differentiability and gradient-based optimization. In the future, we plan to further investigate simulation-driven optimization methods in the context of robotics and control of highly complex multi-body systems that combine rigid and flexible elements.

ACKNOWLEDGMENTS

This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 866480), the ETH Zürich Postdoctoral Fellowship 18-1 FEL-09, and the Personalized Health and Related Technologies (PHRT) SwissHeart grant. We thank our colleagues Pol Banzet, Dongho Kang, and Simon Dünser for their help with various experiments.

REFERENCES

- Jérémie Allard, François Faure, Hadrien Courtecuisse, Florent Falipou, Christian Duriez, and Paul G. Kry. 2010. Volume Contact Constraints at Arbitrary Resolution. *ACM Trans. Graph.* 29, 4 (2010).
- Brandon Amos and J. Zico Kolter. 2017. OptNet: Differentiable Optimization as a Layer in Neural Networks. (2017), 136–145.
- M. Anitescu and F. A. Potra. 1997. Formulating Dynamic Multi-rigid-body Contact Problems with Friction as Solvable Linear Complementarity Problems. *NONLINEAR DYNAMICS* 14 (1997), 231–247.

- David Baraff. 1994. Fast Contact Force Computation for Nonpenetrating Rigid Bodies. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '94)*. 23–34. <https://doi.org/10.1145/192161.192168>
- James Bern, Pol Banzet, Roi Poranne, and Stelian Coros. 2019. Trajectory optimization for cable-driven soft robot locomotion. *Robotics: Science and Systems* (2019).
- Florence Bertails-Descoubes, Florent Cadoux, Gilles Daviet, and Vincent Acary. 2011. A Nonsmooth Newton Solver for Capturing Exact Coulomb Friction in Fiber Assemblies. *ACM Trans. Graph.* 30, 1 (2011).
- Andrew Bradley. 2013. *PDE-constrained optimization and the adjoint method*. Technical Report. Stanford University. https://cs.stanford.edu/~ambrad/adjoint_tutorial.pdf
- Bernard Brogliato. 1999. *Nonsmooth Mechanics: Models, Dynamics and Control*. Springer.
- George E. Brown, Matthew Overby, Zahra Forootaninia, and Rahul Narain. 2019. Accurate dissipative forces in optimization integrators. *ACM Transactions on Graphics* 37, 6 (2019). <https://doi.org/10.1145/3272127.3275011>
- Erin Catto. 2011. Reinventing the spring. https://box2d.org/files/ErinCatto_SoftConstraints_GDC2011.pdf
- Desai Chen, David I. W. Levin, Wojciech Matusik, and Danny M. Kaufman. 2017b. Dynamics-aware numerical coarsening for fabrication design. *ACM Transactions on Graphics* 36, 4 (2017), 1–15. <https://doi.org/10.1145/3072959.3073669>
- Xiaozhong Chen, Sheldon Andrews, Derek Nowrouzezahrai, and Paul Kry. 2017a. Ballistic Shadow Art. *Proceedings of Graphics Interface 2017* Edmonton (2017), 190–198. <https://doi.org/10.20380/GI2017.24>
- Stelian Coros, Sebastian Martin, Bernhard Thomaszewski, Christian Schumacher, Robert Sumner, and Markus Gross. 2012. Deformable Objects Alive! *ACM Trans. Graph.* 31, 4 (2012).
- Filipe de Avila Belbute-Peres, Kevin Smith, Kelsey Allen, Josh Tenenbaum, and J. Zico Kolter. 2018. End-to-End Differentiable Physics for Learning and Control. In *Advances in Neural Information Processing Systems 31*. 7178–7189.
- Jonas Degraeve, Michiel Hermans, Joni Dambre, and Francis Wyffels. 2019. A Differentiable Physics Engine for Deep Learning in Robotics. *Frontiers in Neurorobotics* 13 (2019), 6.
- Ounan Ding and Craig Schroeder. 2020. Penalty Force for Coupling Materials with Coulomb Friction. *IEEE Transactions on Visualization and Computer Graphics* 26, 7 (2020), 2443–2455. <https://doi.org/10.1109/tvcg.2019.2891591>
- Christian Duriez, Frederic Dubois, Abderrahmane Kheddar, and Claude Andriot. 2006. Realistic Haptic Rendering of Interacting Deformable Objects in Virtual Environments. *IEEE Transactions on Visualization and Computer Graphics* 12, 1 (2006).
- T. Erez and E. Todorov. 2012. Trajectory optimization for domains with contacts using inverse dynamics. In *2012 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*.
- Kenny Erleben. 2007. Velocity-Based Shock Propagation for Multibody Dynamics Animation. *ACM Trans. Graph.* 26, 2 (2007), 12.
- K. Erleben, M. Macklin, S. Andrews, and P. G. Kry. 2019. The Matchstick Model for Anisotropic Friction Cones. *Computer Graphics Forum* (2019).
- Guillermo Gallego and Anthony Yezzi. 2015. A compact formula for the derivative of a 3-D rotation in exponential coordinates. *Journal of Mathematical Imaging and Vision* 51, 3 (2015), 378–384.
- Moritz Geilinger, Roi Poranne, Ruta Desai, Bernhard Thomaszewski, and Stelian Coros. 2018. Skaterbots: optimization-based design and motion synthesis for robotic creatures with legs and wheels. *ACM Transactions on Graphics* 37, 4 (2018), 1–12.
- David Hahn, Pol Banzet, James M. Bern, and Stelian Coros. 2019. Real2Sim: visco-elastic parameter estimation from dynamic motion. *ACM Transactions on Graphics* 38, 6 (2019), 1–13. <https://doi.org/10.1145/3355089.3356548>
- David Harmon, Etienne Vouga, Breannan Smith, Rasmus Tamstorf, and Eitan Grinspun. 2009. Asynchronous Contact Mechanics. In *ACM SIGGRAPH 2009 Papers (SIGGRAPH '09)*. Article 87. <https://doi.org/10.1145/1576246.1531393>
- Philipp Holl, Nils Thuerey, and Vladlen Koltun. 2020. Learning to Control PDEs with Differentiable Physics. In *Int. Conf. on Learning Representations*. (to appear).
- Shayan Hoshayari, Hongyi Xu, Espen Knoop, Stelian Coros, and Moritz Bächer. 2019. Vibration-Minimizing Motion Retargeting for Robotic Characters. *ACM Transactions on Graphics* 38, 4 (2019).
- Yuanming Hu, Luke Anderson, Tzu-Mao Li, Qi Sun, Nathan Carr, Jonathan Ragan-Kelley, and Frédo Durand. 2019a. DiffTaichi: Differentiable Programming for Physical Simulation. *International Conference on Learning Representations (ICLR)* (2019).
- Yuanming Hu, Jiancheng Liu, Andrew Spielberg, Joshua B Tenenbaum, William T Freeman, Jiajun Wu, Daniela Rus, and Wojciech Matusik. 2019b. ChainQueen: A Real-Time differentiable Physical Simulator for Soft Robotics. *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)* (2019).
- Sumit Jain and C. Karen Liu. 2011. Controlling Physics-Based Characters Using Soft Contacts. *ACM Transactions on Graphics* 30, 6 (2011).
- Danny M. Kaufman, Timothy Edmunds, and Dinesh K. Pai. 2005. Fast Frictional Dynamics for Rigid Bodies. *ACM Transactions on Graphics* 24, 3 (2005), 946–956.
- Danny M. Kaufman, Shinjiro Sueda, Doug L. James, and Dinesh K. Pai. 2008. Staggered Projections for Frictional Contact in Multibody Systems. *ACM Transactions on Graphics (SIGGRAPH Asia 2008)* 27, 5 (2008), 164:1–164:11.
- Diederik P. Kingma and Jimmy Ba. 2014. ADAM: A Method for Stochastic Optimization. (2014). [arXiv:cs.LG/http://arxiv.org/abs/1412.6980v9](https://arxiv.org/abs/1412.6980v9)
- Jie Li, Gilles Daviet, Rahul Narain, Florence Bertails-Descoubes, Matthew Overby, George E. Brown, and Laurence Boissieux. 2018. An Implicit Frictional Contact Solver for Adaptive Cloth Simulation. *ACM Trans. Graph.* 37 (2018), 15.
- Junbang Liang, Ming Lin, and Vladlen Koltun. 2019. Differentiable Cloth Simulation for Inverse Problems. In *Advances in Neural Information Processing Systems* 32. 772–781.
- C. Karen Liu and Sumit Jain. 2012. A quick tutorial on multibody dynamics. http://fab.cba.mit.edu/classes/865.18/design/optimization/dynamics_1.pdf
- Guillaume Loubet, Nicolas Holzschuch, and Wenzel Jakob. 2019. Reparameterizing Discontinuous Integrands for Differentiable Rendering. *ACM Trans. Graph.* 38, 6 (2019).
- Mickaël Ly, Romain Casati, Florence Bertails-Descoubes, Mélina Skouras, and Laurence Boissieux. 2018. Inverse Elastic Shell Design with Contact and Friction. *ACM Trans. Graph.* 37, 6 (2018).
- Miles Macklin, Kenny Erleben, Matthias Müller, Nuttapon Chentanez, Stefan Jeschke, and Viktor Makovychuk. 2019. Non-smooth Newton Methods for Deformable Multi-body Dynamics. *ACM Transactions on Graphics* 38, 5 (2019).
- Aron Monszpart, Nils Thuerey, and Niloy J. Mitra. 2016. SMASH: Physics-Guided Reconstruction of Collisions from Videos. *ACM Trans. Graph.* 35, 6 (2016).
- Jorge Nocedal. 1980. Updating quasi-Newton matrices with limited storage. *Math. Comp.* 35, 151 (1980), 773–773. <https://doi.org/10.1090/s0025-5718-1980-0572855-7>
- Keith O'Hara. 2019. OptimLib. <https://www.kthohr.com/optimlib.html>. 2019-11-28.
- Miguel A. Otaduy, Rasmus Tamstorf, Denis Steinemann, and Markus Gross. 2009. Implicit Contact Handling for Deformable Objects. *Computer Graphics Forum* 28, 2.
- Simon Pabst, Bernhard Thomaszewski, and Wolfgang Straundföner. 2009. Anisotropic Friction for Deformable Surfaces and Solids. In *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 149–154.
- Dinesh K. Pai, Kees van den Doel, Doug L. James, Jochen Lang, John E. Lloyd, Joshua L. Richmond, and Som H. Yau. 2001. Scanning physical interaction behavior of 3D objects. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques - SIGGRAPH '01*.
- Zherong Pan and Dinesh Manocha. 2018. Position-Based Time-Integrator for Frictional Articulated Body Dynamics. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. <https://doi.org/10.1109/iro.2018.8593817>
- Mark Pauly, Dinesh K. Pai, and Leonidas J. Guibas. 2004. Quasi-Rigid Objects in Contact. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics Symposium on Computer Animation (SCA '04)*. Eurographics Association, 109–119.
- Jovan Popović, Steven M. Seitz, Michael Erdmann, Zoran Popović, and Andrew Witkin. 2000. Interactive manipulation of rigid body simulations. In *Proceedings of the 2000 annual conference on Computer graphics and interactive techniques - SIGGRAPH '00*.
- Yixuan Qiu. 2019. LBFGS++. <http://yixuan.cos.name/LBFGSpp>. 2019-03-12.
- Connor Schenck and Dieter Fox. 2018. SPNets: Differentiable Fluid Dynamics for Deep Neural Networks. In *Proceedings of The 2nd Conference on Robot Learning (Proceedings of Machine Learning Research)*, Vol. 87. PMLR, 317–335.
- Tamar Shinar, Craig A. Schroeder, and Ronald Fedkiw. 2008. Two-way Coupling of Rigid and Deformable Bodies. In *Proceedings of the 2008 Eurographics/ACM SIGGRAPH Symposium on Computer Animation (SCA '08)*. 95–103.
- David E. Stewart. 2000. Rigid Body Dynamics with Friction and Impact. *SIAM Rev.* (2000), 3–39.
- David E. Stewart and J. C. Trinkle. 1996. An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction. *Internat. J. Numer. Methods Engrg.* 39, 15 (1996), 2673–2691.
- Jie Tan, Greg Turk, and C. Karen Liu. 2012. Soft Body Locomotion. *ACM Trans. Graph.* 31, 4 (2012).
- E. Todorov. 2011. A convex, smooth and invertible contact model for trajectory optimization. In *2011 IEEE International Conference on Robotics and Automation*. 1071–1076.
- E. Todorov, T. Erez, and Y. Tassa. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- Marc Toussaint, Kelsey Allen, Kevin Smith, and Joshua Tenenbaum. 2018. Differentiable Physics and Stable Modes for Tool-Use and Manipulation Planning. In *Robotics: Science and Systems XIV*. Robotics: Science and Systems Foundation.
- Christopher D. Twigg and Doug L. James. 2007. Many-Worlds Browsing for Control of Multibody Dynamics. *ACM Transactions on Graphics* 26, 3 (2007).
- Ying Wang, Nicholas J. Weidner, Margaret A. Baxter, Yura Hwang, Danny M. Kaufman, and Shinjiro Sueda. 2019. RedMax. *ACM Transactions on Graphics* 38, 4 (2019). <https://doi.org/10.1145/3306346.3322952>
- Peter Wriggers. 2006. *Computational Contact Mechanics*. Springer-Verlag GmbH.
- K. Yunt and C. Glocker. 2006. Trajectory optimization of mechanical hybrid systems using SUMT. In *9th IEEE International Workshop on Advanced Motion Control*.
- Simon Zimmermann, Roi Poranne, James M. Bern, and Stelian Coros. 2019b. Puppet-Master: robotic animation of marionettes. *ACM Transactions on Graphics* 38, 4 (2019). <https://doi.org/10.1145/3306346.3323003>
- Simon Zimmermann, Roi Poranne, and Stelian Coros. 2019a. Optimal control via second order sensitivity analysis. *CoRR* (2019). <http://arxiv.org/abs/1905.08534>