Learning to Singulate Objects using a Push Proposal Network

Andreas Eitel, Nico Hauff and Wolfram Burgard

Abstract Learning to act in unstructured environments such as cluttered piles of objects poses a substantial challenge for manipulation robots. We present a novel neural network-based approach that separates unknown objects in clutter by selecting favourable push actions. Our network is trained from data collected through autonomous interaction of a PR2 robot with randomly organized tabletop scenes. The model is designed to propose meaningful push actions based on over-segmented RGB-D images. We evaluate our approach by singulating up to 8 unknown objects in clutter. We demonstrate that our method enables the robot to perform the task with a high success rate and a low number of required push actions. Our results based on real-world experiments show that our network is able to generalize to novel objects of various sizes and shapes as well as to arbitrary object configurations. Videos of our experiments can be viewed at http://robotpush.cs.uni-freiburg.de

1 Introduction

Robot manipulation tasks such as tidying up a room or sorting piles of objects are a substantial challenge for robots, especially in scenarios with unknown objects. The objective of object singulation is to physically separate a set of cluttered objects through manipulation, a capability regarded as relevant for service robots operating in unstructured household environments. Further, the ability to separate unknown objects from surrounding objects provides great benefit for object detection, which still remains an open problem for overlapping objects. A key challenge of object singulation is the required interaction of multiple capabilities such as perception, manipulation and motion planning. Perception may fail because of occlusions from objects but also from occlusion by the manipulator itself. Manipulation actions such as grasping may also fail when attempting to manipulate objects that the perception

All authors are with the Department of Computer Science, University of Freiburg, Germany. Corresponding author's email: eitel@cs.uni-freiburg.de

Andreas Eitel, Nico Hauff and Wolfram Burgard



Fig. 1: Our approach consists of a convolutional neural network that ranks a set of potential push actions based on an over-segmented input image in order to clear objects in clutter. The proposed pushes are executed using motion planning. Our push proposal network is trained in an iterative manner from autonomous interaction with cluttered object scenes.

systems fails to correctly segment. Motion planning is particularly prone to errors when applied in unknown, unstructured environments.

Previous approaches for object singulation strongly incorporate the concept of object segments or even complete objects. Push action strategies are selected based on the current belief of the configuration of all segments or objects in the scene, an assumption that is not robust, since segments might merge and object detection might fail. We apply a different strategy that aims to relax the concept of segments and objects for solving the task at hand, making task execution less prone to modelling errors from the perception modules. Similar to recent approaches that operate directly on image inputs [19, 22] we aim for a learning-based action selection approach that takes as input an RGB-D image of a tabletop scene and proposes a set of meaningful push actions using a convolutional neural network, as shown in Fig. 1.

Our primary contributions are: 1) a push proposal network (Push-CNN), trained in an iterative manner to detect push candidates from RGB-D images in order to separate cluttered objects, 2) a method that samples potential push candidates based on depth data, which are ranked with our Push-CNN, and 3) real-world experiments with a PR2 robot in which we compare the performance of our method against a strong manually-designed baseline.

We quantitatively evaluate our approach on 4 sets of real-robot experiments. These experiments involve singulation tasks of increasing difficulty ranging from 4 unknown objects up to configurations with 8 unknown objects. In total the robot executed over 400 push actions during the evaluation and achieved an object separation success rate of up to 70%, which shows that our Push-CNN is able to generalize to previously unseen object configurations and object shapes. Our Freiburg Robotpush dataset is publicly available at http://robotpush.cs.uni-freiburg.de.

2 Related Work

Standard model-based approaches for object singulation require knowledge of object properties to perform motion planning in a physics simulator and to choose push actions accordingly [6, 7]. However, estimating the location of objects and other properties of the physical environment in order to plan push actions can be subject to errors, especially for unknown objects and cluttered scenes [21, 30, 31]. Interactive model-free methods have been applied to solve the task by accumulating evidence of singulated items over a history of interactions including push and grasp primitives [5]. Difficulties arise when objects have to be tracked after each interaction step, requiring small motion pushes due to partial observations and occlusions. We follow a model-free approach that encodes evidence of singulated objects in the learned feature space of the network. Hermans et al. [13] perform object singulation using several push primitives similar to ours. Their method is based on object edges to detect splitting locations between potential objects to apply pushes in those regions respectively, but does not include learned features and does not take into account stacked objects. Katz et al. [14] present an interactive segmentation algorithm to singulate cluttered objects using pushing and grasping. Similar to them, we also create object hypothesis from over-segmenting objects into surface facets but use a different method based on the work of [26]. They perform supervised learning with manual features to detect good manipulation actions such as push, pull and grasp. We take the idea of learning a step further by directly learning from over-segmented images and therefore removing the need for manual feature design. Boularias et al. [3] learn favourable push actions to improve grasping of objects in clutter in a reinforcement learning setting, but use manual features and do not show the applicability of their method for more than 2 objects. Gupta et al. [10] present an approach to sort small cluttered objects using a set of motion primitives, but do not show experiments with objects of various sizes and shapes. Laskey et al. [18] leverage learning from human demonstrations to learn control policies for object singulation, but thus far only considered singulating a single object from a cluttered scene, while our method can singulate up to 6 objects and achieve a higher success rate. Our approach is related to robot learning from physical interaction [24], including

Our approach is related to robot learning from physical interaction [24], including approaches that learn to predict how rigid objects behave if manipulated [17, 12, 32] and self-supervised learning methods for grasping [25, 20]. Gualtieri *et al.* [9] present a grasp pose detection method in which they detect grasp candidates using a convolutional neural network. They follow a similar methodology but for a different task and goal. Recent methods learn the dynamics of robot-object interaction for the task of pushing objects to a target location [8, 1], but do not evaluate on scenes with many touching objects and do not follow an active singulation strategy. Byravan *et al.* [4] learn to predict rigid body motions of pushed objects using a neural network, but to not demonstrate results for multiple objects in a real-world scenario. Object separation is used as a source to leverage perceptual information from interaction following the paradigm of interactive perception [2]. As demonstrated in our experiments, our singulation method can be used to generate informative sensory signals that lower the scene complexity for interactive object segmentation [27, 11, 29, 16].



Fig. 2: For training the push proposal CNN each push action (small arrow) from the collected dataset is labeled as positive or negative by a user who assesses the outcome of the action. The criteria for a positive label are: the object got singulated, the push did not move multiple objects, the object was pushed close to its center of mass and the pushed object was not already singulated.

3 Learning to Singulate Objects

We consider the problem of learning good push actions to singulate objects in clutter. Let **o** be an image with height *H* and width *W* taken from an RGB-D camera with a known camera intrinsics matrix **K**. Further, let $\mathbf{a} = (\mathbf{c}, \alpha)$ be a push proposal with start position $\mathbf{c} = (x, y)$ and push angle α both specified in the image plane. We aim to learn a function $p = F(\mathbf{o}, \mathbf{a}; \theta)$ with learned parameters θ that takes as input the image together with a push proposal and outputs a probability of singulation success.

3.1 Data

The prediction function F is trained in an iterative manner (we use 2 training iterations F_1, F_2) on experiences that the robot collects trough interaction with the environment. At the first iteration, we gather a dataset \mathcal{D} of push proposals together with images from randomly pushing objects in simulation. Then, we train a classifier F_1 that best mimics the labels from an expert user in terms of successful and unsuccessful singulation actions. In the next iteration, we use F_1 to collect more data and add this data to \mathcal{D} . The next classifier F_2 is trained to mimic the expert labels

Learning to Singulate Objects using a Push Proposal Network

on the whole dataset \mathscr{D} . We refer to F_1 as a vanilla network that is trained solely on experience collected from random interaction and F_2 as an aggregated network. The training examples $\mathscr{D} = \{(\mathbf{o}^1, \mathbf{a}^1, y^1), \dots, (\mathbf{o}^N, \mathbf{a}^N, y^N)\}$ are labeled by an expert user, where y^i is a corresponding binary label. Fig. 2 depicts the labels given by an expert user for one exemplary singulation trial. The function is trained in a supervised manner using a negative log likelihood loss $\min_{\theta} \sum_{i=1}^{N} \mathscr{L}(p, y^i)$ and a sigmoid function.

3.2 Approach

Our approach is divided into three modules, as shown in Fig. 1: 1) a sampling module that generates a set of push proposals, 2) a neural network module that classifies the set of push proposals and ranks them accordingly, and 3) an action execution module that computes arm motion plans for the ranked pushes and then executes the first push proposal that has both a high probability of singulation success and a successful motion plan.

3.2.1 Push Proposal Sampling

Our push proposal sampling method is designed to generate a set of push proposals $\{\mathbf{a}^1, \ldots, \mathbf{a}^M\}$. First, we sample 3D push handles $\{\mathbf{h}^1, \ldots, \mathbf{h}^M\}$ represented as point normals pointing parallel to the xy-plane of the robot's odometry frame $\mathbf{h}^m = (x, y, z, \mathbf{n}_h)$, $\mathbf{n}_h = (-n_x, -n_y, 0)$. Specifically, given the raw depth image, we first apply a surface-based segmentation method [26] (we only use the presegmentation and surface-detection steps of the approach) to obtain a set of segments $\{\mathbf{s}^1, \ldots, \mathbf{s}^L\}$ together with a surface normals map and a table plane. Second, we sample for each segment \mathbf{s}^l a fixed number of push handles and remove push handles below the table plane. We assume a fixed push length $l_a = 0.2$ for all handles and compute the convex hull for the table in order to filter out push handles that have push endpoints outside of the table. Finally, we obtain a set of push proposals by transforming the push handles into the camera frame using the transformation matrix **C** from the odometry to the camera frame together with the camera intrinsics $\mathbf{a} = \mathbf{KCh}$.

3.2.2 Push Proposal Network and Input

We parametrize the predictor $F(\mathbf{0}, \mathbf{a}; \theta)$ using a deep convolutional neural network, denoted as a push proposal CNN in Fig. 1. The distinction between **o** and **a** as input is somewhat artificial, because both inputs are combined in a new image **o**_{res} which constitutes the final input of the network. To predict the probability of success of a push proposal for the given singulation task the most important feature is the rela-



Fig. 3: We encode $\mathbf{0}$ and \mathbf{a} together in a push proposal image $\mathbf{0}_{res}$ that maps both the configuration of the objects in the scene and the push proposal direction into a 'push centric' frame.



Fig. 4: Our push proposal network architecture to predict a probability of singulation success for a push proposal image \mathbf{o}_{res} . The input to the network is a single channel grayscale version of the transformed image \mathbf{o}_{res} with a resolution of 64 × 64 pixels. The network consists of 5 convolutional layers and 2 fully connected layers.

tion of the candidate segment that we want to push and its relation to neighbouring segments. We propose to fuse the segmented camera image with a push proposal using rigid image transformations \mathbf{T}_t , \mathbf{T}_r , see Fig. 3. First, we transform the image by a rigid translation $\mathbf{T}_t(\mathbf{c})$ according to the start position \mathbf{c} of the push proposal. Second, we apply image rotation $\mathbf{T}_r(\alpha)$ according to the push angle α . The resulting image is $\mathbf{o}_{res} = \mathbf{T}_r \mathbf{T}_t \mathbf{o}$, as depicted in the right image of Fig. 3. At test time we forward the set of transformed input images $\{\mathbf{o}_{res}^1, \dots, \mathbf{o}_{res}^M\}$ into the push proposal network to predict a probability of success for each push proposal in the set $\{\mathbf{a}^1, \dots, \mathbf{a}^M\}$. The network architecture is depicted in Fig. 4.

We train the network from scratch using random uniform weight initialization and Adam [15]. We performed network architecture search with 10 different configurations and found that for our task the depicted network architecture yielded best performance on 2 separate validation sets. We experienced a drop in classification performance when increasing the image size, introducing dropout, or removing pooling operations, and did not experience improvement for adding xavier weight initialization or other activation functions besides ReLUs. Finally, we train our model for 25 epochs with a batch size of 64, using the default parameters for Adam (learning rate= 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\varepsilon = 1e - 08$, decay= 0.0).

3.2.3 Push Motion Planning

To find a motion plan for each push proposal we use the LBKPIECE motion planning algorithm provided by an out-of-the-box framework [28]. Our arm motion strategy consists of two steps. In the first step, we plan to reach the target push proposal \mathbf{a}^m . In a second step, given that the gripper reached the desired push start position, the robot performs a straight line push with fixed push length l_a . We compute motion plans for both steps before execution to avoid that the robot executes the reaching plan and then fails to find a plan for the straight line push. We found our two-step procedure to be more stable than executing a reach and push trajectory at once, due to arm controller errors.

4 Experiments

We conduct both qualitative and quantitative experiments on a real PR2 robot to benchmark the performance of our overall approach. Hereby, we aim to answer the following questions: 1) Can our push proposal network predict meaningful push actions for object shapes and configurations it has not seen during training? 2) Can our network trained in simulation generalize to a real-world scenario? 3) Can our model reason about scene ambiguities including multiple object clusters and focus its attention to objects that are more cluttered than others? In order to answer question (1) we perform real-world experiments with objects that the network has not seen during training and compare against a manual model-free baseline that reasons



Fig. 5: Starting configurations of the 25 test trials with 4 objects.

about available free space in the scene and keeps a history of previous actions. To answer question (2) we test a network that we train solely with data from simulation (denoted as vanilla network). Finally, three challenging experiments with 6 and 8 objects aim to answer question (3), where the learned model has to trade-off between further separating isolated objects to create more space and split isolated object clusters to solve the task. The overall intent of our experiments is to show that a learned model is able to directly map from sensor input to a set of possible actions, without prior knowledge about objects, spatial relations or physics simulation.

4.1 Experimental Setup

Our PR2 robot is equipped with a Kinect 2 camera mounted on its head that provides RGB-D images with a resolution of 960×540 . We use both arms of the robot, each with 7-DOF to perform the pushing actions. The task is designed as follows: a trial consists of a fixed number of objects from a test set of everyday objects with different shapes that the robot has never seen before. In the starting configuration for a trial all objects are placed or stacked on a table to be in collision. The starting configurations are saved and reproduced for all methods in a manual manner using point cloud alignment. We conduct different experiments of increasing difficulty ranging from 4 to 8 objects, in which for each experiment the number of objects stays fixed. The 25 starting configurations for our quantitative evaluation with 4 objects are depicted in Fig. 5.

We report the number of successful singulation trials. A singulation trial is considered successful if all objects are separated by a minimum distance of 3cm. Further, we report results for singulation success after every action, to show which methods can successfully execute the task with less actions. For each trial the robot performs a fixed set of actions. Given the number of objects n_o for each experiment the maximum number of pushes is $n_{pushes} = \lfloor 1.3 \cdot n_o \rfloor + 1$. The robot is allowed to stop before, if it reasons that all objects are singulated. This is implicitly encoded into our method. If the robot does not find a motion plan for the set of push proposals that the network ranked as positive it automatically terminates the trial.

To train our push proposal network we use an iterative training procedure. First, the vanilla network is trained on labeled data from a total of 2,486 (243 positives, 2,243 negatives) random interactions performed in simulation. Then the aggregated network is trained on additional 970 interactions (271 positives, 699 negatives), which we collected using the vanilla network in both simulation and real-world, resulting in a training dataset size of 3,456 push interactions. We use object meshes from the Freiburg Spatial Relations Dataset for our simulation experiments [23].

4.2 Baseline Method

We provide a quantitative comparison against a manually designed baseline method, to evaluate whether or not the predictions of the push proposal network lead to improved object singulation results. We do not provide knowledge about the objects and their physical properties to our model, correspondingly our baseline method is model-free and follows an interactive singulation strategy that is reasonably effective. The method which we denote as 'free space+tracking' works as follows:

- 1. Given a set of segments $\{s^1, ..., s^L\}$ from our object over-segmentation approach we construct a graph that includes all segments, where each segment is a node in the graph and the edges correspond to distances between the segments. We represent each segment as an axis-aligned bounding box (AABB) and compute an edge between 2 segments by means of the Manhattan distance of the 2 AABBs.
- 2. We compute 2 features for scoring push proposals with the baseline method. The first feature is a predictive feature that reasons about the resulting free space if a segment would be pushed to some target location. To compute the feature, we predict a straight line motion of the respective segment to which the push proposal is assigned to, according to the push direction and the length. Next, we compute the Manhattan distance between the resulting transformed AABB and all other segments in the scene, which we assume will remain static. The free space feature f_s is the minimum distance between the transformed AABB and the other AABBs. If the predicted final position of a segment would lead to collision with another segment the free space feature is zero.
- 3. The second feature includes the push history that we store for each segment. It follows the intuition that the same segment should not be pushed too often, regardless of the predicted free space around it. To retrieve the push history over all segments, we follow a segment-based tracking approach, which aligns the centroid and the principal components of two segments from the current set of segments and the set of segments from the last interaction. We match a set of segments using a weighted average of the principal components and the segment centroid distances $d(\mathbf{s}^l, \mathbf{s}^m) = 0.6 \cdot d_{pca}(\mathbf{s}^l, \mathbf{s}^m) + 0.4 \cdot d_c(\mathbf{s}^l, \mathbf{s}^m)$. To punish multiple pushes *r* of a segment throughout a trial we map the push motion history into a normalized feature using an exponential decay function $f_h = exp(-r)$.
- 4. Accordingly, a push proposal \mathbf{a}^m receives the score $p^m = 0.5 \cdot f_s + 0.5 \cdot f_h$.

4.3 Quantitative Comparisons

We extensively evaluate our method in real-robot object singulation experiments with 4,6 and 8 unknown objects. Additionally we perform an experiment with 6 identical objects to further challenge our approach. Our results with 4 objects (experiment 1), shown in Table 1, indicate that our method is able to improve over the performance of the manual baseline. The success rate of our vanilla network is 68%,

Experiment No.	Number of ob- jects	Method	Success/total trials	Success rate	Mean number of pushes
1	4 4	Free space + tracking Vanilla network	14/25 17/25	56% 68%	5.4±1.68 4.96±1.84
2	6 6 6	Free space + tracking Vanilla network Aggregated network	4/10 4/10 7/10	40% 40% 70%	7.8±1.61 8.1±1.37 6.7±2.0
3	6 (identical objects)	Aggregated network	10/20	50%	7.6±1.76
4	8	Aggregated network	4/10	40%	11.1±1.60

Table 1: Object singulation results for our four experiments of increasing difficulty.

which suggests that the model is making meaningful predictions about which push actions to perform in order to singulate all 4 objects. Fig. 6 provides a more finegrained evaluation, showing the success rate with respect to the number of pushes. Note that the network requires less push actions to complete the task. Interestingly, the baseline method performs on par with the network after the robot has executed five push actions, but then it does not further improve after 6 executed pushes. We noted that the baseline method sometimes fails to singulate the last 2 objects of the scene and instead will choose an object that might already be singulated because it lacks a good attention mechanism.

The task is more complex with 6 or 8 objects, due to additional space constraints and formation of multiple object clusters. When looking at the scene in Fig. 1 one sees that 4 out of the 6 objects are located very closely on the table (green drink bottle, blue box, white round bowl, white coffee bag). Although there are many possible push actions that we would consider reasonable, only pushing the coffee bag or the round bowl to the left side of the table can clear the scene. Accordingly, we



Fig. 6: Success rate with respect to number of pushes (max. 6) required to clear 4 objects (experiment 1).



Fig. 7: Success rate with respect to number of pushes (max. 8) required to clear 6 objects (experiment 2).

find that the performance of the vanilla network drops with respect to the previous experiment with 4 objects. During training it has only seen a small amount of scenes where a random baseline would have cleared all but 2 objects and even less likely has seen examples where the robot by chance chose the right action to clear the 2 remaining objects. Therefore, in this scenario the vanilla network and the baseline method perform on par with a success rate of 40%. When comparing the average number of pushes the baseline performs slightly better, see Fig. 7. Results show that the aggregated network clearly outperforms the other two methods, winning seven out of ten trials with an average of 6.7 push actions needed to singulate all 6 objects, as shown in Table 1 (experiment 2). To get an intuition about the numerical results we refer the reader to the performance reported by Hermans *et al.* [13], who evaluated a very similar task with 6 objects. They report a success rate of 20% with twelve average number of push actions.

We perform an additional experiment with 6 identical objects (sponges) using the aggregated network and reach a success rate of 50%. The most challenging experiment is conducted with 8 objects and even though the network was not trained with such large amount of objects on the table our approach performs well and is able to generalize. Both experiments are depicted in Table 1 (experiment 3 and 4).

4.4 Qualitative Results

Fig. 8 gives insights about the singulation trials from the perspective of the Kinect camera for experiment 1. The first and the third column show the current scene, while the second and fourth column show the robot performing a push action. The trial evolves from top to bottom of the figure. Fig. 9 shows qualitative end results of singulation runs with 6 identical objects (experiment 3) and 8 objects (experiment 4).



Fig. 8: On the left, we see a successful trial. Note that the robot tries several times to clear a small plastic bottle which is stuck under a blue box and manages to singulate the objects in the very last action. On the right we see a run that fails because of two typical reasons: First, the robot is not able to perform a save push action (fourth column, fifth row), which results in 2 objects moving more closely together (blue box, cereal box). Second, the network is not able to draw its attention to the same 2 objects that are now touching and instead proposes an action that moves the green bottle further away from the scene (lower right image).

5 Conclusions

We presented a novel neural network-based approach to singulate unknown objects in clutter by means of pushing actions and showed that it is possible to perform challenging perceptual manipulation tasks by learning rich feature representations from visual input. Unlike traditional methods we do not manually design features for the task at hand. Instead, we train a high-capacity convolutional neural network that we incorporate into a novel system of motion planning and action selection, leading to good generalization performance, while requiring a reasonable amount of labeled training data. We tested our method in extensive real-robot experiments using a PR2 robot and showed the ability of our method to achieve good performance for the challenging task of singulating up to 8 objects. In the future it would be interesting to train the network in a self-supervised manner, which poses the challenge to auLearning to Singulate Objects using a Push Proposal Network



Fig. 9: The first row depicts a successful singulation run with six identical objects (experiment 3). The left image depicts the starting configuration and the right image the end result after manipulation. The second row shows a failed trial. The third row shows a trial where the robot managed to singulate 8 objects (experiment 4). The last row depicts a trial that failed because the object cluster at the bottom is not cleared.

tomatically generate labels from ambiguous visual data. Furthermore, the network could be extended to a multi-class model that predicts different push lengths, which would require collecting additional training data with varying push lengths.

Acknowledgements This work was partially funded by the German Research Foundation under the priority program Autonomous Learning SPP 1527 and under grant number EXC 108. We thank

Seongyong Koo for advice with the baseline method. We further thank Sudhanshu Mittal, Oier Mees and Tim Welschehold for their help and ideas.

References

- Agrawal, P., Nair, A.V., Abbeel, P., Malik, J., Levine, S.: Learning to poke by poking: Experiential learning of intuitive physics. In: Advances in Neural Information Processing Systems (NIPS) (2016)
- Bohg, J., Hausman, K., Sankaran, B., Brock, O., Kragic, D., Schaal, S., Sukhatme, G.: Interactive perception: Leveraging action in perception and perception in action. arXiv preprint arXiv:1604.03670 (2016)
- Boularias, A., Bagnell, J.A.D., Stentz, A.T.: Learning to manipulate unknown objects in clutter by reinforcement. In: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI) (2015)
- Byravan, A., Fox, D.: Se3-nets: Learning rigid body motion using deep neural networks. arXiv preprint arXiv:1606.02378 (2016)
- Chang, L., Smith, J.R., Fox, D.: Interactive singulation of objects from a pile. In: Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA) (2012)
- Cosgun, A., Hermans, T., Emeli, V., Stilman, M.: Push planning for object placement on cluttered table surfaces. In: Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS) (2011)
- 7. Dogar, M., Srinivasa, S.: A framework for push-grasping in clutter. Robotics: Science and systems VII 1 (2011)
- Finn, C., Levine, S.: Deep visual foresight for planning robot motion. arXiv preprint arXiv:1610.00696 (2016)
- Gualtieri, M., ten Pas, A., Saenko, K., Platt, R.: High precision grasp pose detection in dense clutter. In: Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS). IEEE (2016)
- Gupta, M., Müller, J., Sukhatme, G.S.: Using manipulation primitives for object sorting in cluttered environments. IEEE transactions on Automation Science and Engineering 12(2) (2015)
- Hausman, K., Balint-Benczedi, F., Pangercic, D., Marton, Z.C., Ueda, R., Okada, K., Beetz, M.: Tracking-based interactive segmentation of textureless objects. In: Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA) (2013)
- Hermans, T., Li, F., Rehg, J.M., Bobick, A.F.: Learning contact locations for pushing and orienting unknown objects. In: Proc. of the IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids) (2013)
- Hermans, T., Rehg, J.M., Bobick, A.: Guided pushing for object singulation. In: Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS) (2012)
- 14. Katz, D., Venkatraman, A., Kazemi, M., Bagnell, J.A., Stentz, A.: Perceiving, learning, and exploiting object affordances for autonomous pile manipulation. Autonomous Robots (2014)
- 15. Kingma, D., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Koo, S., Lee, D., Kwon, D.S.: Unsupervised object individuation from rgb-d image sequences. In: Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS) (2014)
- Kopicki, M., Zurek, S., Stolkin, R., Mörwald, T., Wyatt, J.: Learning to predict how rigid objects behave under simple manipulation. In: Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA) (2011)
- Laskey, M., Chuck, C., Lee, J., Mahler, J., Krishnan, S., Jamieson, K., Dragan, A., Goldberg, K.: Comparing human-centric and robot-centric sampling for robot deep learning from demonstrations. In: Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA) (2017)

Learning to Singulate Objects using a Push Proposal Network

- Levine, S., Finn, C., Darrell, T., Abbeel, P.: End-to-end training of deep visuomotor policies. Journal of Machine Learning Research 17(39) (2016)
- Levine, S., Pastor, P., Krizhevsky, A., Quillen, D.: Learning hand-eye coordination for robotic grasping with large-scale data collection. In: International Symposium on Experimental Robotics (ISER) (2016)
- Lynch, K.M., Mason, M.T.: Stable pushing: Mechanics, controllability, and planning. The International Journal of Robotics Research 15(6) (1996)
- Mahler, J., Liang, J., Niyaz, S., Laskey, M., Doan, R., Liu, X., Ojea, J.A., Goldberg, K.: Dexnet 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics. arXiv preprint arXiv:1703.09312 (2017)
- Mees, O., Abdo, N., Mazuran, M., Burgard, W.: Metric learning for generalizing spatial relations to new objects. In: Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS) (2017)
- 24. Pinto, L., Gandhi, D., Han, Y., Park, Y.L., Gupta, A.: The curious robot: Learning visual representations via physical interactions. arXiv preprint arXiv:1604.01360 (2016)
- Pinto, L., Gupta, A.: Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. In: Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA) (2016)
- Richtsfeld, A., Mörwald, T., Prankl, J., Zillich, M., Vincze, M.: Segmentation of unknown objects in indoor environments. In: Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS) (2012)
- Schiebener, D., Ude, A., Morimoto, J., Asfour, T., Dillmann, R.: Segmentation and learning of unknown objects through physical interaction. In: Proc. of the IEEE-RAS Int. Conf. on Humanoid Robots (Humanoids) (2011)
- Sucan, I.A., Moll, M., Kavraki, L.E.: The open motion planning library. IEEE Robotics & Automation Magazine 19(4), 72–82 (2012)
- Van Hoof, H., Kroemer, O., Peters, J.: Probabilistic segmentation and targeted exploration of objects in cluttered environments. IEEE Transactions on Robotics 30(5) (2014)
- Yu, K.T., Bauza, M., Fazeli, N., Rodriguez, A.: More than a million ways to be pushed. a high-fidelity experimental dataset of planar pushing. In: Proc. of the IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS) (2016)
- Zhou, J., Paolini, R., Bagnell, J.A., Mason, M.T.: A convex polynomial force-motion model for planar sliding: Identification and application. In: Proc. of the IEEE Int. Conf. on Robotics & Automation (ICRA) (2016)
- Zhu, S., Kimmel, A., Bekris, K.E., Boularias, A.: Model identification via physics engines for improved policy search. arXiv preprint arXiv:1710.08893 (2017)