

Understanding Physical Effects for Effective Tool-use

Zeyu Zhang Ziyuan Jiao Weiqi Wang Yixin Zhu Song-Chun Zhu Hangxin Liu

Abstract—We present a robot learning and planning framework that produces an effective tool-use strategy with the least joint efforts, capable of handling objects different from training. Leveraging a Finite Element Method (FEM)-based simulator that reproduces fine-grained, continuous visual and physical effects given observed tool-use events, the essential physical properties contributing to the effects are identified through the proposed Iterative Deepening Symbolic Regression (IDSR) algorithm. We further devise an optimal control-based motion planning scheme to integrate robot- and tool-specific kinematics and dynamics to produce an effective trajectory that enacts the learned properties. In simulation, we demonstrate that the proposed framework can produce more effective tool-use strategies, drastically different from the observed ones in two exemplar tasks.

Index Terms—Tool use; Symbolic regression; Finite element method

I. INTRODUCTION

A Robot extends its capability to a broader range of tasks by using tools. Unlike treating a tool as a part of the end-effector that commonly appears in industrial settings [1, 2], researchers have proposed various learning-based approaches that empower more adept tool-use behaviors. However, existing learning objectives either focus on low-level motions [3, 4] without an explicit understanding of the tasks or on higher-level concepts with simplified motion patterns [5–7]. As a result, robots are still far from producing situational tool-use strategies: Given a set of objects (typical tools or canonical objects), which one would be the best to accomplish the task? Once an object is chosen as the tool, how to efficiently use it given robot- and tool-specific kinematics and dynamics?

To tackle these challenges, we propose an integrated learning and planning framework wherein robots understand and produce effective tool-use strategies by reasoning about the essential physical properties that contribute to the success of the task. Fig. 1 shows an overview of our integrated framework.

Manuscript received February 24, 2022; accepted: June 26, 2022. Date of publication July xx, 2022; date of current version July xx, 2022. This letter was recommended for publication by Editor Hanna Kurniawati upon evaluation of the Associate Editor and Reviewers’ comments. (Zeyu Zhang and Ziyuan Jiao contributed equally to this work.) (Corresponding author: Hangxin Liu.)

Zeyu Zhang, Ziyuan Jiao, and Song-Chun Zhu are with the Beijing Institute for General Artificial Intelligence, (BIGAI), Beijing 100080, China, and with the Center for Vision, Cognition, Learning, and Autonomy (VCLA), Statistics Department, UCLA (e-mail: zeyuzhang@ucla.edu; zyjiao@ucla.edu; sczhu@stat.ucla.edu).

Weiqi Wang is with the Center for Vision, Cognition, Learning, and Autonomy (VCLA), Statistics Department, UCLA (e-mail: weiqi.wang@ucla.edu).

Yixin Zhu is with the Institute for Artificial Intelligence, Peking University, Beijing 100871, China (e-mail: yixin.zhu@pku.edu.cn).

Song-Chun Zhu is also with the Institute for Artificial Intelligence, Peking University, Beijing 100871, China.

Hangxin Liu is with the Beijing Institute for General Artificial Intelligence, (BIGAI), Beijing 100080, China (e-mail: liuhx@bigai.ai).

This letter has supplementary downloadable material available at <https://doi.org/XXXX/XXXXXXX>, provided by the authors.

Digital Object Identifier (DOI): see top of this page.

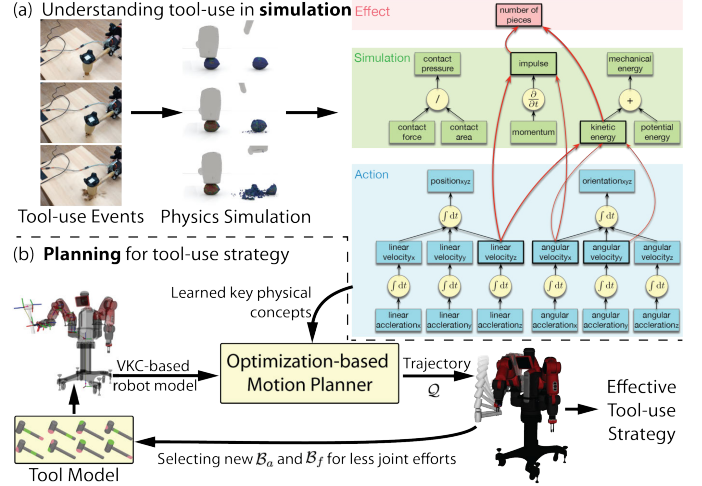


Fig. 1: **Overview of the proposed framework.** (a) After observing tool-use events, we learn the essential physical properties involved in the processes from the effects reproduced by physics-based simulation. (b) The learned results are formulated into a motion planning scheme to produce various strategies to use an object, and the most effective strategy with minimal joint efforts among others is selected.

Compared to prior arts in robotics literature, our framework identifies the invariant learning objective of tool-uses at a more fundamental level; instead of using pure vision-based methods [8, 9], our framework focuses on the physical effects produced by the tool and learns to recognize the essential physical properties in accomplishing the task. Specifically, we adopt a state-of-the-art Finite Element Method (FEM) [10] to simulate how both visual and physical effects evolve over time (e.g., stress, energy, contact) in a continuous manner. A symbolic regression-based Iterative Deepening Symbolic Regression (IDSR) algorithm is devised to trace the set of physical properties produced by the simulator and to efficiently identify how much each property contributes to the effect.

Next, we formulate the learned results into an optimal control-based motion planning scheme that allows the robot to generate various tool-use strategies whose efficiency is evaluated by joint efforts. To ease the motion planning problem and make the scheme more generic (i.e., handle robots with different morphology, tools in diverse shapes, and various ways to operate tools), we introduce a Virtual Kinematic Chain (VKC) perspective [11, 12] that treats the tool as an additional link of the robot and integrates their kinematic and dynamic properties as a whole in motion planning.

In two exemplar tasks—cracking walnut and cutting carrot, we demonstrate that the proposed learning and planning framework can (i) identify the essential physical properties significant to the success of the task and (ii) produce an effective tool-use strategy that emulates the essential properties while minimizing joint efforts using seen and unseen objects as tools. As a result, the proposed framework allows the robot

to better understand the physical environment by leveraging physics-based simulations and become more competent in bootstrapping novel (*i.e.*, not observed) tool-use strategies.

A. Related Work

Learning tool-use involves several cognitive and intelligent processes, challenging even for humans. Replicating such a skill set at the full spectrum is thus difficult, and existing literature mainly focuses at one of three different levels. Low-level **planning and control** methods track desired tool-use trajectories with impedance control [1], alter force and motion constraints at different stages [2], or apply learning-based control [3,4]; robust execution is of the central interest. At mid-level, various **intermediate representations** are identified for better understanding tool-uses, such as keypoints [6,7], primitive parts [13–16], and kinematic models [12,17]. Although introducing these representations facilitates learning more diverse tool-use skills, they are still restricted to the geometric association between shapes and task specifications. To capture high-level **concepts** embedded in tool-uses, researchers adopt task and motion planning [18], functionality and affordance [5,19,20], causality [21], and common-sense [22,23], achieving better generalization capabilities. Empowered by physics-based simulation, we advance this line of work by taking all three views into account: (i) learning related physical properties as the concepts from the tasks at the high-level, (ii) integrating tool's properties to robots by adopting VKC as the intermediate representation at the mid-level, and (iii) planning tool-use strategies via optimal control at the low-level.


Recently, physics-based **simulation** significantly facilitates various robotics tasks, *e.g.*, Liu *et al.* simulate forces to bridge human and robot's embodiments [24], Kennedy *et al.* plan liquid pouring [25], Matl *et al.* infer granular materials' properties [26], Hahn *et al.* approximate soft objects' motions by estimating visco-elastic parameters [27], Geilinger *et al.* develop simulation framework for rigid and soft bodies with fictional contact to promote robot locomotion [28], Li *et al.* improve UAV designs [29], and Heiden *et al.* optimize robot's cutting and slicing motions [30]. Though sharing a similar spirit, the FEM simulator adopted in the paper [10] is designed to produce a wider range of physical properties for robot learning instead of optimizing for dedicated applications.

B. Overview


The remainder of this paper is organized as follows. Section II formally presents the problem definition and our optimal planning framework. Section III describes the simulation setup and the learning of physical properties. We demonstrate our framework's efficacy in Section IV and conclude the paper with discussions in Section V.

II. PROBLEM DEFINITION

We define a tool-use strategy $\mathcal{S} = (\mathcal{B}_a, \mathcal{B}_f, \mathcal{Q})$ by (i) an affordance basis \mathcal{B}_a to be grasped by the robot gripper, (ii) a functional basis \mathcal{B}_f to act on the target object, and (iii) a trajectory \mathcal{Q} directing the functional basis to move towards the target object. Given a tool partitioned into a set of sub-meshes $\{\mathcal{M}_i\}$, a sampling process assigns one sub-mesh as



(a) Various partitions of a hammer. The green regions denote affordance bases \mathcal{B}_a as potential areas for grasps, whereas the red regions denote functional bases \mathcal{B}_f as a candidate area to act on the target object. Surface normals calculated at the regions' center are the directions to grasp or act.



(b) VKCs can be constructed by consolidating the kinematic models with dynamic information of the robot and the tool.

Fig. 2: A VKC perspective that promotes motion planning. (a) Given a sampled bases combination (highlighted in red box) of \mathcal{B}_a and \mathcal{B}_f , (b) a VKC is constructed by assigning a virtual joint between the robot's gripper and the \mathcal{B}_a , and \mathcal{B}_f becomes the new end-effector. This VKC conversion and construction supports efficient and optimal motion planning to produces proper tool-use trajectories by taking both kinematic and dynamic factors into account.

\mathcal{B}_a and another as \mathcal{B}_f , as illustrated in Fig. 2a. The surface normal vector \mathbf{n} at the center of the corresponding sub-mesh indicates the direction for the robot's gripper to approach or for the tool to act on the target object. Assuming the robot can firmly grasp the tool at \mathcal{B}_a , generating a tool-use strategy \mathcal{S} can be formulated as a motion planning problem that finds a collision-free trajectory $\mathcal{Q} = \mathbf{q}_{1:T}$ given \mathcal{B}_a and \mathcal{B}_f .

A. VKC for Motion Planning

The theory of body schema [31] suggests that humans can extend the body's representation to incorporate an external object and treat it as part of their limb for efficient motions and manipulations, which plays a significant role in tool-use [32]. This idea has been introduced to the robotics community to represent robot structures and guide robot's behaviors [33]. Recent modeling approaches adopting VKC [11,12] provide an effective means to model robot tool-uses: By inserting a virtual joint between robot end-effector and tool's \mathcal{B}_a , the kinematics and dynamics of the robot and the tool are integrated, and their motions are planned collectively, resulting in more coordinated motion and higher planning success rate [11,12].

We first adopt an articulated body algorithm [34] to compute the forward dynamics analytically for the constructed VKC. Next, the objective of the motion planning for robot tool-use is formulated by optimal control:

$$\min_{x,u,T} \int_0^T L(x(t), u(t)) dt + \phi(x(T)) \quad (1)$$

$$L(x(t), u(t)) = \dot{q}^\top W_{\dot{q}} \dot{q} + u^\top W_u u, \quad (2)$$

$$\phi(x(T)) = T, \quad T \in \mathbb{R}^+, \quad (3)$$

where $W_{\dot{q}}$ and W_u are weight matrices for joint velocities and joint torques, $u : \mathbb{R} \rightarrow \mathbb{R}^n$ the control input consisted of joint torques, $\phi(x(T))$ measures the quality of the terminal state, particularly, we penalize the total elapsed time T . $x : \mathbb{R} \rightarrow \mathbb{R}^{2n+2m+1}$ is the state variable, which includes (i) joint positions q and velocities \dot{q} of a manipulator with n Degree of Freedom (DoF), (ii) q and \dot{q} of underactuated joints in a tool, and (iii) the virtual joint at the grasp point with a total of m DoF. Eq. (1) penalizes the weighted quadratic cost on joint velocity and torques for the entire trajectory and the total elapsed time.

During the motion planning, we further impose several safety constraints:

$$\dot{x}(t) = f(x(t), u(t)), \quad t \in [0, T] \quad (4)$$

$$g(x(t), u(t)) = 0, \quad t \in [0, T] \quad (5)$$

$$x_{lb} \leq x(t) \leq x_{ub}, \quad t \in [0, T] \quad (6)$$

$$u_{lb} \leq u(t) \leq u_{ub}, \quad t \in [0, T] \quad (7)$$

where Eq. (4) is the system dynamics, Eq. (5) is a task-dependent constraint for tool-use, Eq. (6) and Eq. (7) are safety constraints that bound the robot workspace and control limit.

B. Goal Specification

Formally, the goal for a tool-use is expressed as:

$$f_{\text{task}}(n_T(\mathcal{G}), \text{VKC}) \Rightarrow g(\cdot), \quad (8)$$

where $n_T(\mathcal{G})$ is a set of physical properties that are essential to the task, to be detailed in Section III. f_{task} maps these physical properties and VKCs (as constructed in Fig. 2b) to a constraint function g for motion planning. The intuition is for the robot to emulate those essential physical properties in execution while considering the robot and tool's kinematics and dynamics.

To be more specific, let us take the walnut cracking task as an example. Given the goal position where the contact occurs p_g , the tool should act on the target object with a velocity vector \mathbf{v}_{tool} and the tool's orientation \mathbf{d}_{tool} (to be detailed in Section III-D), both represented in world frame. Eq. (9) first finds a possible robot goal pose q_g through solving inverse kinematics to regulate the tool's orientation when contacting the target object:

$$\frac{f_z(q_g) \cdot \mathbf{v}_{\text{tool}}}{\|f_z(q_g)\| \cdot \|\mathbf{v}_{\text{tool}}\|} = \cos(\mathbf{d}_{\text{tool}}), \quad (9)$$


where $f_z : \mathbb{R}^n \rightarrow \mathbb{R}^3$ finds the surface normal of \mathcal{B}_f . Then, the goal joint velocities are computed by:

$$\dot{q}_g = J_{\text{VKC}}^\top (J_{\text{VKC}} J_{\text{VKC}}^\top)^{-1} \mathbf{v}_{\text{tool}}, \quad (10)$$

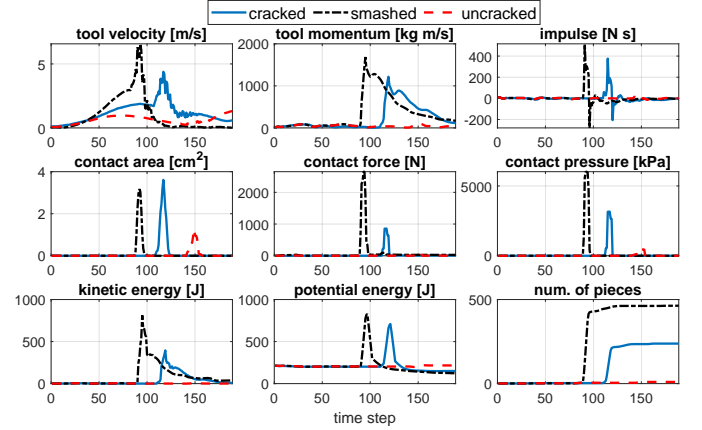
where J_{VKC} is the geometric Jacobian from the robot's base frame to the tool's functional basis at the joint position q_g . Finally, Eq. (8) can be expressed in terms of joint velocity w.r.t. two constraint functions q_g and \dot{q}_g :

$$g_q(x(t), u(t)) = x_q(T) - q_g = 0 \quad (11)$$

$$g_{\dot{q}}(x(t), u(t)) = x_{\dot{q}}(T) - \dot{q}_g = 0 \quad (12)$$



(a) Comparisons between the effect produced by experiments (left) and simulations (right) with different fracturing limits and Young's modulus.



(b) Fine-grained physical properties evolved in time reflect the different physical effects among uncracked, cracked, and smashed.

Fig. 3: **Examples of qualitative and quantitative results produced by the FEM-based simulation.** (a) We qualitatively choose the parameters (in red) that best match the final effect of observed tool-use events. (b) By adopting an FEM-based simulator, the data collection process records physical properties evolved in time.

III. SIMULATION AND LEARNING

This section starts with the technical background of physics-based simulation, followed by how it reproduces fine-grained physical properties and helps understand tool-uses events.

A. Background

Solid simulation approximates objects' physical status. It is oftentimes formulated with the Finite Element Method (FEM) [35], which discretizes each object into small elements with a discrete set of sample points as the degree-of-freedom. Mass and momentum conservation equations are discretized on the mesh and integrated overtime to capture the dynamics. This paper utilizes an Incremental Potential Contact (IPC) handling method [10, 36, 37], a state-of-the-art FEM-based simulator, to address the difficulty of simulating non-smooth contacts between a tool and a target object. To further support object fracture during tool-uses, our simulator measures the displacement of every pair of points that both connect to all the nodes of a triangle on the mesh. If the displacement relative to their original distance exceeds a certain strain threshold, we mark the triangle in-between as separated.

B. Reproducing Effect

To produce similar effects in the simulation that sufficiently match those in the physical world, some parameters governing an object's material property need to be appropriately set. In particular, Young's modulus reflects the object's stiffness—the stiffer the object is, the harder for it to deform or fracture, and fracture limit determines the number and the size of segments

a large piece will fracture into. Fig. 3a qualitatively shows how the resulting effects vary given different Young's modulus and fracture limit. These two parameters are calibrated such that the simulated effects match the observation in physical world.

We use two VIVE Trackers to capture the tool-use events. One to track the movement of the tool (e.g., a hammer), and another placed on the table serving as the reference point for the target object (e.g., a walnut). Both VIVE trackers are calibrated such that their relative poses and captured trajectories are expressed in the same coordinate frame, with a time step of the inverse of their sampling frequency. The meshes of the target object and the tool are pre-scanned using an RGB-D camera. Combining scanned meshes and captured trajectories, we can fully reconstruct an observed event in both space and time and further simulate the effects of the target object both visually and physically. Examples of keyframes of the collected data with corresponding simulated results are shown in Fig. 1a. Fig. 3b visualizes the continuous numerical values of some notable physical properties obtained from simulation. Of note, capturing how the object's status changes and its physical properties evolve over time is highly challenging, if not impossible, using visual information alone.

C. Learning Essential Physical Properties

We quantize the space of physical properties into three different levels; see Fig. 1a for an illustration. (i) **Action** (in blue) includes the trajectory data (position and orientation) directly observed in tool-use events and its velocity and acceleration calculated by finite-difference; these properties are usually controllable by robots. (ii) **Simulation** (in green) includes the physical properties estimated by the simulation given the observed Action. (iii) **Effect** (in red) includes the physical properties representing the tool-use effect. In the case of cracking and cutting tasks, we represent the effect by the number of pieces the target object transforms into.

Given various physical properties estimated and reproduced by the simulation, a robot has to learn how much these properties contribute to the success of the task and distill knowledge at all three levels, such that it can plan its motion in new and even unseen scenarios. To encode the connections across all three levels of physical properties, we propose to learn a **Physical Relation Graph (PRG)** representation through symbolic regression [38, 39]. Specifically, setting the Effect as the target variable y , the symbolic regression is tasked to find a valid expression of y using the set of given variables \mathbf{x} in Simulation and Action: $y = f(\mathbf{x})$. To prevent over-fitting, we further balance the expression's complexity (i.e., how many physical properties are involved) and accuracy (i.e., how well it expresses the target variable). As such, the relations in PRG is sparse and only involve a small subset of the variables that succinctly express the target variable.

Typical symbolic regression problems oftentimes have a large search space. To tackle it, we devise an IDSR algorithm, a variant of symbolic regression, that utilizes the hierarchical information among physical properties at each level to prune the searching space. Specifically, as illustrated in Fig. 4a, typical symbolic regression algorithms directly explore the entire domain with all variables, whereas the proposed IDSR would iteratively deepen the domain based on the hierarchy among

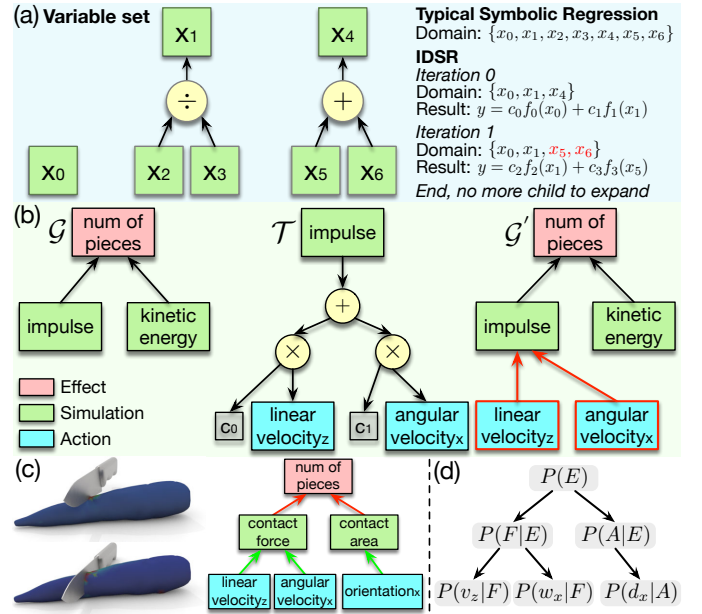


Fig. 4: **Learning relations among physical properties using IDSR.** (a) An example of deepening the variable domain. Since x_4 is not included in the resulted expression in the iteration 0, it is thus removed, and its children are added to the domain in the next iteration. (b) An example of constructing PRG. \mathcal{G}' is the updated graph after inserting the expression \mathcal{T} into the previous graph \mathcal{G} ; newly added nodes and edges are highlighted in red. (c) The PRG constructed for the cutting task. (d) Inferring necessary values at the Action level for the goal specification in planning.

them. If one variable is not selected after an iteration, the domain will replace it with its child variables and reiterate the algorithm, and the resulting expression will only be updated if those child variables play a more significant role. This process continues until all the variables from one level in the domain are selected, or non-selected variables have no child. Algorithm 1 outlines the procedure.

In the case of cracking a walnut (see Fig. 4b), only after the set of relations between *Effect* and *Simulation* is explored would the algorithm subsequently identify the set of relations between *Simulation* and *Action*, expanding the PRG. As a result, this algorithm design saves the memory compared to conventional symbolic regression algorithms while preserving the full capability of distilling the essential relations among variables. The sub-graph highlighted in red in Fig. 1a shows the learned PRG of cracking a walnut, wherein the edge thicknesses are proportional to the physical properties' contribution to the effect. In another task of cutting a carrot by half using a knife (see Fig. 4c), the IDSR algorithm identifies the *contact area* governed by the *orientation* as an essential physical property, since the deviation from a proper orientation range may lead to the increment of contact area.

D. Reasoning about Goal Specification

The \mathcal{G} identified by IDSR is still insufficient to support the proposed planning scheme because it only deduces the relation among those physical properties in a symbolic level, i.e. *velocity* for the task of cracking a walnut, and both *velocity* and *orientation* \mathbf{d}_{tool} for cutting as shown in Fig. 4bc. The

corresponding values of \mathbf{v}_{tool} and/or \mathbf{d}_{tool} applicable for robot planning is not determined yet.

To address this issue, we devise a sequential inference pipeline based on learned \mathcal{G} . As illustrated in Fig. 4d, by modeling the values of observed effect as a Gaussian distribution $P(E)$, a Gaussian Mixture Model (GMM) is learned to capture the joint probability between the effect and an identified physical property according to \mathcal{G} , e.g. $P(E, F)$ for effect and contact force in Fig. 4cd, using the EM algorithm [40]. Specifically, the mixture models are fitted on the data obtained from the simulator that reproduces human demonstrations. Next, given a desired effect, inferring specific values of *contact force* is performed by drawing samples from the distribution $P(F|E) = P(E, F)/P(E)$ [41], and a velocity in z direction v_z is subsequently obtained by sampling from $P(v_z|F)$ following the same protocol. Eventually, this process produces the necessary values at the Action level (\mathbf{v}_{tool} and \mathbf{d}_{tool}) as goal specifications for Eqs. (9) and (10).

IV. EXPERIMENTS

We conduct three sets of studies regarding different types of manipulators with various settings. Using a human arm model [42], we first validate that our planning scheme produces a feasible tool-use strategy identical to human choices; see Section IV-A. Next, we show that our proposed framework produces diverse tool-use strategies for Baxter arm and UR5 manipulator under different scenarios; the most effective ones in terms of least joint effort are demonstrated in Section IV-B. Finally, the produced strategies are fed to simulations for robot planning and execution; see Section IV-C. Experimental results verify that the framework indeed captures the essential physical properties, capable of converting these learned relations into goal specification, resulting in the success of motion planning and task completion.

In all experiments, we solve the motion planning problem defined in Section II by CasADi [43] with the OpenOCL [44] support. A tool-use is considered invalid if the planner cannot produce a feasible solution. We assume the manipulators' bases are fixed. The tool structures are scanned by an RGB-D camera and reconstructed into watertight meshes, and the

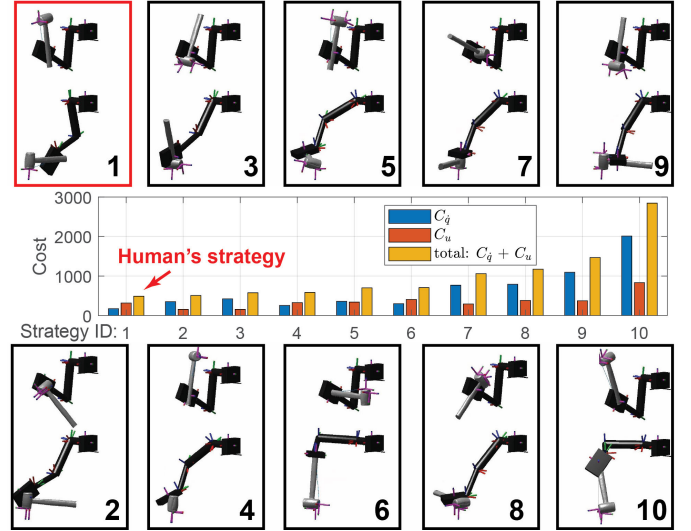


Fig. 5: **Different strategies of tool-use using an approximated human arm model.** \mathcal{B}_a s and \mathcal{B}_f s are sampled from partitioned regions on the hammer, and the trajectory \mathcal{Q} is produced by the optimal control using VKCs. The optimal strategy (in red) indeed follows human intuition of operating a hammer. C_q is the trajectory smoothness cost, and C_u is the joint torque effort cost.

tool's material is homogeneous. For fair comparisons, the target object (e.g., walnut) is placed at the same location within the operational space for each type of manipulator, and the initial pose of the manipulator is identical across all trials. In each trial, the target object has 1229 mesh vertices and is simulated for 200 time steps. The simulation runs on a 16-core AMD Ryzen 9 5950X machine and the average run time for one trial is 77.08 minutes with parallelization for the linear system computations [10]. Algorithm-wise parallelization for FEM still remains an open problem.

A. Validating Optimal Planning by Task Efficiency

In this experiment, we evaluate whether the optimal control-based planning scheme is effective by comparing the produced tool-use strategies with that of human's rational choice, which should be regarded as near-optimal. Using the human arm model [42] that consists of 7 DoFs (3 for shoulder, 2 for elbow, and 2 for wrist) with corresponding arm's physical properties (i.e., mass, inertia) measured by human subjects, we sample various combinations of \mathcal{B}_a and \mathcal{B}_f and produce the corresponding tool-use trajectories. Fig. 5 shows initial and final arm postures and their computed costs of replicated human tool-uses and nine examples of alternative solutions.

Our results show that Strategy 1 is the most efficient one. Compared with a conventional swinging action, holding hammerhead reduces the inertia compensated by actions, resulting in a lower joint torque effort costs C_u s in Strategy 2 and 3. However, the trajectory smoothness costs C_q s are higher as a larger acceleration is required to reach the goal velocity, making their total costs higher than the cost in Strategy 1. Since both Strategy 4 and 5 start from a similar \mathcal{B}_a as in Strategy 1 followed by a swinging trajectory, their C_u s are similar to that of Strategy 1; however, their C_q s are higher since their \mathcal{B}_f s do not well aligned with arm postures. Strategy 6 to 10 are some less typical examples with high C_u s and

Algorithm 1: IDSR

Data: Data samples: \mathcal{D} . Target variable: v_t . Variable set: \mathcal{V}
Result: Best matched expression tree: \mathcal{T}

```

1 Domain  $\leftarrow \{\text{AllRoots}(\mathcal{V})\}$  while not terminate do
2   terminate  $\leftarrow \text{True};$ 
   //Symbolic regression on Domain
3    $\mathcal{T} \leftarrow \text{SR}(\mathcal{D}, v_t, \text{Domain});$ 
4   diff  $\leftarrow \text{Domain} \setminus \mathcal{T}.\text{leaveSymbols}();$ 
   //Deepening the searching domain
5   foreach  $v$  in diff do
6     if  $v$  has child then
7       Domain.add( $v.\text{children}()$ );
8       Domain.remove( $v$ );
9       terminate  $\leftarrow \text{False};$ 
10    end
11  end
12 end
13 return  $\mathcal{T}$  //Return the latest  $\mathcal{T}$ 

```

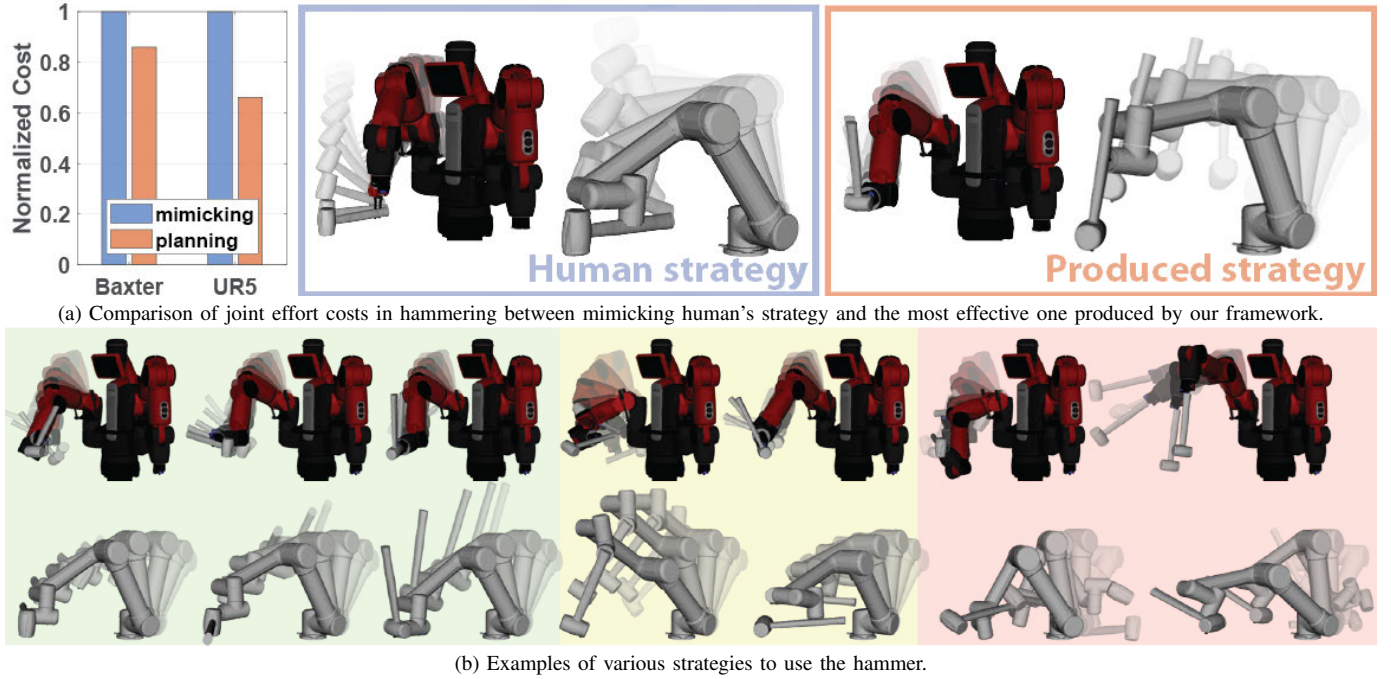


Fig. 6: **Generated various strategies in using a hammer.** (a) Given an inferred velocity vector acting on the walnut, the best tool-use strategy for each robot found by our framework is more efficient than simply mimicking human's strategy, indicated by lower cost. (b) Other strategies found by the proposed framework: low cost (in green), high cost (in yellow), and invalid with violation of constraints (in red).

C_q s; we seldom observe these strategies in real life. Together, these results indicate that our planning scheme can produce an efficient tool-use trajectory with underlying rationales akin to human tool-use behaviors, and thus we expect it to uncover similar insights into robot tool-uses.

B. Effective Tool-Uses

After validating our optimal control-based planning scheme, we test the efficacy of tool-use strategies derived from learned physical properties using two different robots (a Baxter robot and a UR5 manipulator) in two tasks (cracking and cutting).

Due to significant differences in kinematic structures, the observed human strategy of tool-uses may not be ideal for robots. In Fig. 6a, two robots first mimic human's strategy. Specifically, the robots select the observed human's B_a and B_f and mimic the observed trajectory Q by inverse kinematics to operate the hammer. The resulting costs are higher than those of the best strategy found by our framework; the ones found by the proposed framework are dramatically different but more effective for the robots. Fig. 6b further displays some other tool-use strategies with low-cost (effective), high-cost (ineffective), or are invalid by violating constraints.

Our framework is generic and generalizable to more challenging cases. It can further generate effective strategies using unconventional daily objects. The costs of operating those objects are ranked from low to high in Fig. 7a (Baxter) and Fig. 7b (UR5). The experiment reveals some objects (piler and wrench for Baxter, and axe and pan for UR5) are surprisingly more handy for robots compared with the hammer (indicated by the black bar). We further visualize the executed trajectories in Fig. 7c. Of note, the same pan is more suitable for UR5 as the cost of operating it is lower than using a

hammer but not that effective for Baxter. In comparison, the efficiency of using the rock and the toy (Psyduck) are similar for both robots. These results demonstrate that our learning and planning framework enables a situational tool-use skill for various robots.

In another task of cutting carrot, both robots do not perform well if concerned only about the *velocity* as they did in walnut-cracking; the target object will not align with the knife's blade properly as illustrated in Fig. 8a. By incorporating tool's *orientation* as uncovered in Fig. 4c, the robots overcome this deficiency and produce desired effects successfully; see Fig. 8b. Compared with the walnut-cracking task, the cutting task poses greater challenges in selecting unseen objects as tools since not all objects can lead to task completion; *i.e.*, one cannot use a hammer to successfully cut a carrot as a knife does. Yet in Fig. 8c, the result still demonstrates the robots' reasonable efforts in this difficult situation by choosing a sharp edge to contact with the object, showing that our framework successfully captures the essential physics in tool-uses and leverages them in producing its own strategies.

C. Testing Robot Tool-use in Simulation

Finally, we evaluate how well the best strategy found by the proposed framework (*e.g.*, produced strategies in Fig. 6a) can be executed in the simulator. This step is crucial as it separates the proposed framework from purely vision-based methods.

Since no existed work can solve the proposed task, we design a kinematic-based motion planner as a baseline that accounts for the physical properties involved in the task. In the case of the walnut-cracking task, the baseline needs to plan a trajectory that moves the functional basis of the tool to the center of the walnut while keeping its surface normal aligned

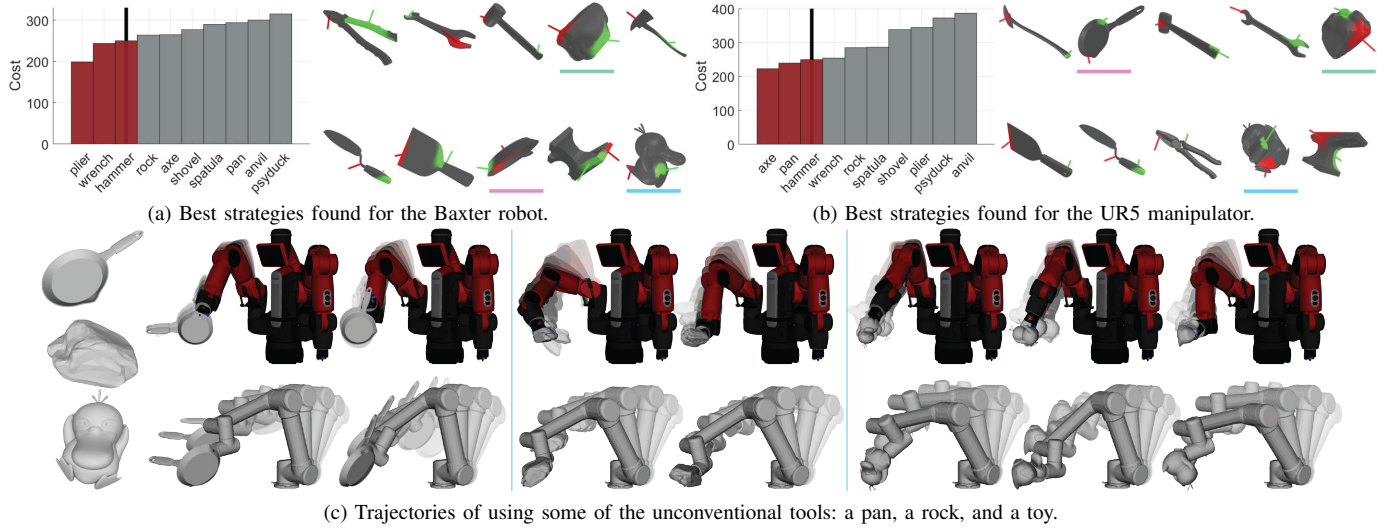


Fig. 7: **Effective tool-uses with unseen objects for the walnut-cracking task.** (a)(b) The best strategies (least cost) for ten different objects to crack a walnut use a Baxter robot and a UR5 manipulator, respectively. (c) Examples of valid trajectories of the Baxter robot (upper) and a UR5 manipulator (lower) using a pan (left), a piece of rock (middle), and a Psyduck toy (right).

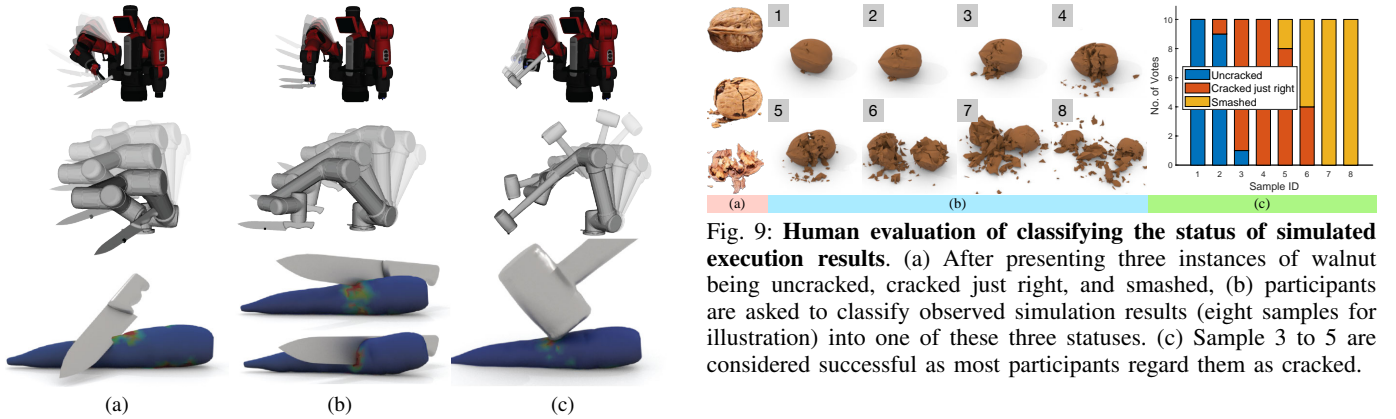


Fig. 8: **Tool-use strategies for cutting the carrot.** (a) Robots fail to accomplish the task without incorporating a tool's orientation. (b) The successful use of a knife requires incorporating orientation properties as learned in Fig. 4c. (c) Even using an object (a hammer) unsuitable for this task, our framework still produces an effective strategy by finding a tool orientation that minimize contact.

with the gravity direction. Fifty trials are simulated for both the Baxter robot and the UR5 manipulator using trajectories produced by the baseline and the proposed framework, and the parameters governing walnut's fracturing properties in each trial are set based on the values shown in Fig. 3a with a randomness of 10% for variations.

Due to the lack of quantitative evaluation of the performance of the walnut cracking task, we conducted a human study to compare the results between the baseline and the proposed framework. Ten participants were recruited online and asked to classify the total of 200 simulated execution results into one of the three statuses based on three instances shown in Fig. 9a. An execution is considered successful if more participants regard the walnut's status as cracked. Fig. 9bc show eight examples of the results based on the human study. The success rates are shown in Table I, demonstrating the necessity of understanding the physics in tool-use. Together,




Fig. 9: **Human evaluation of classifying the status of simulated execution results.** (a) After presenting three instances of walnut being uncracked, cracked just right, and smashed, (b) participants are asked to classify observed simulation results (eight samples for illustration) into one of these three statuses. (c) Sample 3 to 5 are considered successful as most participants regard them as cracked.

the results show the proposed framework indeed enables a better understanding of complex physical events that occurred during the tool-uses and successful productions of tool-use behaviors for robots.

V. CONCLUSION AND DISCUSSION

We presented a learning and planning framework for robots to understand the physics behind tool-use events and generate tool-use strategies suitable for the robots' own kinematics and dynamics. A physics-based FEM simulator was developed to generate physical properties in a continuous manner, from which we devised an IDSR algorithm to learn the essential properties critical to the success of the task. By formulating the learned properties into an optimal control-based motion planning scheme, our experiments demonstrated that the proposed framework allows robots to find tool-use strategies different from human demonstrations when handling seen and unseen objects, with better efficiency measured by least joint efforts.

TABLE I: Success Rate in Cracking Walnut in Simulator

Robot Type	Baseline	Proposed
Baxter	14%	62%
UR5	16%	52%

While our work is conducted in simulation, our planning scheme outputs torque commands that are possible for deployment on physical robots in the future. As grasping remains an unsolved problem, we plan to incorporate more sophisticated methods (e.g., [45]) to generate firm grasp configurations on the tool, such that we can produce more realistic and adaptive tool-uses. The reality gap is another major challenge to realize the physical deployment of the proposed framework. Physics-based simulation is difficult to tune or match the real world precisely. However, it is still a powerful tool for robot understanding and uncovering the task goal.

Acknowledgement: The authors would like to thank Dr. Minchen Li and Dr. Chenfanfu Jiang from the UCLA Mathematics Department for the simulation setup.

REFERENCES

- [1] H. Asada and Y. Asari, "The direct teaching of tool manipulation skills via the impedance identification of human motions," in *International Conference on Robotics and Automation (ICRA)*, 1988.
- [2] R. Holladay, T. Lozano-Pérez, and A. Rodríguez, "Force-and-motion constrained planning for tool use," in *International Conference on Intelligent Robots and Systems (IROS)*, 2019.
- [3] K. Kawaharazuka, K. Okada, and M. Inaba, "Adaptive robotic tool-tip control learning considering online changes in grasping state," *IEEE Robotics and Automation Letters (RA-L)*, 2021.
- [4] N. Saito, T. Ogata, S. Funabashi, H. Mori, and S. Sugano, "How to select and use tools?: Active perception of target objects using multimodal deep learning," *IEEE Robotics and Automation Letters (RA-L)*, 2021.
- [5] S. Akizuki and Y. Aoki, "Tactile logging for understanding plausible tool use based on human demonstration," in *British Machine Vision Conference (BMVC)*, 2018.
- [6] Z. Qin, K. Fang, Y. Zhu, L. Fei-Fei, and S. Savarese, "Keto: Learning keypoint representations for tool manipulation," in *International Conference on Robotics and Automation (ICRA)*, 2020.
- [7] D. Turpin, L. Wang, S. Tsogkas, S. Dickinson, and A. Garg, "Gift: Generalizable interaction-aware functional tool affordances without labels," in *Robotics: Science and Systems (RSS)*, 2021.
- [8] Y. Liu, P. Wei, and S.-C. Zhu, "Jointly recognizing object fluents and tasks in egocentric videos," in *International Conference on Computer Vision (ICCV)*, 2017.
- [9] H. Tang, P. Wei, H. Li, and N. Zheng, "Inferring tasks and fluents in videos by learning causal relations," in *International Conference on Pattern Recognition (ICPR)*, 2020.
- [10] M. Li, Z. Ferguson, T. Schneider, T. Langlois, D. Zorin, D. Panozzo, C. Jiang, and D. M. Kaufman, "Incremental potential contact: Intersection-and inversion-free, large-deformation dynamics," *ACM Transactions on Graphics (TOG)*, 2020.
- [11] Z. Jiao, Z. Zhang, X. Jiang, D. Han, S.-C. Zhu, Y. Zhu, and H. Liu, "Consolidating kinematic models to promote coordinated mobile manipulations," in *International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [12] Z. Jiao, Z. Zhang, W. Wang, D. Han, S.-C. Zhu, Y. Zhu, and H. Liu, "Efficient task planning for mobile manipulation: a virtual kinematic chain perspective," in *International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [13] L. Nair, J. Balloch, and S. Chernova, "Tool macgyvering: Tool construction using geometric reasoning," in *International Conference on Robotics and Automation (ICRA)*, 2019.
- [14] L. Nair, N. S. Srikanth, Z. M. Erickson, and S. Chernova, "Autonomous tool construction using part shape and attachment prediction," in *Robotics: Science and Systems (RSS)*, 2019.
- [15] H. Wicaksono and C. Sammut, "A learning framework for tool creation by a robot," in *Proceedings of ACRA*, 2015.
- [16] H. Wicaksono and C. Sammut, "A cognitive robot equipped with autonomous tool innovation expertise," *International Journal of Electrical & Computer Engineering*, vol. 10, no. 2, 2020.
- [17] K. Takahashi, K. Kim, T. Ogata, and S. Sugano, "Tool-body assimilation model considering grasping motion through deep learning," *Robotics and Autonomous Systems*, vol. 91, pp. 115–127, 2017.
- [18] M. A. Toussaint, K. R. Allen, K. A. Smith, and J. B. Tenenbaum, "Differentiable physics and stable modes for tool-use and manipulation planning," in *Robotics: Science and Systems (RSS)*, 2018.
- [19] Y. Zhu, Y. Zhao, and S.-C. Zhu, "Understanding tools: Task-oriented object modeling, learning and recognition," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [20] M. Leviñ and H. Christensen, "Using environment objects as tools in unknown environments," in *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2015.
- [21] J. Brawer, M. Qin, and B. Scassellati, "A causal approach to tool affordance learning," in *International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [22] K. R. Allen, K. A. Smith, and J. B. Tenenbaum, "Rapid trial-and-error learning with simulation supports flexible tool use and physical reasoning," *Proceedings of the National Academy of Sciences (PNAS)*, vol. 117, no. 47, pp. 29302–29310, 2020.
- [23] S. Tuli, R. Bansal, R. Paul, et al., "Tango: Commonsense generalization in predicting tool interactions for mobile manipulators," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2021.
- [24] H. Liu, C. Zhang, Y. Zhu, C. Jiang, and S.-C. Zhu, "Mirroring without overimitation: Learning functionally equivalent manipulation actions," in *AAAI Conference on Artificial Intelligence (AAAI)*, 2019.
- [25] M. Kennedy, K. Schmeckpeper, D. Thakur, C. Jiang, V. Kumar, and K. Daniilidis, "Autonomous precision pouring from unknown containers," *IEEE Robotics and Automation Letters (RA-L)*, 2019.
- [26] C. Matl, Y. Narang, R. Bajcsy, F. Ramos, and D. Fox, "Inferring the material properties of granular media for robotic tasks," in *International Conference on Robotics and Automation (ICRA)*, 2020.
- [27] D. Hahn, P. Banzet, J. M. Bern, and S. Coros, "Real2sim: Viscoelastic parameter estimation from dynamic motion," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 6, pp. 1–13, 2019.
- [28] M. Geilinger, D. Hahn, J. Zehnder, M. Bächer, B. Thomaszewski, and S. Coros, "Add: Analytically differentiable dynamics for multi-body systems with frictional contact," *ACM Transactions on Graphics (TOG)*, vol. 39, no. 6, pp. 1–15, 2020.
- [29] X. Li, J. McWilliams, M. Li, C. Sung, and C. Jiang, "Soft hybrid aerial vehicle via bistable mechanism," in *International Conference on Robotics and Automation (ICRA)*, 2021.
- [30] E. Heiden, M. Macklin, Y. Narang, D. Fox, A. Garg, and F. Ramos, "Disect: A differentiable simulation engine for autonomous robotic cutting," in *Robotics: Science and Systems (RSS)*, 2021.
- [31] S. Gallagher, *How the body shapes the mind*. Clarendon Press, 2006.
- [32] N. P. Holmes and C. Spence, "Beyond the body schema: Visual, prosthetic, and technological contributions to bodily perception and awareness," in *Human body perception from the inside out: Advances in visual cognition*, pp. 15–64, Oxford University Press, 2006.
- [33] M. Hoffmann, H. Marques, A. Arieta, H. Sumioka, M. Lungarella, and R. Pfeifer, "Body schema in robotics: a review," *Transactions on Autonomous Mental Development*, vol. 2, no. 4, pp. 304–324, 2010.
- [34] R. Featherstone, *Rigid body dynamics algorithms*. Springer, 2014.
- [35] O. C. Zienkiewicz and R. L. Taylor, *The finite element method*, vol. 2. Butterworth-Heinemann, 2000.
- [36] M. Li, M. Gao, T. Langlois, C. Jiang, and D. M. Kaufman, "Decomposed optimization time integrator for large-step elastodynamics," *ACM Transactions on Graphics (TOG)*, vol. 38, no. 4, 2019.
- [37] M. Li, D. M. Kaufman, and C. Jiang, "Codimensional incremental potential contact," *ACM Transactions on Graphics (TOG)*, vol. 40, no. 4, 2021.
- [38] M. Schmidt and H. Lipson, "Distilling free-form natural laws from experimental data," *Science*, vol. 324, no. 5923, pp. 81–85, 2009.
- [39] S.-M. Udrescu and M. Tegmark, "Ai feynman: A physics-inspired method for symbolic regression," *Science Advances*, vol. 6, no. 16, 2020.
- [40] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *Journal of the Royal Statistical Society: Series B (Methodological)*, vol. 39, no. 1, 1977.
- [41] C. M. Bishop and N. M. Nasrabadi, "Conditional Gaussian distributions," in *Pattern recognition and machine learning*, ch. 2.3.1, pp. 85–88. Springer, 2006.
- [42] G. Wu, F. C. Van der Helm, H. D. Veeger, M. Makhssous, P. Van Roy, C. Anglin, et al., "Isb recommendation on definitions of joint coordinate systems of various joints for the reporting of human joint motion—part ii: shoulder, elbow, wrist and hand," *Journal of Biomechanics*, vol. 38, no. 5, pp. 981–992, 2005.
- [43] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "Casadi: a software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, 2019.
- [44] J. Koenemann, G. Licitra, M. Alp, and M. Diehl, "Openoc—open optimal control library," 2017.
- [45] T. Liu, Z. Liu, Z. Jiao, Y. Zhu, and S.-C. Zhu, "Synthesizing diverse and physically stable grasps with arbitrary hand structures using differentiable force closure estimator," *IEEE Robotics and Automation Letters (RA-L)*, 2022.