

Towards Automated Chicken Deboning via Learning-based Dynamically-Adaptive 6-DoF Multi-Material Cutting

Zhaodong Yang¹, Ai-Ping Hu^{1,2}, Harish Ravichandar¹
¹Georgia Institute of Technology, ²Georgia Tech Research Institute
{halyang, ahu6, harish.ravichandar}@gatech.edu

Abstract—Automating chicken shoulder deboning requires precise 6-DoF cutting through a partially occluded, deformable, multi-material joint, since contact with the bones presents serious health and safety risks. Our work makes both systems-level and algorithmic contributions to train and deploy a reactive force-feedback cutting policy that dynamically adapts a nominal trajectory and enables full 6-DoF knife control to traverse the narrow joint gap while avoiding contact with the bones. First, we introduce an open-source custom-built simulator for multi-material cutting that models coupling, fracture, and cutting forces, and supports reinforcement learning, enabling efficient training and rapid prototyping. Second, we design a reusable physical testbed to emulate the chicken shoulder: two rigid “bone” spheres with controllable pose embedded in a softer block, enabling rigorous and repeatable evaluation while preserving essential multi-material characteristics of the target problem. Third, we train and deploy a residual RL policy, with discretized force observations and domain randomization, enabling robust zero-shot sim-to-real transfer and the first demonstration of a learned policy that debones a real chicken shoulder. Our experiments in our simulator, on our physical testbed, and on real chicken shoulders show that our learned policy reliably navigates the joint gap and reduces undesired bone/cartilage contact, resulting in up to a 4x improvement over existing open-loop cutting baselines in terms of success rate and bone avoidance. Our results also illustrate the necessity of force feedback for safe and effective multi-material cutting. The project website is at <https://star-lab.cc.gatech.edu/papers/Yang-automated-deboning>.

I. INTRODUCTION

Many tasks in food processing, healthcare, and manufacturing require separating or removing specific parts from objects made of multiple materials. Chicken deboning is a particularly impactful and challenging example, in which a knife must precisely cut through the shoulder joint of the bird (see Fig. 1A) to separate the chicken’s wing from its body without cutting into the shoulder bones or cartilage that are concealed beneath soft tissue and skin. A failure to avoid the bones or cartilage can introduce bone/cartilage chips into meat, introducing a foreign material into the consumer meat product. Today, this demanding task is performed by expert humans in grueling conditions — each processing plant requires this precise cut be manually repeated approximately 300,000 chickens each day [1] while standing, while temperatures are held between 40 to 50°F.

Despite decades of progress in agricultural automation [1], dating back to the 1990s [2], reliable and high-throughput automated chicken deboning has remained largely elusive due to three crucial and related gaps in the relevant literature:

Precise 6-DoF cutting: The chicken shoulder is comprised of the humerus and coracoid bones (roughly spherical; see Fig. 1A) connected by tendons, muscle and skin. To cut through the spacing between the humerus and scapula, the knife must not only translate but also rotate to align with the 3D cutting plane that is perpendicular to the line connecting the centers of the two spheres. This requires motion planning and control in a full 6-DoF task space, whereas many prior studies on robot cutting have considered cutting only in two-dimensional spaces (see [3], [4], [5], [6]). *Dynamic adaptability in a partially observable environment:* Before the knife enters the shoulder region, the bones are concealed beneath soft tissue, preventing precise trajectory planning in advance. The robot must therefore dynamically adapt and determine its motion based on online force feedback in order to avoid undesirable bone contacts. In contrast, prior work on cutting single-material objects has often assumed open-loop trajectories without online adaptation ([1], [7], [8], [9], [10]). *Multi-material interaction:* Errors in estimating the shoulder joint spacing could result in accidental cutting through cartilage/bone, and existing approaches to robot cutting typically do not model the deformations and fracture of multiple interacting deformable materials [3].

Our work makes both systems-level and algorithmic contributions to train and deploy a reactive cutting policy that can dynamically adapt a given nominal trajectory based on force feedback to achieve precise 6 DoF knife control, with a particular emphasis on chicken deboning (see Fig. 1).

Our first contribution involves the design and implementation of a **customized open-source multi-material cutting simulator** that abstracts chicken deboning and supports policy learning. Our simulator offers a lower-fidelity approximation, capturing only the most-impactful features of the deboning process (e.g., relative cutting forces, coupling and fracture of heterogeneous materials), and is integrated with modern reinforcement learning (RL) libraries for efficient training. Our design strategically avoids a compute-intensive full-fidelity model that renders learning intractable.

While our custom simulator enables low-cost training, design, and evaluation, simulation results might not always translate to reality. At the same time, evaluating on real-world chickens is neither scalable nor repeatable. A variety of practical reasons prohibit careful and controlled experiments: chickens exhibit natural variability, the ground truth configuration of chicken shoulder bones cannot be known

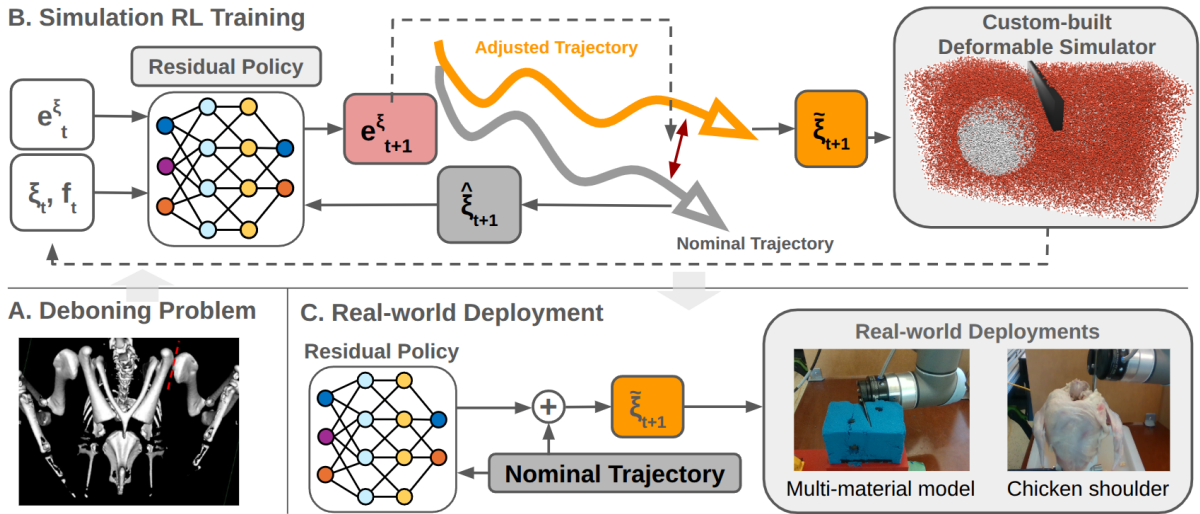


Fig. 1: Overview of this paper. A. MRI scan of a chicken. The red line indicates the shoulder joint position and orientation, where we need to precisely cut through. B. We first train the residual policy in the simulator. For each time step, the residual policy observes the current knife pose ξ_t and knife force sensor reading f_t from the simulator. It also observes residual pose e_t^ξ from the previous time step, and one step future knife pose $\hat{\xi}_{t+1}$ retrieved from the nominal trajectory. The residual policy eventually produces the next residual pose: e_{t+1}^ξ . And $\hat{\xi}_{t+1} = \xi_{t+1} + e_{t+1}^\xi$ is the adjusted pose of the nominal trajectory, also the next knife target pose. We use different nominal trajectories for different cutting tasks. Nominal trajectory is an interpolated trajectory of several keyframe poses, which are either fixed or mapped from the bone position. C. We deploy the residual policy on two challenging real-world multi-material cutting tasks.

apriori, and their configuration changes during cutting. To enable rigorous study and repeatable experimentation in the real-world, our second contribution involves the design and implementation of **a simplified reusable physical testbed for multi-material cutting**, modeled after the chicken shoulder. In this model, the bones are represented as two Reusable Clay spheres with a fixed relative position, ensuring the existence of a valid cutting trajectory, while the surrounding muscle is approximated as a cuboid of Kinetic Sand embedding the two spheres. These two materials exhibit distinct hardness levels, causing the knife to experience different cutting forces when interacting with them. This set-up allows one to precisely control the bone position and orientation in real-world experiments while preserving the essential multi-material characteristics of the bird shoulder deboning task.

Our third contribution involves the design and training of **a dynamically adaptive 6 DoF cutting policy** via RL-based residual policy training. During training and evaluation, our residual policy adaptively adjusts the given nominal knife trajectory based on force feedback to avoid cutting into the two bone spheres whose ground truth positions are not accurately known. To facilitate successful sim-to-real transfer, we discretize the force signals and incorporate it into the policy’s observation space. We find that our residual cutting policy can successfully learn to cut through the spacing between the two bone spheres while avoiding collision with the spheres in simulation. When combined with domain randomization, we show that our adaptive cutting policy can be trained entirely in sim and deployed zero-shot both in our physical testbed and for real chicken deboning.

We conducted rigorous experiments in both our custom-

built simulator and our simplified physical model. Further, to ensure that our contributions ultimately translate to the target setting, we also conducted experiments on real chicken shoulders using our residual policy training and sim2real transfer pipeline. Our results demonstrate the need for force feedback and suggest that our method significantly outperforms previous open-loop automation approaches [10] by up to 4x in terms of success rate and bone avoidance.

In summary, our primary contributions include:

- The first customized open-source simulator for multi-material cutting that supports coupling and fracture of heterogeneous materials, and reinforcement learning, enabling rapid prototyping and efficient training.
- The first real-world reusable multi-material cutting testbed modeled after the chicken shoulder, enabling controlled and repeatable experiments.
- The first learning-based approach to 6 DoF multi-material cutting that transfers zero-shot to both our physical testbed and real chicken shoulders.

II. RELATED WORKS

A. Deformable Cutting Simulator

While several robotic learning simulators support the deformation of soft materials ([11], [12], [13], [14], [15], [16], [17], [18]), most lack the capability to simulate fracture. Approaches to fracture modeling are typically divided into mesh-based and mesh-free methods. Mesh-based techniques like the Finite Element Method (FEM) require complex and costly re-meshing at the fracture interface, often restricting their application to simple planar cuts ([7], [19]). In contrast, mesh-free methods such as the Material Point Method

(MPM) ([20], [21], [22]), Position-based Dynamics (PBD) ([23], [24], [25]) and Smoothed Particle Hydrodynamics (SPH) ([26]) represent objects as particle aggregates, which more naturally accommodate arbitrary fracture surfaces.

However, a critical limitation of most mesh-free simulators for robotics is the absence of force feedback. Notable exceptions exist, but they retain significant constraints. For instance, DiSECT [7] computes interaction forces but is limited to planar cuts, while RoboNinja [3] provides force readings for arbitrary cuts but is restricted to simulating simple bi-material objects consisting of one deformable material and another rigid material. To address these gaps, our developed simulator leverages MPM with the CPIC [20] method. This allows us to simulate arbitrary cutting trajectories for complex objects composed of multiple deformable materials, all while providing the crucial force feedback necessary for robotic learning tasks.

B. Robotic Single-Material Cutting

Robotic cutting has been explored across a diverse range of applications, including food preparation tasks like processing vegetables ([4], [27], [28], [29], [5], [9]), fruits ([30], [31], [32]), dough [33], and meat [8]. Its application extends to specialized industrial and medical domains, from hot-blade cutting of polystyrene [34] to precise surgical procedures ([35], [36], [37]). To enhance robotic perception and control, researchers have developed sensorized tools, such as smart knives that provide force and depth feedback [38], and have utilized multi-modal haptic data to improve robustness [39].

However, a common thread in these prior works is the focus on cutting objects composed of a single, relatively uniform material, often executed through pre-defined open-loop trajectories or simple, repetitive planar cuts. In contrast, our research addresses the more complex challenge of cutting objects composed of multiple distinct deformable materials requiring complex spatial knife paths. This task necessitates real-time, multi-modal sensory feedback to dynamically adapt the cutting strategy to, for instance, precisely incise one material without damaging an adjacent one.

C. Robotic Multi-Material Cutting

While research into robotic cutting of multi-material objects is limited, a few notable studies exist. Early work in this area tackled the challenging problem of chicken deboning by developing highly specialized, automated machinery ([1], [10]). However, these systems were rigid in their design, lacking the versatility to adapt to natural variations in chicken size or bone positioning.

More recent studies have employed learning-based approaches with sensory feedback. For instance, one system successfully scooped the flesh from a grapefruit without piercing the peel by using a trained classifier and a closed-loop controller to manage the 1D cutting depth [6]. Similarly, the RoboNinja framework demonstrated impressive adaptability in cutting multi-material fruits like avocados by using a state estimator and an adaptive policy to control a 2D cutting plane [3]. Despite its success on objects with varied

core geometries, its approach is fundamentally constrained to a 2D workspace.

A critical limitation of these recent advancements is that they address the problem assuming low-dimensional action spaces (1D or 2D). This is insufficient for complex, three-dimensional tasks like chicken deboning, which requires navigating intricate geometries and dynamically adapting the cutting tool’s full 6D pose. Therefore, the challenge of versatile, high-degree-of-freedom robotic cutting for multi-material objects remains an open problem.

III. METHODOLOGY

In this work, we utilize residual policy for trajectory adaptation and achieve successful sim2real transfer on applications such as the chicken deboning problem. We first implement a multi-material deformable cutting simulator that supports coupling and fracture of heterogeneous materials, as well as realistic modeling of cutting forces. Using this simulator, we construct a chicken shoulder cutting environment corresponding to our simplified real world model and integrate it with Stable-Baselines3, a reinforcement learning (RL) framework. Within this environment, we train a reactive residual cutting policy, where the policy adaptively adjusts the knife’s motion in response to force feedback. To promote robust sim-to-real transfer, we discretize force feedback and incorporate it into the policy’s observation space. Finally, we evaluate the trained residual cutting policy on a reusable real world multi-material cutting model.

A. Multi-Material Cutting Simulator

We develop a multi-material cutting simulation environment that supports both the modeling of deformable materials and the coupling between them. In our scene, bones and meat are represented using an elastoplastic continuum model simulated with MLS-MPM [20], governed by the von Mises yield criterion. MPM-based methods naturally capture large deformations and topological changes, making them well-suited for cutting tasks. To differentiate materials and maintain simulation stability, we assign the following distinct physical parameters: $\lambda_m = 27.78Pa$, $\mu_m = 41.67Pa$, $\rho_m = 10^3kg/m^3$, $\lambda_b = 222.22Pa$, $\mu_b = 333.33Pa$, $\rho_b = 2.819 \times 10^3kg/m^3$, where λ_m , μ_m , λ_b , and μ_b are Lamé parameters for meat material and bone material, respectively, and ρ_m and ρ_b the density for meat material and bone material.

The knife is represented as a time-varying signed distance field (SDF) and approximated as a thin cuboid to accelerate computation. Contact between the deformable materials and the rigid knife is modeled by computing the cuboid surface normals and applying Coulomb friction. Material separation during cutting is handled directly by MLS-MPM, which inherently supports sharp and clean splitting of material points in soft media.

Our simulator is implemented in Python and Taichi [40], and the environment inherits from the OpenAI Gym interface to ensure compatibility with Stable-Baselines3 [41]. This allows reinforcement learning policies to be trained directly using algorithms provided by Stable-Baselines3. To ensure

stable and efficient training, we further apply large force and velocity clipping, which prevents non-converging simulations and guarantees smooth RL rollouts.

B. Residual Cutting Policy Learning

The residual cutting policy is trained within our customized simulator. We let the knife follow a nominal trajectory determined based on estimated bone positions. At each simulation timestep, the residual policy adjusts the 6D (position and orientation) nominal trajectory. The policy’s observation space includes the current knife 6D pose, the next nominal pose (one timestep ahead), the measured 3D cutting force, and the previous 6D residual action. Based on these inputs, the policy outputs a 6D residual adjustment, which is added to the nominal trajectory to generate the executed action. To ensure stability and prevent unrealistic motions, the total residual action is clipped before being applied. To bridge the simulation and reality gap, especially the cutting force, we discretize each component of the 3D force and use the discretized 3D force in the residual policy observation.

Observation Space: At each timestep, the residual policy observes the following information from the environment: the current knife pose ξ_t ; one step future knife pose $\hat{\xi}_{t+1}$, retrieved from the nominal trajectory; current knife force sensor reading f_t ; and residual pose e_t^ξ .

Action Space: The residual policy produces a delta residual pose as the action a_t , which is added to the residual pose e_t^ξ to produce the next residual pose: $e_{t+1}^\xi = e_t^\xi + \eta a_t$, where η is an action scaling factor. e_{t+1}^ξ will be added to the one-step future knife pose $\hat{\xi}_{t+1}$ to adjust the nominal trajectory: $\tilde{\xi}_{t+1} = \hat{\xi}_{t+1} + e_{t+1}^\xi$. And $\tilde{\xi}_{t+1}$ will be the next knife target pose. The initial value of the residual pose e_0^ξ is $\mathbf{0}$.

Nominal Trajectory: Nominal trajectory is interpolated from several waypoints which are determined based on estimated bone positions (refer to [10] for details on anatomical chicken shoulder joint prediction). For example, previous studies utilizes a shoulder prediction model parameterized by visual bird features to obtain the waypoints [10].

Reward: Since we already have a nominal trajectory to roughly command the motion of the knife, we design a reward function to guide the residual policy to learn to adjust the nominal trajectory to avoid cutting into the bones, while also avoiding unnecessary adjustment: $R = R_{bone} + R_{action}$. R_{bone} is the bone avoidance reward. $R_{bone} = C - \alpha_1 b$. C is a constant value, α_1 is the scaling parameter and $b \in \{0, 1\}$ is a binary value. $b = 0$ when the knife cuts the bone at the current time step, while $b = 1$ when the knife does not cut the bone at the current time step. This reward term is crucial for eliciting desired behavior. R_{action} is the action penalty, which regularizes the behavior of the residual policy. $R_{action} = \alpha_2 \|e_t^\xi\|_2^2$, where α_2 is a scaling parameter and e_t^ξ is the residual pose.

C. Reusable Real-World Model and Testbed

We designed a simplified, reusable chicken shoulder model to enable quantitative and controllable study of the deboning problem. In this model, the bones are represented as two




Fig. 2: The setup process of our reusable physical testbed.

spheres with a fixed relative position, ensuring the existence of a valid cutting trajectory, while the surrounding soft tissue is approximated as a cuboid embedding the two spheres (Fig 2). We use Kinetic Sand to represent the soft tissue because it has stable physical properties [3] and requires a relatively small force to be cut through. For the bone/cartilage material, which should allow fracture for simulating failure cases and require a larger cutting force to fracture than the soft tissue, we tested the cutting force of three different materials: Reusable Clay, Air-dry Clay and Play-Doh. The average cutting force of the Re-usable Clay and the Air-dry clay are larger than that of Kinetic Sand, while the cutting force of Play-Doh is close to cutting Kinetic Sand. Therefore, Play-Doh is not a proper proxy for the bone material. We eventually chose Reusable Clay over Air-dry Clay for better efficiency, since Air-dry Clay requires over 24 hours of drying time and cannot be reused after fracturing.

We design and 3D-print a base and support for the deformable model (Fig. 2). The base can be clamped to a table to maintain a fixed model position during cutting, while the support consists of two parallel cylindrical columns mounted on the base. Reusable clay spheres, crafted with a spherical mold, are placed on the columns to represent the shoulder bones with fixed relative positions. The remaining volume is filled with Kinetic Sand using a box mold, representing the chicken muscle meat.

D. Sim2Real Transfer

Due to the difficulty in resetting the real-world cutting environment and obtaining the environment states, we are unable to directly train RL policy in the real world. Instead, we utilize our custom-built simulator to train a policy and then deploy it to the real world. To zero-shot deploy the trained residual policy in the real world, we apply domain randomization. Specifically, we add noise into the proprioceptive observations of the residual policy while training in simulation. We add Gaussian noise $\mathcal{N}(0, 0.01)$ at the current knife position and add Gaussian noise $\mathcal{N}(0, 0.1)$ at the current knife orientation quaternion. We also observe that previous studies utilize discretization for successful sim2real transfer of tactile information [42] on in-hand manipulation tasks. While our simulator captures the relative force trends necessary to distinguish between different materials during

multi-material cutting, accurately modeling the absolute contact dynamics remains challenging, resulting in an inherent domain gap in the force profiles. To further bridge the gap between simulation force observation and real world force observation, we adopt discretization by normalizing and discretizing the force to make each orthogonal force component fall within the set $\{0, 0.1, 0.2, \dots, 0.9, 1.0\}$. Crucially, our results demonstrate that by relying on this discretized force abstraction alongside proprioceptive domain randomization, a policy trained on simplified simulation dynamics can still successfully generalize to complex real-world environments.

IV. EXPERIMENTAL EVALUATION

We conducted three experiments (one in simulation, one using our physical model, and one using real chicken shoulders) in order to answer the following questions:

- **Q1:** How does our approach improve the cutting performance of nominal trajectory?
- **Q2:** How important is force feedback for adaptation?
- **Q3:** Can residual policy be successfully transferred to the real-world?

Baselines: We evaluated against two baselines that challenge our core designs.

- *Nominal:* This baseline tracks the precomputed nominal trajectory to cut the deformable model.
- *Adaptive w/o Force:* This baseline utilizes both nominal trajectory and residual policy. But the residual policy does not observe knife force.
- *Adaptive (Ours):* *Adaptive* utilizes both nominal trajectory and residual policy mentioned in Section III-B.

Note that RoboNinja [3] is not a suitable baseline for our approach given that it addresses a different and simpler problem. Nevertheless, we experimented with extending RoboNinja’s adaptive cutting policy to our physical testbed. As expected, it could not complete our cutting task successfully due to three main reasons: (1) it relies on the detections of collisions to estimate the shape of the rigid core but such collisions can severely damage non-rigid bones (see Fig.3), (2) it can only adjust the trajectory to one direction, while our cutting task requires adjustment in arbitrary directions and orientation, and (3) it can produce large adjustments that miss the spacing between the two bone spheres.

Policy Learning: We used the PPO [43] implementation from StableBaselines3 to train our residual policy. Both policy and value functions are parameterized via a MLP network with 2 hidden layers. The size of the hidden layers for both policy and value functions are (64, 64). The activation functions are all set as hyperbolic tangent. We use the same PPO hyper-parameters for all the baselines and our method (γ : 0.99, λ : 0.95, learning rate: 0.0003, clip range: 0.2, and batch size: 64). We train the policies on a computer with a single NVIDIA RTX 4090 GPU.

Environment Setup: As detailed in Sec. III-A, the physical simulation parameters remain identical when training the cutting policies for both the real-world proxy model and the actual chicken experiments.




Fig. 3: Roboninja’s adaptive cutting policy qualitative evaluation results. The left picture shows that RoboNinja relies on knife-bone collision for interactive state estimation, which does crucial damage to non-rigid core in our task. The right is a failure case of RoboNinja when the knife first cuts into the right bone. Because RoboNinja can only adjust its trajectory to one direction (right), it keeps cutting into the right bone deeper and deeper.



Fig. 4: We use a cone which is clamped on the table to fix the chicken while conducting the cutting experiment.

In real-world model, the radius of the two bone spheres is $0.021m$ and the distance between the center of the two spheres is $0.05m$, hence leaving a $0.008m$ gap between the two sphere. To establish scenarios with varying estimation errors, we varied the 3D positions of the two spheres by $[\delta x, \delta y, \delta z]$, with sampling ranges δx : $[-0.005m, 0.005m]$, δy : $[-0.01m, 0.01m]$ and δz : $[-0.01m, 0.01m]$. The relative position of the two spheres remains constant, guaranteeing the existence of a valid cutting trajectory between the gap that avoids sphere (bone) contact. To maintain numerical stability, the simulated environment is scaled up by a factor of 5 relative to the physical real-world dimensions. For actual chicken experiment, the setup of chicken shoulder cutting experiment is shown in Fig. 4.

Real-world Robot Setup: Our real-world experiments are conducted on a UR5 robotic arm platform. A Robotiq FT 300 force–torque sensor is mounted between the robot wrist and a knife blade tool, providing force feedback during cutting. The sensor has a measurement range of 0–300 N, sufficient for capturing the interaction forces encountered in the multi-material cutting task.

Nominal Trajectory: In simulation and physical model experiments, we generated the nominal trajectory by interpolating three waypoints. We held the three waypoints constant, and systematically varied the ground truth bone positions to reflect the fact that nominal trajectories are computed offline based solely on noisy estimates of bone positions. For cutting real chicken shoulders, we used the nominal




Fig. 5: Simulation qualitative evaluation result. The first two rows show typical failure cases of the *Nominal* and *Adaptive w/o Force*. Due to lack of adaptivity, *Nominal* rarely avoids the bone in most of the cases. Though *Adaptive w/o Force* has the ability to adjust its trajectory, it can not perform well because it lacks necessary force feedback. The third row demonstrates a typical successful example of *Adaptive*, which has both the ability of adaptation and receiving necessary force feedback.

trajectory interpolated from waypoints (determined based on a shoulder prediction model parameterized by visual bird features [10]). We used cubic splines to interpolate the translation of the waypoints and spherical linear interpolation to interpolate between 3D rotations.

A. Simulation Experiment

We first evaluated each method on a set of 50 cuts in our simulator with randomly sampled bone positions.

Metrics:

- *Violation Duration*: To quantify bone avoidance, we count the number of timesteps during which the knife is cutting into the bones during each cut as the duration of violation.
- *Success Rate*: We consider an episode “successful” if there are no bone cuts in the episode. We consider it to be a bone cut if the knife is cutting into the bone during at least two consecutive environment steps.

Results: We provide qualitative results in Fig.5 and quantitative results in TABLE I. As shown, Open-loop *Nominal* only results in 0.2 success rate and highest violation duration, due to its lack of adaptivity, underscoring its sensitivity to partial observability and the need for adaptation in multi-material cutting tasks. While *Adaptive w/o Force* is able to perform better than *Nominal* likely by improving the nominal trajectory while training in RL, it still lacks crucial force feedback. As expected, *Adaptive* achieves the highest success rate and lowest bone cut count among all the baselines, demonstrating the effectiveness of residual policy training and the need for force-driven adaptivity.

B. Real-world Model Experiment

We next evaluate each method over 20 cuts using our reusable physical model described in Section III-C (see Section III-D our sim2real transfer approach).

TABLE I: Simulation Evaluation Results

Methods	Avg. Violation Duration (%) (\downarrow)	Success Rate (\uparrow)
<i>Nominal</i>	18.1 ± 11.88	0.20
<i>Adaptive w/o Force</i>	9.04 ± 10.76	0.54
<i>Adaptive</i>	4.74 ± 9.66	0.80

Metrics:

- *Bone Cut Mass*: The weight of reusable clay bone fragment(s) resulting from an unsuccessful cut, as measured by a jewelry scale with a resolution of 0.001 grams.
- *Success Rate*: We consider a trial to be “successful” if no part of bones was cut during the trial.

Results: Please refer to Fig.6 for qualitative results, and TABLE II for quantitative results. Similar to simulation results, our approach (*Adaptive*) achieves the highest success rate and a substantially lower average bone cut mass than the two baselines. These results, not only reinforce the effectiveness of force-driven adaptation, but also demonstrate the impressive sim2real transfer performance of our learned residual policy. Further, taking advantage of our access to ground-truth bone position when using the simplified model, we studied the relation between the extent of estimation error and bone cut mass (see Fig.7). Here, we approximate estimation error using the amount by which we shift the true bone position along the y axis (roughly perpendicular to the cutting direction), away from the estimated position. We observe that, in contrast to the two baselines struggle, *Adaptive* still performs well even when the estimated bone positions are far from ground truth.

C. Real Chicken Deboning Experiment

Finally, we evaluated each method on cutting real chicken shoulder joints. We first trained both learning-based methods in our custom-built simulator, while randomizing the position

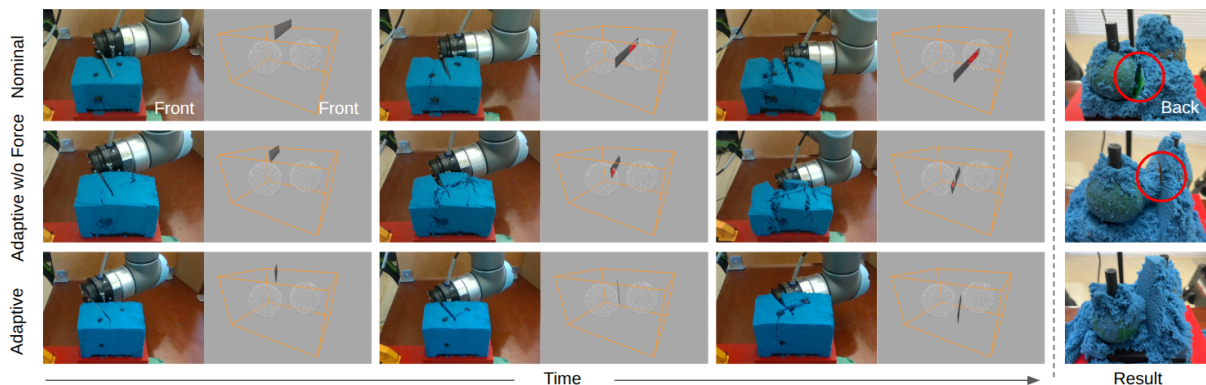


Fig. 6: Real world model testbed qualitative evaluation result. To be able to “see through” the deformable model, we record robot end-effector trajectory and bone sphere position and visualize the cutting results using Taichi. The knife-bone intersection area will turn red in the visualization when a collision occurs. We also tear down the model after each cut to expose the bone and take a picture from the back. The pictures are shown in the last column. The red circle indicates the parts of the bone spheres that are cut off.




Fig. 7: Bone cut mass a function of estimation error.

Method	Avg. Bone Mass (g)	Success Rate
<i>Nominal</i>	3.032 ± 3.090	0.05
<i>Adapt. w/o f.</i>	0.730 ± 0.884	0.35
<i>Adapt.</i>	0.065 ± 0.130	0.75

TABLE II: Results on Real-World Testbed.




Fig. 8: Top: Knife trajectory for cutting the chicken shoulder. Bottom left: typical cutting result of the *Nominal* baseline, in which one large piece of the shoulder bone is cut off. Bottom right: example of utilizing the *Adaptive*. The bone is successfully avoided and the knife precisely cuts between the two shoulder bones.

of the two simulated bone spheres according to the distribution of real chicken shoulder bone positions. We then used the sim2real transfer approach described in Section III-D.

Metrics: We use the same evaluation metrics as in Section IV-B. In addition, we also compute the ratio of cuts that bone fragments weighing under 0.05 grams since such small fragments are likely to be less dangerous.

Baselines: We use the nominal trajectory proposed in a previous robotic chicken deboning study [10] as a baseline.

Results: We provide the qualitative and quantitative results of cutting real chicken shoulder in Fig. 8 and TABLE III, respectively. Even on the challenging real-world chicken

deboning task, our *Adaptive* approach is able to considerably improve upon the nominal trajectory. These results also demonstrate our overall approach’s ability to overcome dramatic sim2real gaps (simulated particles with computed forces to real chicken deboning with a force-torque sensor mounted on to a physical robot).

TABLE III: Real Chicken Cutting Experiment Results

Methods	Avg. Bone Mass (g) (↓)	Success Rate (↑)	Bone mass < 0.05g (↑)
<i>Nominal</i>	0.339 ± 0.316	0.10	0.20
<i>Adaptive</i>	0.121 ± 0.239	0.50	0.70

V. CONCLUSIONS

In this paper, we present the first multi-material deformable simulator that supports coupling and fracturing of heterogeneous materials and reinforcement learning, and the first real-world reusable multi-material cutting testbed modeled after the chicken shoulder. We use learning-based approach to tackle 6 DoF multi-material cutting tasks and achieve successful zero-shot transfer to both our physical testbed and real chicken shoulders. While our approach has demonstrated necessary adaptivity and transferability, it still has certain limitations for future extension. First, we have not tried our approach on more diverse bone shapes. Second, we rely on normalization and discretization of the force signal for sim2real transfer. We need to collect force data in advance and find the maximum value for manual normalization. Future research could consider cutting tasks with more arbitrary trajectory and more diverse bone shapes, especially more anatomically correct bone geometry. Other sim2real transfer techniques, such as Rapid Motor Adaptation can be applied to automatically adapt the trained policy to cutting tasks of different materials. Besides, we can improve the fidelity of our multi-material deformable simulator and even include more evaluation metrics for the multi-material cutting task.

VI. ACKNOWLEDGMENT

This work is supported by the USDA National Institute of Food and Agriculture.

REFERENCES

- [1] D. Zhou, J. Holmes, W. Holcombe, and G. McMurray, "Automation of the bird shoulder joint deboning," in *2007 IEEE/ASME international conference on advanced intelligent mechatronics*, 2007, pp. 1–6.
- [2] W. Daley, T. He, K.-M. Lee, and M. Sandlin, "Modeling of the natural product deboning process using biological and human models," in *1999 IEEE/ASME International Conference on Advanced Intelligent Mechatronics (Cat. No. 99TH8399)*. IEEE, 1999, pp. 49–54.
- [3] Z. Xu, Z. Xian, X. Lin, C. Chi, Z. Huang, C. Gan, and S. Song, "RoboNinja: Learning an Adaptive Cutting Policy for Multi-Material Objects," in *Proceedings of Robotics: Science and Systems*, Daegu, Republic of Korea, July 2023.
- [4] C. C. Beltran-Hernandez, N. Erbeti, and M. Hamaya, "Sliceit! - a dual simulator framework for learning robot food slicing," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 4296–4302.
- [5] X. Mu, Y. Xue, and Y.-B. Jia, "Dexterous robotic cutting based on fracture mechanics and force control," *IEEE Transactions on Automation Science and Engineering*, vol. 21, no. 4, pp. 5198–5215, 2024.
- [6] A. Straižys, M. Burke, and S. Ramamoorthy, "Surfing on an uncertain edge: Precision cutting of soft tissue using torque-based medium classification," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4623–4629.
- [7] E. Heiden, M. Macklin, Y. S. Narang, D. Fox, A. Garg, and F. Ramos, "DiSECT: A Differentiable Simulation Engine for Autonomous Robotic Cutting," in *Proceedings of Robotics: Science and Systems*, Virtual, July 2021.
- [8] P. Long, A. A. Moughibay, W. Khalil, and P. Martinet, "Robotic Meat Cutting," in *ICT-PAMM Workshop*, Kumamoto, Japan, Nov. 2013.
- [9] X. Mu, Y. Xue, and Y.-B. Jia, "Robotic cutting: Mechanics and control of knife motion," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 3066–3072.
- [10] A.-P. Hu, J. Bailey, M. Matthews, G. McMurray, and W. Daley, "Intelligent automation of bird deboning," in *IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*, 2012.
- [11] J. Liang, V. Makoviychuk, A. Handa, N. Chentanez, M. Macklin, and D. Fox, "GPU-accelerated robotic simulation for distributed reinforcement learning," in *CoRL*, 2018.
- [12] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012, pp. 5026–5033.
- [13] G. Authors, "Genesis: A generative and universal physics engine for robotics and beyond," December 2024. [Online]. Available: <https://github.com/Genesis-Embodied-AI/Genesis>
- [14] C. Lin, J. Fan, Y. Wang, Z. Yang, Z. Chen, L. Fang, T.-H. Wang, Z. Xian, and C. Gan, "UBSoft: A simulation platform for robotic skill learning in unbounded soft environments," in *8th Annual Conference on Robot Learning*, 2024.
- [15] S. Li, Z. Huang, T. Chen, T. Du, H. Su, J. B. Tenenbaum, and C. Gan, "Dexdeform: Dexterous deformable object manipulation with human demonstrations and differentiable physics," in *The Eleventh International Conference on Learning Representations*, 2023.
- [16] Z. Xian, B. Zhu, Z. Xu, H.-Y. Tung, A. Torralba, K. Fragkiadaki, and C. Gan, "Fluidlab: A differentiable environment for benchmarking complex fluid manipulation," in *International Conference on Learning Representations*, 2023.
- [17] X. Lin, Y. Wang, J. Olkin, and D. Held, "Softgym: Benchmarking deep reinforcement learning for deformable object manipulation," in *Conference on Robot Learning*, 2020.
- [18] F. Xiang, Y. Qin, K. Mo, Y. Xia, H. Zhu, F. Liu, M. Liu, H. Jiang, Y. Yuan, H. Wang, L. Yi, A. X. Chang, L. J. Guibas, and H. Su, "SAPIEN: A simulated part-based interactive environment," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [19] L. Jeřábková and T. Kühlen, "Stable cutting of deformable objects in virtual environments using xfem," *IEEE Computer Graphics and Applications*, vol. 29, no. 2, pp. 61–71, 2009.
- [20] Y. Hu, Y. Fang, Z. Ge, Z. Qu, Y. Zhu, A. Pradhana, and C. Jiang, "A moving least squares material point method with displacement discontinuity and two-way rigid body coupling," *ACM Trans. Graph.*, vol. 37, no. 4, Jul. 2018.
- [21] S. Wang, M. Ding, T. F. Gast, L. Zhu, S. Gagniere, C. Jiang, and J. M. Teran, "Simulation and visualization of ductile fracture with the material point method," *ACM Comput. Graph. Interact. Tech.*, 2019.
- [22] J. Wolper, Y. Fang, M. Li, J. Lu, M. Gao, and C. Jiang, "Cdmpm: continuum damage material point methods for dynamic fracture animation," *ACM Trans. Graph.*, vol. 38, no. 4, Jul. 2019.
- [23] Y. Han, F. Liu, and M. C. Yip, "A 2d surgical simulation framework for tool-tissue interaction," *arXiv preprint arXiv:2010.13936*, 2020.
- [24] I. Berndt, R. Torchelsen, and A. Maciell, "Efficient surgical cutting with position-based dynamics," *IEEE Computer Graphics and Applications*, vol. 37, no. 3, pp. 24–31, 2017.
- [25] F. Liu, Z. Li, Y. Han, J. Lu, F. Richter, and M. C. Yip, "Real-to-sim registration of deformable soft tissue with position-based dynamics for surgical robot autonomy," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 12 328–12 334.
- [26] K. Raveendran, C. Wojtan, and G. Turk, "Hybrid smoothed particle hydrodynamics," in *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ser. SCA '11. New York, NY, USA: Association for Computing Machinery, 2011, p. 33–42.
- [27] H.-S. Fang, H. Fang, Z. Tang, J. Liu, C. Wang, J. Wang, H. Zhu, and C. Lu, "Rh20t: A comprehensive robotic dataset for learning diverse skills in one-shot," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 653–660.
- [28] J. Grannen, Y. Wu, B. Vu, and D. Sadigh, "Stabilize to act: Learning to coordinate for bimanual manipulation," in *7th Annual Conference on Robot Learning*, 2023.
- [29] I. Mitsioni, P. Tajvar, D. Kragic, J. Tumova, and C. Pek, "Safe data-driven model predictive control of systems with complex dynamics," *IEEE Transactions on Robotics*, vol. 39, no. 4, pp. 3242–3258, 2023.
- [30] I. Mitsioni, Y. Karayiannidis, and D. Kragic, "Modelling and learning dynamics for robotic food-cutting," in *2021 IEEE 17th International Conference on Automation Science and Engineering (CASE)*, 2021, pp. 1194–1200.
- [31] P. Jamdagni and Y.-B. Jia, "Robotic cutting of solids based on fracture mechanics and fem," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 8252–8257.
- [32] I. Mitsioni, Y. Karayiannidis, J. A. Stork, and D. Kragic, "Data-driven model predictive control for the contact-rich task of food cutting," in *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*, 2019, pp. 244–250.
- [33] H. Shi, H. Xu, S. Clarke, Y. Li, and J. Wu, "Robocook: Long-horizon elastoplastic object manipulation with diverse tools," in *7th Annual Conference on Robot Learning*, 2023.
- [34] A. Søndergaard, J. Feringa, T. Nørbjerg, K. Steenstrup, D. Brander, J. Graversen, S. Markvorsen, A. Bærentzen, K. Petkov, J. Hattel, K. Clausen, K. Jensen, L. Knudsen, and J. Kortbek, *Robotic Hot-Blade Cutting*. Springer International Publishing, 2016.
- [35] A. Straižys, M. Burke, and S. Ramamoorthy, "Learning robotic cutting from demonstration: Non-holonomic dmps using the udwadia-kalaba method," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 5034–5040.
- [36] L. Han, H. Wang, Z. Liu, W. Chen, and X. Zhang, "Vision-based cutting control of deformable objects with surface tracking," *IEEE/ASME Transactions on Mechatronics*, vol. 26, no. 4, pp. 2016–2026, 2021.
- [37] B. Thananjeyan, A. Garg, S. Krishnan, C. Chen, L. Miller, and K. Goldberg, "Multilateral surgical pattern cutting in 2d orthotropic gauze with deep reinforcement learning policies for tensioning," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 2371–2378.
- [38] A. Mason, D. Romanov, L. E. Cordova-Lopez, and O. Korostynska, "Smart knife: Integrated intelligence for robotic meat cutting," *IEEE Sensors Journal*, vol. 22, no. 21, pp. 20 475–20 483, 2022.
- [39] K. Zhang, M. Sharma, M. Veloso, and O. Kroemer, "Leveraging multimodal haptic sensory data for robust cutting," in *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*, 2019, pp. 409–416.
- [40] Y. Hu, T.-M. Li, L. Anderson, J. Ragan-Kelley, and F. Durand, "Taichi: a language for high-performance computation on spatially sparse data structures," *ACM Trans. Graph.*, vol. 38, no. 6, Nov. 2019.
- [41] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, "Stable-baselines3: Reliable reinforcement learning implementations," *Journal of Machine Learning Research*, 2021.
- [42] H. Qi, B. Yi, S. Suresh, M. Lambeta, Y. Ma, R. Calandra, and J. Malik, "General in-hand object rotation with vision and touch," in *7th Annual Conference on Robot Learning*, 2023.
- [43] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017.