

Multi-Modal Geometric Learning for Grasping and Manipulation

David Watkins-Valls, Jacob Varley, and Peter Allen

Abstract—This work provides an architecture that incorporates depth and tactile information to create rich and accurate 3D models useful for robotic manipulation tasks. This is accomplished through the use of a 3D convolutional neural network (CNN). Offline, the network is provided with both depth and tactile information and trained to predict the object’s geometry, thus filling in regions of occlusion. At runtime, the network is provided a partial view of an object. Tactile information is acquired to augment the captured depth information. The network can then reason about the object’s geometry by utilizing both the collected tactile and depth information. We demonstrate that even small amounts of additional tactile information can be incredibly helpful in reasoning about object geometry. This is particularly true when information from depth alone fails to produce an accurate geometric prediction. Our method is benchmarked against and outperforms other visual-tactile approaches to general geometric reasoning. We also provide experimental results comparing grasping success with our method.

I. INTRODUCTION

Robotic grasp planning based on raw sensory data is difficult due to occlusion and incomplete information regarding scene geometry. Often, one sensory modality does not provide enough context to enable reliable planning. For example, a single depth sensor image cannot provide information about occluded regions of an object, and tactile information is incredibly sparse. This work uses a 3D convolutional neural network to enable stable robotic grasp planning by incorporating both tactile and depth information to infer occluded geometries. This multi-modal system uses both tactile and depth information to form a more complete model of the space the robot can interact with and also to provide a complete object model for grasp planning.

At runtime, a point cloud of the visible portion of the object is captured, and multiple guarded moves are executed in which the hand is moved towards the object, stopping when contact with the object occurs. The newly acquired tactile information is combined with the original partial view, voxelized, and sent through the CNN to create a hypothesis of the object’s geometry.

Depth information from a single point of view often does not provide enough information to accurately predict object geometry. There is often unresolved uncertainty about the geometry of the occluded regions of the object. To alleviate this uncertainty, we utilize tactile information to generate a new, more accurate hypothesis of the object’s 3D

This work is supported by NSF Grant CMMI 1734557. We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research. Authors are with Columbia University, (davidwatkins, jvarley, allen)@cs.columbia.edu

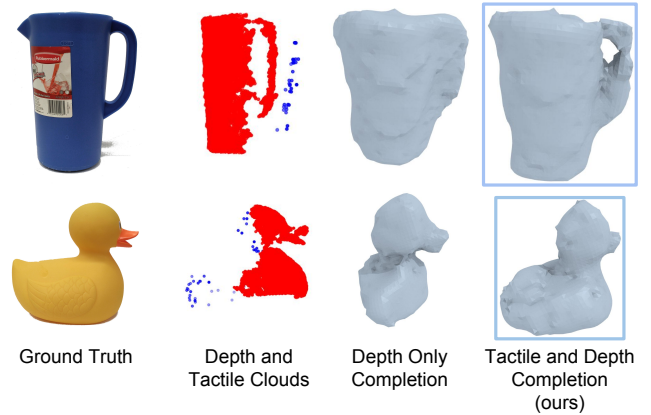


Fig. 1: Completion example from tactile and depth data. Small amounts of additional tactile data can significantly improve the systems ability to reason about 3D geometry. The Depth Only Completion for the pitcher does not capture the handle well, whereas the tactile information gives a better geometric understanding. The additional tactile data allowed the CNN to correctly identify a handle in the completion mesh and similar completion improvement was found for the novel rubber duck not in the training set.

geometry, incorporating both visual and tactile information. Fig. 1 demonstrates an example where the understanding of the object’s 3D geometry is significantly improved by the additional sparse tactile data collected via our framework. An overview of our sensory fusion architecture is shown in Fig. 2.

This work is differentiated from others [1] in that our CNN is acting on both the depth and tactile as input information fed directly into the model rather than using the tactile information to update the output of a CNN not explicitly trained on tactile information. This enables the tactile information to produce non-local changes in the resulting mesh. In many cases, depth information alone is insufficient to differentiate between two potential completions, for example a pitcher vs a rubber duckie. In these cases, the CNN utilizes sparse tactile information to affect the entire completion, not just the regions in close proximity to the tactile glance. If the tactile sensor senses the occluded portion of a drill, the CNN can turn the entire completion into a drill, not just the local portion of the drill that was touched.

The contributions of this work include: 1) a framework for integrating multi-modal sensory data to holistically reason about object geometry and enable robotic grasping, 2) an open source dataset for training a shape completion

system using both tactile and depth sensory information, 3) open source code for alternative visual-tactile general completion methods, 4) experimental results comparing the completed object models using depth only, the combined depth-tactile information, and various other visual-tactile completion methods, and 5) real and simulated grasping experiments using the completed models. This dataset, code, and extended video are freely available at <http://crlab.cs.columbia.edu/visualtactilegrasping/>.

II. RELATED WORK

The idea of incorporating sensory information from vision, tactile and force sensors is not new [2]. Despite the intuitiveness of using multi-modal data, there is still no consensus on which framework best integrates multi-modal sensory information in a way that is useful for robotic manipulation tasks. While prior work has been done to complete geometry using depth alone, none of these works consider tactile information[3][4]. In this work, we are interested in reasoning about object geometry, and in particular, creating models from multi-modal sensory data that can be used for grasping and manipulation.

Several recent uses of tactile information to improve estimates of object geometry have focused on the use of Gaussian Process Implicit Surfaces (GPIS) [5]. Several examples along this line of work include [6][7] [8][9][10][11][12]. This approach is able to quickly incorporate additional tactile information and improve the estimate of the object’s geometry local to the tactile contact or observed sensor readings. There has additionally been several works that incorporate tactile information to better fit planes of symmetry and superquadrics to observed point clouds [13][14][15]. These approaches work well when interacting with objects that conform to the heuristic of having clear detectable planes of symmetry or are easily modeled as superquadrics.

There has been successful research in utilizing continuous streams of visual information similar to Kinect Fusion [16] or SLAM [17] in order to improve models of 3D objects for manipulation, an example being [18][19]. In these works, the authors develop an approach to building 3D models of unknown objects based on a depth camera observing the robot’s hand while moving an object. The approach integrates both shape and appearance information into an articulated ICP approach to track the robot’s manipulator and the object while improving the 3D model of the object. Similarly, another work [20] attaches a depth sensor to a robotic hand and plans grasps directly in the sensed voxel grid. These approaches improve their models of the object using only a single sensory modality but from multiple points in time.

In previous work [21], we created a shape completion method using single depth images. The work provides an architecture to enable robotic grasp planning via shape completion, which was accomplished through the use of a 3D CNN. The network was trained on an open source dataset of over 440,000 3D exemplars captured from varying viewpoints. At runtime, a 2.5D point cloud captured from a single point of view was fed into the CNN, which fills in the

occluded regions of the scene, allowing grasps to be planned and executed on the completed object. The runtime of shape completion is rapid because most of the computational costs of shape completion are borne during offline training. This prior work explored how the quality of completions vary based on several factors. These include whether or not the object being completed existed in the training data, how many object models were used to train the network, and the ability of the network to generalize to novel objects, allowing the system to complete previously unseen objects at runtime. The completions are still limited by the training datasets and occluded views that give no clue to the unseen portions of the object. From a human perspective, this problem is often alleviated by using the sense of touch. In this spirit, this paper addresses this issue by incorporating sparse tactile data to better complete the object models for grasping tasks.

III. VISUAL-TACTILE GEOMETRIC REASONING

Our framework utilizes a trained CNN to produce a mesh of the target object, incorporating both depth and tactile information. We utilize the same architecture as found in [21]. The model was implemented using the Keras [22] deep learning library. Each layer used rectified linear units as nonlinearities except the final fully connected (output) layer which used a sigmoid activation to restrict the output to the range $[0, 1]$. We used the cross-entropy error $E(y, y')$ as the cost function with target y and output y' :

$$E(y, y') = -(y \log(y') + (1 - y) \log(1 - y'))$$

This cost function encourages each output to be close to either 0 for unoccupied target voxels or 1 for occupied target voxels. The optimization algorithm Adam [23], which computes adaptive learning rates for each network parameter, was used with default hyperparameters ($\beta_1=0.9$, $\beta_2=0.999$, $\epsilon=10^{-8}$) except for the learning rate, which was set to 0.0001. Weights were initialized following the recommendations of [24] for rectified linear units and [25] for the logistic activation layer (batch size=32).

IV. COMPLETION OF SIMULATED GEOMETRIC SHAPES

Three networks with the architecture from [21] were trained on a simulated dataset of geometric shapes (Fig. 3) where the front and back were composed of two differing shapes. Sparse tactile data was generated by randomly sampling voxels along the occluded side of the voxel grid. We trained a network that only utilized tactile information. This performed poorly due to the sparsity of information. A second network was given only the depth information during training and performed better than the tactile-only network did. It still encountered many situations where it did not have enough information to accurately complete the obstructed half of the object. A third network was given depth and tactile information which successfully utilized the tactile information to differentiate between plausible geometries of occluded regions.

The Jaccard similarity improved from 0.890 in the depth only network to 0.986 in the depth and tactile network. This

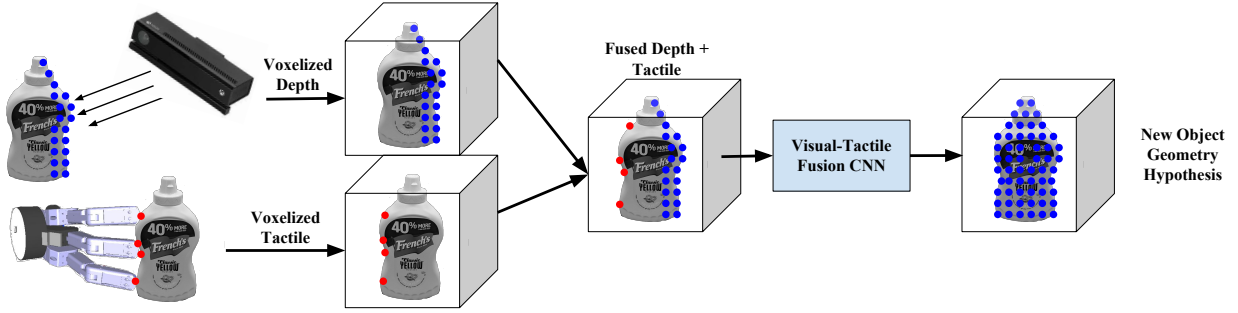


Fig. 2: Both tactile and depth information are independently captured and voxelized into 40^3 grids. These are merged into a shared occupancy map which is fed into a CNN to produce a hypothesis of the object’s geometry.

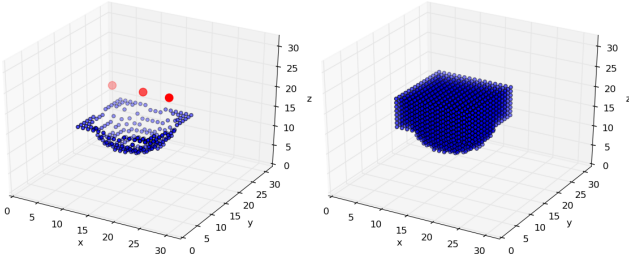


Fig. 3: Example training pair from the geometric shape dataset. For the left, red dots represent tactile readings and blue dots represent the the depth image. The blue points on the right are the ground truth 3D geometry.

Algorithm 1 Simulated YCB/Grasp Tactile Data Generation

```

1: procedure SAMPLE_TACTILE(vox_gt)
2:   grid_dim = 40 // resolution of voxel grid
3:   npts = 40 // num locations to check for contact
4:   vox_gt_cf = align_gt.to_depth_frame(vox_gt)
5:   xs = rand_ints(start=0, end=grid_dim-1, size=npts)
6:   ys = rand_ints(start=0, end=grid_dim-1, size=npts)
7:   tactile_vox = []
8:   for x, y in xs, ys do
9:     for z in range(grid_dim-1, -1, -1) do
10:      if vox_gt_cf[x, y, z] == 1 then
11:        tactile_vox.append(x, y, z)
12:        continue
13:   tactile_points = vox2point cloud(tactile_vox)
14:   return tactile_points

```

task demonstrated that a CNN can be trained to leverage sparse tactile information to decide between multiple object geometry hypotheses. When the object geometry had sharp edges in its occluded region, the system would use tactile information to generate a completion that contained similar sharp edges in the occluded region. This completion is more accurate not just in the observed region of the object but also in the unobserved portion of the object.

V. COMPLETION OF YCB/GRASP DATASET OBJECTS

We used the dataset from [21] to create a new dataset consisting of half a million triplets of oriented voxel grids:

depth, tactile, and ground truth. Depth voxels are marked as occupied if visible to the camera. Tactile voxels are marked occupied if tactile contact occurs within the voxel. Ground truth voxels are marked as occupied if the object intersects a given voxel, independent of perspective. The point clouds for the depth information were synthetically rendered in the Gazebo [26] simulator. This dataset consists of 608 meshes from both the Grasp [27] and YCB [28] datasets. 486 of these meshes were randomly selected and used for a training set and the remaining 122 meshes were kept for a holdout set.

The synthetic tactile information was generated according to Algorithm 1. In order to generate tactile data, the voxelization of the ground truth high resolution mesh (vox_gt) (Alg.1:L1) was aligned with the captured depth image (Alg.1:L4). 40 random (x, y) points were sampled in order to generate synthetic tactile data (Alg.1:L5-6). For each of these points (Alg.1:L7), a ray was traced in the $-z$, direction and the first occupied voxel was stored as a tactile observation (Alg.1:L11). Finally this set of tactile observations was converted back to a point cloud (Alg.1:L13).

Two identical CNNs were trained where one CNN was provided only depth information (**Depth Only**) and a second was provided both tactile and depth information (**Tactile and Depth**). During training, performance was evaluated on simulated views of meshes within the training data (*Training Views*), novel simulated views of meshes in the training data (*Holdout Views*), novel simulated views of meshes not in the training data (*Holdout Meshes*), and real non-simulated views of 8 meshes from the YCB dataset (*Holdout Live*).

The *Holdout Live* examples consist of depth information captured from a real Kinect and tactile information captured from a real Barrett Hand attached to a Staubli Arm. We used depth filtering to mask out the background of the captured depth cloud. The object was fixed in place during the tactile data collection process. While collecting the tactile data, the arm was manually moved to place the end effector behind the object and 6 exploratory guarded motions were made where the fingers closed towards the object. Each finger stopped independently when contact was made with the object, as shown in Fig. 4.

Fig. 5 demonstrates that the difference between the **Depth Only** CNN completion and the **Tactile and Depth** CNN completion becomes larger on more difficult completion




Fig. 4: Barrett hand showing contact with a fixed object. (a) The hand is manually brought to an approach position, (b) approaches the object, and (c) the fingers are curled to contact the object and collect tactile information. This process is repeated 6 times over the occluded surface of the object.

Completion Method	Train View(Sim)	Holdout View(Sim)	Holdout Model(Sim)	Holdout (Live)
Partial	0.01	0.02	0.01	0.01
Convex Hull	0.50	0.51	0.46	0.43
GPIS	0.47	0.45	0.35	0.48
Depth CNN	0.68	0.65	0.65	0.37
Ours	0.69	0.66	0.65	0.64

TABLE I: **Jaccard similarity results**, measuring the intersection over union of two voxelized meshes, as described in Section VI. (Larger is better)

Completion Method	Train View(Sim)	Holdout View(Sim)	Holdout Model(Sim)	Holdout (Live)
Partial	19.9mm	21.1mm	16.6mm	18.6mm
Convex Hull	13.9mm	16.1mm	14.1mm	10.5mm
GPIS	17.1mm	16.0mm	21.3mm	20.8mm
Depth CNN	12.1mm	13.7mm	12.4mm	22.9mm
Ours	7.7mm	13.9mm	13.6mm	6.2mm

TABLE III: **Pose error results** from simulated grasping experiments. This is the average L2 difference between planned and realized grasp pose averaged over the 3 finger tips and the palm of the hand, in millimeters. (Smaller is better)

Completion Method	Train View(Sim)	Holdout View(Sim)	Holdout Model(Sim)	Holdout (Live)
Partial	7.8	7.0	7.6	11.9
Convex Hull	32.7	45.1	49.1	11.6
GPIS	59.9	79.2	118.0	17.9
Depth CNN	6.5	6.9	6.5	16.5
Ours	5.8	5.8	6.2	7.4

TABLE II: **Hausdorff distance results**, measuring the mean distance in millimeters from points on one mesh to points on another mesh, as described in Section VI. (Smaller is better)

Completion Method	Train View(Sim)	Holdout View(Sim)	Holdout Model(Sim)	Holdout (Live)
Partial	8.19°	6.71°	8.78°	7.67°
Convex Hull	3.53°	4.01°	4.59°	3.77°
GPIS	4.65°	4.79°	4.95°	5.92°
Depth CNN	3.09°	3.56°	4.52°	6.83°
Ours	2.48°	3.41°	4.95°	2.43°

TABLE IV: **Joint error results** from simulated grasping experiments. This is the mean L2 distance between planned and realized grasps in degrees averaged over the hand’s 7 joints. Our method’s smaller error demonstrates a more accurate geometry reconstruction. (Smaller is better)

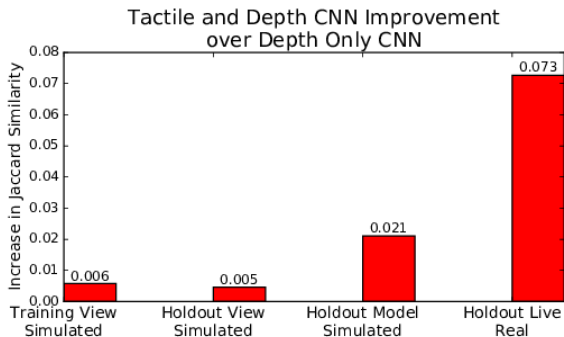


Fig. 5: As the difficulty of the data splits increase, the delta between the **Depth Only** CNN completion accuracy and the **Tactile and Depth** CNN completion accuracy increases. The additional tactile information is more useful on more difficult completion problems.

problems. The performance of the **Depth Only** CNN nearly matches the performance of the **Tactile and Depth** CNN on the training views. Because these views are used during training, the network is capable of generating reasonable completions. Moving from *Holdout Views* to *Holdout Meshes* to *Holdout Live*, the completion problems move further away from the examples experienced during training. As the problems become harder, the **Tactile and Depth** network

outperforms the **Depth Only** network by a greater margin, as it is able to utilize the sparse tactile information to differentiate between various possible completions. This trend shows that the network is able to make more use of the tactile information when the depth information alone is insufficient to generate a quality completion. We generated meshes from the output of the combined tactile and depth CNN using a marching cubes algorithm. We also preserve the density of the rich visual information and the coarse tactile information by utilizing the post-processing from [21].

VI. COMPARISON TO OTHER COMPLETION METHODS

In this work we benchmarked our framework against the following general visual tactile completion methods.

Partial Completion: The set of points captured from the Kinect is concatenated with the tactile data points. The combined cloud is run through marching cubes, and the resulting mesh is then smoothed using Meshlab’s [29] implementation of Laplacian smoothing. These completions are incredibly accurate where the object is directly observed but make no predictions in unobserved areas of the scene.

Convex Hull Completion: The set of points captured from the Kinect is concatenated with the tactile data points. The combined cloud is run through QHull to create a convex hull. The hull is then run through Meshlab’s implementation




Fig. 6: The entire *Holdout Live* dataset. These completions were all created from data captured from a real Kinect and a real Barrett Hand attached to a Staubli Arm. The **Depth and Tactile Clouds** have the points captured from a Kinect in red and points captured from tactile data in blue. Notice many of the **Depth Only** completions do not extend far enough back but instead look like other objects that were in the training data (ex: cell phone, banana). Our method outperforms the **Depth Only**, **Partial**, and **Convex Hull** methods in terms of Hausdorff distance and Jaccard similarity. Note that the **GPIS** completions form large and inaccurate completions for the Black and Decker box and the Rubbermaid Pitcher, whereas our method correctly bounds the end of the box and finds the handle of the pitcher.

of Laplacian smoothing. These completions are reasonably accurate near observed regions. However, a convex hull will fill regions of unobserved space.

Gaussian Process Implicit Surface Completion (GPIS): Approximated depth cloud normals were calculated using PCL's KDTree normal estimation. Approximated tactile

cloud normals were computed to point towards the camera origin. The depth point cloud was downsampled to size M and appended to the tactile point cloud. We used a distance offset d to add positive and negative observation points along the direction of the surface normal. We then sampled the Gaussian process using [30] with a n^3 voxel

Completion Method	Partial	Convex Hull	GPIS	Depth CNN	Ours
Lift Success (%)	62.5%	62.5%	87.5%	75.0%	87.5%
Joint Error (°)	6.37°	6.05°	10.61°	5.42°	4.67°
Time (s)	1.533s	0.198s	45.536s	3.308s	3.391s

TABLE V: **Lift Success** is the percentage of successful lift executions. **Joint Error** is the average error per joint in degrees between the planned and executed grasp joint values. While GPIS and our method have the same lift success, our method is 1340% faster and has 41% of the joint error, making the process more reliable. (Smaller is better). **Average time to complete a mesh** using each completion method. While the convex hull completion method is fastest, ours has a superior tradeoff between speed and quality.

grid and a noise parameter s to create meshes from the point cloud. We empirically determined the values of M, s, n, d by sampling the Jaccard similarity of GPIS completions where $M = [200, 300, 400]$, $s = [0.001, 0.005]$, $n = [40, 64, 100]$, and $d = [0.005, 0.0005]$. We found $M = 300$ to be a good tradeoff between speed and completion quality. Additionally we used $s = 0.001$, $d = 0.0005$, and $n = 100$.

In prior work [21] the Depth Only CNN completion method was compared to both a RANSAC based approach [31] and a mirroring approach [32]. These approaches make assumptions about the visibility of observed points and do not work with data from tactile contacts that occur in unobserved regions of the workspace.

A. Geometric Comparison Metrics

The Jaccard similarity was used to compare 40^3 CNN outputs with the ground truth. We also used this metric to compare the final resulting meshes from several completion strategies. The completed meshes were voxelized at 80^3 and compared with the ground truth mesh. The results are shown in Table I. Our method results in higher similarity to the ground truth meshes than do all other described approaches.

The Hausdorff distance metric computes the average distance from the surface of one mesh to the surface of another. A symmetric Hausdorff distance was computed with Meshlab’s Hausdorff distance filter in both directions. Table II shows the mean values of the symmetric Hausdorff distance for each completion method. In this metric, our tactile and depth CNN mesh completions are significantly closer to the ground truth compared to the other approaches’ completions.

Both the partial and Gaussian process completion methods are accurate when close to the observed points but fail to approximate geometry in occluded regions. We found that in our training, the Gaussian Process completion method would often create a large and unruly object if the observed points were only a small portion of the entire object or if no tactile points were observed in simulation. Using a neural network has the added benefit of abstracting object geometries, whereas the alternative completion methods fail to approximate the geometry of objects which do not have points bounding their geometry.

B. Grasp Comparison in Simulation

In order to evaluate our framework’s ability to enable grasp planning, the system was tested in simulation using the same set of completions. The use of simulation allowed for the quick planning and evaluation of 7900 grasps. GraspIt! was used to plan grasps on all of the completions of the objects by uniformly sampling different approach directions. These grasps were then executed not on the completed object but on the ground truth meshes in GraspIt!. In order to simulate a real-world grasp execution, the completion was removed from GraspIt! and the ground truth object was inserted in its place. Then the hand was placed 20 cm away from the ground truth object along the approach direction of the grasp. The spread angle of the fingers was set, and the hand was moved along the approach direction of the planned grasp either until contact was made or a maximum approach distance was traveled. Then fingers closed to the planned joint values and each finger continued to close until either contact was made with the object or the joint limits were reached.

Table III shows the average difference between the planned and realized Cartesian finger tip and palm poses. Table IV shows the difference in pose of the end effector between the planned and realized grasps averaged over the 7 joints of the hand. Using our method, the end effector ended up closer to its intended location in both joint space and the palm’s Cartesian position versus other completion methods’ grasps.

C. Live Grasping Results

To further test our network’s efficacy, the grasps were planned and executed on the Holdout Live views using a Staubli arm with a Barrett Hand. The grasps were planned using meshes from the different completion methods described above. For each of the 8 objects, we ran the arm once using each completion method. The results are shown in Fig. 6 and Table V. Our method enabled an improvement over the other visual-tactile shape completion methods in terms of grasp success rate and resulted in executed grasps closer to the planned grasps, as shown by the lower average joint error (and much faster than GPIS).

VII. CONCLUSION

Our method provides an open source novel visual-tactile completion method which outperforms other general visual-tactile completion methods in completion accuracy, time of execution, and grasp posture utilizing a dataset which is representative of household and tabletop objects. We demonstrated that even small amounts of additional tactile information can be incredibly helpful in reasoning about object geometry. Experimental results verified that utilizing both vision and tactile was superior to using depth alone. In the future we hope to relax the fixed object assumption by using novel tactile sensors [33] we are developing with fast dynamic response that allow object contact without motion. We are also interested in using Next-Best-View planning [34], [35] for refining and optimizing the tactile exploration strategy to recover unseen geometry.

REFERENCES

- [1] S. Wang, J. Wu, X. Sun, W. Yuan, W. T. Freeman, J. B. Tenenbaum, and E. H. Adelson, "3d shape perception from monocular vision, touch, and shape priors," *arXiv:1808.03247*, 2018.
- [2] P. K. Allen, A. Miller, B. Leibowitz, and P. Oh, "Integration of vision, force and tactile sensing for grasping," *Int. Journal of Intelligent Mechatronics*, vol. 4, no. 1, pp. 129–149, 1999.
- [3] A. Dai, C. R. Qi, and M. Nießner, "Shape completion using 3d-encoder-predictor cnns and shape synthesis," in *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2017.
- [4] C. S. Weerasekera, T. Dharmasiri, R. Garg, T. Drummond, and I. D. Reid, "Just-in-time reconstruction: Inpainting sparse maps using single view depth predictors as priors," *CoRR*, vol. abs/1805.04239, 2018. [Online]. Available: <http://arxiv.org/abs/1805.04239>
- [5] O. Williams and A. Fitzgibbon, "Gaussian process implicit surfaces," *Gaussian Proc. in Practice*, pp. 1–4, 2007.
- [6] S. Caccamo, Y. Bekiroglu, C. H. Ek, and D. Kragic, "Active exploration using gaussian random fields and gaussian process implicit surfaces," in *IROS*. IEEE, 2016, pp. 582–589.
- [7] Z. Yi, R. Calandra, F. Veiga, H. van Hoof, T. Hermans, Y. Zhang, and J. Peters, "Active tactile object exploration with gaussian processes," in *IROS*. IEEE, 2016, pp. 4925–4930.
- [8] M. Bjorkman, Y. Bekiroglu, V. Hogman, and D. Kragic, "Enhancing visual perception of shape through tactile glances," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 2013, pp. 3180–3186.
- [9] S. Dragiev, M. Toussaint, and M. Gienger, "Gaussian process implicit surfaces for shape estimation and grasping," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 2845–2850.
- [10] N. Jamali, C. Ciliberto, L. Rosasco, and L. Natale, "Active perception: Building objects' models using tactile exploration," in *Humanoid Robots (Humanoids), 2016 IEEE-RAS 16th International Conference on*. IEEE, 2016, pp. 179–185.
- [11] N. Sommer, M. Li, and A. Billard, "Bimanual compliant tactile exploration for grasping unknown objects," in *ICRA*. IEEE, 2014, pp. 6400–6407.
- [12] J. Mahler, S. Patil, B. Kehoe, J. van den Berg, M. Ciocarlie, P. Abbeel, and K. Goldberg, "GP-GPIS-OPT: Grasp planning with shape uncertainty using gaussian process implicit surfaces and sequential convex programming," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2015.
- [13] J. Ilonen, J. Bohg, and V. Kyrki, "Three-dimensional object reconstruction of symmetric objects by fusing visual and tactile sensing," *The International Journal of Robotics Research*, vol. 33, no. 2, pp. 321–341, 2014.
- [14] —, "Fusing visual and tactile sensing for 3-d object reconstruction while grasping," in *Robotics and Automation (ICRA), 2013 IEEE International Conference on*. IEEE, 2013, pp. 3547–3554.
- [15] A. Bierbaum, I. Gubarev, and R. Dillmann, "Robust shape recovery for sparse contact location and normal data from haptic exploration," in *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*. IEEE, 2008, pp. 3200–3205.
- [16] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohi, J. Shotton, S. Hodges, and A. Fitzgibbon, "Kinectfusion: Real-time dense surface mapping and tracking," in *Mixed and augmented reality (ISMAR), 2011 10th IEEE international symposium on*. IEEE, 2011, pp. 127–136.
- [17] S. Thrun and J. J. Leonard, "Simultaneous localization and mapping," in *Springer handbook of robotics*. Springer, 2008, pp. 871–889.
- [18] M. Krainin, P. Henry, X. Ren, and D. Fox, "Manipulator and object tracking for in-hand 3d object modeling," *The International Journal of Robotics Research*, vol. 30, no. 11, pp. 1311–1327, 2011.
- [19] M. Krainin, B. Curless, and D. Fox, "Autonomous generation of complete 3d object models using next best view manipulation planning," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 5031–5037.
- [20] A. Hermann, F. Mauch, S. Klemm, A. Roennau, and R. Dillmann, "Eye in hand: Towards gpu accelerated online grasp planning based on pointclouds from in-hand sensor," in *Humanoids*. IEEE, 2016, pp. 1003–1009.
- [21] J. Varley, C. DeChant, A. Richardson, A. Nair, J. Ruales, and P. Allen, "Shape completion enabled robotic grasping," in *Intelligent Robots and Systems (IROS), IEEE/RSJ 2017 International Conference on*, extended version preprint at arXiv:1609.08546.
- [22] F. Chollet, "Keras," <https://github.com/fchollet/keras>, 2015.
- [23] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [24] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human-level performance on imagenet classification," *arXiv preprint arXiv:1502.01852*, 2015.
- [25] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *Proceedings of the 13th AISTATS*, 2010, pp. 249–256.
- [26] N. Koenig and A. Howard, "Design and use paradigms for Gazebo, an open-source multi-robot simulator," in *IROS*, vol. 3. IEEE, 2004, pp. 2149–2154.
- [27] D. Kappler, J. Bohg, and S. Schaal, "Leveraging big data for grasp planning," in *ICRA*. IEEE, 2015, pp. 4304–4311.
- [28] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The ycb object and model set: Towards common benchmarks for manipulation research," in *Advanced Robotics (ICAR), 2015 International Conference on*. IEEE, 2015, pp. 510–517.
- [29] P. Cignoni, M. Corsini, and G. Ranzuglia, "Meshlab: an open-source 3d mesh processing system," *Ercim news*, vol. 73, pp. 45–46, 2008.
- [30] M. P. Gerardo-Castro, T. Peynot, F. Ramos, and R. Fitch, "Robust multiple-sensing-modality data fusion using gaussian process implicit surfaces," in *Information Fusion (FUSION), 2014 17th International Conference on*. IEEE, 2014, pp. 1–8 <https://github.com/marcospaul/GPIS>.
- [31] C. Papazov and D. Burschka, "An efficient ransac for 3d object recognition in noisy and occluded scenes," in *Asian Conference on Computer Vision*. Springer, 2010, pp. 135–148.
- [32] J. Bohg, M. Johnson-Roberson, B. León, J. Felip, X. Gratal, N. Bergström, D. Kragic, and A. Morales, "Mind the gap-robotic grasping under incomplete observation," in *ICRA*. IEEE, 2011, pp. 686–693.
- [33] I. Kyriassis, C. Yu, Y. Hsu, P. Piacenza, E. Hannigan, M. Ciocarlie, and P. Allen, "Sheet-based flexible technologies for mechanical sensing," in *2018 International Flexible Electronics Technology Conference (IFETC)*, Aug 2018, pp. 1–1.
- [34] W. R. Scott, G. Roth, and J.-F. Rivest, "View planning for automated three-dimensional object reconstruction and inspection," *ACM Comput. Surv.*, vol. 35, no. 1, pp. 64–96, Mar. 2003.
- [35] P. S. Blaer and P. K. Allen, "View planning and automated data acquisition for three-dimensional modeling of complex sites," *Journal of Field Robotics*, vol. 26, no. 1112, pp. 865–891, 2009.