Vision-Based Contact Localization Without Touch or Force Sensing

Anonymous Author(s)
Affiliation
Address
email

Abstract: Contacts play a critical role in most manipulation tasks, yet robots today do not have any approach to reliably sense contacts in general settings. Force torque and touch sensing are limited by how little of the world they can sense and by sensor drift, so that today's contact perception approaches built on them require restrictive assumptions or prior knowledge of object geometries and frequent re-calibration. We propose a challenging vision-based extrinsic contact localization task: with only a single RGB-D camera view of a robot workspace, identify when and where an object held by the robot contacts the rest of the environment. We show that careful task-attuned design is critical for a neural network trained in simulation to discover solutions that transfer well to a real robot. Our final approach `im2contact` demonstrates the promise of versatile general-purpose contact perception from vision alone, performing well for localizing various contact types (point, line, or planar; sticking, sliding, or rolling; single or multiple), and even under occlusions in its camera view. Appendix and video results can be found at: https://sites.google.com/view/im2contact/home

Keywords: contact perception, manipulation, vision-based

1 Introduction

Perceiving and reacting to contact is critical for performing manipulation tasks [1–5]. Consider what happens when a person puts a book on a crowded shelf: they hold the book and aim for a gap until it meets resistance, jostle the book to make room, then press sideways to line up the book against its neighbor, and finally slide the book snugly into place. Throughout, the key events they must track all have to do with the physical contacts between various surfaces: when and where they are made, broken, and transition from sticking to sliding. What means does a robot have today to sense such contacts between its body, its tools, and external objects?

Contact perception techniques for robots today operate mainly from force torque and touch sensing. Force torque (F/T) sensors [6] located at the robot's joints can inform model-based contact estimation techniques. However, this estimation problem is under-determined. To see this, consider a two-fingered manipulator in two contact configurations: applying a 10 N force to one finger, versus applying a 5 N force to each finger. A wrist F/T sensor located directly behind gripper's midpoint, is aliased: it senses identical forces in both cases. For such direct contacts with the robot, aliasing can be partially resolved with proximal touch sensors [7–10] applied to contacting surfaces on the robot, but over only a limited area. Further, consider as in our introductory example, a robot holding a book that contacts a bookshelf. Such "extrinsic" contact is still only sensed indirectly at the fingers, and remains impossible to resolve. To overcome these issues, today's contact estimation techniques reduce the number of unknowns by operating under restrictive assumptions, such as about the number and types of contacts, and the shapes of the various objects involved (Sec 2). Finally, even when operating within these assumptions, F/T and touch sensors are often expensive, and drift or deteriorate quickly over use [10–13] requiring frequent and cumbersome re-calibration.

Submitted to the 7th Conference on Robot Learning (CoRL 2023). Do not distribute.

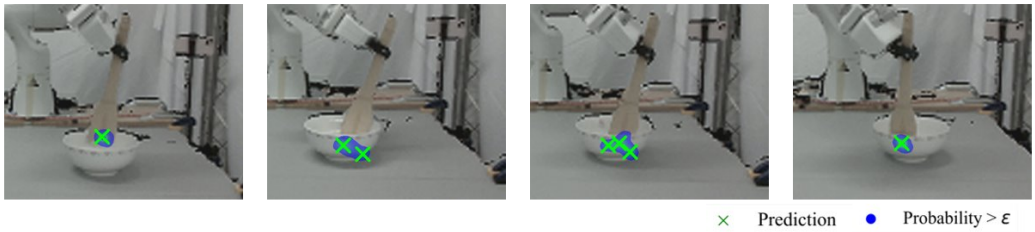

Figure 1: Sequence of predictions of our method `im2contact` for localizing extrinsic contact in an image between a grasped object (such as a spatula) and the environment (such as a bowl)

To address both the lack of global information available to F/T and touch sensors, and the sensor drift issue, we consider instead operating with camera images. Cameras can see the entire scene, and function without deterioration for long durations. Since contacts are a function of the shapes and movements of the objects and the environment, contact estimation would be easy if we could first get perfect 3D shape and pose estimates from vision. However, we argue that contact from 3D would be a poor solution. First, given occlusions around the critical contact regions and arbitrary-shaped objects, near-perfect shape and pose estimates would require impractically many camera views. Next, the often-crucial binary distinction between near-contacts and true contacts can be hard to capture in a finite-resolution image.

Rather than infer contacts from intermediate visual representations such as shapes and poses, we propose `im2contact`, a more direct approach for data-driven, model-free visual contact location estimation at the outputs of a deep neural network. To get contact labels, we train inside a simulator, but this introduces an inevitable domain gap. After finding that merely training with depth inputs does not suffice to overcome this gap, we propose three carefully constructed fixes: we crop out irrelevant portions of the image, specify the grasped object through an additional reference image input, and make motion information available by feeding optical flow as another input. In our zero-shot transfer evaluations on a real robot, `im2contact` models demonstrate the possibility and promise of versatile general-purpose contact perception from vision alone, performing well for localizing various contact types (point, line, or planar; sticking, or sliding; single or multiple), and even under occlusions in its camera view.

2 Related Work

The robotics community has long recognized the crucial role of external contact sensing in manipulation, particularly of contacts between the manipulator and environment (e.g. [14] [15] [16]), coined by Bicchi et al. [17] as "intrinsic contact". However, "extrinsic contact," e.g. between a held object and the environment, is similarly useful, but notably more difficult. Prior work on extrinsic contact has typically found success through strong assumptions which ultimately limit the potential scope of the results. For example, related work has used specialized tactile sensors and the assumption and/or enforcement of a known contact configuration [18–20], or specialized sensors and a limited set of possible held objects [21]. Other approaches require use of exploratory motions designed to collect information [22] or privileged access to the geometry of the held object [23], [24].

Recent work, integrating a high-cost depth scanner together with force-torque sensing, has demonstrated the utility of extrinsic contact sensing for a scraping task, but restricted to a grasped object class of spatulas and line contacts with a planar surface [25].

In order to enable robots to freely extend their affordances using the tools they may encounter in unstructured and unmodeled settings, many of the above assumptions must be lifted. We take a step towards that goal via the choice of vision-based sensing with no explicit assumptions made on the number or types of contact, shapes of objects, and minimal access to privileged knowledge of the object or environment.

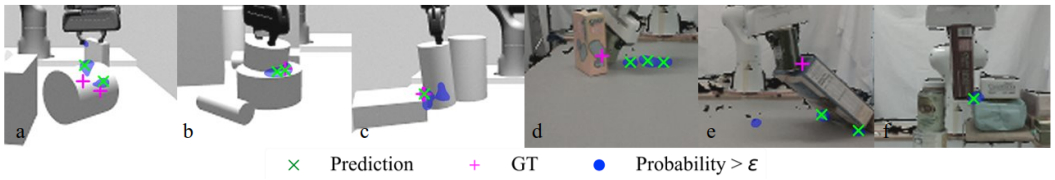

Figure 2: U-Net Depth generally performs well in sim but poorly on real in challenging cases such as: **(a,d)** occluded contact **(b,e)** ambiguous grasped object geometry **(c,f)** near contact with background object

3 Visually Localizing Extrinsic Contacts In General Manipulation Scenes

For a robot arm holding a grasped object, we are interested in localizing contact between the held object and the rest of the environment, as the robot arm moves in its workspace. Unlike much of the prior work described above, we do not assume access to prior information about the held object or the environment, other than the table.

As inputs at each time instant, we assume a HxW view of the robot workspace from a single fixed RGB-D camera as shown in Fig 3, alongside proprioceptive sensing of the robot state from joint encoders. This is a minimal sensing setup for vision-based robot manipulation, chosen to maximize the scope of our problem formulation.

As output, we would like our system to generate a HxW map of estimated contact locations, that can be overlaid on the camera image. To achieve this, we will train neural networks with standard pixel-wise binary cross-entropy classification objectives. This treats the the output contact map values at each pixel as contact probabilities, and maximizes the likelihood under the model of the annotated ground truth contact locations in the training data.

3.1 Simulated Training Data

It is not practical to obtain ground truth contact location annotations from real video, but fortunately simulators provide this information. To generate target contact maps for training, we project 3D contact points from the Gazebo simulator [26] into the camera frame.

Next, we generate training data in an episodic fashion. In each 15-second episode, we randomize the shapes and sizes, masses, friction coefficients, and initial poses of grasped and environmental objects. Shapes are chosen to be cylinders, spheres, and cuboids with random parameters. The grasped object is rigidly attached to the robot arm throughout the episode, as the robot end-effector moves to randomly set targets with a low-impedance controller to generate rich interaction data. We provide more details on our website. Our 4500 episodes (675000 frames at 10 fps) of training data span many types of extrinsic contacts: point, line, and plane contacts, instantaneous collisions and sustained sliding or rolling contacts along with simultaneous contact with multiple bodies.

3.2 A Simple Baseline To Illustrate The Difficulty of Sim-To-Real Transfer

To motivate our final method, we first showcase the difficulty of sim-to-real transfer with a baseline, U-Net Depth. U-Net Depth builds on the widely used U-Net [27, 28] architecture for image-to-image problems. In the U-Net, an encoder first assimilates information from over the entire 240x320 input image into a "bottleneck" representation of size 15x20x1024. Then, a decoder iteratively spatially upsamples this representation with the aid of skip connections from intermediate encoder layers, to finally produce an output over the original 240x320 input dimensions. To facilitate sim-to-real transfer and access 3D information, U-Net Depth uses depth images [29] as input rather than color images. We train using pixel-wise cross entropy.

U-Net Depth trains to near-zero training losses, and performs very well on held-out data in simulation. However, on real data, its performance deteriorates significantly. In Fig 2, we anecdotally

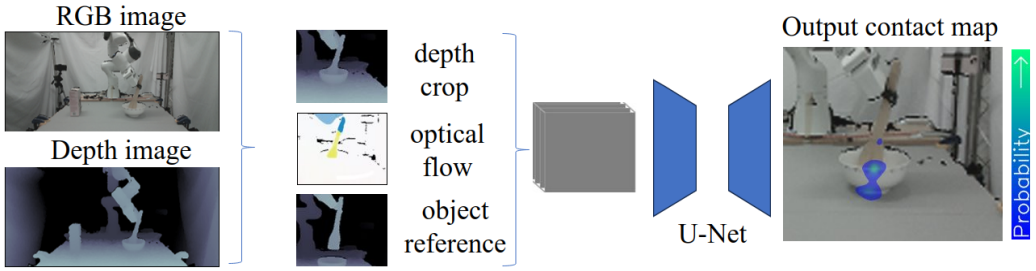

Figure 3: `im2contact` architecture. Depth and flow are both cropped and concatenated before being passed through the U-Net. The object reference is passed through a separate encoder whose output is concatenated at the bottleneck of the U-Net. The final output of the U-Net is a probability map over the original image dimensions.

note a few challenging scenarios where this baseline struggles in the real world: (a) under significant occlusion, it misses the occluded contact and predicts irrelevant false positives on the table, (b) when the geometry of the grasped object is ambiguous, it predicts contact between an extraneous box and table, and (c) it produces far more false positives in near-contact scenarios than in sim. We more thoroughly evaluate U-Net Depth in Sec 4.

3.3 Facilitating Generalizable Contact Localization With `im2contact`

To mitigate the failures of U-Net Depth on real data, we carefully reason about potential causes of poor generalization. Machine learning systems commonly over rely on spurious correlates of the target labels in the training data [30–32], which fails to generalize under distribution shift. During training, spurious correlates could distract the optimizer [32] from finding true "causes"; this can be mitigated by anticipating such distracting but irrelevant correlates in advance and removing them from the learner's inputs. Inputs might be also incomplete, lacking information on the true underlying cause of the outputs. We propose three simple improvements to U-Net Depth, one to remove distracting correlates, and two to add missing causal information into the inputs.

- **Depth Cropping (+ crop):** Assuming grasped objects that are not very large, all extrinsic contacts will occur close to the end-effector. With a calibrated camera and proprioceptive knowledge, it is easy to locate the end-effector in the image. We propose to focus the network on the most relevant regions by cropping the depth images into a 90x110 box around this end-effector location at each time, before feeding them to the U-Net. In addition, we add three channels for coordinate convolution as proposed by Liu et al. [33] to provide the network with the spatial relationship of the cropped image with respect to the original image.
- **Grasped Object Reference (+ obj-ref):** We have thus far reduced contact localization to an instantaneous task, performed from the current sensory observations. However, consider the scenario when the grasped object is occluded by the environment: if we do not know its spatial extent, it may be impossible to know whether there is any extrinsic contact. In other scenarios, during close contact, its can be difficult, from the depth map, to see where one object ends and the other begins, which is critical for predicting contact. We propose to provide such missing information through an additional input: a single reference depth image of the grasped object in the robot gripper, partially specifying the shape of the grasped object.
- **Optical Flow (+ flow):** From a single image, it may be difficult, or even impossible, to differentiate between contact and near contact. For example, a gap between the grasped object and the environment of a few millimeters is likely imperceptible without a perfect vantage. However, contact forces often induce motion in the environment, causing environmental objects to slide or roll. To capture these cues, we propose to use optical flow computed from the camera RGB images.

In the rest of the paper, we use the name `im2contact` for the combined approach: U-Net Depth + crop + obj-ref + flow. Fig 3 schematically depicts the network with all inputs and outputs.

Reducing To Discrete Contact Locations. Our training procedure generates contact probability maps, but it is convenient for evaluation and useful for downstream tasks to identify discrete contact locations. To do this, we adapt the greedy non-maximum suppression (NMS) technique [34–36], commonly used in object detection. To identify spatially separated peaks in the contact probability score map, we first reject all pixels scoring below ϵ=.01, then sort pixels by descending score, and iteratively reject pixels within a 5 px neighborhood of higher-scoring pixels. See website for details.

Implementation Details. The pixel-wise cross-entropy loss treats each pixel independently of all others and hence penalizes a 1 px deviation in predicted location equivalently to a 100 px deviation. To mitigate this, we spatially blur ground-truth contact maps before computing our training target, and find that this accelerates training. We monitor performance on held-out simulation data to implement early termination. We will publicly release all code and models via our website after the review period. For computing optical flow, we use the off-the-shelf RAFT [37] model.

4 Experiments

We test `im2contact` in simulation and in real to evaluate performance in realistic tabletop manipulation settings. We further test, via ablation studies, whether, and under what settings, our proposed changes (crop, obj-ref, and flow) do improve real performance. We include additional anecdotal experiments which push `im2contact` beyond its training settings, evaluating generalization.

Real Robot Data, Annotations, and Performance Metrics. We perform experiments in a tabletop environment with a Franka Emika Panda robot arm and an Intel Realsense L515 RGB-D camera, with reasonably well-matched simulation and real robot setups, as shown earlier in Sec 3. To facilitate some quantitative comparison, we manually annotate contact locations for 30 episodes (12,362 frames, of which 1/3 involve at least one contact) of real robot interaction data, spanning variations in the grasped object, environment objects, robot movements, and contact scenarios. Recall that for simulation experiments, annotated contacts are readily available as in Sec 3; we use 500 episodes (75,000 frames) for computing simulation performance metrics.

Our main metrics are precision, recall, and their harmonic mean, the F1 score. All three must be in $[0, 1]$, and higher is better. To compute them, we first follow the procedure from Sec 3.3 for reducing network output contact maps to discrete locations. We then match each ground truth contact point to a prediction with the Hungarian algorithm, and drop predictions where the ground truth is over 15px away as false positives. For our metrics, ground truth contacts with matches are true positives, and those without are false negatives. We run 5 random seeds for each method and report means and standard errors.

Quantitative Results on Simulation and Real Data. Fig 4 plots contact localization precision and recall in real (hollow circles) and simulated (solid circles) test data, for `im2contact`, the baseline U-Net Depth (Sec 3.2), and leave-one-out ablations of `im2contact` that in turn drop crop, obj-ref, and flow. First, for the simple U-Net Depth baseline, these results clearly validate our qualitative observations from Sec 3: precision and recall both deteriorate dramatically from sim to real. Next, U-Net Depth and `im2contact` are both among the best performing methods on sim data, but `im2contact` only drops marginally in performance on real data, so it remains among the best-performing methods in real data, while U-Net Depth performs worst.

Moving on to the leave-one-out ablations, "w/o crop" deteriorates nearly as much as U-Net Depth, while "w/o obj-ref" and "w/o flow" degrade more gracefully upon transfer to real data, crop contributes the most among the three components of `im2contact`. Fig 4 (middle) plots standard error for real data and the legend (right) lists F1 scores. From this initial coarse analysis with limited quantitative metrics, losing obj-ref marginally hurts precision, and losing flow does not significantly affect these scores. Note that these aggregate scores might not reflect performance in rare scenarios, of which there are many. We show full tables on our website.

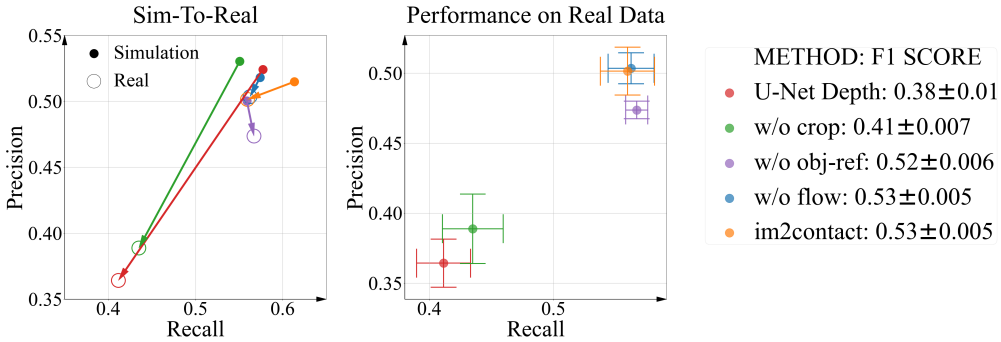

Figure 4: **(Left)** Degradation under sim-to-real transfer **(Middle)** Precision-Recall on real data **(Right)** Legend along with F1 scores on real data.

Qualitative Analysis. Armed with the initial evidence from the coarse quantitative results, we performed a thorough qualitative analysis over 30 real robot episodes (approximately 13 mins). Figs 5 and 6 present selected video reels. Fig 7 compares the outputs of all methods at some selected frames in the data. More illustrative video results are on the website. These examples are selected to illustrate some key insights from our more comprehensive analysis:

- As our metrics above suggested, crop is indeed the most important contributor to `im2contact` performance. In Fig 7, both `im2contact` "w/o crop" and U-Net Depth perform similarly poorly, demonstrating the criticality of focusing the model on relevant regions for finding good solutions.
- The grasped object reference frame (obj-ref) frequently improves performance in occluded contact situations (Fig 5, left), or when the grasped object moves coherently with others, creating a misleading flow field (Fig 5, right and Fig 7 (d)). This is consistent with our arguments in Sec 3.3; in such situations, knowing the shape of the grasped object is critical for localizing extrinsic contacts.
- Consistent with the metrics above, the effects of optical flow are less obvious, and w/o flow is very similar to `im2contact` in Fig 7. We have however observed that it improves performance in some

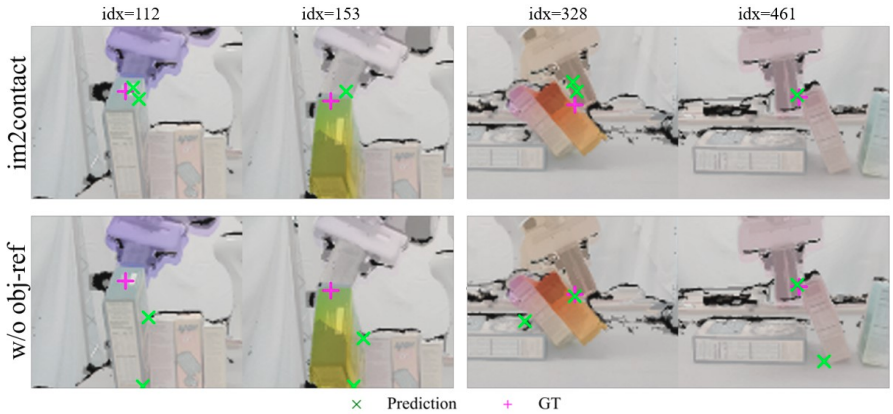

Figure 5: `im2contact` **w/o obj-ref**: obj-ref provides useful information about the grasped object's shape, which is not otherwise available under occlusions or misleading flow fields. Optical flow fields overlaid on all images. **(Left)** The grasped can behind the cereal box, slides up its side to topple the box. Here, the grasped object is occluded. **(Right)** The robot pushes two boxes to topple them, then pushes one of them back upright. These sustained pushes lead to coherent flow fields that do not help to separate a grasped object from a neighbor. In both cases, "w/o obj-ref" errs by treating the can-cum-boxes collective as a single grasped object, generating false positives at locations where the boxes contact their neighbors.

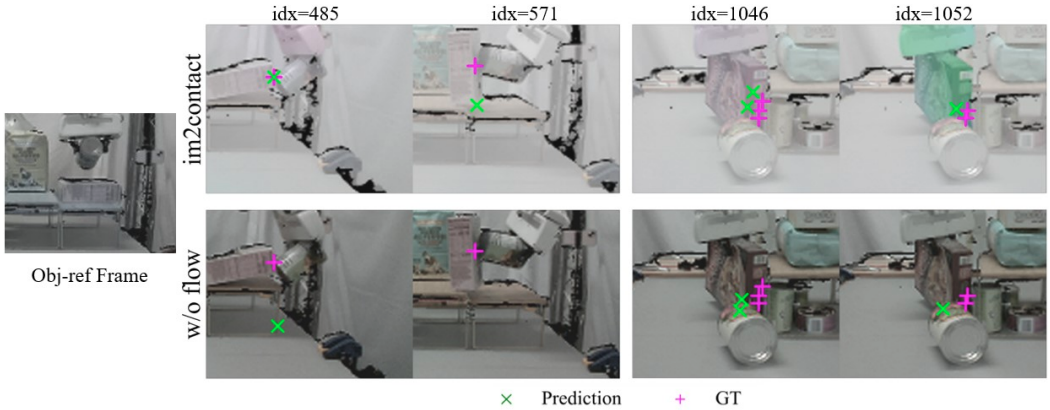

Figure 6: `im2contact` **w/o flow** (Optical flow fields overlaid only on `im2contact` images): **(Left)** The robot uses a grasped can to press against and pivot a box into an upright position. The reference frame does not clearly show the can's shape, so flow must separate the grasped object from others. Sure enough, in frame 485 during the motion, "w/o flow" predicts irrelevant contacts at the foot of the shelf, but `im2contact` performs correctly. Later, in frame 571 immediately after the motion, `im2contact` no longer has the flow cue, and fails. **(Right)** The robot uses a large box to sweep a number of cans under a shelf. In this highly cluttered scene, "w/o flow" consistently predicts a false contact between a stationary can in the foreground rather than the true contact happening in the background underneath the shelf. `im2contact` on the other hand uses the flow on the moving background can to localize the true contacts.

cases. In Fig 6 (a), flow suppresses a false positive in a near-contact situation and helps identify a contact with a can within clutter in Fig 6 (b).

Out-of-Distribution evaluation of `im2contact` We conclude with a brief highlight reel in Fig 8 visualizing the output of `im2contact` on out-of-distribution objects that are difficult to annotate manually. We find that despite training only on rigid grasps of primitive rigid bodies in simulation, `im2contact` produces reasonable estimates of extrinsic contact.

5 Limitations and Future Work

While it is surprising how much contact perception can be achieved from just a single camera and proprioception, our approach nevertheless only partially addresses the contact perception problem: we only localize contacts in the 2D image. We do not localize contacts in 3D, perceive contact forces, or classify modes of contact. Our deliberately simple sensing setup might be fundamentally limited for solving these broader problems. For example, predictions over the RGB-D image may be backprojected into 3D spatial coordinates using the depth map. However, single-viewpoint depth images are only a 2.5D representation of geometry, so we are unable to localize full 3D coordinates of contact points during occluded contacts. Multiple cameras may help address this issue.

In addition, our cropping method assumes some upper bound on the dimensions of the grasped object. In future work, integrating an off-the-shelf object detection method may be able to provide more adaptive bounding boxes without having to make explicit assumptions.

Lastly, our method leverages the global information provided by vision over more local tactile sensing. In doing so, we sacrifice the precision of our method which is prone to false positives during near-contact and occlusions, even with the addition of optical flow. In future work, we hope to incorporate the advantages of both sensing modalities.

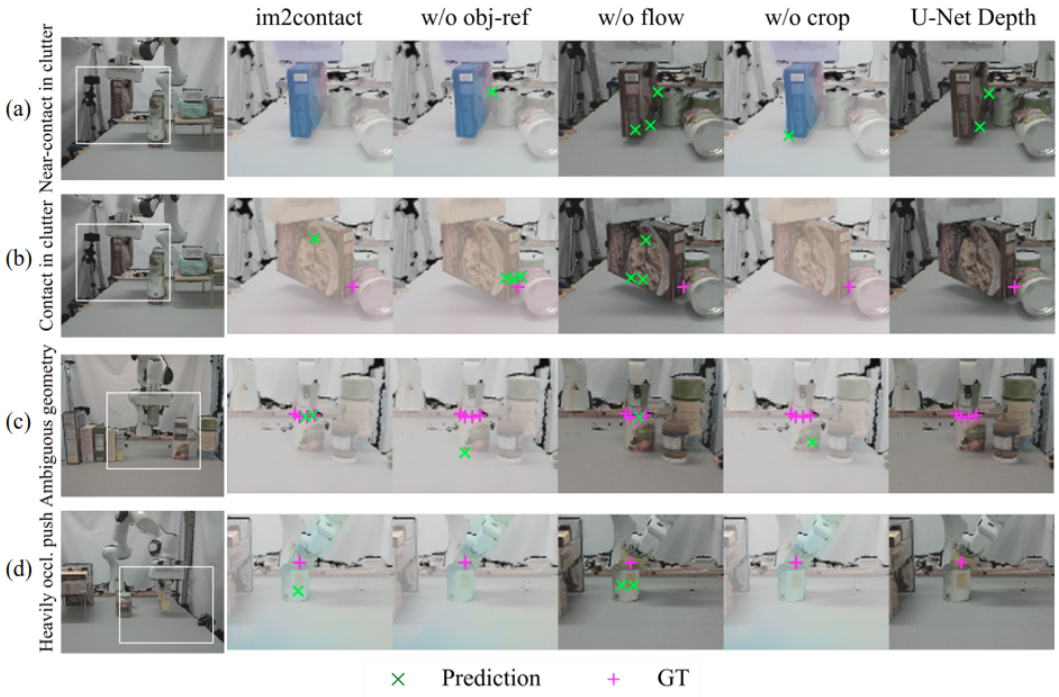

Figure 7: Our full ablations on four illustrative samples (flow visualized on the models that use them): **(a)** a grasped cereal box encounters numerous near-contacts in a highly cluttered scene. **(b)** From the same episode, the grasped cereal box comes into contact with a foreground can in the cluttered scene. **(c)** A can is stacked on top of another in a cluttered scene **(d)** A significantly occluded grasped sugar box pushes a can across the table

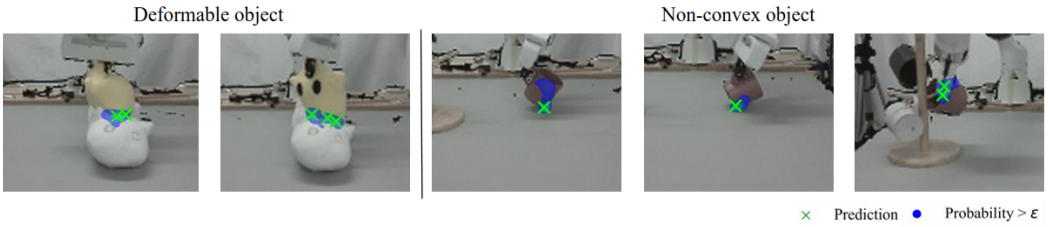

Figure 8: `im2contact` predictions on **(Left)** interactions between two moderately deformable plushies **(Right)** an in-hand re-orientation of a non-convex mug, followed by hanging on a rack

6 Conclusion

We present a method that enables extrinsic contact localization from vision alone with minimal prior knowledge or explicit assumptions on the grasped object and environment. By incorporating well-motivated inputs to our model in simulation, we show successful sim-to-real transfer of our model whereas a naive baseline fares poorly. The method also shows promise on our chosen out-of-distribution objects which include deformable and non-convex objects. In addition to addressing our methods limitations, in the future we hope to further demonstrate the utility of our extrinsic contact estimator through the coupling of a downstream control task involving tool-use or assembly.

References

[1] R. S. Johansson and J. R. Flanagan. Coding and use of tactile signals from the fingertips in object manipulation tasks. *Nature Reviews Neuroscience*, 10(5):345–359, 2009.

[2] R. D. Howe. Tactile sensing and control of robotic manipulation. *Advanced Robotics*, 8(3): 245–261, 1993.

[3] S. Tian, F. Ebert, D. Jayaraman, M. Mudigonda, C. Finn, R. Calandra, and S. Levine. Manipulation by feel: Touch-based control with deep predictive models. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 818–824. IEEE, 2019.

[4] A. Aydinoglu, P. Sieg, V. M. Preciado, and M. Posa. Stabilization of complementarity systems via contact-aware controllers. *IEEE Transactions on Robotics*, 38(3):1735–1754, 2021.

[5] J. M. Romano, K. Hsiao, G. Niemeyer, S. Chitta, and K. J. Kuchenbecker. Human-inspired robotic grasp control with tactile sensing. *IEEE Transactions on Robotics*, 27(6):1067–1079, 2011.

[6] M. Y. Cao, S. Laws, and F. R. y Baena. Six-axis force/torque sensors for robotics applications: A review. *IEEE Sensors Journal*, 21(24):27238–27251, 2021.

[7] J. A. Fishel and G. E. Loeb. Sensing tactile microvibrations with the biotac—comparison with human sensitivity. In *2012 4th IEEE RAS & EMBS international conference on biomedical robotics and biomechatronics (BioRob)*, pages 1122–1127. IEEE, 2012.

[8] W. Yuan, S. Dong, and E. H. Adelson. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors*, 17(12):2762, 2017.

[9] M. Lambeta, P.-W. Chou, S. Tian, B. Yang, B. Maloon, V. R. Most, D. Stroud, R. Santos, A. Byagowi, G. Kammerer, et al. Digit: A novel design for a low-cost compact high-resolution tactile sensor with application to in-hand manipulation. *IEEE Robotics and Automation Letters*, 5(3):3838–3845, 2020.

[10] R. Bhirangi, T. Hellebrekers, C. Majidi, and A. Gupta. Reskin: versatile, replaceable, lasting tactile skins. *arXiv preprint arXiv:2111.00071*, 2021.

[11] B. Wu, Z. Wu, and F. Shen. Research on calibration system error of 6-axis force/torque sensor integrated in humanoid robot foot. In *2010 8th World Congress on Intelligent Control and Automation*, pages 6878–6882. IEEE, 2010.

[12] F. J. A. Chavez, G. Nava, S. Traversaro, F. Nori, and D. Pucci. Model based in situ calibration with temperature compensation of 6 axis force torque sensors. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 5397–5403. IEEE, 2019.

[13] M. Reinvee, K. Jansen, et al. Utilisation of tactile sensors in ergonomic assessment of hand–handle interface: a review. *Agronomy Research*, 12(3):907–914, 2014.

[14] A. De Luca, A. Albu-Schaffer, S. Haddadin, and G. Hirzinger. Collision detection and safe reaction with the DLR-III lightweight manipulator arm. In *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1623–1630. doi:10.1109/IROS.2006.282053. ISSN: 2153-0866.

[15] 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. ISBN 978-1-5090-3762-9. doi:10.1109/IROS.2016.7759743. URL `http://ieeexplore.ieee.org/document/7759743/`.

[16] P. Grady, J. A. Collins, S. Brahmbhatt, C. D. Twigg, C. Tang, J. Hays, and C. C. Kemp. Visual pressure estimation and control for soft robotic grippers. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3628–3635. doi:10.1109/IROS47612.2022.9982073. ISSN: 2153-0866.

[17] A. Bicchi, J. K. Salisbury, and D. L. Brock. Contact sensing from force measurements. 12(3): 249–262. ISSN 0278-3649. doi:10.1177/027836499301200304. URL https://doi.org/10.1177/027836499301200304. Publisher: SAGE Publications Ltd STM.

[18] K.-T. Yu and A. Rodriguez. Realtime state estimation with tactile and visual sensing for inserting a suction-held object. doi:10.1109/IROS.2018.8594077.

[19] 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. doi:10.1109/ICRA48506.2021.9561781. ISSN: 2577-087X.

[20] S. Kim, D. K. Jha, D. Romeres, P. Patre, and A. Rodriguez. Simultaneous tactile estimation and control of extrinsic contact. URL http://arxiv.org/abs/2303.03385.

[21] A. Molchanov, O. Kroemer, Z. Su, and G. S. Sukhatme. Contact localization on grasped objects using tactile sensing. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 216–222. doi:10.1109/IROS.2016.7759058. ISSN: 2153-0866.

[22] Y. Karayiannidis, C. Smith, F. E. Vina, and D. Kragic. Online contact point estimation for uncalibrated tool use. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2488–2494. IEEE. ISBN 978-1-4799-3685-4. doi:10.1109/ICRA.2014.6907206. URL http://ieeexplore.ieee.org/document/6907206/.

[23] C. Higuera, S. Dong, B. Boots, and M. Mukadam. Neural contact fields: Tracking extrinsic contact with tactile sensing. URL http://arxiv.org/abs/2210.09297.

[24] A. Sipos and N. Fazeli. Simultaneous contact location and object pose estimation using proprioception and tactile feedback. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3233–3240. doi:10.1109/IROS47612.2022.9981762. ISSN: 2153-0866.

[25] M. Van der Merwe, D. Berenson, and N. Fazeli. Learning the dynamics of compliant tool-environment interaction for visuo-tactile contact servoing. URL http://arxiv.org/abs/2210.03836.

[26] N. Koenig and A. Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566)*, volume 3, pages 2149–2154. IEEE, 2004.

[27] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.

[28] N. Siddique, S. Paheding, C. P. Elkin, and V. Devabhaktuni. U-net and its variants for medical image segmentation: A review of theory and applications. *Ieee Access*, 9:82031–82057, 2021.

[29] M. Danielczuk, M. Matl, S. Gupta, A. Li, A. Lee, J. Mahler, and K. Goldberg. Segmenting unknown 3d objects from real depth images using mask r-cnn trained on synthetic data. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 7283–7290. IEEE, 2019.

[30] H. Ye, C. Xie, T. Cai, R. Li, Z. Li, and L. Wang. Towards a theoretical framework of out-of-distribution generalization. *Advances in Neural Information Processing Systems*, 34:23519–23531, 2021.

[31] R. Geirhos, J.-H. Jacobsen, C. Michaelis, R. Zemel, W. Brendel, M. Bethge, and F. A. Wichmann. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673, 2020.

[32] V. Nagarajan, A. Andreassen, and B. Neyshabur. Understanding the failure modes of out-of-distribution generalization. *arXiv preprint arXiv:2010.15775*, 2020.

[33] R. Liu, J. Lehman, P. Molino, F. Petroski Such, E. Frank, A. Sergeev, and J. Yosinski. An intriguing failing of convolutional neural networks and the coordconv solution. *Advances in neural information processing systems*, 31, 2018.

[34] G. Burel and D. Carel. Detection and localization of faces on digital images. *Pattern Recognition Letters*, 15(10):963–967, 1994.

[35] P. Felzenszwalb, D. McAllester, and D. Ramanan. A discriminatively trained, multiscale, deformable part model. In *2008 IEEE conference on computer vision and pattern recognition*, pages 1–8. Ieee, 2008.

[36] R. Girshick, J. Donahue, T. Darrell, and J. Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.

[37] Z. Teed and J. Deng. RAFT: Recurrent all-pairs field transforms for optical flow. URL http://arxiv.org/abs/2003.12039.