



HAL
open science

Data-efficient learning of object-centric grasp preferences

Yoann Fleytoux, Anji Ma, Serena Ivaldi, Jean-Baptiste Mouret

► **To cite this version:**

Yoann Fleytoux, Anji Ma, Serena Ivaldi, Jean-Baptiste Mouret. Data-efficient learning of object-centric grasp preferences. ICRA 2022 - IEEE International Conference on Robotics and Automation, May 2022, Philadelphia, United States. hal-03466780v2

HAL Id: hal-03466780

<https://hal.univ-lorraine.fr/hal-03466780v2>

Submitted on 24 Jun 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Data-efficient learning of object-centric grasp preferences

Yoann Fleytoux¹, Anji Ma^{1,2}, Serena Ivaldi¹, Jean-Baptiste Mouret¹

Abstract—Grasping made impressive progress during the last few years thanks to deep learning. However, there are many objects for which it is not possible to choose a grasp by only looking at an RGB-D image, might it be for physical reasons (e.g., a hammer with uneven mass distribution) or task constraints (e.g., food that should not be spoiled). In such situations, the preferences of experts need to be taken into account.

In this paper, we introduce a data-efficient grasping pipeline (Latent Space GP Selector — LGPS) that learns grasp preferences with only a few labels per object (typically 1 to 4) and generalizes to new views of this object. Our pipeline is based on learning a latent space of grasps with a dataset generated with any state-of-the-art grasp generator (e.g., Dex-Net). This latent space is then used as a low-dimensional input for a Gaussian process classifier that selects the preferred grasp *among those proposed by the generator*.

The results show that our method outperforms both GR-ConvNet and GG-CNN (two state-of-the-art methods that are also based on labeled grasps) on the Cornell dataset, especially when only a few labels are used: only 80 labels are enough to correctly choose 80% of the grasps (885 scenes, 244 objects). Results are similar on our dataset (91 scenes, 28 objects).

I. INTRODUCTION

Robots are often tasked with grasping objects from a box or a conveying belt, especially in industrial settings [1]. Thanks to the recent advances in deep learning, they can now grasp unknown objects with success rates that exceed 90% [2]. To do so, they learn the relationship between the shape of the objects and the most appropriate grasp.

Nevertheless, the best grasp according to the shape is not always the grasp that should be favored. For instance, a hammer should usually be taken by the head because the mass distribution is not uniform, whereas the shape suggests that it should be taken by the handle. Similarly, a sharp knife should usually be grasped by the handle to avoid damaging the gripper; but it might be necessary to grasp it by the blade if the robot needs to give it to an operator. For some objects, it might be legally forbidden to grasp them in some way, for example in the food industry; in some other cases, the objects might break or be scratched when grasped in the wrong way, for example, by grasping prescription glasses by a lens, or the objects require taking specific properties into account, like fresh paint on a sub-part. Importantly, the grasp that




Fig. 1. Grasp generated by Dex-Net 4.0 [3] for a toy hammer and with our method with only 2 labels (one positive label, in green, and one negative label, in red). See section V-A.

should be preferred does not always help grasping; instead, it is often a constraint of the task.

Hence, expert knowledge for grasping specific objects is required in many situations; but labeling thousands of images is time-consuming and would need to be performed for each application (e.g., for a hammer factory, then for a knife warehouse, etc.). Our grasp pipeline aims at learning object-centric grasping preferences with very few examples (typically 1 to 4 per object) and generalizing to other views of the same object. In this work, we target single objects on a simple background, for instance, objects on a conveyor belt in industrial settings.

Our main insight is that we can leverage generic grasp generators to make learning preferences data-efficient in two ways: (1) they can be run on a large dataset of images to learn a latent space for grasps, and (2) they can be used to generate grasp candidates so that a preference-based classifier only has to choose among good grasps. In other words, a grasp generator takes the needles out of the haystack, and a classifier only has to choose the best needle in a low-dimensional space.

Our main assumption is that we have access to a set of passively-obtained RGB-D images on which a shape-based grasp generator can be run; we consider this dataset as “large and cheap”. We represent the generated grasps using rotated and adjusted image patches that embed the grasp and its context in a single input [4], which makes it possible to learn a low-dimensional representation of both the grasps and their context with a Variational Auto-Encoder (VAE) [5], [6]. We then use this low-dimensional representation to train/query a Gaussian Process (GP) classifier [7], [8] that *filters* the grasps generated by a grasp generator, which are often already effective, according to the preferences of the expert. Importantly, the expert can give both “positive labels”, that is, grasps that should be favored (e.g., “this is how this should be done”), and “negative labels”, that is,

*This work was supported by the CHIST-ERA grant HEAP (CHIST-ERA-17-ORMR-003).

¹Université de Lorraine, CNRS, Inria, F-54600, France

²School of Mechatronic Engineering, Beijing Institute of Technology, Beijing 100081, China

jean-baptiste.mouret@inria.fr

grasps that should be avoided (e.g., “do not do this”). We call our method “Latent Space GP Selector” (LGPS).

Once trained, our pipeline generates grasp candidates, encodes them in the latent space, then queries a Gaussian process classifier to know the preference of the expert; the robot executes this grasp with a standard planning algorithm [9]. For the preference training, the grasp selected by the expert is encoded to the same latent space and the Gaussian process classifier is updated.

While VAE and GP have been combined together in different fields (e.g., [10] for videos), our main contribution is the combination of grasp generators (which can, for instance, be based on deep learning) with an image-based grasp representation to learn a latent space of grasps in an unsupervised way. We show that our pipeline makes it possible to learn grasps that are about 80% consistent with the expert labels with less than one example per object on the Cornell dataset [11] (885 scenes, 244 objects) and on our own dataset (91 scenes, 28 objects).

II. RELATED WORK

Vision-based robotic grasp planning methods can be classified into analytical [12] and data-driven approaches [13]. Analytical approaches require information about the physical properties of the manipulated objects (shape, mass, center of mass, friction coefficient), which is possible only on well-controlled manufacturing scenarios. By contrast, data-driven approaches attempt to generalize to unknown objects by being trained on datasets of grasp examples.

Data-driven approaches mainly differ in the kind of dataset they use. A recently successful idea is to create large synthetic databases from 3D models and simulations [4], [14]–[18]. The strength of these datasets is their size (millions of objects), but they currently only take into account the 3D shape (via depth data): they assume a uniform mass distribution of objects. As they strongly rely on simulation, they often need adaptation methods to be effective with real robots [19]; on the other hand, automatically gathering sufficient training data through trial and error with real robots [20], [21] is highly time-consuming and does not necessarily provide better results (e.g., 50,000 trials and 700 hours of robot use in [20], [21]). Human feedback was used [22], [23] such that the robot only performs hundreds of grasps during task learning, but these manipulations can cause safety concerns when dealing with particular objects where certain grasps must be avoided.

The second kind of data-driven approach learns from a dataset in which objects have been labeled by experts to specify where/how to grasp objects. Several studies have been proposed to exploit data from captured expert manipulations such as [24], which learns interactions from videos of experts, or [25], [26], which uses custom handheld devices to collect grasping demonstration. Nevertheless, many recent studies [27]–[29] use the Cornell dataset [11] or similarly acquired datasets [30], [31] which provide thousands of grasp locations labeled by humans.

These grasp demonstrations have been used to infer an evaluation function that ranks grasp candidates according to the expert specifications [20], [32], [33]. Recently Generative Grasping CNN (GG-CNN) [28], Generative Residual Convolutional Neural Network (GR-ConvNet) [29] and other work use hand-designed labels to generate pixel-wise grasp affordance map. The grasping task is then similar to semantic segmentation, in which some parts of the object are deemed graspable whereas some others have to be avoided.

While generic grasping of unknown objects needs to use generic datasets, labeling specific datasets for specific objects is highly time-consuming. A few papers therefore focus on learning grasp preferences with as few “demonstrations” as possible [34]. In particular, [35] proposes a CNN, pixel-wise segmentation pipeline that predicts authorized grasping locations from depth image while avoiding prohibited locations. The demonstrations are gathered by applying colored tape to the operator’s fingers, who thus can demonstrate parallel grasps. This pipeline was evaluated with 5 types of industrial objects (socket wrenches, pliers, light bulbs, cups, and screws), using between 1 and 3 demonstrations for training. To allow CNN to be trained with so few labels, the authors apply heavy data augmentation and train on each type of object separately, and that allows them to reach between 70% and 90% success depending on the object. However, they report that their pipeline can at best grasp 2 objects with 2 different grasping strategies with the same neural network.

III. PROBLEM STATEMENT

Our main assumptions are: (1) the robot has an RGB-D camera; (2) all the objects can be grasped from the top; (3) we have access to a large unlabeled dataset of RGB-D images (**D**); (4) we have access to a small labeled dataset (≤ 4 per object) of grasps as either positive (good grasp) or negative (grasp to avoid) (dataset **E**).

Our main objective is to learn to reproduce the grasps that the expert labeled as good for new views of objects that have been seen previously. Our secondary objective is to generalize to objects that have never been seen but that are close to those already labeled.

We evaluate the performance using the *rectangle metric* [11], which makes it possible to compare to previous works using published datasets [11], [36]. For this metric, grasps are represented as rectangles centered on the gripper’s center position (x, y) , rotated according to the orientation $\theta \in [-\frac{\pi}{2}, +\frac{\pi}{2}]$, with a width that is equaling the gripper opening l , and a height that encodes the tolerance. Two grasps are compared by looking at how much the two corresponding rectangles overlap. More precisely, given a proposed grasp GC and a ground truth grasp GT , the Intersection over Union (IoU, also called the Jaccard index) corresponds to the normalized area of intersection:

$$IoU(GT, GC) = \frac{|GT \cap GC|}{|GT \cup GC|} \quad (1)$$

The rectangle metric RM is equal to 1 if this overlapping area is above 0.25 and the angular difference is below $\frac{\pi}{6}$:

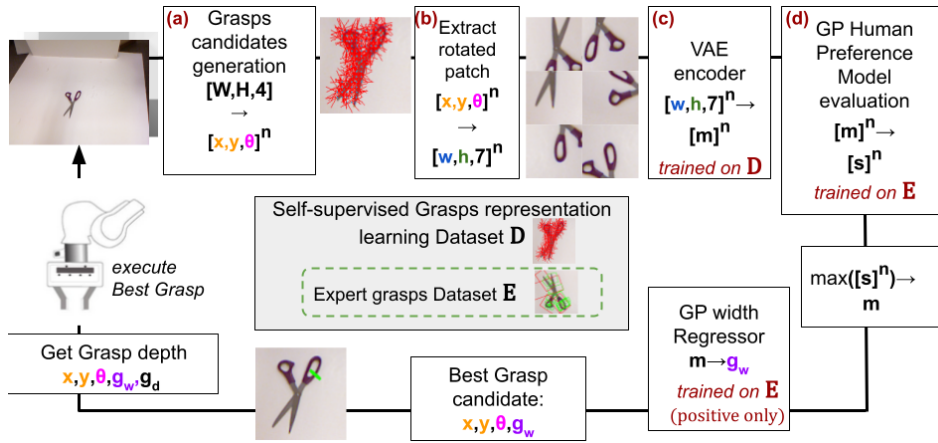


Fig. 2. The LGPS grasping pipeline, assuming that the VAE has been trained before (Sec. IV-C). From an RGB-D image, a grasp generator (Sec. IV-A) creates grasp candidates as segments. These grasp candidates are represented as rotated patches centered on the middle of the segment (Sec. IV-B). Each of them is fed to a VAE to get their latent representation, which is, in turn, the input of the GP classifier (Sec. IV-D) to obtain the estimated probability of being selected by the expert; the grasp with the highest probability is selected. A second GP is queried to get the width of the gripper for the selected grasp. The “depth” of the grasp (the z-position of the gripper) is computed using the depth image.

$$RM(GT, GC) = \begin{cases} 1, & \text{if } IoU(GT, GC) > 0.25 \\ & \text{and } |GT.\theta - GC.\theta| \leq \frac{\pi}{6} \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

We compare our results to GR-Convnet [29] and GC-CNN [28] with the rectangle metric on the Cornell Dataset [11] and our own dataset, which is focused on meaningful object-specific labels (scissors, hammer, etc.). Our method does not learn a height (tolerance) for the rectangles, therefore we have it set to 38 pixels.

While the process described in the present paper is offline — learning is performed with a labeled dataset (or a subset of this dataset) —, we envision future *online* deployments in which an expert corrects the robot online only when it is wrong.

IV. METHOD

Our full algorithmic pipeline is illustrated in (Fig. 2). The following subsections provide details about the respective components (a-d) of the pipeline.

A. Grasp candidate generation (a)

Given an RGB-D image, grasp candidates can be generated using computer vision techniques [37], or by randomly sampling them, or by using methods based on deep learning like Dex-Net [38], Generative Grasping CNN (GG-CNN) [28] or Generative Residual Convolutional Neural Network (GR-ConvNet) [29]. Here we use a grasp generator based on computer vision, which has proved to be both fast and effective in our experiments, once combined with a Gaussian process classifier (Sec. IV-D), although a deep-learning generator could take its place easily.

We first extract the object with the GrabCut algorithm [39] which works like a “magic wand” to separate an object from its background using four classes of pixels: from the objects, probably from the object, not from the object, and

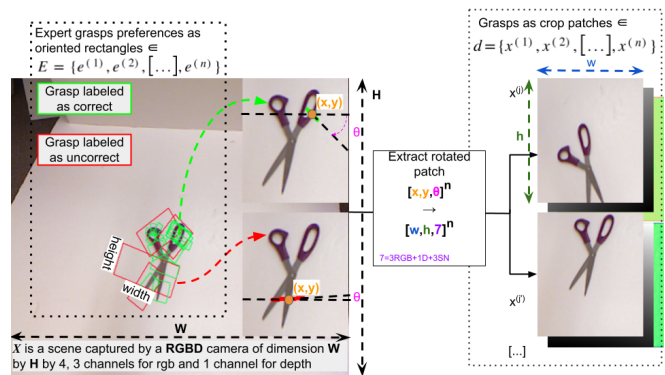


Fig. 3. Grasps are represented as a rotated and cropped patch of the image that is centered on the grasp. This representation makes it possible to feed both grasps and their context to convolutional neural networks.

probably not from the object. We determine the class of each pixel using three methods: HSV-based color segmentation, saliency based Segmentation [40], Gaussian Mixture-based Background/Foreground Segmentation [41]. We run a Canny edge detector [42] to obtain edges and we compute the skeleton [43].

We generate grasps candidates by computing lines that are perpendicular to each edge pixel and to each skeleton pixel [37]. This is achieved by looking at the two nearest neighbors of each point of the edges/skeleton. To keep the number of grasp candidates low, we skip edge/skeleton points if the distance with the previous point is below 4 pixels (Cornell dataset) and 8 pixels (our dataset). Finally, we add a random angle (± 0.4 rad, Gaussian distribution) to each candidate to increase the grasp diversity. At that stage, grasps have no gripper width since they are used to (1) generate patches (Section IV-B) which have a fixed size (therefore only the position and orientation matters), and (2) to query a GP, which will predict the gripper width using the latent representation of the patch.

These n points and angles $[x, y, \theta]^n$ are the list of grasps candidates for a specific RGB-D image.

B. Grasp candidate representation (b)

We encode grasps as 7-channel *image patches* (Fig. 3), like in [20] and [4], which centers each image on its respective grasp. The main benefit of this representation is that it combines a specification of the grasp (orientation, position relative to the object) with an image of the object to recognize it. These image patches are easily fed to convolutional neural networks, by contrast to a coordinate-based or feature-based representation, which would need to be associated with the right object.

To create these patches from $[x, y, \theta]$ (section IV-A), we: (1) rotate the image to correspond to the orientation of the segment, (2) translate the image to be centered on the center of the segment, (3) crop the image to 128×128 pixels. The width of the gripper is ignored for the patch extraction stage. Our image patches are 7-channel: 3 channels for the RGB image, 1 channel for the depth, and 3 channels for the surface normal image, generated from the depth gray-scale image.

C. Latent space for grasp (c)

We here assume that we have access to a large dataset of patches (called d) that is not labelled. We train a convolutional β -Variational Auto-Encoder [6], [44] using a large number of patches (at least 40,000, depending on the dataset) generated from RGB-D images. The last layer of the decoder uses a tanh activation function because we normalize our input to $[-1, 1]$. Given the trained VAE, the extraction of grasp patches and their encoding into the latent space takes less than 2 seconds per scene on average on our computer.

D. Expert preference model learning (d) with Gaussian processes

We now assume that we have access to a second dataset, for which a small set of patches have been labeled either as positive (selected by the expert) or negative (grasp to avoid). For the positive examples, the dataset also contains the gripper opening width selected by the expert.

We train a Gaussian processes classifier [8] that takes as input the latent code for a patch, that is, a grasp candidate, and outputs a score s between 0 and 1 that describes the probability that the grasp would be chosen by the expert. Using the positive example, we also train Gaussian process regressor [8] that takes the same input and outputs a probability distribution of the width of the gripper. Compared to neural networks, Gaussian processes classifiers are more accurate when there are few data [8], at the expense of a longer query/training time (query is $O(n^2)$ with n the number of samples). In our tests, we observed a query time of less than 0.06 second per scene.

For each generated grasp (Sec. IV-A), we first generate the corresponding patch, encode it to the latent space using the β -VAE decoder, then query the GP to obtain its score. We select the grasp candidate with the highest score. To scale the GP classifier to many samples without compromising on

query time, we use ‘‘Scalable Variational Gaussian Process classifiers’’ [7], which is based on a variational inducing point framework, although any GP classifier could be used instead (for less than about 500 labels, a standard GP classifier works well). The gripper opening width g_w is selected by querying the GP regressor with the selected patch (a standard GP is used for this as we only need a single query, compared to many queries for the classifier).

To use the oriented rectangle representation (see Fig. 3), we use a fixed *height* value. When executing the selected grasp with the robot, the depth is computed by using the depth data from the RGB-D camera: we extract an oriented cropped patch of the depth point cloud with the width of the selected grasp and a fixed height (5 pixels), and we use the closest point to the gripper g_d (that is, the highest point of the object) as the z -reference.

V. EXPERIMENTAL EVALUATION

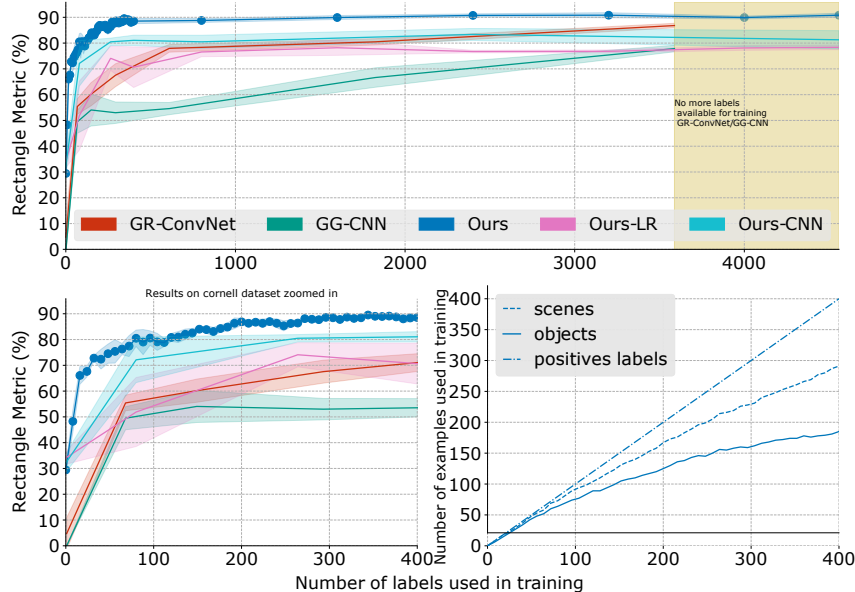
We evaluate our approach on the Cornell dataset [45] (885 scenes, 244 objects, 5055 positive label, 2822 negative labels) and on a custom dataset for which humans typically have preferences (91 scenes, 28 objects, 447 positive labels, 145 negative labels). For the Cornell dataset, the VAE is trained on the 885 scenes (244 objects). For our dataset, images are obtained with an Intel RealSense D415 Depth Camera mounted on the gripper of a Franka-Emika Panda robot (Fig. 1) that is positioned 65 centimeters above the objects. The objects are mostly from the YCB dataset [46]. We used a 32-dimensional latent space (output of the VAE and input of the GP).

We focus on the performance with very few labels (fewer than 4000, ideally fewer than 50) because we envision interactive or semi-interactive scenarios in which an expert does not want to spend much time in labeling. In addition, we are interested in generalizing to new views of objects for which we have labels, and not necessarily to new objects that have never been seen. As a consequence, we make no effort to split the training set and the learning set into disjoint sets of objects; on the contrary, we typically expect our algorithm to select the right grasp with 1-2 grasp examples of the same object from different views. For instance, we do not expect our algorithm to know how to grasp a hammer if it has never seen a hammer, but we want it to learn how to grasp a hammer from any point of view once it has been explained how to do so once or twice.

We compare our method to two state-of-the-art grasping methods that are based on expert labels: Generative Grasping CNN (GG-CNN) [28] and Generative Residual Convolutional Neural Network (GR-ConvNet) [29] using the rectangle metric (Sec. III). We were unable to run Dex-Net 4.0 [3] on the Cornell dataset because the objects are too far from the camera (Dex-Net requires 0.5 to 0.7m) and the dataset does not give the camera intrinsic parameters. GG-CNN and GR-Convnet are two recent algorithms that use a fully-convolutional neural network to generate grasp quality and grasp poses/width (x, y, θ, w) at every pixel from RGB-D

Results on the Cornell dataset (885 scenes, 244 objects)

IW split training 231 objects 641 scenes
3621 labels+, 2111 labels-



Results on our dataset (91 scenes, 28 objects)

IW split training 26 objects 63 scenes
304 labels+, 100 labels-

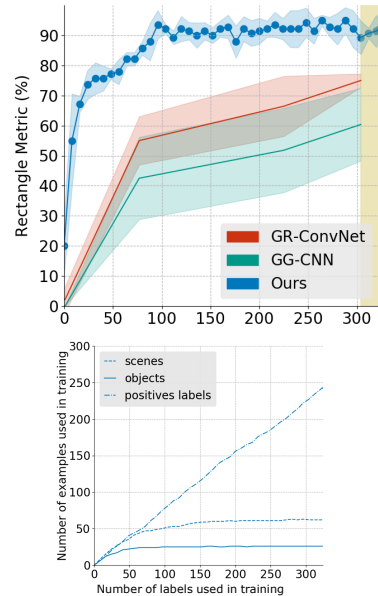


Fig. 4. Rectangle metric for the Cornell dataset and our dataset. The x -axis corresponds to the number of labels used for training (training+validation data, 641 scenes) and the y -axis to the rectangle metric in percent of the number of scenes in the test set (244 scenes). The rectangle metric counts how often the selected grasp matches at least one of the positive labels. For the Cornell dataset, the top plot shows the rectangle metric from 0 to all the available labels (GR-ConvNet and GG-CNN do not use the negative labels, therefore they cannot use all the labels), the bottom left plot corresponds to the same data but focuses on the results from 0 to 400 labels, and the bottom right plot reports the number of scenes and objects with regards to the number of labels used for training. To better understand the performance of our pipeline, we compare two “ablations”: switching the GP classifier with a logistic regression model for binary classification (“Ours-LR”) and using a CNN binary classifier instead of the VAE+GP combination (“Ours-CNN”).

scenes and to sample grasp candidates. GR-ConvNet reports a state-of-the-art accuracy of 97.7% on the Cornell dataset.

For the following experiments, we are interested in comparing our method’s performances (in terms of following the preferences) according to the number of available labels. To split our dataset image-wise, we order the RGB-D scenes by objects and randomly pick one scene per object for the testing data set, on which every algorithm is tested.

We generate the training and validation sets for a specified number of labeled grasps by randomly selecting data points outside of the test set. These labels are then split into 5 groups in order to perform 5-fold cross validation to train GG-CNN and GR-ConvNet. Each method is trained with the same labels available for training, with this difference: GG-CNN and GR-ConvNet discard negative grasps for training (in these methods, everything that is not positive is considered as negative), and our method discards the validation data (since we do not need it for the GP). For instance, for 100 randomly selected labels (e.g., 60 positives and 40 negatives), each fold contains 80 labels for training and 20 for validation; our method is trained with 80 labels whereas GG-CNN and GR-ConvNet would use on average 48 labels (36 labels for training and 12 labels for validation). Please note that we report data using the number of used labels (training and validation), not the number of randomly selected labels, so that the comparison is as fair as possible.

We run our method 5 times for each number of labels (each run is independent); we test with 0, 10, 20, \dots , 500 la-

bels and 1000, 2000, \dots , 4558 labels. Overall, we therefore launch 280 ($56 \times 5 = 280$) independent runs of our learning algorithm (and therefore 280 tests). For GR-ConvNet and GG-CNN, we test with 0, 100, 300, 400, 1000, 3000, and 3588 labels (these methods only use the positive labels) and 5 replicates (total 45 runs for each baseline). We train the VAE part of our method once for each dataset \mathcal{D} , using a 70%-30% training and validation split. The VAE network is trained for 48 hours on a Nvidia GTX 1080 Ti using the Keras framework and the Adam optimizer with the following parameters: $|B| = 64, \eta = 0.001, w = 128, h = 128, \beta = 1.0, m = 32$. GG-CNN and GR-ConvNet are both trained on Nvidia GTX 1080 with the provided settings and data augmentations suited for the Cornell Dataset.

A. Qualitative results

We first checked that our method fits our expectations on an object with a clear choice (Fig. 1). In this experiment, a (toy) hammer has to be taken from the top of the handle (close to the head), and not from the bottom (far from the head). We have chosen this example because a hammer has a very non-uniform distribution of mass that cannot be deduced from the shape alone: grasps that are not close to the head are unlikely to be successful.

For this preliminary experiment, we used a single image with two labels: a positive label for the top part of the handle and a negative label for the bottom part (Fig. 1). We then captured 21 additional scenes of the same hammer

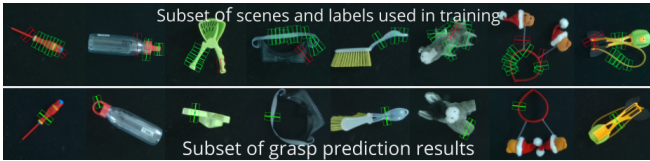


Fig. 5. Typical results from our dataset: the top row shows the objects in scenes used for training (additional scenes of these objects were also used), and the bottom row shows some of our predicted grasps on unseen scenes of the same object. For benchmarking data-efficiency, a random subset of labels is used (e.g., only 2 labels for each object).

in different positions and evaluated how often our method chose to grasp the hammer from the top part (to perform this evaluation, we labeled the 21 scenes, but did not use these labels in training).

The results show that our method selects the right grasp in 19/21 scenes and only fails in the 2 scenes in which the hammer is upright on the table (for which the handle is not accessible with a top-down grasp).

B. Quantitative results

After about 16 example scenes from the Cornell dataset (16 random labels, which correspond to about 13 positive examples for about 14 objects), our method selects the right grasp (according to the rectangle metric) for 66% of the test scenes (244 scenes) (Fig. 4) that include many unseen objects (244 objects total, whereas at best 16 different objects have been seen). As a comparison point, randomly selecting the grasp among those generated leads to a score of 29% (this corresponds to the “0 label” value on the plot), which means that the GP classifier is learning useful knowledge. After 68 labels, our method accurately predicts the grasps for 76% of the scenes, whereas GR-ConvNet selects the right grasp in only 55% of the test scenes, and GG-ConvNet only 49%. Our method exceeds 80% with 80 labels (63 different objects, 71 scenes, 54 positive labels) and 89.5% with 344 labels. By contrast, the best score reached by GR-ConvNet after 3588 labels¹ is 86.8% and by GG-CNN is 77.9%. Overall, our method clearly outperforms the best score of all the baselines after only 344 labels (versus more than 3000 for the baselines), and only a few hundred labels are enough to get more than 80% of the grasps right.

The results are similar for our dataset (Fig. 5), which is smaller but similar to the Cornell dataset (single object with labels) : our method reaches a score of 74% with only 24 labels (from 22 scenes, 15 different objects, 20 of them being positive), and 93.57% with 96 labels. To provide a reference point, with 300 labels, GR-ConvNet obtains a score of 75% and GG-CNN a score of 60%.

C. Ablation study

Our pipeline combines a VAE to learn a generic latent space with a GP to classify grasp candidates in this la-

¹Please note that the authors GR-ConvNet [29] report an accuracy of 97.7% with all the available labels, but we were unable to obtain the same score with their code (they do not report any score with a subset of the Cornell dataset, as we do here). Our results are consistent with the results reported in the GG-CNN literature [29].

tent space. To understand the contribution of the GP, we replaced it with a logistic regression classifier (from Scikit-Learn). The results (Fig. 4, “Ours-LR”) show that the logistic regression leads to scores similar to those obtained by GR-ConvNet, which shows that the GP classifier is a key component. To evaluate the contribution of the VAE trained on many unlabeled images, we then trained a CNN classifier with the same architecture as the VAE (it takes the same patches as input) but that outputs directly the class of the grasp candidate (correct or incorrect). This CNN is trained only on the labeled patches. The results show (Fig. 4, “Ours-CNN”) that this simple approach can lead to competitive results but is clearly outperformed by our pipeline. Overall, both training on unlabeled images and using a GP are necessary to reach a good rectangle metric score.

D. Tests with a Franka-Emika Panda robot

The supplementary video (<https://youtu.be/2-ncHGVChe0>) shows a Franka-Emika Panda robot that performs the grasps learned with our pipeline and our dataset (Fig. 1). For each of the 28 objects, the camera on the gripper is positioned 65 cm above the table, our pipeline selects the grasp, and the robot executes it using the planning algorithms of the MoveIt library [47].

The objective of our method is to follow the expert preferences, not to achieve the most successful grasps: the expert might require a grasp that is much harder than other grasps for their own reason (which is what the rectangle metrics evaluate). In these robot experiments, 17% of the grasps did not match what the expert expected and 20% of the grasps were unsuccessful (these two sets of failure do not fully overlap). Most failures came from issues with the grasp generation step (e.g., the preferred grasp was not proposed by the generator or no good grasp was proposed).

VI. CONCLUSIONS

Our pipeline is highly data-efficient because it uses a shape-based grasp generator as a prior, instead of learning grasps only from labels. A direct consequence is that our pipeline is at least as good as the grasp generator that is used: it always chooses among the proposed grasps. When more effective grasp generators are developed, they can be directly leveraged to both create the latent space and the candidates, and our pipeline becomes more effective.

This paper focuses on offline training in order to compare it to the state-of-the-art on well-defined datasets. However, our pipeline would fit well in online scenarios during which a “supervisor” corrects the robot only whenever it is wrong, that is, when the shape-based grasp generator is wrong. The supervisor would then act like a teacher with a competent student, with minimal supervision. In future work, we will design human-in-the-loop studies in conjunction with recent grasp generators to assess the performance on such an online learning process.

REFERENCES

- [1] D. Prattichizzo and J. C. Trinkle, “Grasping,” in *Springer handbook of robotics*. Springer, 2016, pp. 955–988.

- [2] G. Du, K. Wang, S. Lian, and K. Zhao, "Vision-based robotic grasping from object localization, object pose estimation to grasp estimation for parallel grippers: a review," *Artificial Intelligence Review*, vol. 54, no. 3, pp. 1677–1734, 2021.
- [3] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, "Learning ambidextrous robot grasping policies," *Science Robotics*, vol. 4, no. 26, 2019.
- [4] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," in *Robotics: Science and Systems XIII*, 2017.
- [5] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in *International Conference on Learning Representations (ICLR)*, 2014.
- [6] C. P. Burgess, I. Higgins, A. Pal, L. Matthey, N. Watters, G. Desjardins, and A. Lerchner, "Understanding disentangling in β -vae," *CoRR*, vol. abs/1804.03599, 2018. <http://arxiv.org/abs/1804.03599>
- [7] J. Hensman, A. G. de G. Matthews, and Z. Ghahramani, "Scalable variational gaussian process classification," in *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics, AISTATS*, vol. 38, 2015.
- [8] C. E. Rasmussen and C. K. I. Williams, *Gaussian processes for machine learning*. MIT Press, 2006.
- [9] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, A. Y. Ng *et al.*, "ROS: an open-source robot operating system," in *ICRA workshop on open source software*, 2009.
- [10] Y. Yoo, S. Yun, H. Jin Chang, Y. Demiris, and J. Young Choi, "Variational autoencoded regression: high dimensional regression of visual data on complex manifold," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [11] Y. Jiang, S. Moseson, and A. Saxena, "Efficient grasping from RGBD images: Learning using a new rectangle representation," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2011.
- [12] C. Ferrari and J. F. Canny, "Planning optimal grasps," in *IEEE International Conference on Robotics and Automation (ICRA)*, 1992, pp. 2290–2295.
- [13] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis — a survey," *IEEE Transactions on Robotics*, vol. 30, no. 2, pp. 289–309, 2013.
- [14] E. Johns, S. Leutenegger, and A. J. Davison, "Deep learning a grasp function for grasping under gripper pose uncertainty," in *IEEE/RSJ International Conference on Intelligent Robots and Systems, (IROS)*, 2016.
- [15] M. Gualtieri and R. P. Jr., "Learning 6-dof grasping and pick-place using attention focus," in *2nd Annual Conference on Robot Learning (CoRL)*, vol. 87. PMLR, 2018, pp. 477–486.
- [16] A. Zeng, S. Song, S. Welker, J. Lee, A. Rodriguez, and T. A. Funkhouser, "Learning synergies between pushing and grasping with self-supervised deep reinforcement learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [17] A. Mousavian, C. Eppner, and D. Fox, "6-DOF GraspNet: Variational grasp generation for object manipulation," in *IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 2901–2910.
- [18] U. Viereck, A. ten Pas, K. Saenko, and R. P. Jr., "Learning a visuomotor controller for real world robotic grasping using simulated depth images," in *Conference on Robot Learning (CoRL)*, 2017.
- [19] K. Bousmalis, A. Irpan, P. Wohlhart, Y. Bai, M. Kecey, M. Kalakrishnan, L. Downs, J. Ibarz, P. Pastor, K. Konolige, S. Levine, and V. Vanhoucke, "Using simulation and domain adaptation to improve efficiency of deep robotic grasping," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 4243–4250.
- [20] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.
- [21] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [22] C. Daniel, O. Kroemer, M. Viering, J. Metz, and J. Peters, "Active reward learning with a novel acquisition function," *Autonomous Robots*, vol. 39, no. 3, pp. 389–405, 2015.
- [24] A. Migliozi, G. Laudante, P. Falco, and B. Siciliano, "Vision-based
- [23] R. Pinsler, R. Akrouf, T. Osa, J. Peters, and G. Neumann, "Sample and feedback efficient hierarchical reinforcement learning from human preferences," in *IEEE International Conference on Robotics and Automation, (ICRA)*, 2018, pp. 596–601.
- grasp learning of an anthropomorphic hand-arm system in a synergy-based control framework," *Science Robotics*, vol. 4, no. 26, 2019.
- [25] P. Praveena, G. Subramani, B. Mutlu, and M. Gleicher, "Characterizing input methods for human-to-robot demonstrations," in *14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2019.
- [26] S. Song, A. Zeng, J. Lee, and T. A. Funkhouser, "Grasping in the wild: Learning 6DoF closed-loop grasping from low-cost demonstrations," *CoRR*, vol. abs/1912.04344, 2019. <http://arxiv.org/abs/1912.04344>
- [27] F. Chu, R. Xu, and P. A. Vela, "Real-world multiobject, multigrasp detection," *IEEE Robotics and Automation Letters*, 2018.
- [28] D. Morrison, J. Leitner, and P. Corke, "Closing the loop for robotic grasping: A real-time, generative grasp synthesis approach," in *Robotics: Science and Systems XIV*, 2018.
- [29] S. Kumra, S. Joshi, and F. Sahin, "Antipodal robotic grasping using generative residual convolutional neural network," *CoRR*, vol. abs/1909.04810, 2019. <http://arxiv.org/abs/1909.04810>
- [30] A. Zeng, S. Song *et al.*, "Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [31] E. De Coninck, T. Verbelen, P. Van Molle, P. Simoens, and B. Dhoedt, "Learning robots to grasp by demonstration," *Robotics and Autonomous Systems*, vol. 127, p. 103474, 2020.
- [32] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *The International Journal of Robotics Research*, vol. 34, no. 4-5, pp. 705–724, 2015.
- [33] D. Park and S. Y. Chun, "Classification based grasp detection using spatial transformer network," *arXiv preprint arXiv:1803.01356*, 2018.
- [34] P. Van Molle, T. Verbelen, E. De Coninck, C. De Boom, P. Simoens, and B. Dhoedt, "Learning to grasp from a single demonstration," *arXiv preprint arXiv:1806.03486*, 2018.
- [35] F. H el enon, L. Bimont, E. Nyiri, S. Thiery, and O. Gibaru, "Learning prohibited and authorised grasping locations from a few demonstrations," in *29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2020, pp. 1094–1100.
- [36] A. Depierre, E. Dellandr ea, and L. Chen, "Jacquard: A large scale dataset for robotic grasp detection," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [37] M. Vohra, R. Prakash, and L. Behera, "Real-time grasp pose estimation for novel objects in densely cluttered environment," *CoRR*, vol. abs/2001.02076, 2020. <http://arxiv.org/abs/2001.02076>
- [38] J. Mahler, F. T. Pokorny, B. Hou, M. Roderick, M. Laskey, M. Aubry, K. Kohlhoff, T. Kr oger, J. J. Kuffner, and K. Goldberg, "Dex-net 1.0: A cloud-based network of 3d objects for robust grasp planning using a multi-armed bandit model with correlated rewards," in *IEEE International Conference on Robotics and Automation, (ICRA)*, 2016.
- [39] C. Rother, V. Kolmogorov, and A. Blake, "'grabcut': interactive foreground extraction using iterated graph cuts," *ACM Transactions on Graphics*, vol. 23, no. 3, pp. 309–314, 2004.
- [40] X. Hou and L. Zhang, "Saliency detection: A spectral residual approach," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007.
- [41] P. KaewTraKulPong and R. Bowden, "An improved adaptive background mixture model for real-time tracking with shadow detection," in *Proc. 2nd European Workshop on Advanced Video Based Surveillance Systems*, 2002.
- [42] J. F. Canny, "A computational approach to edge detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 8, no. 6, pp. 679–698, 1986.
- [43] T. Lee, R. L. Kashyap, and C. Chu, "Building skeleton models via 3-d medial surface/axis thinning algorithms," *Graphical Models and Image Processing*, vol. 56, no. 6, pp. 462–478, 1994.
- [44] I. Higgins, L. Matthey, A. Pal, C. Burgess, X. Glorot, M. Botvinick, S. Mohamed, and A. Lerchner, "Beta-vae: Learning basic visual concepts with a constrained variational framework," in *5th International Conference on Learning Representations (ICLR)*, 2017.
- [45] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," in *Robotics: Science and Systems IX*, 2013.
- [46] B.  alli, A. Walsman, A. Singh, S. S. Srinivasa, P. Abbeel, and A. M. Dollar, "Benchmarking in manipulation research: The YCB object and model set and benchmarking protocols," *CoRR*, vol. abs/1502.03143, 2015. <http://arxiv.org/abs/1502.03143>
- [47] S. Chitta, I. Sucas, and S. Cousins, "MoveIt!" *IEEE Robotics & Automation Magazine*, vol. 19, no. 1, pp. 18–19, 2012.