

Learning to Slide Unknown Objects with Differentiable Physics Simulations

Changkyu Song and Abdeslam Boularias
 Computer Science Department, Rutgers University, New Jersey, USA
 {cs1080,abdeslam.boularias}@cs.rutgers.edu

Abstract—We propose a new technique for pushing an unknown object from an initial configuration to a goal configuration with stability constraints. The proposed method leverages recent progress in differentiable physics models to learn unknown mechanical properties of pushed objects, such as their distributions of mass and coefficients of friction. The proposed learning technique computes the gradient of the distance between predicted poses of objects and their actual observed poses, and utilizes that gradient to search for values of the mechanical properties that reduce the reality gap. The proposed approach is also utilized to optimize a policy to efficiently push an object toward the desired goal configuration. Experiments with real objects using a real robot to gather data show that the proposed approach can identify mechanical properties of heterogeneous objects from a small number of pushing actions.

I. INTRODUCTION

Nonprehensile manipulation of objects is a practical skill used frequently by humans to displace objects from an initial configuration to a desired final one with a minimum effort. In robotics, this type of manipulation can be more advantageous than the traditional pick-and-place when an object cannot be easily grasped by the robot, due to the design of the end-effector and the size of the object, or the obstacles surrounding the manipulated object. For example, combined pushing and grasping actions have been shown to succeed where traditional grasp planners fail, and to work well under uncertainty [1–7].

The mechanics of planar pushing was extensively explored in the past [8]. Large datasets of images of planar objects pushed by a robot on a flat surface were also recently presented [9, 10]. Recent techniques for planar sliding mechanics focus on learning data-driven models for predicting the motions of the pushed objects in simulation. While this problem can be solved to a certain extent by using generic end-to-end machine learning tools such as neural networks [10], model identification methods that are explicitly derived from the equations of motion are generally more efficient [11]. A promising new direction is to directly differentiate the prediction error with respect to the model of the object’s mechanical properties, such as its mass and friction distributions, and to use standard gradient descent algorithms to search for values of the properties that reduce the gap between simulated motions and observed ones. Unfortunately, most popular physics engines do not natively provide the derivatives of the predicted poses [12–16], and the only way to differentiate them is through numerical finite differences, which are expensive computationally.

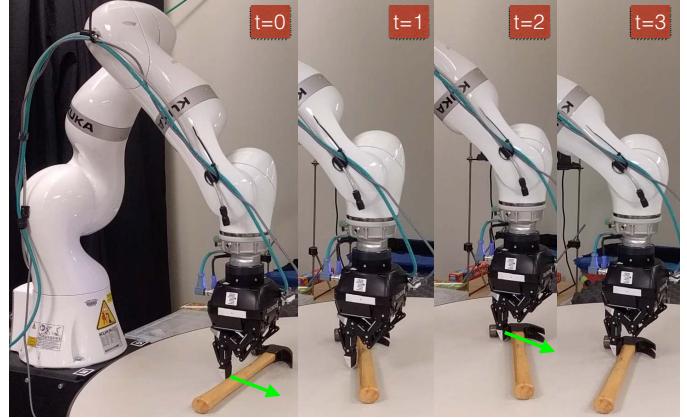


Fig. 1. Robotic setup used for pre-grasp sliding experiments. In this example, the robot pushes a hammer into a desired final pose on the edge of the table through consecutive pushes. The choice of contact points, pushing directions and a stable goal configuration is non-trivial as it depends on the mass distribution and the friction coefficients of different parts of the object, which are unknown *a priori*. This work proposes a method for identifying these parameters from a small number of observed motions of the object.

In the present work¹, we propose a method for safely sliding an unknown object from an initial configuration to a target one. The proposed approach integrates a model identification algorithm with a planner. To account for non-uniform surface properties and mass distributions, the object is modeled as a large set of small cuboids that may have different material properties and that are attached to each other with fixed rigid joints. A simulation error function is given as the distance between the centers of the objects in simulated trajectories and the true observed ones. The gradient of the simulation error is used to search for the object’s mass and friction distributions.

The *main contribution* is the derivation of the analytical gradient of the simulation error with respect to the mass and friction distributions using the proposed cuboid representation of objects. The *second contribution* is the use of the derived analytical gradients for identifying models of unknown objects by using a robotic manipulator, and demonstrating the computational and data efficiency of the proposed approach. The *third contribution* is the use of the proposed integrated method for pre-grasp sliding manipulation of thin unknown objects.

A video of the data collection process, the pre-grasp sliding and grasping experiments using the robotic setup shown in Figure 1 can be viewed at <https://bit.ly/371Y6Y1>.

¹This work was supported by NSF awards 1734492, 1723869 and 1846043.

II. RELATED WORK

Algorithms for *Model-based reinforcement learning* explicitly learn the unknown dynamics, often from scratch, and search for an optimal policy accordingly [17–22]. The unknown dynamics are often modeled using an off-the-shelf statistical learning algorithm, such as a *Gaussian Process* (GP) [23–29], or a neural network [30–33]. This approach was recently used to collect images of pushed objects and build models of their motions [10, 23, 34, 35]. While the proposed method belongs to the category of model-based RL, it differs from most related methods by the explicit use of the dynamics equations, which drastically improves its data-efficiency. The identified mass distribution can also be used to predict the balance and stability of the object in new configurations that are not covered in the training data. For instance, a GP or a neural net cannot predict if an object remains stable when pushed to the edge of a support table, unless such an example is included in the training data, with the risk dropping the object and losing it.

The *mechanics of pushing* was explored in several past works [8, 36–45], from both a theoretical and algorithmic point of view. Notably, Mason [8] derived the voting theorem to predict the rotation and translation of an object pushed by a point contact. A strategy for stable pushing when objects remain in contact with an end-effector was also proposed in [36]. Yoshikawa and Kurisu [43] proposed a regression method for identifying the support points of a pushed object by dividing the support surface into a grid and approximating the measured frictional force and torque as the sum of unknown frictional forces applied on the grid’s cells. A similar setup was considered in [37] with a constraint to ensure positive friction coefficients. The *limit surface* is a convex set of all friction forces and torques that can be applied on an object in quasi-static pushing. The limit surface is often approximated as an ellipsoid [38], or a higher-order convex polynomial [40–42]. An ellipsoid approximation was also used to simulate the motion of a pushed object to perform a *push-grasp* [39]. In contrast with our method, these works identify only the friction parameters, and assume that the mass distribution is known or irrelevant in a quasi-static regime.

There has been a recent surge of interest in developing *natively differentiable physics engines* [46, 47]. A combination of a learned and a differentiable simulator was used to predict effects of actions on planar objects [48], and to learn fluid parameters [49]. Differentiable physics simulations were also used for manipulation planning and tool use [50]. Recently, it has been observed that a standard physical simulation, formulated as a Linear Complementary Problem (LCP), is also differentiable and can be implemented in PyTorch [51]. In [52], a differentiable contact model was used to allow for optimization of several locomotion and manipulation tasks.

III. PROBLEM SETUP AND NOTATION

We consider the problem of displacing a rigid object on a flat homogeneous surface from an initial pose x_0 to a desired final pose x_T^d . The object has an unknown shape and

material properties. We assume that a depth-sensing camera provides a partial 3D view of the object. The partial view contains only the object’s upper surface. A 3D shape is then automatically constructed by assuming that the occluded bottom side is flat. Most of the objects used in our experiments are not flat. However, the autonomously learned friction model simply assigns near-zero friction coefficients to the regions of the bottom surface that do not actually touch the tabletop. Thus, learned near-zero friction forces compensate for wrongly presumed flat regions in the occluded bottom part of the object.

We approximate the object as a finite set of small cuboids. The object is divided into large number of connected cells $1, 2, \dots, n$, using a regular grid structure. Each cell i has its own local mass and coefficient of friction that can be different from the other cells. The object’s pose x_t at time $t \in [0, T]$ is a vector in $[SE(2)]^n$ corresponding to the translation and rotation in the plane for each of the n cells. In other terms, $x_t = [p_{x,t}^1, p_{y,t}^1, \theta_t^1, \dots, p_{x,t}^n, p_{y,t}^n, \theta_t^n]^T$, where $(p_{x,t}^i, p_{y,t}^i)$ is the i^{th} cell’s 2D position on the surface, and θ_t^i is its angle of rotation. Similarly, we denote the object’s generalized velocity (a *twist*) at time t by $\dot{x}_t = [\dot{p}_{x,t}^i, \dot{p}_{y,t}^i, \dot{\theta}_t^i]_{i=0}^n$, where $(\dot{p}_{x,t}^i, \dot{p}_{y,t}^i)$ is the i^{th} cell’s linear velocity on the surface, and $\dot{\theta}_t^i$ is its angular velocity. The object’s mass matrix \mathcal{M} is a diagonal $3n \times 3n$ matrix, where the diagonal is $[\mathcal{I}_1, \mathcal{M}_1, \mathcal{M}_1, \mathcal{I}_2, \mathcal{M}_2, \mathcal{M}_2, \dots, \mathcal{I}_n, \mathcal{M}_n, \mathcal{M}_n]$, \mathcal{I}_i is the moment of inertia of the i^{th} cell of the object, and \mathcal{M}_i is its mass. $\mathcal{I}_i = \frac{1}{6}\mathcal{M}_i w^2$ where w is the width of a cuboid. μ is a $3n \times 3n$ diagonal matrix, where the diagonal is $[\mu_1, \mu_1, \mu_1, \mu_2, \mu_2, \mu_2, \dots, \mu_n, \mu_n, \mu_n]$. μ_i is the coefficient of friction between the i^{th} cell of the object and the support surface. We assume that: $\forall i \in \{1, \dots, n\} : \mu_i \in [0, \mu_{max}]$, where μ_{max} is a given upper bound. An external generalized force (a *wrench*) denoted by F is an $1 \times 3n$ vector $[f_x^1, f_y^1, \tau^1, \dots, f_x^n, f_y^n, \tau^n]^T$, where $[f_x^i, f_y^i]$ and τ^n are respectively the force and torque applied on cell i . External forces are generated from the contact between the object and a fingertip of the robotic hand used to push the object. We assume that at any given time t , at most one cell of the object is in contact with the fingertip. Therefore, $F = [0, 0, 0, \dots, f_x^{c(t)}, f_y^{c(t)}, \tau^{c(t)}, \dots, 0, 0, 0]^T$ where $c(t) \in \{0, \dots, n\}$ is the index of the contacted cell at time t .

A ground-truth trajectory \mathcal{T}^g is a state-action sequence $(x_0^g, \dot{x}_0^g, F_0, \dots, x_{T-1}^g, \dot{x}_{T-1}^g, F_{T-1}, x_T^g, \dot{x}_T^g)$, wherein (x_t^g, \dot{x}_t^g) is the observed pose and velocity of the pushed object, and F_t is the external force applied at time t , as defined above. A corresponding simulated trajectory \mathcal{T} is obtained by starting at the same initial state \hat{x}_0 in the corresponding real trajectory, i.e., $\hat{x}_0 = x_0$, and applying the same control sequence $(F_0, F_1, \dots, F_{T-1})$. Thus, the simulated trajectory \mathcal{T} results in a state-action sequence $(x_0^g, \dot{x}_0^g, F_0, x_1, \dot{x}_1, F_1, \dots, x_{T-1}, \dot{x}_{T-1}, F_{T-1}, x_T, \dot{x}_T)$, where $x_{t+1} = x_t + \dot{x}_t dt$ is the predicted next pose. Velocity \dot{x}_t is a vector corresponding to translation and angular velocities in the plane for each of the n cells, it is predicted in simulation as $\dot{x}_{t+1} = V(x_t, \dot{x}_t, F_t, \mathcal{M}, \mu)$. The goal is to identify mass distribution \mathcal{M} and friction map μ that result in simulated trajectories that are as close as possible to the real observed ones. Therefore, the objective is to solve the

$$\boxed{\begin{bmatrix} \mathbf{0} \\ \mathbf{0} \\ \rho \\ \xi \end{bmatrix} + \begin{bmatrix} \mathcal{M} & \mathcal{J}_e^T(x_t) & \mathcal{J}_f^T(\dot{x}_t) & \mathbf{0} \\ \mathcal{J}_e(x_t) & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathcal{J}_f(\dot{x}_t) & \mathbf{0} & \mathbf{0} & -I \\ \mathbf{0} & \mathbf{0} & I & \mathbf{0} \end{bmatrix} \begin{bmatrix} -\dot{x}_{t+dt} \\ \lambda_e \\ \lambda_f \\ \gamma \end{bmatrix} = \begin{bmatrix} -\mathcal{M}\dot{x}_t - dt\mathbf{F}_t \\ \mathbf{0} \\ \mathbf{0} \\ \mu \text{diag}(\mathcal{M}) \end{bmatrix} \text{ s.t. } \begin{bmatrix} \rho \\ \xi \end{bmatrix} \geq \mathbf{0}, \begin{bmatrix} \lambda_f \\ \gamma \end{bmatrix} \geq \mathbf{0}, \begin{bmatrix} \rho \\ \xi \end{bmatrix} \begin{bmatrix} \lambda_f \\ \gamma \end{bmatrix} = \mathbf{0}}$$

Fig. 2. Equations of Motion as a Linear Complementarity Problem (LCP)

following optimization problem,

$$\begin{aligned} (\mathcal{M}^*, \mu^*) &= \arg \min_{\mathcal{M}, \mu} \text{loss}(\mathcal{M}, \mu), \\ \text{loss}(\mathcal{M}, \mu) &\stackrel{\text{def}}{=} \sum_{t=0}^{T-2} \|x_{t+2}^g - (x_{t+1}^g + V(x_t^g, \dot{x}_t^g, F_t, \mathcal{M}, \mu)dt)\|_2. \end{aligned} \quad (1)$$

Since x_t is a vector containing all cells' positions, the loss is the sum of distances between each cell's ground-truth pose and its predicted pose, which is equivalent to the average distance (ADD) metric as proposed in [53]. In the following, we explain how velocity function V is computed.

IV. FORWARD SIMULATION

We adopt here the formulation presented in [51, 54, 55]. We adapt and customize the formulation to exploit the proposed grid-structure representation, and we extend it to include frictional forces between a pushed object and a support surface. The transition function is given as $\dot{x}_{t+1} = x_t + \dot{x}_t dt$ where dt is the duration of a constant short time-step. Velocity \dot{x}_t is a function of force F_t and mechanical parameters \mathcal{M} and μ . To find \dot{x}_{t+1} , we solve the system of equations of motion that we present in Figure 2, where x_t and \dot{x}_t are inputs, $[\rho, \xi]$ are slack variables, $[\dot{x}_{t+dt}, \lambda_e, \lambda_f, \gamma]$ are unknown vectors, and $[\mathcal{M}, \mu]$ are hypothesized mass and friction matrices. $\text{diag}(\mathcal{M})$ is a $1 \times 3n$ vector corresponding to the main diagonal of \mathcal{M} .

$\mathcal{J}_e(x_t)$ is a global Jacobian matrix of all the adjacency constraints in the grid structure. These constraints ensure that the different cells of the object move together with the same velocity. $\mathcal{J}_e(x_t)$ is an $m \times n$ matrix where n is the number of cells, and m is the number of pairs of adjacent cells.

$$\mathcal{J}_e(x_t) = \begin{bmatrix} \mathcal{J}_e^{1,1}(x_t) & \dots & \mathcal{J}_e^{1,n}(x_t) \\ \vdots & \dots & \vdots \\ \mathcal{J}_e^{m,1}(x_t) & \dots & \mathcal{J}_e^{m,n}(x_t) \end{bmatrix},$$

If cell i , whose four sides have length l , is one of the two adjacent cells in the pair indexed by k , then

$$\mathcal{J}_e^{k,i}(x_t) = \begin{bmatrix} -\frac{l}{2} \sin \theta_t^i & 1 & 0 \\ \frac{l}{2} \cos \theta_t^i & 0 & 1 \end{bmatrix}. \text{Else, } \mathcal{J}_e^{k,i}(x_t) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}.$$

λ_e is a $2m \times 1$ variable vector that is multiplied by the Jacobian $\mathcal{J}_e(x_t)$ to generate the vector of impulses $\mathcal{J}_e^T(x_t)\lambda_e$, which are time-integrals of internal forces that preserve the rigid structure of the object.

$\mathcal{J}_f(\dot{x}_t)$ is an $n \times n$ Jacobian matrix related to the frictional forces between the object's cells and the support surface, and the corresponding constraints. The main block-diagonal of

$\mathcal{J}_f(\dot{x}_t)$ is $[\mathcal{J}_f^1(\dot{x}_t), \mathcal{J}_f^2(\dot{x}_t), \dots, \mathcal{J}_f^n(\dot{x}_t)]$, wherein

$$\mathcal{J}_f^i(\dot{x}_t) = \begin{bmatrix} -\text{sign}(\dot{\theta}_t^i) & 0 & 0 \\ 0 & \frac{\dot{p}_{x,t}^i}{\sqrt{(\dot{p}_{x,t}^i)^2 + (\dot{p}_{y,t}^i)^2}} & \frac{\dot{p}_{y,t}^i}{\sqrt{(\dot{p}_{x,t}^i)^2 + (\dot{p}_{y,t}^i)^2}} \end{bmatrix},$$

and the remaining entries of $\mathcal{J}_f(\dot{x}_t)$ are all zeros. λ_f is a $2n \times 1$ variable vector. $\mathcal{J}_f(\dot{x}_t)$ is multiplied by λ_f to generate a vector of the frictional forces and torques between the support surface and each cell of the object. $\mathcal{J}_f(\dot{x}_t)$ defines the direction of the frictional forces and torques as the opposite of its current velocity $\dot{x}_t = [\dot{\theta}_t^i, \dot{p}_{x,t}^i, \dot{p}_{y,t}^i]_{i=1}^n$, whereas λ_f defines the scalar magnitudes of the frictional forces and torques.

The friction terms have complementary constraints, stated in Fig. 2. These constraints are used to distinguish between the cases when the object is moving and friction magnitudes λ_f are equal to $\mu \text{diag}(\mathcal{M})$, and the case when the object is stationary and the friction magnitudes λ_f are smaller than $\mu \text{diag}(\mathcal{M})$. When the object moves, and assuming that the change in the direction of motion happens smoothly, we have $\mathcal{J}_f(\dot{x}_t)\dot{x}_{t+dt} < 0$. Therefore, $\gamma > 0$ because of the constraints $\rho = \mathcal{J}_f(\dot{x}_t)\dot{x}_{t+dt} + \gamma I$ and $\rho \geq 0$ and $\gamma \geq 0$. Then $\xi = 0$ because of the constraint $\gamma\xi = 0$. We conclude that $\lambda_f = \mu \text{diag}(\mathcal{M})$ from the constraint $\xi + \lambda_f = \mu \text{diag}(\mathcal{M})$. Similarly, one can show that $\lambda_f < \mu \text{diag}(\mathcal{M})$ if $\dot{x}_{t+dt} = 0$.

To simulate a trajectory $(x_0, \dot{x}_0, F_0, x_1, \dot{x}_1, F_1, \dots)$, we iteratively find velocities \dot{x}_{t+dt} by solving the equations in Fig. 2 where $(x_t, \dot{x}_t, F_t, \mu, \mathcal{M})$ are fixed inputs and the remaining variables are unknown. The solution is obtained, after an initialization step, by iteratively minimizing the residuals from the equations in Fig. 2, using the convex optimizer of [56].

V. MASS AND FRICTION GRADIENTS

To obtain material parameters $[\mathcal{M}, \mu]$, a gradient descent on the loss function in Equation 1 is performed. A first approach to compute the gradient is to use the *Autograd* library for automatic derivation in Python. We propose here a second simpler and faster approach based on deriving analytically the closed forms of the gradients $\frac{\partial \text{loss}(\mathcal{M}, \mu)}{\partial \mu}$ and $\frac{\partial \text{loss}(\mathcal{M}, \mu)}{\partial \mathcal{M}}$.

Let us denote by $(\dot{x}_{t+1}^*, \lambda_e^*, \lambda_f^*, \gamma^*)$ the solutions for $(\dot{x}_{t+1}, \lambda_e, \lambda_f, \gamma)$ in the system in Fig. 2. Let us also use $D(x)$ to denote a matrix that contains a vector x as a main diagonal and zeros elsewhere. Finally, let $\text{diag}(\mathcal{M})$ refer to the main diagonal of \mathcal{M} . In other terms, $\text{diag}(\mathcal{M}) = [\mathcal{I}_1, \mathcal{M}_1, \mathcal{M}_1, \dots, \mathcal{I}_n, \mathcal{M}_n, \mathcal{M}_n]$. Then,

$$\begin{cases} -\mathcal{M}\dot{x}_{t+1}^* + \mathcal{J}_e^T(x_t)\lambda_e^* + \mathcal{J}_f^T(\dot{x}_t)\lambda_f^* + \mathcal{M}\dot{x}_t + F_t dt = 0, \\ -\mathcal{J}_e(x_t)\dot{x}_{t+1}^* = 0, \\ D(\lambda_f^*)(-\mathcal{J}_f(\dot{x}_t)\dot{x}_{t+1}^* - \gamma^*) = 0, \\ D(\gamma^*)(\lambda_f^* - \mu \text{diag}(\mathcal{M})) = 0. \end{cases}$$

The differentials of the system are given as

$$\begin{cases} -\partial\mathcal{M}\dot{x}_{t+1}^* - \mathcal{M}\partial\dot{x}_{t+1} + \mathcal{J}_e^T(x_t)\partial\lambda_e + \mathcal{J}_f^T(\dot{x}_t)\partial\lambda_f + \partial\mathcal{M}\dot{x}_t = 0, \\ -\mathcal{J}_e(x_t)\partial\dot{x}_{t+1} = 0, \\ D(-\mathcal{J}_f(\dot{x}_t)\dot{x}_{t+1}^* - \gamma^*)\partial\lambda_f + D(\lambda_f^*)(-\mathcal{J}_f(x_t)\partial\dot{x}_{t+1} - \partial\gamma) = 0, \\ D(\lambda_f^* - \mu \text{diag}(\mathcal{M}))\partial\gamma + D(\gamma^*)(\partial\lambda_f - \partial\mu \text{diag}(\mathcal{M}) - \mu\partial\text{diag}(\mathcal{M})) = 0, \end{cases}$$

wherein ∂x_t , $\partial\dot{x}_t$, $\partial\mathcal{J}_e^T(x_t)$, $\partial\mathcal{J}_f^T(\dot{x}_t)$ are all zero matrices and vectors because x_t and \dot{x}_t are fixed and treated as a constant since they are set to x_t^g and \dot{x}_t^g in Equation 1. Also, $\partial F_t = \mathbf{0}$ because the applied force at time t is given as a constant in the identification phase. The differentials can be arranged in the following matrix form: $GX = Y$,

$$G = \begin{bmatrix} \mathcal{M} & \mathcal{J}_e^T(x_t) & \mathcal{J}_f^T(\dot{x}_t) & \mathbf{0} \\ \mathcal{J}_e(x_t) & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ D(\lambda_f^*)\mathcal{J}_f(\dot{x}_t) & \mathbf{0} & D(-\mathcal{J}_f(\dot{x}_t)\dot{x}_{t+1}^* - \gamma^*) & -D(\lambda_f^*) \\ \mathbf{0} & \mathbf{0} & D(\gamma^*) & D(\lambda_f^* - \mu \text{diag}(\mathcal{M})) \end{bmatrix}$$

$$X = [-\partial\dot{x}_{t+1}, \partial\lambda_e, \partial\lambda_f, \partial\gamma]^T,$$

$$Y = [\partial\mathcal{M}(\dot{x}_{t+1}^* - \dot{x}_t), \mathbf{0}, \mathbf{0}, D(\gamma^*)\partial(\mu \text{diag}(\mathcal{M}))]^T.$$

Also,

$$\begin{aligned} \frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}(-\partial\dot{x}_{t+1}) &= \left[\frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}, \mathbf{0}, \mathbf{0}, \mathbf{0} \right] X \\ &= \left[\frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}, \mathbf{0}, \mathbf{0}, \mathbf{0} \right] G^{-1}Y = [\alpha_1, \alpha_2, \alpha_3, \alpha_4]Y, \end{aligned}$$

wherein $[\alpha_1, \alpha_2, \alpha_3, \alpha_4]$ is defined as $[\frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}, \mathbf{0}, \mathbf{0}, \mathbf{0}]G^{-1}$. We use the blockwise matrix inversion to compute G^{-1} ,

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix}^{-1} = \begin{bmatrix} (A - BD^{-1}C)^{-1} & -(A - BD^{-1}C)^{-1}BD^{-1} \\ g(A, B, C, D) & h(A, B, C, D) \end{bmatrix}$$

where A and D are square matrices, and D and $(A - BD^{-1}C)$ are invertible. Notice that to compute $[\alpha_1, \alpha_2, \alpha_3, \alpha_4]$, we only need the upper quarter of G^{-1} , because the remaining rows will be multiplied by $\mathbf{0}$. Consequently, terms $g(A, B, C, D)$ and $h(A, B, C, D)$ do not matter here, and we only need the terms $(A - BD^{-1}C)^{-1}$ and $-(A - BD^{-1}C)^{-1}BD^{-1}$. The first term $(A - BD^{-1}C)^{-1}$ corresponds to

$$\left(\begin{bmatrix} \mathcal{M} & \mathcal{J}_e^T(x_t) \\ \mathcal{J}_e(x_t) & \mathbf{0} \end{bmatrix} - \begin{bmatrix} \mathcal{J}_f^T(\dot{x}_t) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} D(-\mathcal{J}_f(\dot{x}_t)\dot{x}_{t+1}^* - \gamma^*) & -D(\lambda_f^*) \\ D(\gamma^*) & D(\lambda_f^* - \mu \text{diag}(\mathcal{M})) \end{bmatrix} \right)^{-1} \begin{bmatrix} D(\lambda_f^*)\mathcal{J}_f(\dot{x}_t) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}^{-1}.$$

In the model identification phase, we only utilize data points where the object actually moves when pushed by the robot. Thus, $-\mathcal{J}_f(\dot{x}_t)\dot{x}_{t+1}^* - \gamma^* = \mathbf{0}$ and $\lambda_f^* - \mu \text{diag}(\mathcal{M}) = \mathbf{0}$, and

$$(A - BD^{-1}C)^{-1} = \begin{bmatrix} \mathcal{M} & \mathcal{J}_e^T(x_t) \\ \mathcal{J}_e(x_t) & \mathbf{0} \end{bmatrix}^{-1} = \begin{bmatrix} X_{1,1} & X_{1,2} \\ X_{2,1} & X_{2,2} \end{bmatrix}.$$

Using the blockwise matrix inversion, we find that

$$X_{1,1} = \mathcal{M}^{-1} + \mathcal{M}^{-1}\mathcal{J}_e^T(x_t)(-\mathcal{J}_e(x_t)\mathcal{M}^{-1}\mathcal{J}_e^T(x_t))^{-1}\mathcal{J}_e(x_t)\mathcal{M}^{-1}.$$

We will see in the following that the remaining matrices, $X_{1,2}$, $X_{2,1}$ and $X_{2,2}$, will not be needed. Similarly, the top right of G^{-1} is $-(A - BD^{-1}C)^{-1}BD^{-1}$. It is given as

$$\begin{aligned} & - \left(\begin{bmatrix} X_{1,1} & X_{1,2} \\ X_{2,1} & X_{2,2} \end{bmatrix} \begin{bmatrix} \mathcal{J}_f^T(\dot{x}_t) & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{0} & D^{-1}(\gamma^*) \\ -D^{-1}(\lambda_f^*) & \mathbf{0} \end{bmatrix} \right) \\ & = \begin{bmatrix} \mathbf{0} & X_{1,1}\mathcal{J}_f^T(\dot{x}_t)D^{-1}(\gamma^*) \\ \mathbf{0} & X_{2,1}\mathcal{J}_f^T(\dot{x}_t)D^{-1}(\gamma^*) \end{bmatrix}. \end{aligned}$$

Therefore,

$$G^{-1} = \begin{bmatrix} X_{1,1} & X_{1,2} & \mathbf{0} & X_{1,1}\mathcal{J}_f^T(\dot{x}_t)D^{-1}(\gamma^*) \\ X_{2,1} & X_{2,2} & \mathbf{0} & X_{2,1}\mathcal{J}_f^T(\dot{x}_t)D^{-1}(\gamma^*) \\ X_{3,1} & X_{3,2} & X_{3,3} & X_{3,4} \\ X_{4,1} & X_{4,2} & X_{4,3} & X_{4,4} \end{bmatrix}.$$

The first term in $[\alpha_1, \alpha_2, \alpha_3, \alpha_4]$ is then

$$\alpha_1 = \frac{\partial\text{loss}}{\partial(-\dot{x}_{t+dt})}X_{1,1},$$

and the fourth term is

$$\alpha_4 = \frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}(-X_{1,1}\mathcal{J}_f^T(\dot{x}_t)D^{-1}(\gamma^*)).$$

Since $\frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}(-\partial\dot{x}_{t+1}) = [\alpha_1, \alpha_2, \alpha_3, \alpha_4][\partial\mathcal{M}(\dot{x}_{t+1}^* - \dot{x}_t), \mathbf{0}, \mathbf{0}, D(\gamma^*)\partial(\mu \text{diag}(\mathcal{M}))]^T$, then

$$\begin{aligned} \frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}(-\partial\dot{x}_{t+1}) &= \frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}(X_{1,1}\partial\mathcal{M}(\dot{x}_{t+1}^* - \dot{x}_t) \\ &\quad + (-X_{1,1}\mathcal{J}_f^T(\dot{x}_t)D^{-1}(\gamma^*))D(\gamma^*)\partial(\mu \text{diag}(\mathcal{M}))) \\ &= \frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}(X_{1,1}\partial\mathcal{M}(\dot{x}_{t+1}^* - \dot{x}_t) + (-X_{1,1}\mathcal{J}_f^T(\dot{x}_t))\partial(\mu \text{diag}(\mathcal{M}))) \end{aligned}$$

In the following, we show how to use the equation above to derive $\frac{\partial\text{loss}}{\partial\mathcal{M}}$ and $\frac{\partial\text{loss}}{\partial\mu}$ and use them in a *coordinate descent* algorithm to identify (\mathcal{M}^*, μ^*) from data.

A. Mass Gradient

We calculate $\frac{\partial\text{loss}}{\partial\mathcal{M}}$ while setting $\partial\mu = \mathbf{0}$.

$$\begin{aligned} \frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}(-\partial\dot{x}_{t+1}) &= \frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}(X_{1,1}\partial\mathcal{M}(\dot{x}_{t+1}^* - \dot{x}_t) \\ &\quad + (-X_{1,1}\mathcal{J}_f^T(\dot{x}_t))\mu\partial\text{diag}(\mathcal{M})) \end{aligned}$$

Then,

$$\frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}\left(\frac{\partial\dot{x}_{t+1}}{\partial\mathcal{M}}\right) = \frac{\partial\text{loss}}{\partial(\dot{x}_{t+1})}X_{1,1}(\mathcal{J}_f^T(\dot{x}_t)\mu - D(\dot{x}_{t+1}^* - \dot{x}_t))$$

From the definition of the loss function in Equation 1, we can see that $\frac{\partial\text{loss}}{\partial(\dot{x}_{t+1})} = 2dt \sum_{t=1}^{T-1} D(x_{t+1} - x_{t+1}^g)$, wherein x_{t+1}^g is the observed ground-truth pose of the object and x_{t+1} is its predicted pose, computed as $x_{t+1} = x_t^g + \dot{x}_t^g dt$. Finally,

$$\frac{\partial\text{loss}}{\partial\mathcal{M}} = 2dt \sum_{t=1}^{T-1} (D(x_{t+1} - x_{t+1}^g)X_{1,1}(\mathcal{J}_f^T(\dot{x}_t)\mu - D(\dot{x}_{t+1}^* - \dot{x}_t))).$$

B. Friction Gradient

We calculate $\frac{\partial\text{loss}}{\partial\mu}$ while setting $\partial\mathcal{M} = \mathbf{0}$.

$$\frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}(-\partial\dot{x}_{t+1}) = \frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}(-X_{1,1}\mathcal{J}_f^T(\dot{x}_t)\partial\mu\text{diag}(\mathcal{M}))$$

Then,

$$\frac{\partial\text{loss}}{\partial(-\dot{x}_{t+1})}\left(\frac{\partial\dot{x}_{t+1}}{\partial\mu}\right) = \frac{\partial\text{loss}}{\partial(\dot{x}_{t+1})}(-X_{1,1}\mathcal{J}_f^T(\dot{x}_t)\mathcal{M})$$

Finally,

$$\frac{\partial\text{loss}}{\partial\mu} = 2dt \sum_{t=1}^{T-1} (D(x_{t+1} - x_{t+1}^g)X_{1,1}\mathcal{J}_f^T(\dot{x}_t)\mathcal{M})$$

Algorithm 1: Learning Mass and Friction with Differentiable Physics Simulations

Input: A set of real-world trajectories $\mathcal{D} = \{(x_t^g, \dot{x}_t^g, F_t)_{t=0}^T\}$; a learning rate α_{rate} ; and initial mass matrix \mathcal{M} and friction map μ ; Maximum mass and friction $m_{\max}, u_{\max} \in \mathbb{R}$;

Output: Updated mass matrix \mathcal{M} and friction map μ

```

1 repeat
2   foreach  $\mathcal{T}^g = (x_t^g, \dot{x}_t^g, F_t)_{t=0}^T \in \mathcal{D}$  do
3      $(x_0, \dot{x}_0) \leftarrow (x_0^g, \dot{x}_0^g); x_1 \leftarrow x_1^g;$ 
4     for  $t \in \{0, T-2\}$  do
5        $\dot{x}_{t+1} \leftarrow V(x_t, \dot{x}_t, F_t, \mathcal{M}, \mu)$ ; ▷ Solving the LCP
6        $x_{t+2} \leftarrow x_{t+1} + \dot{x}_{t+1} dt$ ; ▷ Predicting next pose
7        $X \leftarrow \mathcal{M}^{-1} +$ 
8        $\mathcal{M}^{-1} \mathcal{J}_e^T(x_t) (-\mathcal{J}_e(x_t) \mathcal{M}^{-1} \mathcal{J}_e^T(x_t))^{-1} \mathcal{J}_e(x_t) \mathcal{M}^{-1}$ 
9        $W_\mu \leftarrow X \mathcal{J}_f^T(\dot{x}_t) \mathcal{M};$ 
10       $W_M \leftarrow X (\mathcal{J}_f^T(\dot{x}_t) \mu - D(\dot{x}_{t+1} - \dot{x}_t));$ 
11       $\mu \leftarrow \mu - \alpha_{\text{rate}} D(x_{t+1} - x_{t+1}^g) W_\mu;$ 
12       $\mathcal{M} \leftarrow \mathcal{M} - \alpha_{\text{rate}} D(x_{t+1} - x_{t+1}^g) W_M;$ 
13       $\mu \leftarrow \arg \min_{\mu' \in [0, I u_{\max}]} \|\mu' - \mu\|_\infty;$ 
14       $\mathcal{M} \leftarrow \arg \min_{\mathcal{M}' \in [0, I m_{\max}]} \|\mathcal{M}' - \mathcal{M}\|_\infty$ ; ▷ Projecting the gradients
15   end
16 end
17 until timeout;

```

C. Mass and Friction Identification Algorithm

The gradient of the loss function with respect to the mass and friction matrices \mathcal{M} and μ are used in Algorithm 1 to search for the ground-truth mass and friction. We present here the stochastic version where the gradient is computed for each trajectory separately. The gradient can also be computed in a batch mode from all trajectories. $\{(x_t^g, \dot{x}_t^g, F_t)_{t=0}^T\}$ of data collected by the robot. The mass and friction are increased or decreased depending on the signs in the error vector $(x_{t+1} - x_{t+1}^g)$, which corresponds to the reality gap. The main computational bottleneck is in computing weights W_M and W_μ , which is linear in the number of cells n because \mathcal{M} is a diagonal matrix, and $\mathcal{J}_e(x_t)$ is block diagonal thanks to the regular grid structure of the cells. Finally, we project the updated mass and friction matrices by rounding their values down to upper limits $m_{\max}, u_{\max} \in \mathbb{R}$. The provided upper limits are the same for every cell in the object, whereas the mass and friction distributions learned by the algorithm are highly heterogenous, as will be shown in the experiments.

VI. POLICY GRADIENT

After identifying the mass distribution \mathcal{M} and friction map μ using Algorithm 1, we search for a new sequence of forces $(F_t)_{t=0}^{T-1}$ to push the object toward a desired terminal goal configuration x_T^d . Algorithm 2 summarizes the main steps of this process. We start by creating a rapidly exploring random tree (RRT^*) to find the shortest path from x_0 to x_T^d . While searching for the shortest path, we eliminate from the tree object poses that are unstable (based on the identified mass distribution \mathcal{M}) or that are in collision with other objects. RRT^* returns a set of waypoints $\mathcal{X}_{\text{waypoint}}$. At each iteration

of the main loop of the algorithm, we find the nearest waypoint in $\mathcal{X}_{\text{waypoint}}$ and search for actions that would push the object toward it. A pushing force is parameterized by a contact point, a direction and a magnitude, as discussed in Section III. We focus here on optimizing the contact point, and we keep the magnitude constant. The direction of the force is chosen to be always horizontal. It is given as the opposite of the surface normal of the object at the contact point, projected down on the 2D plane of the support surface. This choice is made to avoid slippages and changes in contact points during a push.

A contact point is always located on the outer side of a cuboid (cell). Therefore, we limit the search to the outer cells of the grid. The objective of this search is to select a contact point that reduces the gap between the predicted pose of the object after pushing it, and the nearest waypoint x_{target} that has not been reached yet. We select the initial contact point as the outer cell that is most aligned to the axis $\hat{x} - x_{\text{target}}$, where \hat{x} is the estimated center of mass. The gradient of the gap with respect to the contact point is computed by using the finite-difference method. The contact point is moved in the direction that minimizes the gap until a local optimum is reached. Force F_t is then defined based on the selected contact point. The pose and velocity of the object are replaced by the predicted ones that result from applying force F_t . This process is repeated until the object reaches the desired goal configuration. The time duration of each pushing action is also optimized by using finite differences. This part is omitted for simplicity's sake.

The surface of the object often contains non-differentiable parts where the analytical gradient with respect to the contact point is undefined. Even on smooth parts, there is no clear advantage of computing the analytical gradient here, because the space of contact points is uni-dimensional, in contrast to the high-dimensional space of non-uniform mass and friction distributions. In low-dimensional search spaces, finite-difference methods are computationally efficient.

VII. MAIN ALGORITHM

Algorithm 3 summarizes the main steps of the proposed approach and the protocol followed in the experiments. In summary, the robot first “plays” with the unknown object by applying random short-lasting horizontal forces for a safe and local exploration. A mass and friction model is then inferred from the gathered data by using Algorithm 1. Based on the inferred mass distribution, a safe goal configuration is sampled from a desired goal region. For example, a pregrasp sliding manipulation can be used to grasp a thin object that cannot be directly grasped from a flat surface. In [57], a known object is pushed to the edge of a table and then grasped from there. Pushing an unknown object to the edge of a table results often in losing the object. Our method avoids this issue by sampling a goal configuration that allows a sufficient part of the object to be graspable, while keeping the object balanced on the edge thanks to the identified mass distribution. Once a goal is selected, Algorithm 2 is used to generate a sequence of actions to push the object to the goal.

Algorithm 2: Planning and Control

Input: Identified mass and friction \mathcal{M} and μ , and presumed 3D shape \mathcal{S} of the object; initial ground-truth configuration x_0 ; desired final configuration x_T^d ;

Output: A sequence of pushing forces $(F_t)_{t=0}^{T-1}$;

1 Find a set $\mathcal{X}_{waypoint}$ of waypoints by calling $RRT^*(\mathcal{S}, x_0, x_T^d)$;

2 $\dot{x}_0 \leftarrow \mathbf{0}$; $t \leftarrow 0$;

3 **while** $\mathcal{X}_{waypoint} \neq \emptyset$ **do**

4 **repeat**

5 $\hat{x} \leftarrow \frac{\sum_{i \in \{1, \dots, n\}} (\mathcal{M}_i)(p_{x,t}^i, p_{y,t}^i)}{\sum_{i \in \{1, \dots, n\}} \mathcal{M}_i}$; ▷ center of mass

6 $x_{target} \leftarrow \arg \min_{x \in \mathcal{X}_{waypoint}} \|x - x_t\|_2$;

7 $i^* \leftarrow \arg \max_{i \in \{1, \dots, n\}} \frac{(\hat{x} - x_{target})((p_{x,t}^i, p_{y,t}^i) - \hat{x})^T}{\|\hat{x} - x_{target}\|_2 \|(p_{x,t}^i, p_{y,t}^i) - \hat{x}\|_2}$;

8 ▷ Initializing the contact point

9 improvement \leftarrow true;

10 **repeat**

11 Let i_{left}^* and i_{right}^* be two cells that are adjacent to i^* and that are also on the outer envelope of the object;

12 **for** $i \in \{i^*, i_{left}^*, i_{right}^*\}$ **do**

13 Choose the direction of the force (f_x^i, f_y^i) as the opposite of the object's surface normal at cell i .

14 $F^i \leftarrow [0, 0, 0, \dots, f_x^i, f_y^i, 0, \dots, 0, 0, 0]$;

15 $gap_i \leftarrow \|x_t + \dot{x}_t dt + V(x_t, \dot{x}_t, F^i, \mathcal{M}, \mu) dt - x_{target}\|_2$;

16 **end**

17 $i^* \leftarrow \arg \min_{i \in \{i^*, i_{left}^*, i_{right}^*\}} gap_i$;

18 $F_t \leftarrow [0, 0, 0, \dots, f_x^{i^*}, f_y^{i^*}, 0, \dots, 0, 0, 0]$;

19 $\dot{x}_{t+1} \leftarrow V(x_t, \dot{x}_t, F^{i^*}, \mathcal{M}, \mu)$; $x_{t+1} \leftarrow x_t + \dot{x}_t dt$;

20 **if** $i^* \neq i_{left}^* \wedge i^* \neq i_{right}^*$ **then**

21 improvement \leftarrow false;

22 **end**

23 **until** improvement $=$ false;

24 $t \leftarrow t + 1$;

25 **until** $\min_{x \in \mathcal{X}_{waypoint}} \|x - x_t\|_2 \leq \epsilon$;

26 $\mathcal{X}_{waypoint} \leftarrow \mathcal{X}_{waypoint} \setminus \{\arg \min_{x \in \mathcal{X}_{waypoint}} \|x - x_t\|_2\}$;

27 **end**

VIII. EVALUATION

We report here the results of three sets of experiments to evaluate the proposed approach. The most important set is the one related to mass and friction identification.

A. Experimental Setup

The experiments are performed on both simulated and real robot and objects. In the real robot setup, a *Robotiq* 3-finger hand mounted on a *Kuka* robot is repeatedly moved to collide with a rigid object that is set on a table-top and to push it forward, as shown as Figure 1. The initial, final and intermediate point clouds of the object are recorded. The simulation experiments are performed using the physics engine *Bullet* and models of the robot and objects. The experiments are performed on five real objects: a book, a hammer, a snack, a toolbox and a spray gun, and eight simulated objects: a box, a hammer, a book, a crimp, a snack, a ranch, a spray gun and a toothpaste. The number of cells per object varies from 28 to 88 depending on the size of the object.

Algorithm 3: Learning to Slide Unknown Objects with Differentiable Physics Simulations

Input: Point cloud of an object's upper surface; Maximum mass and friction $m_{max}, u_{max} \in \mathbb{R}$; Desired goal region \mathcal{G} ;

1 Create a 3D shape \mathcal{S} of the object by projecting its upper surface down on the support surface;

2 Decompose the 3D shape into a regular grid of n small cuboids;

3 Let x_0^g be the vector of ground-truth positions and rotations of the n cuboids at time 0; $\dot{x}_0^g = \mathbf{0}$;

4 Initialize mass matrix \mathcal{M} and friction map μ uniformly with a random value;

5 Sample a small number $T_{\text{exploration}}$ of contact points and horizontal pushing directions ; ▷ For a safe local exploration

6 Use the robot's end-effector to apply forces on the sampled contact points and along the sampled directions for short periods of time dt , and record the resulting trajectories in $\mathcal{D} = \{(x_t^g, \dot{x}_t^g, F_t)_{t=0}^{T_{\text{exploration}}}\}$;

7 Use Algorithm 1 to identify \mathcal{M} and μ from dataset \mathcal{D} ;

8 Sample a goal configuration $x^d \in \mathcal{G}$ where the object is predicted to remain stable under gravity, according to the identified mass distribution \mathcal{M} ;

9 Use Algorithm 2 to obtain a new sequence of forces $(F_t)_{t=0}^{T_{\text{execution}}}$ to push the object to x^d ;

10 Execute the sequence $(F_t)_{t=0}^{T_{\text{execution}}}$ with the robot's end-effector;

B. Tasks

Model Identification. Each real and simulated object is pushed by the robot randomly 10 times on the table. Half of the recorded trajectories are used for learning a mass matrix \mathcal{M} and friction map μ . The other half is used for testing the identified models. Since the ground-truth values of mass and friction are unknown, the identified models are evaluated in terms of the accuracy of the predicted pose of each cell after applying the sequence of actions provided in the test set. The experiments on the real objects are repeated with 25 randomized splits into training and testing sets. The simulation experiments are repeated with 10 different models per object.

Planning and Control. For each one of the eight objects in simulation, we randomly sample 10 values for their mass and friction matrices, and 10 random goal configurations in a disk of a radius of 1m around the initial configuration. The rotations of the goal configurations are also selected randomly. The number of settings is then $8 \times 10 \times 10$. The task is to generate, in each setting, a sequence of forces that pushes the object from the initial configuration to the goal. The objective is to assess the computational efficiency of Algorithm 2.

Pre-grasp Sliding Manipulation. Finally, we evaluate the entire system (Algorithm 3) on the task of sliding an object from a random initial pose on a table to a desired goal region at the edge of the table where the object can be grasped. The object cannot be directly grasped from a flat surface, the goal is to push it to the edge where part of it sticks out of the table and becomes graspable. Since the mass distribution of the object is highly heterogeneous and unknown, the object often becomes unbalanced at the edge and falls from the table if the identified model is incorrect. We report here the percentage of experiments where the object is successfully pushed to

the edge and grasped without losing it. The experiments on this task are performed using the real Kuka robot and a real hammer. The exploration phase contains only 5 random pushing actions that are used for model identification. The reported results are averaged over 16 independent runs, with a different initial pose in each run.



Fig. 3. Learned mass distributions. Red color means higher mass value while blue color means lower mass value. The middle column shows the initial uniform mass distributions provided to Algorithm 1 and the right column shows the mass distributions returned by the algorithm after only five iterations of gradient descent.

C. Compared Methods

Model Identification. Algorithm 1 is compared against the following methods. *Random search* is a baseline method that repeatedly samples random values of the mass and friction matrices and returns the best sampled values that minimize $loss(\mathcal{M}, \mu)$. *Weighted sampling search* generates random values uniformly in the first iteration, and then iteratively generates normally distributed random values around the best parameter obtained in the previous iteration. The standard deviation of the random values is gradually reduced over time, to focus the search on the most promising region of the search space. *The finite differences gradient* is an approximation of the analytical gradient. We add or subtract a small amount to the current parameter values and simulate the trajectories using the neighboring parameter values to approximate the derivatives $\frac{\partial loss}{\partial \mathcal{M}}$ and $\frac{\partial loss}{\partial \mu}$. Because a large number of simulations is required to compute the gradient for all the cells, the parameters of each cell are updated using the coordinate gradient descent. We also compare the proposed method with two

black-box optimization methods: *CMA-ES* [58] and *Nelder-Mead*, and an automatic differentiation of the LCP solver using the *Autograd* function of PyTorch [59]. The same minimum and maximum bounds of mass and friction are provided to all methods and are also used for all cells of objects.

Planning and Control. We perform an ablation study where we substitute the finite-difference gradient in Algorithm 2 with an *exhaustive search* of the optimal contact point.

Pre-grasp Sliding Manipulation. We compare Algorithm 3 to two alternatives. The first one assumes a uniform and homogenous mass and friction values, and uses directly Algorithm 2 to push the object to the goal *without model identification*. The second alternative is identical to Algorithm 3, except that *no upper limit* on the friction coefficients is used.

D. Results

Model Identification. Figure 3 shows mass distributions identified using Algorithm 1. The mass distributions of the book, the hammer, the snack, the toolbox and the spray gun are all identified using the real objects and robot. The experiments on the remaining objects were performed in simulation because they were too thin for a safe robotic manipulation. The results show that the identified models quickly converge to the ground-truth models. Figure 4 (a)-(m) shows the average distance between the predicted cell positions and the ground-truth ones in the test data as a function of the number forward physics simulations used by the different identification methods. In our method, the number of physics simulations is the same as the number of steps of the gradient-descent, since one simulation is performed after each update of \mathcal{M} and μ . Note that the physics simulations dominate the computation time, with $1.3448(\pm 0.5604)$ second per a simulation, while the computation of the gradient using the proposed algorithm takes only $0.0069(\pm 0.0024)$ second. The results demonstrate that global optimization methods suffer from the curse of dimensionality due to the combinatorial explosion in the number of possible parameters for all cells. The results also demonstrate that the proposed method can estimate the parameters within a surprisingly small number of gradient-descent steps and a short computation time (30 seconds). The average error of the predicted cell positions using the identified mass and friction distributions is less than 1.53cm in simulation and 2.27cm in the real experiments. Note that the error in real experiments is higher due to sensing and control errors. The finite differences approach also failed to converge to an accurate model due to the high computational cost of the gradient computation, as well as the sensitivity of the computed gradients to the choice the grid size. Figure 4 (o) shows that the number of training actions improves the accuracy of the model learned by Algorithm 1. Increasing the number of training actions allows the robot to uncover properties of different parts of the object more accurately.

Planning and Control. Figure 4 (p) shows the number of physics simulations used to optimize the contact location. The proposed policy gradient algorithm requires a small fraction of the exhaustive search's computation time, while both methods

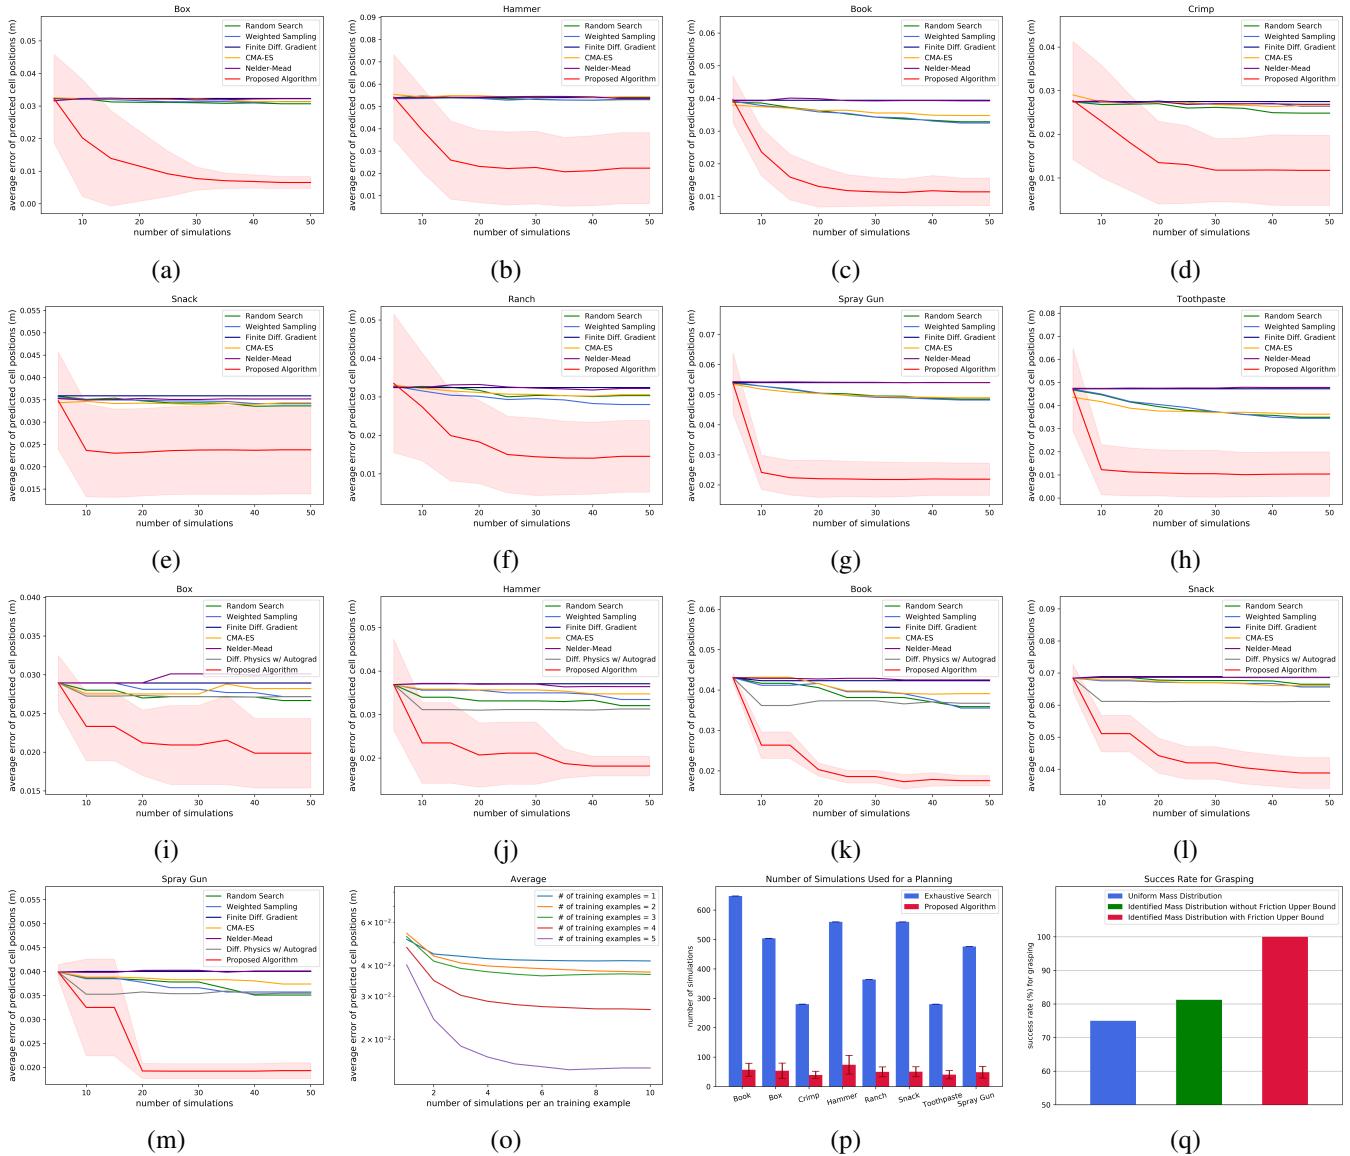


Fig. 4. Average predicted cell position error (in meters) for each object as a function of the number of performed physics simulations. Figures (a)-(h) correspond to the simulation experiments and (i)-(m) correspond to the real experiments. Figure (o) shows the average predicted cell position error with different numbers of training examples. Figure (p) shows the predicted cell position error with different numbers of training examples. Figure (q) shows the success rate for the sliding and grasping task.

attained a 100% success rate in finding a sequence of pushes that reaches the goal in simulation.

Pre-grasp Sliding Manipulation. This task integrates the two previous tasks and evaluates the entire proposed system. Figure 4 (q) shows that in all of the 16 trials with random initial poses, the robot successfully identified the hammer’s mass and friction distributions using Algorithm 1, selected a physically stable goal configuration in the desired goal region based on the identified model, planned and executed a sequence of pushing actions using Algorithm 2, and grasped the object from the part pushed out of the table. Only 12 trials resulted in successful grasps when the mass distribution was not learned and was assumed to be uniform. The upper bound on the dynamic coefficient of friction is set to 1.00 in Algorithm 1. The success rate drops to 13/16 when no upper

bound limit is set on the friction in Algorithm 1.

IX. FINAL REMARK

The decline in the success rate in Figure 4 (q) when no upper limit on the friction is used is an important observation. The algorithm attributed most of the rotations to unrealistically high frictions in certain regions, instead of an uneven mass distribution. Despite identifying a wrong model of mass and friction, the predicted motion of the object was accurate as long as the object was entirely on the table. But once pushed to the edge, and its heavy side is not anymore supported by the table’s surface, the object falls. This clearly demonstrates the importance of identifying not only the friction, but also the accurate mass distribution of objects for a safe manipulation. The upper bound can be seen as an *inductive bias* that is necessary for learning from limited data.

REFERENCES

- [1] Mehmet Remzi Dogar and Siddhartha S. Srinivasa. A framework for push-grasping in clutter. In *Robotics: Science and Systems*, 2011.
- [2] Mehmet R. Dogar and Siddhartha S. Srinivasa. A planning framework for non-prehensile manipulation under clutter and uncertainty. *Autonomous Robots*, 33(3):217–236, Oct 2012.
- [3] Jennifer E. King, Joshua A. Haustein, Siddhartha S. Srinivasa, and Tamim Asfour. Nonprehensile whole arm rearrangement planning on physics manifolds. In *IEEE International Conference on Robotics and Automation, ICRA 2015, Seattle, WA, USA, 26-30 May, 2015*, pages 2508–2515. IEEE, 2015. ISBN 978-1-4799-6923-4. doi: 10.1109/ICRA.2015.7139535. URL <https://doi.org/10.1109/ICRA.2015.7139535>.
- [4] Jennifer E. King, Marco Cognetti, and Siddhartha S. Srinivasa. Rearrangement planning using object-centric and robot-centric action spaces. In *2016 IEEE International Conference on Robotics and Automation, ICRA, Stockholm, Sweden, May 16-21, 2016*, pages 3940–3947, 2016.
- [5] J. E. King, V. Ranganeni, and S. S. Srinivasa. Unobservable monte carlo planning for nonprehensile rearrangement tasks. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4681–4688, May 2017. doi: 10.1109/ICRA.2017.7989544.
- [6] Lerrel Pinto, Aditya Mandalika, Brian Hou, and Siddhartha Srinivasa. Sample-efficient learning of nonprehensile manipulation policies via physics-based informed state distributions. *CoRR*, abs/1810.10654, 2018. URL <http://arxiv.org/abs/1810.10654>.
- [7] Weihao Yuan, Johannes A. Stork, Danica Kragic, Michael Yu Wang, and Kaiyu Hang. Rearrangement with nonprehensile manipulation using deep reinforcement learning. In *2018 IEEE International Conference on Robotics and Automation, ICRA 2018, Brisbane, Australia, May 21-25, 2018*, pages 270–277, 2018. doi: 10.1109/ICRA.2018.8462863. URL <https://doi.org/10.1109/ICRA.2018.8462863>.
- [8] Matthew T. Mason. Mechanics and planning of manipulator pushing operations. *The International Journal of Robotics Research*, 5(3):53–71, 1986.
- [9] R. Tedrake N. Fazeli, R. Kolbert and A. Rodriguez. Parameter and contact force estimation of planar rigid-bodies undergoing frictional contact. *IJRR*, 2017.
- [10] Y. Lin T. Lozano-Perez L. Kaelbling P. Isola M. Bauza, F. Alet and A. Rodriguez. Omnipush: accurate, diverse, real-world dataset of pushing dynamics with rgbd images. In *IROS*, 2019.
- [11] Jiaji Zhou, Matthew T Mason, Robert Paolini, and Drew Bagnell. A convex polynomial model for planar sliding mechanics: theory, application, and experimental validation. *The International Journal of Robotics Research*, 37(2-3):249–265, 2018.
- [12] Tom Erez, Yuval Tassa, and Emanuel Todorov. Simulation tools for model-based robotics: Comparison of bullet, havok, mujoco, ODE and physx. In *IEEE ICRA*, 2015.
- [13] DART physics egnine. [Online]. Available: <http://dartsim.github.io>.
- [14] PhysX physics engine. [Online]. Available: www.geforce.com/hardware/technology/physx.
- [15] Bullet physics engine. [Online]. Available: [www.bulletphysics.org](http://bulletphysics.org).
- [16] Open dynamics engine. [Online]. Available: <http://ode.org>.
- [17] Mehmet Dogar, Kaijen Hsiao, Matei Ciocarlie, and Siddhartha Srinivasa. Physics-Based Grasp Planning Through Clutter. In *R:SS*, July 2012.
- [18] K. M. Lynch and M. T. Mason. Stable pushing: Mechanics, control- lability, and planning. *International Journal of Robotics Research*, 18, 1996.
- [19] Jonathan Scholz, Martin Levihn, Charles L. Isbell, and David Wingate. A Physics-Based Model Prior for Object- Oriented MDPs. In *ICML*, 2014.
- [20] Jiaji Zhou, Robert Paolini, J. Andrew Bagnell, and Matthew T. Mason. A convex polynomial force-motion model for planar sliding: Identification and application. In *ICRA*, pages 372–377, 2016.
- [21] Pieter Abbeel, Morgan Quigley, and Andrew Y Ng. Using inaccurate models in reinforcement learning. In *Proc. of ICML*. ACM, 2006.
- [22] Abdeslam Boularias, James Andrew Bagnell, and Anthony Stentz. Learning to manipulate unknown objects in clutter by reinforcement. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, January 25-30, 2015, Austin, Texas, USA.*, pages 1336–1342, 2015. URL <http://www.aaai.org/ocs/index.php/AAAI/AAAI15/paper/view/9360>.
- [23] F. Hogan* M. Bauza* and A. Rodriguez. A data-efficient approach to precise and controlled pushing. In *CoRL*, 2018.
- [24] M. Deisenroth, C. Rasmussen, and D. Fox. Learning to Control a Low-Cost Manipulator using Data-Efficient Reinforcement Learning. In *R:SS*, 2011.
- [25] Roberto Calandra, André Seyfarth, Jan Peters, and Marc P. Deisenroth. Bayesian optimization for learning gaits under uncertainty. *Annals of Mathematics and Artificial Intelligence (AMAI)*, 76(1):5–23, 2016. ISSN 1573-7470.
- [26] Alonso Marco, Felix Berkenkamp, Philipp Hennig, Angela P. Schoellig, Andreas Krause, Stefan Schaal, and Sebastian Trimpe. Virtual vs. real: Trading off simulations and physical experiments in reinforcement learning with bayesian optimization. In *ICRA*, pages 1557–1563, 2017.
- [27] Somil Bansal, Roberto Calandra, Ted Xiao, Sergey Levine, and Claire J Tomlin. Goal-driven dynamics learning via bayesian optimization. In *IEEE CDC*, 2017.
- [28] Rémi Pautrat, Konstantinos Chatzilygeroudis, and Jean-

Baptiste Mouret. Bayesian optimization with automatic prior selection for data-efficient direct policy search. *arXiv preprint arXiv:1709.06919*, 2017.

[29] A. Sintov, A. S. Morgan, A. Kimmel, A. M. Dollar, K. E. Bekris, and A. Boularias. Learning a state transition model of an underactuated adaptive hand. *IEEE Robotics and Automation Letters*, 4(2):1287–1294, 2019.

[30] Junhyuk Oh, Xiaoxiao Guo, Honglak Lee, Richard L. Lewis, and Satinder P. Singh. Action-conditional video prediction using deep networks in atari games. In *NIPS*, pages 2863–2871, 2015.

[31] Silvia Chiappa, Sébastien Racaniere, Daan Wierstra, and Shakir Mohamed. Recurrent environment simulators. *arXiv preprint arXiv:1704.02254*, 2017.

[32] Chelsea Finn and Sergey Levine. Deep visual foresight for planning robot motion. *ICRA 2017*.

[33] Chelsea Finn, Xin Yu Tan, Yan Duan, Trevor Darrell, Sergey Levine, and Pieter Abbeel. Deep spatial autoencoders for visuomotor learning. *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pages 512–519, 2016.

[34] N. Fazeli KT. Yu, M. Bauza and A. Rodriguez. More than a million ways to be pushed: A high-fidelity experimental dataset of planar pushing. In *IROS*, pages 30–37, 2016.

[35] F. Hogan N. Doshi and A. Rodriguez. Hybrid differential dynamic programming for planar manipulation primitives. In *ICRA*, 2020.

[36] Kevin M. Lynch and Matthew T. Mason. Stable pushing: Mechanics, controllability, and planning. *The International Journal of Robotics Research*, 15(6):533–556, 1996.

[37] Kevin M. Lynch. Estimating the friction parameters of pushed objects. In *1993 International Conference on Intelligent Robots and Systems*, 1993 International Conference on Intelligent Robots and Systems, pages 186–193, 12 1993.

[38] Robert D. Howe and Mark R. Cutkosky. Practical force-motion models for sliding manipulation. *The International Journal of Robotics Research*, 15(6):557–572, 1996.

[39] Mehmet Remzi Dogar and Siddhartha S. Srinivasa. Push-grasping with dexterous hands: Mechanics and a method. *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2123–2130, 2010.

[40] Jiaji Zhou, Robert Paolini, J. Andrew Bagnell, and Matthew T. Mason. A convex polynomial force-motion model for planar sliding: Identification and application. In *2016 IEEE International Conference on Robotics and Automation, ICRA 2016, Stockholm, Sweden, May 16-21, 2016*, pages 372–377, 2016.

[41] Jiaji Zhou, Yifan Hou, and Matthew T. Mason. Pushing revisited: Differential flatness, trajectory planning, and stabilization. *I. J. Robotics Res.*, 38(12-13), 2019.

[42] Jiaji Zhou, Matthew T. Mason, Robert Paolini, and Drew Bagnell. A convex polynomial model for planar sliding mechanics: theory, application, and experimental validation. *I. J. Robotics Res.*, 37(2-3):249–265, 2018.

[43] Tsuneo Yoshikawa and Masamitsu Kurisu. Identification of the center of friction from pushing an object by a mobile robot. *Proceedings IROS '91:IEEE/RSJ International Workshop on Intelligent Robots and Systems '91*, pages 449–454 vol.2, 1991.

[44] Shaojun Zhu, Andrew Kimmel, Kostas Bekris, and Abdeslam Boularias. Fast model identification via physics engines for improved policy search. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI), Stockholm, Sweden*, 2018.

[45] Changkyu Song and Abdeslam Boularias. Identifying mechanical models of unknown objects with differentiable physics simulations. In *Proceedings of the 2020 Learning for Dynamics and Control Conference (L4DC), Berkeley, California*, 2020, 2020.

[46] Jonas Degrave, Michiel Hermans, Joni Dambre, and Francis Wyffels. A differentiable physics engine for deep learning in robotics. *CoRR*, abs/1611.01652, 2016.

[47] Rami Al-Rfou, Guillaume Alain, Amjad Almahairi, Christof Angermüller, Dzmitry Bahdanau, Nicolas Ballas, Frédéric Bastien, Justin Bayer, Anatoly Beilikov, Alexander Belopolsky, Yoshua Bengio, Arnaud Bergeron, James Bergstra, Valentin Bisson, Josh Bleecher Snyder, Nicolas Bouchard, Nicolas Boulanger-Lewandowski, Xavier Bouthillier, Alexandre de Brébisson, Olivier Breuleux, Pierre Luc Carrier, Kyunghyun Cho, Jan Chorowski, Paul F. Christiano, Tim Cooijmans, Marc-Alexandre Côté, Myriam Côté, Aaron C. Courville, Yann N. Dauphin, Olivier Delalleau, Julien Demouth, Guillaume Desjardins, Sander Dieleman, Laurent Dinh, Melanie Ducoffe, Vincent Dumoulin, Samira Ebrahimi Kahou, Dumitru Erhan, Ziye Fan, Orhan Firat, Mathieu Germain, Xavier Glorot, Ian J. Goodfellow, Matthew Graham, Çağlar Gülcöhre, Philippe Hamel, Iban Harlouchet, Jean-Philippe Heng, Balázs Hidasi, Sina Honari, Arjun Jain, Sébastien Jean, Kai Jia, Mikhail Korobov, Vivek Kulkarni, Alex Lamb, Pascal Lamblin, Eric Larsen, César Laurent, Sean Lee, Simon Lefrançois, Simon Lemieux, Nicholas Léonard, Zhouhan Lin, Jesse A. Livezey, Cory Lorenz, Jeremiah Lowin, Qianli Ma, Pierre-Antoine Manzagol, Olivier Mastropietro, Robert McGibbon, Roland Memisevic, Bart van Merriënboer, Vincent Michalski, Mehdi Mirza, Alberto Orlandi, Christopher Joseph Pal, Razvan Pascanu, Mohammad Pezeshki, Colin Raffel, Daniel Renshaw, Matthew Rocklin, Adriana Romero, Markus Roth, Peter Sadowski, John Salvatier, François Savard, Jan Schlüter, John Schulman, Gabriel Schwartz, Iulian Vlad Serban, Dmitriy Serdyuk, Samira Shabanian, Étienne Simon, Sigurd Speckermann, S. Ramana Subramanyam, Jakub Sygnowski, Jérémie Tanguay, Gijs van Tulder, Joseph P. Turian, Sebastian Urban, Pascal Vincent, Francesco Visin, Harm de Vries, David Warde-Farley, Dustin J. Webb, Matthew Willson, Kelvin Xu, Lijun Xue, Li Yao, Saizheng Zhang, and Ying Zhang. Theano: A

python framework for fast computation of mathematical expressions. *CoRR*, abs/1605.02688, 2016.

[48] Alina Kloss, Stefan Schaal, and Jeannette Bohg. Combining learned and analytical models for predicting action effects. *CoRR*, abs/1710.04102, 2017.

[49] Connor Schenck and Dieter Fox. Spnets: Differentiable fluid dynamics for deep neural networks. *CoRR*, abs/1806.06094, 2018.

[50] Marc Toussaint, Kelsey R Allen, Kevin A Smith, and Josh B Tenenbaum. Differentiable physics and stable modes for tool-use and manipulation planning. In *Proc. of Robotics: Science and Systems (R:SS 2018)*, 2018.

[51] Filipe de Avila Belbute-Peres and Zico Kolter. A modular differentiable rigid body physics engine. In *Deep Reinforcement Learning Symposium, NIPS*, 2017.

[52] Igor Mordatch, Emanuel Todorov, and Zoran Popović. Discovery of complex behaviors through contact-invariant optimization. *ACM Trans. Graph.*, 31(4):43:1–43:8, July 2012. ISSN 0730-0301. doi: 10.1145/2185520.2185539. URL <http://doi.acm.org/10.1145/2185520.2185539>.

[53] Stefan Hinterstoisser, Vincent Lepetit, Slobodan Ilic, Stefan Holzer, Gary Bradski, Kurt Konolige, and Nassir Navab. Model based training, detection and pose estimation of texture-less 3d objects in heavily cluttered scenes. In Kyoung Mu Lee, Yasuyuki Matsushita, James M. Rehg, and Zhanyi Hu, editors, *Computer Vision – ACCV 2012*, pages 548–562, 2013.

[54] Michael Bradley Cline. *Rigid body simulation with contact and constraints*. PhD thesis, University of British Columbia, 2002.

[55] Filipe de A. Belbute-Peres, Kevin A. Smith, Kelsey R. Allen, Joshua B. Tenenbaum, and J. Zico Kolter. End-to-end differentiable physics for learning and control. In *Proceedings of the 32Nd International Conference on Neural Information Processing Systems, NIPS’18*, pages 7178–7189, USA, 2018. Curran Associates Inc. URL <http://dl.acm.org/citation.cfm?id=3327757.3327820>.

[56] J. Mattingley and S. Boyd. CVXGEN: A code generator for embedded convex optimization. *Optimization and Engineering*, 12(1):1–27, 2012.

[57] Kaiyu Hang, Andrew Morgan, and Aaron Dollar. Pre-grasp sliding manipulation of thin objects using soft, compliant, or underactuated hands. *IEEE Robotics and Automation Letters*, PP:1–1, 01 2019. doi: 10.1109/LRA.2019.2892591.

[58] Nikolaus Hansen. The cma evolution strategy: a comparing review. *Towards a new evolutionary computation*, pages 75–102, 2006.

[59] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.