Learning the Dynamics of Compliant Tool-Environment Interaction for Visuo-Tactile Contact Servoing

Anonymous Author(s)
Affiliation
Address
email

Abstract: Many manipulation tasks require the robot to control the contact between a grasped compliant tool and the environment, e.g. scraping a frying pan with a spatula. However, modeling tool-environment interaction is difficult, especially when the tool is compliant, and the robot cannot be expected to have the full geometry and physical properties (e.g., mass, stiffness, and friction) of all the tools it must use. We propose a framework that learns to predict the effects of a robot's actions on the contact between the tool and the environment given visuo-tactile perception. Key to our framework is a novel *contact feature* representation that consists of a binary contact value, the line of contact, and an end-effector wrench. We propose a method to learn the dynamics of these contact features from real world data that does not require predicting the geometry of the compliant tool. We then propose a controller that uses this dynamics model for visuo-tactile contact servoing and show that it is effective at performing scraping tasks with a spatula, even in scenarios where precise contact needs to be made to avoid obstacles.

Keywords: Contact-Rich Manipulation, Multi-Modal Dynamics Learning

1 Introduction

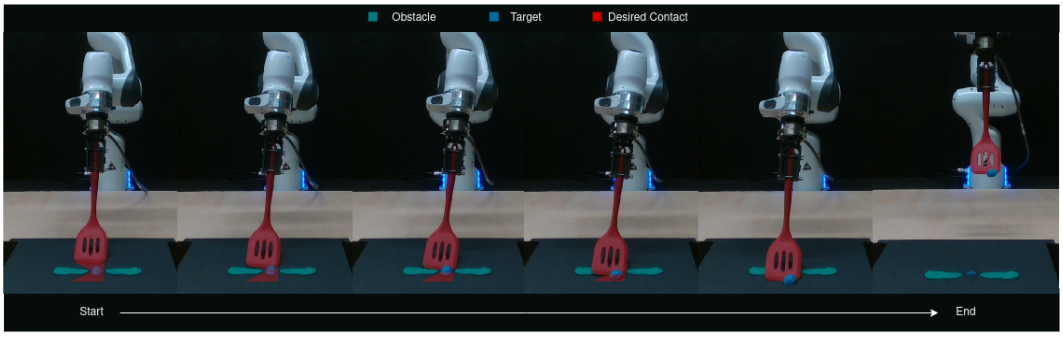

Figure 1: We present a method for *extrinsic contact servoing*, i.e., controlling contact between a compliant tool and the environment. Our method is able to complete the requested contact trajectory, avoiding contact with surface obstacles, and successfully scrape the target object. Note that to do this the spatula must be tilted so that only a corner of it is in contact.

Many manipulation tasks require the robot to control the contact between a grasped tool and the environment. The ability to reason over and control this *extrinsic* contact is crucial to enabling helpful robots that can scrape a frying pan with a spatula, eraser or wipe a surface [1], screw a bottle cap onto a bottle [2], perform peg-in-hole assemblies [3, 4], and perform many other tasks.

Submitted to the 6th Conference on Robot Learning (CoRL 2022). Do not distribute.

In this work, we seek to address the problem of controlling the extrinsic contact between a grasped *compliant* tool (e.g. a spatula) and the environment. In general, the robot cannot expect to have the full geometry and physical properties (e.g., mass, friction, stiffness) of all the tools it must use or the geometries of the environments it must manipulate in. Instead, the robot must utilize multimodal sensory observations, such as pointclouds and tactile feedback, to act on the environment.

In recent years, learning-based methods have become increasingly popular to address the complexities of robotic manipulation, including for contact-rich tasks [5]. These methods can be loosely grouped into model-free methods, that directly learn a policy [3, 2, 6], and model-based methods, that learn system dynamics [7, 8, 9]. By focusing on modeling system dynamics, model-based methods can plan to reach new goals without retraining, and are often more data-efficient [9]. Therefore, we propose learning the dynamics of our system to solve the extrinsic contact servoing task.

It is not obvious which representation to use for these dynamics. Fully recovering tool or environment geometries from visual data [10, 11] and tactile feedback [12] has been widely explored, with recent extensions to compliant geometries [13]; however, even if the system can be fully identified, contact models to resolve interactions can have limited fidelity [14]. On the other hand, learned dynamics representations can be difficult to interpret and require demonstrations or observations from the desired state to specify goals [7, 15]. Instead, we propose a novel *contact feature representation* for our learning method that focuses on tool-environment interaction and bypasses explicitly modeling the whole system. We represent the contact configuration as 1) a **binary contact mode** (indicating if the system is in contact); 2) a **contact geometry** (as a line in 3D space); and 3) an **end-effector wrench**.

We propose a learning architecture to model the dynamics of the proposed contact representation from raw sensory observations over candidate action trajectories. We propose structuring the model as a latent space dynamics model with a decoder that recovers the contact state. We also propose an impedance action offset term in the dynamics that allows us to accurately propagate robot poses, despite robot impedance. To provide labels to our model, we collect self-supervised data on a 7DoF Franka Emika Panda, using sensor data to automatically label the contact state.

We validate our proposed method by completing various desired contact trajectories on the real robot system. We first show that our method can track diverse desired contact trajectories in the absence of obstacles. Next, we demonstrate that we can utilize extrinsic contact servoing to scrape a target object from the table, while handling occlusions and avoiding contact with obstacles.

In summary, we make the following contributions:

- We present a framework for modeling compliant tool-environment contact interactions by learning contact feature dynamics.
- We propose a learned model architecture to capture the dynamics of contact features, trained in a supervised fashion using real world self-supervised data.
- We design and demonstrate a controller using the contact feature dynamics to realize diverse goal trajectories, including in the presence of obstacles.

2 Related Work

Existing research has investigated the task of recovering contact locations. Manuelli et al. [16] localize point contacts on a rigid robot with known geometry by employing a particle filtering approach to update a set of candidate contact locations based on force torque sensing. Kim et al. [4] and Ma et al. [17] model contact between a grasped rigid object and the environment by assuming stationary line contacts and modeling the deformation of a GelSlim gripper. The estimated line contact is then used in a Reinforcement Learning (RL) policy. Neither of these methods extends to compliant tools and neither models the dynamics of the contact configuration.

Other works explore *tactile servoing* methods, where contact at the sensor is driven to a desired configuration. Li et al. [18] use a large tactile pad and define contact configuration features of

objects pressed against the sensor. They manually construct a feedback controller based on these features and use it to drive contacts to desired configurations. Sutanto et al. [19] use a smaller profile tactile sensor and learn the dynamics of a learned latent space. They then employ a MPC scheme to drive contacts to desired configurations on the sensor. Both of these works assume contact is happening at the sensing location. We, on the other hand, seek to servo *extrinsic* contacts, where we do not get direct sensing at the point of contact.

3 Problem Formulation

We parameterize our contact feature as a binary contact indicator $c^b \in \{0, 1\}$, used to indicate whether the tool is in contact, a contact line[1] $\boldsymbol{c}^l \in \mathbb{R}^{2\times 3}$ representing the contact geometry between the tool and the environment, and an end effector wrench $\boldsymbol{c}^w \in \mathbb{R}^6$. The geometry $\boldsymbol{c}^l$ is only active when the tool is in contact $c^b = 1$. The contact representation allows extrinsic contact goals to be expressed as *desired contact trajectories* $G = [\boldsymbol{g}_1, \boldsymbol{g}_2, \ldots, \boldsymbol{g}_L]$, where each $\boldsymbol{g}_i \in \mathbb{R}^{2,3}$ is a desired contact line to reach. We assume that contact should be maintained throughout the task.

We formulate extrinsic contact servoing as a model predictive planning problem, given observations of the current state of the system $\boldsymbol{o}_0$. For a given horizon T, we select the next T desired contact lines $[\boldsymbol{g}_{i+1}, \ldots, \boldsymbol{g}_{i+T}] \subseteq G$ to be our current contact goal sequence. The planning problem is:

$$\min_{\boldsymbol{a}_{0:T-1}} \sum_{t=1}^{T} d(\boldsymbol{c}_t^l, \boldsymbol{g}_{i+t}) \tag{1}$$

$$\text{s.t. } c_t^b = 1, \forall t \in [1, T] \tag{2}$$

$$\{c_{0:T}^b, \boldsymbol{c}_{0:T}^l, \boldsymbol{c}_{0:T}^w\} = g(\boldsymbol{o}_0, \boldsymbol{a}_{0:T-1}) \tag{3}$$

Here g is a model describing the *contact feature dynamics*. The binary constraint ensures that the tool remains in contact while the cost function d measures the distance between the two contact lines, as the average Euclidean distance between the line endpoints. Finally, if $d(\boldsymbol{c}_1^l, \boldsymbol{g}_i + 1) < \epsilon$ we increment i, thus moving to the next sequence of desired contact lines for the next round of planning.

4 Method

4.1 Contact Feature Dynamics Model

To solve our constrained optimization Eq. 1, we require a model g which can map from raw observations $\boldsymbol{o}_0$ and a proposed action trajectory $\boldsymbol{a}_{0:T-1}$ to the resulting contact states $\{c_{0:T}^b, \boldsymbol{c}_{0:T}^l, \boldsymbol{c}_{0:T}^w\}$. We propose modeling the contact feature dynamics as a deep neural network.

We assume access to a pointcloud $\boldsymbol{v}_0$ and input wrench $\boldsymbol{h}_0$ measured at the robot's wrist as our observations, $\boldsymbol{o}_0 = (\boldsymbol{v}_0, \boldsymbol{h}_0)$. Note that end effector wrench is both an input to our method and part of the contact state; predicting future wrench aids the representation learning and provides expected wrenches for planning.

We perform all learning in the local end effector frame. We transform the pointcloud to the end effector frame ${}^{EE_0}\boldsymbol{v}_0$ and clip to a $0.5m^3$ bounding box region around the end effector that contains the contact event. We similarly predict our contact lines in the current end effector frame, ${}^{EE_t}\boldsymbol{c}_t^l, \forall t \in [1, T]$. Learning in the end effector frame provides invariance in the visual domain to translations and rotations of the end effector and removes distractors that do not contribute to the contact state, such as the rest of the robot arm or the scene background.

Our contact feature dynamics model (Figure 2) has three components: an encoder e which maps from raw observations to a learned latent space, a decoder d which maps from the latent space to the contact state, and a dynamics model f which captures dynamics in the latent space. We parameterize the models by a set of learned weights θ.

[1] In this paper we only consider line contacts between the tool and environment, but we believe our method is straightforward to extend to patch contacts by using a richer contact descriptor, e.g. the convex hull of a set of points.

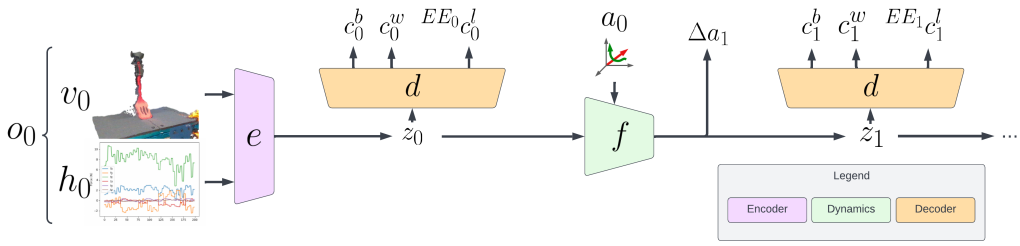

Figure 2: Our proposed contact feature dynamics model. Our architecture embeds raw observations into a latent space where dynamics can be unrolled. We then decode the contact state from the latent space. To account for our impedance controller, we also predict an action offset term which "learns the impedance" in order to relate predicted contact geometries to the world frame.

We start by embedding the current observations into the latent space with our encoder $\hat{\boldsymbol{z}}_0 = e(\boldsymbol{v}_0, \boldsymbol{h}_0)$. We unroll actions in the latent space as the contact state alone has insufficient contextual information (e.g. end-effector pose and local geometry information) to predict the next contact state. Because we predict the tth contact state in the current end effector frame EE_t, an important consideration when designing our dynamics model is being able to accurately recover this frame. Due to the contactful nature of our task, the low-level controller that executes $\boldsymbol{a}_t$ is an impedance controller, meaning the resulting action may different from the commanded action. To account for this, we predict an additional term from our dynamics model $\Delta\hat{\boldsymbol{a}}_{t+1}$, which is an $SE(3)$ transformation that predicts the offset between the commanded and realized next end effector pose. Thus, our dynamics model predicts, $\hat{\boldsymbol{z}}_{t+1}, \Delta\hat{\boldsymbol{a}}_{t+1} = f(\hat{\boldsymbol{z}}_t, \boldsymbol{a}_t)$. This allows us to construct the following recursive estimate of our end effector frame ${}^{W}\hat{T}_{EE_{t+1}} = {}^{W}\hat{T}_{EE_t} T(\boldsymbol{a}_t) T(\Delta\hat{\boldsymbol{a}}_{t+1})$, where ${}^{W}\hat{T}_{EE_t}$ is the $SE(3)$ transformation describing the pose of the end effector at time t. We know the initial transform ${}^{W}\hat{T}_{EE_0}$ from our robot proprioception, $T(\boldsymbol{a}_t)$ provides the transformation for the action command, and $T(\Delta\hat{\boldsymbol{a}}_{t+1})$ for the predicted offset term. To enforce valid $SE(3)$ predictions we predict rotations in the axis-angle representation [20].

Finally, we recover our contact state estimates with our decoder d given the latent state $\hat{\boldsymbol{z}}_t$: $\hat{c}_t^b, {}^{EE_t}\hat{\boldsymbol{c}}_t^l, \hat{\boldsymbol{c}}_t^w = d(\hat{\boldsymbol{z}}_t)$. We can then recover the predicted contact line in the world frame by composing with the estimate of the end effector frame transformation ${}^{W}\hat{T}_{EE_t}$.

An overview of the model architectures is shown in Fig. 2 and full architecture details can be found in Appendix A.

4.1.1 Training Loss

We train our model on rollouts of the system, where a single example is a sequence $[\boldsymbol{v}_t, \boldsymbol{h}_t, \boldsymbol{a}_t, \Delta\boldsymbol{a}_t, \boldsymbol{c}_t^l, \boldsymbol{c}_t^w, c_t^b]_{t=0}^T$. We define the loss over the example as:

$$\mathcal{L}_\theta = \Big(\sum_{t=0}^{T} BCE(\hat{c}_t^b, c_t^b) + \alpha \cdot c_t^b \cdot MSE(\hat{\boldsymbol{c}}_t^l, \boldsymbol{c}_t^l) + \beta \cdot MSE(\hat{\boldsymbol{c}}_t^w, \boldsymbol{c}_t^w)\Big) \tag{4}$$

$$+ \Big(\sum_{t=1}^{T} \rho \cdot MSE(\Delta\hat{\boldsymbol{a}}_t, \Delta\boldsymbol{a}_t) + \gamma \cdot MSE(\hat{\boldsymbol{z}}_t, e(\boldsymbol{v}_t, \boldsymbol{h}_t))\Big) \tag{5}$$

Here BCE is the Binary Cross Entropy classification loss and MSE is the Mean Square Error regression loss. α, β, ρ and γ are loss weighting terms. The first four loss terms are prediction losses over the contact mode, contact geometry, end effector wrench, and impedance offset transformation. The final loss term is a latent consistency loss, which encourages latent rollouts to match the latent state yielded by encoding future observations.

4.2 Extrinsic Contact Servoing Controller

We propose to solve our planning problem using Model Predictive Path Integral (MPPI), which has been shown to be effective for continuous control tasks where sampling is cheap and parallelizable

(e.g., neural network representations) [21]. We convert our binary constraint to a penalty, penalizing a trajectory if it yields actions that lead out of contact. Pairing this with the contact line prediction loss yields the following final cost function:

$$\sum_{t=1}^{T} d({}^{W}\hat{\boldsymbol{c}}_t^l, \boldsymbol{g}_{i+t}) + \phi \cdot \begin{cases} |\hat{c}_t^b - \psi| & \text{if } \hat{c}_t^b < \psi \\ 0 & \text{o.w.} \end{cases} \tag{6}$$

The constraint is violated if the likelihood of binary contact is below the classification threshold ψ, in which case we penalize by the distance to the threshold. ϕ weights the penalty against the contact line loss. With this cost function we apply MPPI to yield the next action and execute it on the robot.

4.3 Extrinsic Contact Dynamics Labeling

Our contact dynamics training loss in Eq. 4 requires ground truth contact state labels $(\boldsymbol{c}_t^l, c_t^b, \boldsymbol{c}_t^w)$ at time t. As accurate simulation of contactful interactions is challenging, we propose a method of data acquisition directly in the real world. To generate contact line labels, we use a high resolution, low frequency scanner, a Photoneo PhoXi 3D Scanner, to generate high quality scans of the contact interaction. Using these scans, we generate contact line labels by filtering points just above the table, clipped to the area around the end effector. We then cluster these points to remove noisy points on the tabletop and generate the contact line $\boldsymbol{c}_t^l$ by selecting the two furthest points in the cluster. See Appendix B for examples of contact labels. We use a force torque sensor to identify the contact state wrench $\boldsymbol{c}_t^w$ and threshold the wrench to identify binary contact c_t^b automatically. This setup is flexible to new robots and tools, so long as a force torque sensor is available on the robot wrist.

5 Results

Our experiments seek to answer three questions: First, can we accurately model the contact feature dynamics when there are no objects on the table? Second, can we perform extrinsic contact servoing in this environment? Third, can our controller be applied to an object scraping task, where it must handle visual occlusions and reaction forces arising from contact with the target object?

5.1 Experimental Setup

We test our method on a Franka Emika Panda with a rigidly-mounted compliant spatula at the end effector (Fig. 1). For our observations $\boldsymbol{o}_0$, we use pointclouds $\boldsymbol{v}_0$ from an Intel Realsense D435 sensor[2] and mount an ATI Gamma Force/Torque sensor between the end effector and tool. We use the last four wrench values received after the previous action completed as the tactile input $\boldsymbol{h}_0$. To collect our dataset, we use a random action policy with a heuristic to encourage contact between the tool and the spatula. No other objects are on the tabletop during data collection to allow proper data supervision, as detailed in Sec. 4.3. We collect 489 random trajectories of data (22005 rollout sequences) to train and test the model. We use 80/10/10 train, validate, test splits of the trajectories.

5.2 Modeling Contact Feature Dynamics

We first investigate the ability of our model to capture the contact feature dynamics exhibited in our dataset. Due to the novelty of our representation, to the authors' knowledge, there are no existing visuo-tactile extrinsic contact feature learning methods that are directly comparable. Thus we focus our comparisons on ablations to understand the importance of the components of our proposed neural network model. We train three variations. First is the full model, as described in Sec. 4.1, hereafter called "Full Model." Second, to understand the importance of modeling the impedance of the robot, we ablate the impedance offset action prediction, thus we propagate the end effector frame only with the commanded action. We call this method "No Impedance." Finally, we investigate our model trained only on visual input data, called "Vision-Only," to demonstrate the drop in performance when we do not use force data.

We train our contact feature dynamics on our dataset with a rollout horizon of $T = 3$. All methods are trained with the Adam optimizer [22] until convergence on the validation set. We set $\alpha = 100.0, \beta = \gamma = \rho = 0.1$ in our loss term in Eq. 4 to balance the scale of the terms.

[2] We don't use the high-fidelity Photoneo scan as it is a very low-frequency scanner.

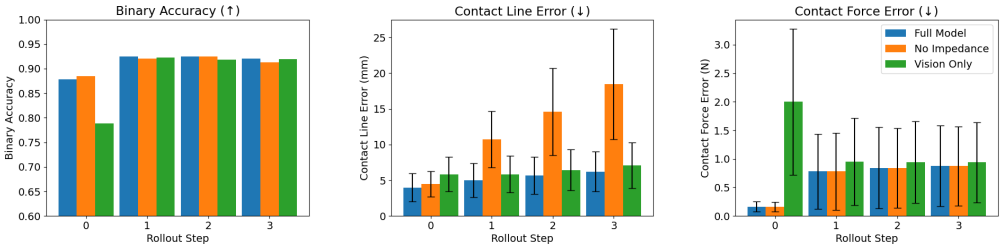

(a) (b) (c)

Figure 3: Contact Feature Dynamics Performance: without our impedance offset term, contact line quickly drifts; without tactile inputs it is difficult to accurately model end effector wrenches.

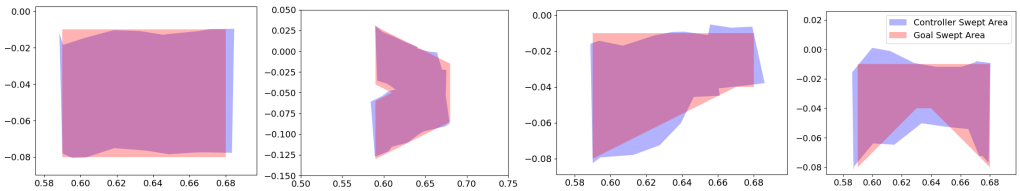

(a) Straight (b) V-Shape (c) Tilt (d) Tilt Back

Figure 4: Qualitative Scraped Areas: our controller (blue) is able to closely match desired contact trajectories (red).

We compare the prediction performance of the models on the test split of the dataset (2250 examples) in Fig. 3 (see Appendix C for torque prediction results, whose results are very similar to force results). Our results show that we can effectively learn contact feature dynamics, with the Full Model achieving around 90% accuracy, about 5mm contact line error, and less than a Newton of contact force error. Fig. 3b shows that without the impedance offset term, the contact line estimate drifts significantly, due to the difference between commanded and realized actions, which makes the estimate of the end effector pose increasingly bad. The Full Model shows slight improvement over the Vision-Only method throughout. The vision-only does especially poor in contact force estimation (Fig. 3c), as it does not get contact wrenches as input.

5.3 Extrinsic Contact Servoing

Next we investigate how our proposed controller performs following specified contact trajectories. Methods that learn latent dynamics for model predictive planning in an unsupervised fashion [7, 9] could be applied to our system to learn the dynamics and potentially plan with them, however, specifying contact trajectory goals in the learned latent spaces would require significant new capabilities for these models. Similarly, methods that recover tool/environment geometries [13, 10] from sensory data would also require significant changes to predict our contact dynamics. We leave these investigations to future work, and focus here on the performance of our proposed methodology.

Obstacle-Free: We start by attempting to servo along four different contact trajectories in the obstacle-free environment. The desired contact trajectories are shown in Fig. 4, and explore translation of contact as well as cases where the robot must tilt the tool to achieve a contact smaller than the width of the tool. We use the same labeling technique introduced in Sec. 4.3 to get ground truth contact trajectories executed by the controller.

We use the controller described in Sec. 4.2, with $\psi = 0.45, \phi = 0.05$. All planning is done with the Full Model. To investigate the planning performance, we run the controller ten times per trajectory. We use the Intersection over Union (IoU) of the desired and swept contact areas as our performance metric. We construct the goal contact area by sweeping the space between the specified goal contact lines and the realized controller swept area by assuming that the space between two consecutive contact states was swept out if the two states were both in contact.

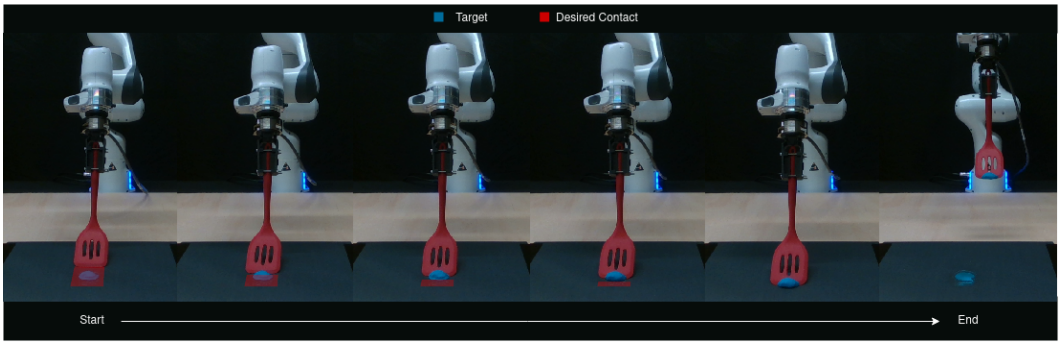

Figure 6: Example of extrinsic contact servoing execution for the "Straight" target obstacle scrape experiment. Our method is able to accurately servo along the desired contact trajectory and successfully scrape the target.

We show qualitative examples of controller realized scrapes compared to the goal scrapes in Fig. 4. We see that the controller is able to closely match the desired swept areas, including in the difficult tilting problems. Fig. 5 shows the average IoU scores over the ten controller rollouts on each desired contact trajectory. For the two trajectories focused on contact line translation, we achieve high IoU (near 80%). For the tilting case, we see a drop in performance, likely because it is more difficult to control transitions and smaller contact shapes.

With Obstacles: We next examine our method's robustness to visual occlusions and reaction forces arising from contact with a target object to be scraped. This task is common in construction, cooking, and cleaning. A deformable and slightly adhesive material (Playdough) is pressed onto the surface and a contact trajectory is specified through the object. In one case, the target object is alone on the tabletop (Fig. 6), and thus we specify a contact trajectory using the full width of the tool. In the second scenario, obstacles are on the table near the target object (Fig. 1), thus we must specify a contact trajectory that avoids them.

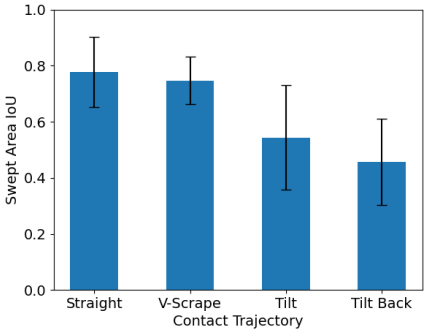

Figure 5: Extrinsic Contact Servoing IoU: our controller realizes diverse contact goal trajectories with high overlap. We see reduced performance for the challenging case of contact trajectories that require tilting.

Besides running our Full Model in these scenarios, we investigated enhancements of the method to aid its performance in the presence of visual occlusions and object reaction forces. First, we applied data augmentation to our training dataset, randomly generating ellipsoids in the pointcloud and using a hidden point removal algorithm [23] to provide corresponding occlusions in the original point cloud. See Appendix B for examples of augmented inputs. We call this method "Full Model + Aug." Second, we investigate using the difference between the predicted and observed wrenches $\Delta \boldsymbol{w} = \hat{\boldsymbol{c}}_t^w - \boldsymbol{c}_t^w$ to derive an action offset to compensate for the extra wrench experienced by the robot. From Δw we derive an action that will counteract this wrench offset $\hat{\boldsymbol{a}}_t = \frac{\Delta \boldsymbol{w}}{\boldsymbol{k}_p}$. $\boldsymbol{k}_p$ are the pose gains of the impedance controller. The offset action is composed with the original action from the controller. We call this method "Full Model + Aug + Wrench Offset."

We use two metrics. First, we measure the approximate mass of the target object to be scraped before and after scraping and determine the percentage of mass successfully removed. Second, we compare the 2D footprint of the material before and after scraping and report the percentage of the footprint successfully removed. The second is a more challenging metric, since even a slightly wrong scrape will leave residue. The quantitative results over 5 runs of each method in each experiment setup are

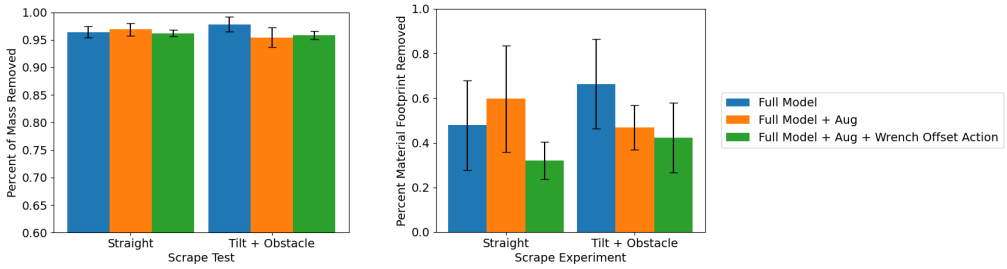

(a) Percent Mass Removed

(b) Percent Footprint Removed

Figure 7: Target Scraping Results: Our Full Model and variations for addressing visual occlusions and reaction forces arising from contact with the target object perform comparably on both metrics over 5 trials on each experiment.

shown in Fig. 7. Examples of scrape executions are shown in Fig. 1 and Fig. 6. See Appendix C for more examples of scrape results.

We see that in each case, all methods were able to remove over 95% of material mass on average and about 40-60% of material's footprint from the tabletop. Surprisingly, we don't see a consistent improvement training our model with visual occlusions or adding action offsets. The Full Model's robustness to occlusions here could be due to the fact that we use a single tool in these experiments, and thus it may be sufficient in most cases to capture the location of the table with respect to the end effector in order to estimate the contact line. Even with visual occlusions near the tool contact, it is likely our method can still recover the relative pose of the tabletop from the surrounding points. The lack of clear improvement from the wrench offset action may be due to the fact that it is sufficient to be able to replan, as we do at every step with our MPPI controller.

6 Limitations and Conclusion

Limitations: A common failure mode for our method is in controlling contact when the tool is tilted, where it is more likely for the method to yield actions that take the robot out of contact. This could, in part, be due to data imbalance. It is difficult to fully cover the space of desired contact configurations. In future work, we are interested in utilizing online learning [9] or curiosity [24] to more effectively cover the space of contacts in our dataset.

There are cases where tool to environment contact is not represented well as a contact line. Consider, for instance, the task of removing a nail with a hammer or putting the cap on a bottle. Extending our contact feature dynamics to these methods will requiring expanding our representation learning method to consider these more diverse contact specifications and consider more complex contact modes. One approach to handling these more complex interactions is to leverage collocated tactile sensing as the robot makes contact with the environment [25, 26]. The resulting tactile cues can provide sufficient information to characterize these more complex interactions.

Finally, our method relies upon supervision. For future contact-rich tasks of interest, the need for labels could become more costly. We hope to investigate how we can remove reliance upon supervision by exploring few/zero shot generalization [27, 28] and domain randomization techniques [29].

Conclusion: Our approach simplifies contact rich interactions for compliant tool manipulation, by avoiding the necessity for full system identification while maintaining interpretability and accuracy by explicitly modeling the *contact state* of the system and how it evolves. In the future, we wish to investigate our method's applicability to other tasks where full state estimation is difficult, but the contact state is crucial, such as wiping with a cloth. Additionally, we wish to investigate our methods ability to generalize to multiple tools and handle more severe occlusions and forceful interactions.

References

[1] R. Martín-Martín, M. A. Lee, R. Gardner, S. Savarese, J. Bohg, and A. Garg. Variable impedance control in end-effector space: An action space for reinforcement learning in contact-rich tasks. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1010–1017. IEEE, 2019.

[2] S. Levine, C. Finn, T. Darrell, and P. Abbeel. End-to-end training of deep visuomotor policies. *The Journal of Machine Learning Research*, 17(1):1334–1373, 2016.

[3] M. A. Lee, Y. Zhu, K. Srinivasan, P. Shah, S. Savarese, L. Fei-Fei, A. Garg, and J. Bohg. Making sense of vision and touch: Self-supervised learning of multimodal representations for contact-rich tasks. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 8943–8950. IEEE, 2019.

[4] S. Kim and A. Rodriguez. Active extrinsic contact sensing: Application to general peg-in-hole insertion. *arXiv preprint arXiv:2110.03555*, 2021.

[5] O. Kroemer, S. Niekum, and G. D. Konidaris. A review of robot learning for manipulation: Challenges, representations, and algorithms. *Journal of machine learning research*, 22(30), 2021.

[6] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *The International journal of robotics research*, 37(4-5):421–436, 2018.

[7] W. Yan, A. Vangipuram, P. Abbeel, and L. Pinto. Learning predictive representations for deformable objects using contrastive estimation. In *Conference on Robot Learning*, pages 564–574. PMLR, 2021.

[8] P. Mitrano, D. McConachie, and D. Berenson. Learning where to trust unreliable models in an unstructured world for deformable object manipulation. *Science Robotics*, 6(54):eabd8170, 2021.

[9] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson. Learning latent dynamics for planning from pixels. In K. Chaudhuri and R. Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2555–2565. PMLR, 09–15 Jun 2019. URL https://proceedings.mlr.press/v97/hafner19a.html.

[10] L. Mescheder, M. Oechsle, M. Niemeyer, S. Nowozin, and A. Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4460–4470, 2019.

[11] C. B. Choy, D. Xu, J. Gwak, K. Chen, and S. Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *European conference on computer vision*, pages 628–644. Springer, 2016.

[12] D. Watkins-Valls, J. Varley, and P. Allen. Multi-modal geometric learning for grasping and manipulation. In *2019 International conference on robotics and automation (ICRA)*, pages 7339–7345. IEEE, 2019.

[13] Y. Wi, P. Florence, A. Zeng, and N. Fazeli. Virdo: Visio-tactile implicit representations of deformable objects. *arXiv preprint arXiv:2202.00868*, 2022.

[14] N. Fazeli, S. Zapolsky, E. Drumwright, and A. Rodriguez. Fundamental limitations in performance and interpretability of common planar rigid-body contact models. In *Robotics Research*, pages 555–571. Springer, 2020.

[15] L. Manuelli, Y. Li, P. Florence, and R. Tedrake. Keypoints into the future: Self-supervised correspondence in model-based reinforcement learning. *arXiv preprint arXiv:2009.05085*, 2020.

[16] L. Manuelli and R. Tedrake. Localizing external contact using proprioceptive sensors: The contact particle filter. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5062–5069. IEEE, 2016.

[17] D. Ma, S. Dong, and A. Rodriguez. Extrinsic contact sensing with relative-motion tracking from distributed tactile measurements. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11262–11268. IEEE, 2021.

[18] Q. Li, C. Schürmann, R. Haschke, and H. J. Ritter. A control framework for tactile servoing. In *Robotics: Science and systems*. Citeseer, 2013.

[19] G. Sutanto, N. Ratliff, B. Sundaralingam, Y. Chebotar, Z. Su, A. Handa, and D. Fox. Learning latent space dynamics for tactile servoing. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3622–3628. IEEE, 2019.

[20] A. Byravan and D. Fox. Se3-nets: Learning rigid body motion using deep neural networks. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 173–180, 2017. doi:10.1109/ICRA.2017.7989023.

[21] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou. Information theoretic mpc for model-based reinforcement learning. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1714–1721, 2017. doi:10.1109/ICRA.2017.7989202.

[22] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

[23] S. Katz, A. Tal, and R. Basri. Direct visibility of point sets. In *ACM SIGGRAPH 2007 papers*, pages 24–es. 2007.

[24] S. Rajeswar, C. Ibrahim, N. Surya, F. Golemo, D. Vazquez, A. Courville, and P. O. Pinheiro. Haptics-based curiosity for sparse-reward tasks. In *Conference on Robot Learning*, pages 395–405. PMLR, 2022.

[25] I. Taylor, S. Dong, and A. Rodriguez. Gelslim3. 0: High-resolution measurement of shape, force and slip in a compact tactile-sensing finger. *arXiv preprint arXiv:2103.12269*, 2021.

[26] A. Alspach, K. Hashimoto, N. Kuppuswamy, and R. Tedrake. Soft-bubble: A highly compliant dense geometry tactile sensor for robot manipulation. In *2019 2nd IEEE International Conference on Soft Robotics (RoboSoft)*, pages 597–604. IEEE, 2019.

[27] S. Gidaris and N. Komodakis. Dynamic few-shot visual learning without forgetting. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4367–4375, 2018.

[28] F. Zhao, J. Zhao, S. Yan, and J. Feng. Dynamic conditional networks for few-shot learning. In *Proceedings of the European conference on computer vision (ECCV)*, pages 19–35, 2018.

[29] P. Mitrano and D. Berenson. Data augmentation for manipulation. *arXiv preprint arXiv:2205.02886*, 2022.