

GQ-STN: Optimizing One-Shot Grasp Detection based on Robustness Classifier

Alexandre Gariépy, Jean-Christophe Ruel, Brahim Chaib-draa and Philippe Giguère

Abstract—Grasping is a fundamental robotic task needed for the deployment of household robots or furthering warehouse automation. However, few approaches are able to perform grasp detection in real time (frame rate). To this effect, we present Grasp Quality Spatial Transformer Network (GQ-STN), a one-shot grasp detection network. Being based on the Spatial Transformer Network (STN), it produces not only a grasp configuration, but also directly outputs a depth image centered at this configuration. By connecting our architecture to an externally-trained grasp robustness evaluation network, we can train efficiently to satisfy a robustness metric via the backpropagation of the gradient emanating from the evaluation network. This removes the difficulty of training detection networks on sparsely annotated databases, a common issue in grasping. We further propose to use this robustness classifier to compare approaches, being more reliable than the traditional rectangle metric. Our GQ-STN is able to detect robust grasps on the depth images of the Dex-Net 2.0 dataset with 92.4 % accuracy in a single pass of the network. We finally demonstrate in a physical benchmark that our method can propose robust grasps more often than previous sampling-based methods, while being more than 60 times faster.

I. INTRODUCTION

Grasping, corresponding to the task of grabbing an object initially resting on a surface with a robotic gripper, is one of the most fundamental problems in robotics. Its importance is due to the pervasiveness of operations required to seize objects in an environment, in order to accomplish a meaningful task. For instance, manufacturing systems often perform pick-and-place, but rely on techniques such as template matching to locate pre-defined grasping points [1]. In a more open context such as household assistance, where objects vary in shape and appearance, we are still far from a completely satisfying solution. Indeed, in an automated warehouse, it is often one of the few tasks still performed by humans [2].

To perform autonomous grasping, the first step is to take a sensory input, such as an image, and produce a grasp configuration. The arrival of active 3D cameras, such as the *Microsoft Kinect*, enriched the sensing capabilities of robotic systems. One could then use analytical methods [3] to identify grasp locations, but these often assume that we already have a model. They also tend to perform poorly in the face of sensing noise. Instead, recent methods have explored data-driven approaches. Although sparse coding has been used [4], the vast majority of new data-driven grasping approaches employ machine learning, more specifically deep learning [5]–[9]. A major drawback to this is that deep learning approaches require a significant amount of training data. Currently, grasping training databases based on real




Fig. 1. Overview of our method. Classical one-shot methods for grasping supervise a prediction (in red) using geometric supervision from randomly selected ground truth (in green). We instead suggest to use robustness supervision (in cyan) to learn fine-grained adjustments, without requiring a ground truth annotation at this specific grasp location.

data are scant, and generally tailored to specific robotic hardware [10], [11]. Given this issue, others have explored the use of simulated data [12], [13].

Similarly to computer vision, data-driven approaches in grasping can be categorized into *classification* and *detection* methods. In classification, a network is trained to predict if the sensory input (a cropped and rotated part of the image) corresponds to a successful grasp location. For the detection case, the network outputs directly the best grasp configuration for the whole input image. One issue with classification-based approaches is that they require a search on the input image, in order to find the best grasping location. This search can be exhaustive, and thus suffers from the curse of dimensionality [14]. To speed-up the search, one might use informed proposals [6], [12], in order to focus on the most promising parts of the input image. This tends to make the approach relatively slow, depending on the number of proposals to evaluate.

While heavily inspired by computer vision techniques, training a network for detection in a grasping context is significantly trickier. As opposed to classic vision problems, for which detection targets are well-defined instances of objects in a scene, grasping configurations are continuous. This means that there exist a potentially infinite number of successful grasping configurations. Thus, one cannot exhaustively generate all possible valid grasps in an input image. Another issue is that grasping databases are not providing the *absolute best* grasping configuration for a given image of an object, but rather a (limited) number of *valid* grasping configurations.

In this paper, we propose a one-shot grasping detection architecture for parallel grippers, based on deep learning. Importantly, our detection approach on depth images can be trained from *sparse* grasping annotations meant to train

a classifier. As such, it does not require the best grasping location to be part of the training dataset. To achieve this, we leverage a pre-existing grasp robustness classifier, called Grasp Quality CNN (GQ-CNN) [12]. This is made possible by the fact that our network architecture directly outputs an image corresponding to a grasp proposal, allowing it to be fed directly to an image-based grasp robustness classifier. Our architecture makes extensive use of the STN [15], which can learn to perform geometric transformations in an end-to-end manner. Because our network is based on STNs, the gradient generated by the GQ-CNN robustness classifier will propagate throughout our architecture. Our network is thus able to climb the robustness gradient, as opposed to simply regressing towards grasp configurations, which are limited in the training database. In some sense, our network is able to learn from the implicit knowledge of the quality of a grasp, knowledge that was captured by GQ-CNN.

In short, our contributions are the following:

- 1) Describing one of the first techniques to train a one-shot detection network on the detection version of the Dex-Net 2.0 dataset. Our network is based on an attention mechanism, the STN, to perform one-shot grasping detection, resulting in our Grasp Quality Spatial Transformer Network (GQ-STN) architecture;
- 2) Using the Grasp Quality CNN (GQ-CNN) as a supervisor to train this one-shot detection network, thus enabling to learn from a limited number of grasp annotations and to achieve a high robustness classification score;
- 3) Showing that our method generalizes well to real-world conditions in a physical benchmark, where our GQ-STN proposes a high rate of robust grasp.

II. RELATED WORK

Over the years, many network architectures have been proposed to solve the grasping problem. Here, we present them grouped by themes, either based on their overall method of operation or on the type of generated output.

A. Proposal + Classification Approaches

Drawing inspiration from previous data-driven methods [3], some approaches work in a two-stage manner, first by proposing grasp candidates then by choosing the best one via a classification score. Note that this section does not include architecture employing Region Proposal Network (RPN), as these are applied on a fixed-grid pattern, and can be trained end-to-end. They are discussed later.

Early work in applying deep learning on the grasping problem employed such a classification approach. For instance, Lenz *et al.* [14] employed a cascaded approach of two fully-connected neural networks. The first one was designed to be small and fast to evaluate and perform the exhaustive search. The second and larger network then evaluated the best 100 proposals of the previous network. This architecture achieved 93.7% accuracy on the Cornell Grasping Dataset (CGD).

Pinto *et al.* [10] reduced the search space of grasp proposals by only sampling grasp locations x, y and cropping

a patch of the image around this location. To find the grasp angle, the author proposed to have 18 outputs, separating the angle prediction into 18 discrete angles by 10° increments.

The *EnsembleNet* Asif *et al.* [16] worked in a different manner. It trained four distinct networks to propose different grasp representations (regression grasp, joint regression-classification grasp, segmentation grasp, and heuristic grasp). Each of these proposals was then ranked by the *SelectNet*, a grasp robustness predictor trained on grasp rectangles.

To alleviate the issue of small training datasets labelled manually, Mahler *et al.* [12] relied entirely on a simulator setup to generate a large database of grasp examples called Dex-Net 2.0 (see section III-B). Each grasp example was rated using a rule-based grasp robustness metric named *Robust Ferrari Canny*. By thresholding this metric, they trained a deep neural network, dubbed Grasp-Quality CNN (GQ-CNN), to predict grasp success or failure. The GQ-CNN takes as input a 32×32 depth image centered on the grasp point, taken from a top view to reduce the dimensionality of the grasp prediction. For grasp detection in an image, they used an antipodal sampling strategy. This way, 1000 antipodal points on the object surface were proposed and ranked with GQ-CNN. Even though their system is mostly trained using synthetic data, it performed well in a real-world setting. For example, it achieves a 93% success rate on objects seen during the training time and 80% success rate on novel objects on a physical benchmark.

Park *et al.* [6] decomposed the search for grasps in different steps, using STNs. The first STN acted as a proposal mechanism, by selecting 4 crops as candidate grasp locations the image. Then, each of these 4 crops were fed into a single network, comprising a cascade of two STNs: one estimated the grasp angle and the last STN chose the image's scaling factor and crop. The latter crop can be seen as a fine adjustment of the grasping location. The four final images were then independently fed to a classifier, to find the best one. Each component, being the STNs and the classifier, were trained on CGD separately using ground truth data and then fine-tuned together. This is a major distinction from other *Proposal + Classification* approaches, as the others cannot jointly train the *proposal* and *classification* sub-systems.

B. Single-shot Approaches

1) *Regression Approaches*: To eliminate the need to perform the exhaustive search of grasp configurations, Redmon *et al.* [17] proposed the first one-shot detection approach. To this effect, the authors proposed different CNN architectures, in which they always used AlexNet[18] pretrained on ImageNet as the feature extractor. To exploit depth, they fed the depth channel from the RGB-D images into the blue color channel, and fine-tuned. The first architecture, named *Direct Regression*, directly regressed from the input image the best grasp rectangle represented by the tuple $\{x, y, width, height, \theta\}$. The second architecture, *Regression + Classification* added object class prediction to test its regularization effect. Kumra *et al.* [19] further developed this one-shot detection approach by employing the more powerful




Fig. 2. **(Left)** Training example from the Dex-Net 2.0 detection dataset. Notice how there are very few annotations, thus not covering all of the possible grasp positions on the entire object. **(Middle)** Training example from the Cornell Grasping Dataset (CGD). These manually-labeled grasp annotations tend to cover a more important fraction of the object, but for a much more limited number of examples. Figure from [17]. **(Right)** Grasp path proposed by Chen *et al.* [21] to augment the grasp rectangle representation on the CGD. A grasp prediction (green) is projected to a grasp path that lies between two ground-truth annotation. This allows for better evaluation of detection approaches. Figure from [21].

ResNet-50 architecture [20]. They also explored a different strategy to integrate the depth modality, while seeking to preserve the benefits of ImageNet pre-training. As a solution, they introduced the *multi-modal grasp* architecture which separated RGB processing and depth processing in two different ResNet-50 networks, both pre-trained on ImageNet. Their architecture then performed late fusion, before the fully connected layers performed direct grasp regression.

2) *Multibox Approaches*: Redmon *et al.* [17] also proposed a third architecture, *MultiGrasp*, separating the image into a regular grid (dubbed Multi-box). At each grid cell, the network predicted the best grasping rectangle, as well as the probability of this grasp being positive. The grasp rectangle with the highest probability was then chosen. Trottier *et al.* [9] improved results by employing a custom ResNet architecture for feature extraction. Another advantage was the reduced need for pre-training on ImageNet. Chen *et al.* [21] remarked that grasp annotations in grasping datasets are not exhaustive. Consequently, they developed a method to transform a series of discrete grasp rectangles to a continuous *grasp path*. Instead of matching a prediction to the closest ground truth to compute the loss function, they mapped the prediction to the closest grasp path. This means that a prediction that falls directly between two annotated ground truths can still have a low loss value, thus (partially) circumventing the limitations of the Intersection-over-Union (IoU) metric when used with sparse annotation, as long as the training dataset is sufficiently densely labeled (see Figure 2). The authors re-used the *MultiGrasp* architecture from Redmon *et al.* for their experimentation.

3) *Anchor-box Approaches*: Zhou *et al.* [5] introduced the notion of oriented anchor-box, inspired by YOLO9000 [22]. This approach is similar to MultiGrasp (as the family of YOLO object detectors is a direct descendant of Multi-Grasp [17]) with the key difference of predicting offsets to predefined anchor boxes for each grid cell, instead of directly predicting the best grasp at each cell. Chu *et al.* [7] extends MultiGrasp to multiple object grasp detection by using region-of-interest pooling layers [23].

4) *Discrete Approaches*: Johns *et al.* [24] proposed to use a discretization of the space with a granularity of 1 cm and 30° . In a single pass of the network, the model predicts

a score at each grid location. Their method can explicitly account for gripper pose uncertainty. If a grasp configuration has a high score, but the neighboring configurations on the grid have a low score, it is probable that a gripper that has a Gaussian error on its position will fail to grasp at this location. The authors explicitly handled this problem by smoothing the 3D grid (two spatial axis, one rotation axis) by a Gaussian kernel corresponding to the gripper error.

Satish *et al.* [25] introduced a fully-convolutional successor to GQ-CNN. It extends GQ-CNN to a k -class classification where each output is the probability of a good grasp at the angle $180^\circ/k$, similar to [10]. They train their network for this classification task. They then transform the fully-connected layer into a convolutional layer, enabling classification at each location of the feature map. This effectively evaluates each discrete location x, y for graspability.

III. PROBLEM DESCRIPTION

A. One-shot Grasp Detection

Given the depth image of an object on a flat surface, we want to find a grasp configuration that maximizes the probability of lifting the object with a parallel-plate gripper. We aimed at performing this detection in a one-shot manner, i.e. with a single pass of the depth image through our network. As prediction output, we used the 5D grasp representation $\{x, y, z, \theta, w\}$, where x, y, z captures the 3D coordinates of the grasp, θ the angle of the gripper and w its opening. This representation considers grasps taken from above the object, perpendicular to the table’s surface, as in [12], [17]. As our network is trained using both the dataset and the grasp robustness classifier GQ-CNN of Dex-Net 2.0 [12], we detail them below.

B. Dex-Net 2.0 Dataset

Dex-Net 2.0 is a large-scale simulated dataset for parallel-gripper grasping. It contains 6.7 million grasps on pre-rendered depth images of 3D models. These 3D models come from two different sources. 1,371 models come from 3DNet [26], a synthetic model dataset built for classification and pose estimation. The other 129 additional models are laser scans from KIT [27]. All of the 3D models were resized to fit within a 5 cm parallel gripper.

The grasp labels in the Dex-net 2.0 dataset were acquired via random sampling of antipodal grasp candidates. A heuristic-based approach developed in previous work (Dex-Net 1.0[28]) was used to compute a robustness metric. This metric was thresholded to determine the grasp robustness label, i.e. *robust* vs. *non-robust*.

Learning one-shot grasp detection on the Dex-Net 2.0 dataset is in itself a challenging task, because of the few positive annotations per image. Annotations are very sparse compared to Cornell Grasping Dataset (CGD), a standard dataset used in one-shot grasp detection. For instance, it can be seen from Figure 2 that the ground truth annotation of Dex-Net 2.0 is clearly sparser than CGD. This prevents the grasp annotation augmentations method such as grasp path [21] from being employed on the former.

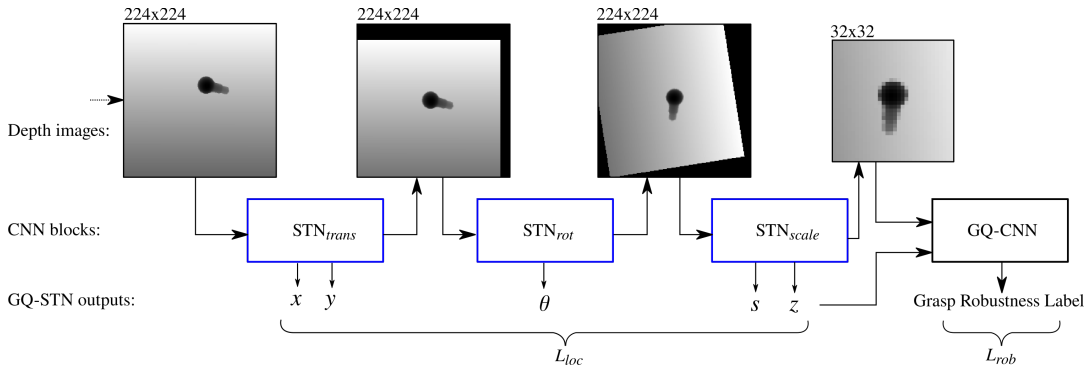


Fig. 3. Our complete one-shot STN-based architecture. The three STNs learn respectively translation to the grasp’s center, rotation to the grasp’s angle and scaling to the grasp’s opening. The intermediary outputs of the STNs are fully observable and are used to determine the grasp location. The last STN feeds into GQ-CNN, which predicts a grasp robustness label. A detailed view of a STN block is depicted in Fig. 4.

There are two available versions of the Dex-net 2.0 dataset. The first version is a classification dataset. It was used by Mahler *et al.* [12] to train GQ-CNN. It contains 32×32 depth images of grasp candidates with associated grasp robustness metrics, which are thresholded to obtain robustness labels. The authors also released a detection version of the dataset. This version contains the centered depth images of the object, at full resolution (400×400).

Please note that in this work, we used the original Dex-Net 2.0 annotations. Recently published work [25] developed a sampling method for generating additional annotations for the Dex-Net 2.0 images. Our approach could potentially benefit from more detection annotations on images contained in the Dex-Net 2.0 dataset. Still, for a given object, there is an infinity of possible grasp configurations which cannot all be annotated. Instead of improving learning at the annotation level, our approach, described in the following section, explicitly handles this inherent constraint.

IV. GQ-STN NETWORK ARCHITECTURE

In this paper, we propose Grasp Quality Spatial Transformer Network (GQ-STN), a neural network architecture for one-shot grasp detection based on the Spatial Transformer Network (STN). This architecture enables us to train directly on a robustness label outputted by GQ-CNN, unlike previous one-shot grasp detection methods that enforce robustness implicitly through geometric regression on annotated locations.

A. Spatial Transformer Network

The main component in our single-shot detection architecture is the Spatial Transformer Network (STN) [15]. In some sense, it acts as an attention mechanism, by narrowing/reorienting objects in a more canonical representation for the task at hand. It is a drop-in block that can be inserted between two feature maps of a Convolutional Neural Network (CNN) to learn a spatial transformation of the input feature map. The Spatial Transformer Network (STN) consists of three parts: a *localization network*, a *grid generator* and a *sampler*. The localization network learns a transformation matrix $\Lambda^{2 \times 3}$ based on the input feature map. The grid generator and the sampler transform the input feature map by the geometric transformation specified by Λ . It does so in a fully differentiable manner, in a process similar to texture

mapping. It can thus stretch, rotate, or skew the input feature map, resulting in a new feature map as output. A pure rotation transformation is illustrated in Figure 4.

A Spatial Transformer Network (STN) can be constrained to only represent specific geometric transformations, instead of freely learning the six elements of Λ . In our approach, we will employ three different transformations matrices:

$$\Lambda_{trans} = \begin{bmatrix} 1 & 0 & x \\ 0 & 1 & y \end{bmatrix}, \quad \Lambda_{scale} = \begin{bmatrix} s & 0 & 0 \\ 0 & s & 0 \end{bmatrix},$$

$$\Lambda_{rot} = \begin{bmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \end{bmatrix}.$$

Λ_{trans} represents a relative translation by a factor of $x, y \in [-0.5, 0.5]$, Λ_{rot} a rotation by an angle θ and Λ_{scale} an isotropic scaling by a factor of s .

B. Full architecture

Instead of predicting all transformations in a single network, we used a cascade of three STN blocks, STN_{trans} , STN_{rot} and STN_{scale} , which are respectively constrained by Λ_{trans} , Λ_{rot} and Λ_{scale} . In other words, STN_{trans} learns the translation x, y to the grasp center, STN_{rot} learns the rotation θ of the gripper and STN_{scale} learns a scaling s representing the opening of the gripper. A motivation behind this architecture is to isolate the regression of the angle θ , which is a challenging task for a one-shot network according to Park *et al.* [29]. All Spatial Transformer Networks (STN) were applied directly to the 1-channel depth map; contrary to Kumra *et al.* [19], we found no benefit in using a 3-channel version pre-trained on ImageNet for the STNs. All STNs also output a depth image, meaning that the communication between blocks of the network is not conducted via high-level feature maps, but via fully-observable depth images.

We used ResNet-34 as localization networks in all three Spatial Transformer Network (STN)s, as in [6]. This yielded slightly better results than the smaller ResNet-18 while maintaining a reasonable training time. Drawing from [17] and [22], the output layers of the ResNet-34 computed the elements of Λ_* as follows:

$$\begin{aligned} x &= \sigma(w_x) - 0.5 & \alpha &= \sigma(w_\alpha) & s &= \gamma e^{w_s} \\ y &= \sigma(w_y) - 0.5 & \beta &= \sigma(w_\beta) & z &= w_z \\ & & \theta &= \text{atan2}(\alpha, \beta)/2 & & \end{aligned}$$




Fig. 4. A Spatial Transformer Network (STN) block performing a rotation of θ on an input depth image, aligning the image to the grasp’s axis. A ResNet-34 localization network predicts the transformation matrix Λ_{rot} . This is the second of the three STNs shown in Fig. 3.

The tuples $\{w_x, w_y\}$, $\{w_\alpha, w_\beta\}$ and $\{w_s, w_z\}$ are the raw outputs of the localization networks of respectively STN_{trans} , STN_{rot} and STN_{scale} . To break the two-fold rotational symmetry of the angle prediction, we predict α, β which are respectively the sine and cosine of twice the angle θ , as in [17]. γ is the mean scaling factor in the training set. In conjunction with the scaling s , the last STN’s localization network also predicts the normalized gripper’s height z .

The input of the complete network, illustrated in Figure 3, is a 224×224 depth image. The translation and rotation STNs both generate a depth image of the same size as the input, while the STN_{scale} generates a depth image at a resolution of 32×32 . STN_{scale} is followed by GQ-CNN. The latter predicts a grasp robustness label given the 32×32 image outputted by STN_{scale} . We use pre-trained weights made available by Mahler *et al.* for GQ-CNN. These weights are frozen throughout training. At evaluation time, GQ-CNN is not required for grasp detection. However, because evaluating a single grasp on GQ-CNN is low-cost, we keep GQ-CNN to avoid a GPU memory transfer cost later if we need a robustness label associated with a detection.

Note that every block in the architecture is fully differentiable, thus allowing us to leverage information from the error on the grasp robustness label, by back-propagating from the latter all the way back to the first STN.

C. Training

At each step of training, we randomly select a ground truth positive grasp example from the Dex-net 2.0, thus obtaining target values for location Λ_{trans}^{gt} , Λ_{rot}^{gt} and Λ_{scale}^{gt} . We train the network using two types of supervision:

- **Localization loss** L_{loc} : the L_2 loss on the predictions of the localization networks of the STNs using Λ_*^{gt} ;
- **Robustness loss** L_{rob} : the cross-entropy loss on the output of GQ-CNN, where the expected value is a positive grasp label.

The total loss L_{tot} is given by:

$$L_{tot} = \xi L_{loc} + (1 - \xi) L_{rob}.$$

The training regimen begins with $\xi = 1$ and we gradually slide the loss mixing parameter toward $\xi = 0$. This way, we bootstrap the learning of our architecture with ground-truth grasp positions. These provide strong cues to the STNs, via the loss L_{loc} . As we reach $\xi = 0$, the network training then focuses on directly improving the grasp quality metric,

irrespective of grasp positions. Importantly, this allows our one-shot detection network to learn from sparsely labeled ground-truth, by eventually strictly focusing on a grasp robustness metric provided by GQ-CNN. The bootstrapping induced by $\xi > 0$ was necessary for the network training to converge, enabling a proper focus on the object. It can be seen in Figure 4 that transformations on the depth image introduce artifacts on the edges. If one would start training with $\xi = 0$, the network would enter a degenerate state where edge artifacts are mistaken for object edges.

During early stages of bootstrapping when $\xi > 0$, training tend to be quite unstable. There is an accumulation of error where, for instance, STN_{scale} cannot provide a good prediction because of errors made by STN_{trans} and STN_{rot} , resulting in a high L_{loc} . We solved this issue by using a *teacher forcing* approach [30] where the STNs are trained in a disjoint manner. Instead of using the Λ_{trans} and Λ_{rot} predicted by the first and second localization networks respectively, we directly transform the images using the ground truth information Λ_{trans}^{gt} , Λ_{rot}^{gt} . Teacher forcing allows the three STN to be trained simultaneously, instead of training them in sequence as proposed in [6], resulting in a shorter training time. Teacher forcing is disabled after $\xi = 0$, allowing a joint training of all parameters on L_{rob} .

V. EXPERIMENTS AND EVALUATION

We compared our architecture against three baselines: the single-shot DirectGrasp and MultiGrasp architectures[17] and the approach based on *Proposal+Classification* from Dex-Net 2.0[12] that we will refer to as Prop+GQ-CNN. For DirectGrasp and MultiGrasp, we replaced the AlexNet feature extractor by a ResNet feature extractor, as seen in Kumra *et al.* [19]. We trained our GQ-STN model and both DirectGrasp and MultiGrasp on 80% of the Dex-Net 2.0 dataset and held 20% in a test set. For the Prop+GQ-CNN approach, we used the pre-trained model made available by the authors.

We further tested GQ-STN, DirectGrasp and MultiGrasp on the Jacquard dataset[31]. Note that because Jacquard does not contain any gripper height information, we could not train the architectures on this dataset, as the gripper height is a required input of GQ-CNN. Therefore, Jacquard is only used here for testing networks that were trained on the Dex-Net 2.0 dataset.

We implemented all the architectures using the Tensorflow library. We trained all models 40 epochs with the Adam Optimizer. For GQ-STN, we had the following scheduling for ξ and the learning rate lr : 6 epochs at $\xi = 1.0, lr = 1 \times 10^{-3}$, 3 epochs at $\xi = 0.5, lr = 2 \times 10^{-4}$, 3 epochs at $\xi = 0.2, lr = 4 \times 10^{-5}$, 9 epochs at $\xi = 0.0, lr = 4 \times 10^{-5}$, and a fine-tuning stage of 19 epochs at $\xi = 0.0, lr = 8 \times 10^{-6}$ using early stopping. Teacher-forcing was turned on for only the first 12 epochs.

For DirectGrasp and MultiGrasp, we employed the same lr schedule, though they converged faster than GQ-STN and the last fine-tuning step with $lr = 8 \times 10^{-6}$ did not improve results. We kept the models that had the highest rectangle

are shown in Figure 6. These auxiliary results show that, especially in the context of sparse grasp annotations such as with the Dex-Net 2.0 dataset, the rectangle metric does not properly represent the performance of a grasping system. This further motivates our choice of evaluating with a robustness classification metric.

C. Metric Results

Table I shows that overall, on the Dex-Net 2.0 dataset, our approach is able to return a significantly higher percentage of high-quality grasps (92.4%) than the one-shot detection approach based on MultiGrasp (30.6%) and DirectGrasp (25.9%). This large performance gap can be explained by the fact that our approach enables us to optimize directly on the robustness classification metric, which is impossible for the two baselines. For all approaches, the rectangle metric tends to under-estimate the performance, which is explainable by sparse grasp annotations of the Dex-Net 2.0 dataset, as discussed in Section V-B.

We also tested the three models on the Jacquard dataset, which, contrary to Dex-Net 2.0, contains dense grasp rectangle annotations. As we can see in Table I, our GQ-STN returns significantly more robust grasp (60.4%) than the best baseline MultiGrasp (34.2%). This shows a good generalization of our method, which is also observed in the physical benchmark (Section V-D).

D. Physical benchmark

We evaluated all three methods in real-world conditions using the physical setup seen in Figure 5. It comprised a *Universal Robots UR5* arm, a *Robotiq 85* gripper and a *Microsoft Kinect* sensor. The *Kinect* sensor was mounted 70 cm perpendicular to the table’s surface. Grasp prediction was based on a single rectified depth image, where we replaced invalid depth pixels using inpainting [24].

We selected 12 household and office objects for testing, shown in Figure 5. We chose objects that have a good variety of shape, material and texture and are similar to the one

TABLE I

COMPARISON OF ONE-SHOT METHODS ON EVALUATION METRICS.

Test Dataset	Model	Precision (%)	
		Rectangle	Robust
Dex-Net 2.0	DirectGrasp	48.1	25.9
	MultiGrasp	48.4	30.6
	GQ-STN (ours)	46.7	92.4
Jacquard (trained on Dex-Net 2.0)	DirectGrasp	67.4	32.7
	MultiGrasp	71.8	34.2
	GQ-STN (ours)	70.8	60.4

TABLE II

COMPARISON OF METHODS ON OUR PHYSICAL BENCHMARK.

Model	Success rate (%)	Robust pred. rate (%)	Grasp detect. time (sec)
MultiGrasp	95	21.7	0.014
GQ-STN (ours)	96.7	61.7	0.024
Prop+GQ-CNN	98.3	48.3	1.5




Fig. 7. Examples of robust and non-robust grasp detection made by GQ-STN and Prop+GQ-CNN in our physical benchmark.

used in [12]. During testing, we placed the target object at a random position near the center of the table, by shaking it under a box to ensure random orientation, as in Mahler *et al.* [12]. We then estimated the grasp configuration with one of the three methods, and used a custom path planner to execute the grasp motion. The gripper default opening was 8.5 cm. It closed on the object until a maximum force feedback is reached. Upon closure, the object was lifted from the table and the success evaluated manually. Each of the 12 objects was tested 5 times, for each compared method. In total, we performed 180 grasp attempts.

We computed three metrics in this physical benchmark:

- 1) **Success rate:** Percentage of the lift attempts that resulted in a success. We execute the detected grasp even if it is not classified robust by the robustness classification metric.
- 2) **Robust prediction rate:** Percentage of the time the detected grasp (or the top grasp candidate for the sampling-based Prop+GQ-CNN) is robust according to the robustness classification metric.
- 3) **Grasp detection time:** Time in seconds between capturing an image and returning a grasp location. Here, we ignore time taken for inpainting.

As we can see in Tab. II, all three methods performed similarly, within the uncertainty of low samples. However, our method returned a robust grasp 61.7% of the time, which is significantly more than MultiGrasp and above Prop+GQ-CNN.

Qualitatively, the approach Prop+GQ-CNN seemed to perform slightly better during real experiments, especially with larger objects such as the red chips clip. In some sense, this is not surprising as it evaluated the grasp quality over 1000 positions. Figure 7 shows examples of grasp detection on our physical benchmark. Even though the methods were trained only on simulated data, its large amount helped generalization to real-world conditions, as noted as well by Mahler *et al.* [12]. Note that no domain-randomization was used here, contrary to Bousmalis *et al.* [13].

In terms of timing, our GQ-STN approach is in the same order of magnitude as the MultiGrasp approach, even though we run an image through three ResNet networks (one per Localization Network inside the STN). The detection time for Prop+GQ-CNN is two orders of magnitude larger than our approach, i.e. around 60 times slower. This limits its ability to perform real-time grasp detection.

Even though GQ-STN returns a single grasp and does so much faster, GQ-STN finds a robust grasp more often than Prop+GQ-CNN's sampling. Considering the high precision of the robustness classification metric, this enables GQ-STN to be used in a framework where we first evaluate the fast GQ-STN then fallback to a slow sampling method if we have not found a robust grasp, improving the overall average planning time.

VI. CONCLUSION

In this paper, we present a novel architecture for one-shot detection of grasp localization, based on the Spatial Transformer Network (STN) architecture. With it, we have demonstrated how one can use supervision from a robustness classifier to train one-shot grasp detection. On the Dex-Net 2.0 dataset, our method returns robust grasps more often than a baseline model that is only trained using the geometric supervision. We showed in a physical benchmark that our method can find robust grasps in real-world conditions more often than sampling methods, while still performing real-time (over 40 Hz), which is greater than frame rate grasp detection on a Kinect.

This speed opens up the possibility of carrying out visual servoing for grasping, for moving objects for instance. If a camera in-hand is used, it makes it possible to explore an object in real-time, similarly to a next-best-view approach, akin to Levine *et al.* [11]. There are other interesting research avenues at the network architecture level for future work. For instance, since all inputs of the Spatial Transformer Network (STN) are similar depth images, one could imagine a parameter sharing mechanism to speed up the training time and reduce the model size.

ACKNOWLEDGEMENTS

This work was financed by the Fonds de Recherche du Québec – Nature et technologies and the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

- [1] J.-P. Mercier, C. Mitash, P. Giguère, and A. Boularias, "Learning object localization and 6d pose estimation from simulation and weakly labeled real images," in *ICRA*, 2019.
- [2] N. Correll, K. E. Bekris, D. Berenson, O. Brock, A. Causo, K. Hauser, K. Okada, A. Rodriguez, J. M. Romano, and P. R. Wurman, "Analysis and observations from the first amazon picking challenge," *T-ASE*, vol. 15, no. 1, pp. 172–188, 2018.
- [3] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis—a survey," *T-RO*, vol. 30, no. 2, pp. 289–309, 2014.
- [4] L. Trottier, P. Giguère, and B. Chaib-draa, "Sparse dictionary learning for identifying grasp locations," in *WACV*, 2017, pp. 871–879.
- [5] X. Zhou, X. Lan, H. Zhang, Z. Tian, Y. Zhang, and N. Zheng, "Fully convolutional grasp detection network with oriented anchor box," in *IROS*, 2018, pp. 7223–7230.
- [6] D. Park and S. Y. Chun, "Classification based grasp detection using spatial transformer network," *CoRR*, 2018.
- [7] F. J. Chu, R. Xu, and P. A. Vela, "Real-world multiobject, multigrasp detection," *RA-L*, vol. 3, no. 4, pp. 3355–3362, Oct. 2018.
- [8] D. Chen, V. Dietrich, Z. Liu, and G. von Wichert, "A probabilistic framework for uncertainty-aware high-accuracy precision grasping of unknown objects," *J. of Intelligent & Robotic Systems*, Oct. 2017.
- [9] L. Trottier, P. Giguère, and B. Chaib-Draa, "Convolutional residual network for grasp localization," in *CRV*, 2017, pp. 168–175.
- [10] L. Pinto and A. Gupta, "Supersizing self-supervision: Learning to grasp from 50K tries and 700 robot hours," in *ICRA*, 2016.
- [11] S. Levine, P. Pastor, A. Krizhevsky, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *I. J. Robotics Res.*, vol. 37, pp. 421–436, 2018.
- [12] 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 *RSS*, 2017.
- [13] K. Bousmalis, A. Irpan, P. Wohlhart, Y. Bai, M. Kelcey, 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 *ICRA*, 2018, pp. 4243–4250.
- [14] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *I. J. Robotics Res.*, vol. 34, no. 4-5, pp. 705–724, 2015.
- [15] M. Jaderberg, K. Simonyan, A. Zisserman, and K. Kavukcuoglu, "Spatial Transformer Networks," in *NIPS*, 2015.
- [16] U. Asif, J. Tang, and S. Harter, "EnsembleNet: Improving grasp detection using an ensemble of convolutional neural networks," in *BMVC*, 2018.
- [17] J. Redmon and A. Angelova, "Real-time grasp detection using convolutional neural networks," in *ICRA*, 2015, pp. 1316–1322.
- [18] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," *NIPS*, 2012.
- [19] S. Kumra and C. Kanan, "Robotic grasp detection using deep convolutional neural networks," in *IROS*, 2017, pp. 769–776.
- [20] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *CVPR*, vol. 7, 2015, pp. 171–180.
- [21] L. Chen, P. Huang, and Z. Meng, "Convolutional multi-grasp detection using grasp path for RGBD images," *Robotics and Autonomous Systems*, vol. 113, Mar. 2019.
- [22] J. Redmon and A. Farhadi, "YOLO9000: Better, faster, stronger," in *CVPR*, 2017, pp. 6517–6525.
- [23] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *NIPS*, pp. 1–10, 2015.
- [24] E. Johns, S. Leutenegger, and A. J. Davison, "Deep learning a grasp function for grasping under gripper pose uncertainty," in *IROS*, 2016.
- [25] V. Satish, J. Mahler, and K. Goldberg, "On-policy dataset synthesis for learning robot grasping policies using fully convolutional deep networks," *RA-L*, vol. 4, no. 2, pp. 1357–1364, 2019.
- [26] W. Wohlkinger, A. Aldoma, R. B. Rusu, and M. Vincze, "3DNet: Large-scale object class recognition from cad models," in *ICRA*, 2012, pp. 5384–5391.
- [27] A. Kasper, Z. Xue, and R. Dillmann, "The KIT object models database: An object model database for object recognition, localization and manipulation in service robotics," *I. J. Robotics Res.*, vol. 31, no. 8, pp. 927–934, 2012.
- [28] J. Mahler, F. T. Pokorny, B. Hou, M. Roderick, M. Laskey, M. Aubry, K. Kohlhoff, T. Kroger, 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 *ICRA*, 2016.
- [29] D. Park, Y. Seo, and S. Y. Chun, "Rotation ensemble module for detecting rotation-invariant features," *CoRR*, 2018.
- [30] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [31] A. Depierre, E. Dellandréa, and L. Chen, "Jacquard: A large scale dataset for robotic grasp detection," in *IROS*, 2018, pp. 3511–3516.
- [32] Y. Jiang, S. Moseson, and A. Saxena, "Efficient grasping from RGBD images: Learning using a new rectangle representation," in *ICRA*, 2011, pp. 3304–3311.
- [33] G. Ghazaei, I. Laina, C. Rupprecht, F. Tombari, N. Navab, and K. Nazarpour, "Dealing with ambiguity in robotic grasping via multiple predictions," in *ACCV*, 2018.