

Dealing with Ambiguity in Robotic Grasping via Multiple Predictions

Ghazal Ghazaei^{1,2}, Iro Laina², Christian Rupprecht²,
Federico Tombari², Nassir Navab², and Kianoush Nazarpour^{1,3}

¹ School of Engineering, Newcastle University, Newcastle, UK

² Technische Universität München, Munich, Germany

³ Institute of Neuroscience, Newcastle University, Newcastle, UK

Abstract. Humans excel in grasping and manipulating objects because of their life-long experience and knowledge about the 3D shape and weight distribution of objects. However, the lack of such intuition in robots makes robotic grasping an exceptionally challenging task. There are often several equally viable options of grasping an object. However, this ambiguity is not modeled in conventional systems that estimate a single, optimal grasp position. We propose to tackle this problem by simultaneously estimating multiple grasp poses from a single RGB image of the target object. Further, we reformulate the problem of robotic grasping by replacing conventional grasp rectangles with grasp belief maps, which hold more precise location information than a rectangle and account for the uncertainty inherent to the task. We augment a fully convolutional neural network with a multiple hypothesis prediction model that predicts a set of grasp hypotheses in under 60 ms, which is critical for real-time robotic applications. The grasp detection accuracy reaches over 90% for unseen objects, outperforming the current state of the art on this task.

Keywords: Robotic Grasping · Deep Learning · Multiple Hypotheses.

1 Introduction

Grasping is a necessary skill for an autonomous agent to interact with the environment. The ability to grasp and manipulate objects is imperative for many applications in the field of personal robotics and advanced industrial manufacturing. However, even under simplified working conditions, robots cannot yet match human performance in grasping. While humans can reliably grasp and manipulate a variety of objects with complex shapes, in robotics this is still an unsolved problem. This is especially true when trying to grasp objects in different positions, orientations or objects that have not been encountered before. Robotic grasping is a highly challenging task and consists of several components that need to take place in real time: perception, planning and control.

In the field of robotic perception, a commonly studied problem is the detection of viable grasping locations. Visual recognition from sensors —such as RGB-D cameras— is required to perceive the environment and transfer candidate grasp points from the image domain to coordinates in the real world. The localization of reliable and effective grasping points on the object surface is a necessary first step for successful manipulation

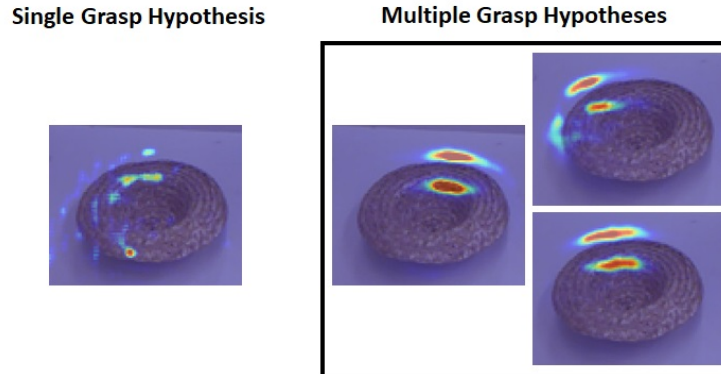


Fig. 1. We propose a model for regressing multiple grasp hypotheses as 2D belief maps, which tackles the ambiguity of grasp detection more effectively than a single grasp detection, in particular for completely unseen shapes, as the one depicted here.

through an end effector, such as a robotic hand or a gripper. The detected target position is then used such that an optimal trajectory can be planned and executed. This visual recognition task has gained great attention in recent years [1, 9, 12, 13, 15, 20, 21, 27, 29, 34, 35, 37, 38, 40] and led to the emergence of benchmark datasets, such as the Cornell grasp detection dataset [20], to evaluate the performance of approaches designed for this specific task.

Early approaches rely on explicitly estimating object geometry to localize grasping points [27, 40]. This tends to slow down the overall run-time and fails in presence of complicated or unseen object shapes. Following the success of deep learning in a wide spectrum of computer vision applications, several recent approaches [9, 15, 20, 24, 29, 37, 38] employed Convolutional Neural Networks (CNNs) [14, 18] to successfully detect grasping points from visual data, typically parametrized by 5-dimensional (5D) grasping representations [12, 20]. It is worth noting that most of these methods rely on depth data, often paired with color information. All these approaches have contributed significantly to improving robotic grasp detection, however they have not exhaustively studied generalization to novel, complex shapes. In particular, although some prior work explicitly aims at grasp estimation for unseen objects from RGB-D/depth data, this aspect is still regarded as an open issue [40]. In this work we propose a novel grasp detection approach from RGB data only. Our method incorporates two measures to explicitly model ambiguity related to the task of robotic grasping. First, we redefine the task of grasp detection as dense belief estimation problem. Thus, instead of the conventional grasp representation based on bounding boxes [12] we model the grasp space with 2D belief maps to be predicted from an input image. This allows the model to predict a grasp distribution with spatial uncertainty that accounts for small-scale ambiguities and exploits the full potential of CNNs in learning spatial representations.

The reformulation of this problem further highlights the inherent ambiguity in grasp detection. Most objects can be gripped in different ways and, although some may be preferable, there is not necessarily a “best” grip. This is also reflected in that current

benchmarks provide multiple grasp rectangles as ground truth for each object. However, aiming for a single output in an ambiguous problem can harm performance as the network typically learns the conditional average of all possible outcomes. To better model larger scale ambiguities, we employ a *multi-grasp prediction* framework and estimate multiple meaningful grasping positions for each input image. This approach allows to better model the output distribution and results in more precise and robust predictions especially in the case of unseen objects. The outcome of our method in comparison to a conventional single-prediction model is depicted in Figure 1. Finally, for the selection of a single grasping position, we propose an additional ranking stage based on Gaussian Mixture Models (GMMs) [25]. This is particularly useful for practical applications of our approach and for fair comparisons with the state of the art. We demonstrate the effectiveness of our approach by evaluating on a common benchmark [20].

2 Related Work

Robotic Grasp Detection Before the immense success of deep learning in computer vision applications, grasp estimation solutions were mostly based on analytic methods [3]. Some of these approaches, such as Graspit! [27], are dependent on the presence of a full 3D model to fit a grasp to it, not feasible for real-time applications. With the improvement of depth sensors, there are also recent methods that leverage geometrical information to find a stable grasp point using single-view point clouds [40].

In addition, the combination of both learning techniques and 3D shape information has led to interesting results. Varley *et al.* [35], use a deep learning based approach to estimate a 3D model of the target object from a single-view point cloud and suggest a grasp using 3D planning methods such as Graspit!. Mahler *et al.* [24] develop a quality measure to predict successful grasp probabilities from depth data using a CNN. Asif *et al.* [1] extract distinctive features from RGB-D point cloud data using hierarchical cascade forests for recognition and grasp detection.

The most recent robotic grasp estimation research is focused solely on deep learning techniques. Lenz *et al.* [20] pioneered the transfer of such techniques to robotic grasping using a two-step cascade system operating on RGB-D input images. A shallow network first predicts high-ranked candidate grasp rectangles, followed by a deeper network that chooses the optimal grasp points. Wang *et al.* [38] followed a similar approach using a multi-modal CNN. Another method [15] uses RGB-D data to first extract features from a scene using a ResNet-50 architecture [11] and then a successive shallower convolutional network applied to the merged features to estimate the optimal point of grasping.

Recent work in robotic grasp detection has also built on object detection methods [30, 31] to directly predict candidate grasp bounding boxes. Redmon *et al.* [29] employ YOLO [30] for multiple grasp detection from RGB-D images. This model produces an output grid for candidate predictions including the confidence of grasp being correct in each grid cell. This MuliGrasp approach improved the state-of-the-art accuracy of grasp detection up to $\sim 10\%$. However, the results are only reported for the best ranked rectangle and the performance of other suggested grasps is not known. Guo *et al.* [9] instead propose a hybrid deep network combining both visual and tactile sensing. The multi-modal data is fed into a visual object detection network [31] and a

tactile network during training and the features of both networks are concatenated as an intermediate layer to be employed in deep visual network during test.

Landmark Localization In our method we define the grasping problem differently. Instead of approaching the task as object detection, *i.e.* detecting grasping rectangles as for example in [9, 29], we express the rectangles as 2D belief maps around the grasping positions. This formulation is inspired by the latest methods in landmark localization, for example in human pose estimation [2, 4, 6, 28, 39], facial keypoint detection [5, 26] and articulated instrument localization [8, 16]. The use of heat maps to represent 2D joint locations has significantly advanced the state of the art in the localization problem. These models are trained so that the output matches the ground truth heat maps, for example through \mathcal{L}_2 regression, and the precise landmark locations can be then computed as the maxima of the estimated heat maps.

Multiple Hypothesis Learning To better model the grasp distribution of varying objects as well as grasp uncertainty, we augment the belief maps along the lines of multiple hypothesis learning [19, 33]. These methods model ambiguous prediction problems by producing multiple possible outcomes for the same input. However, they do not explore the possibility to select the best hypothesis out of the predicted set. The problem of selecting good hypotheses for multi-output methods has been typically addressed by training selection networks [10, 22]. Here, we solve this problem in a task-specific fashion, by scoring the predictions based on their alignment with a parametric Gaussian distribution which was used in training.

3 Methods

In the following, we describe our approach in detail. First, we redefine the problem of robotic grasp detection as prediction of 2D grasp point belief maps (Section 3.1). Specifically, we learn a mapping from a monocular RGB image to grasping confidence maps via CNN regression (Section 3.2). We then introduce our multi-grasp framework to tackle the inherent ambiguity of this problem by predicting multiple grasping possibilities simultaneously (Section 3.3). Finally, we rank all predicted grasps according to GMM likelihood in order to select the top ranked prediction (Section 3.4).

3.1 Grasp Belief Maps

The problem of robotic grasp detection can be formulated as that of predicting the size and pose of a rectangle, which, as suggested by [12], includes adequate information for performing the grasp; that is a 5-dimensional grasp configuration denoted by (x, y, θ, h, w) , where (x, y) is the center of the rectangle and θ is its orientation relative to the horizontal axis. We denote width and height of the bounding box with w and h respectively. These correspond to the length of a grip and the aperture size of the gripper. This representation has been frequently used in prior work [1, 9, 15, 20, 29, 38] as guidance for robotic grippers.




Fig. 2. An illustration of the adaptation of grasp rectangles to their associated grasp belief maps. The belief maps are constructed using the centers of the gripper plates as means for the normal distributions. The variance σ_x is proportional to the gripper height, while σ_y is a chosen constant.




Fig. 3. Samples of rectangle grasps and grasp belief maps shown for the same item.

In this work, we propose an alternative approach to model the detection of a robotic grasp using 2D belief maps. For an N -finger robotic gripper, these belief maps can be represented as a mixture model of N bivariate normal distributions fitted around the finger locations.

For a parallel gripper, the previously used grasping rectangle representation can be encoded in belief maps as follows. The means $\boldsymbol{\mu}^{(n)} = (\mu_x^{(n)}, \mu_y^{(n)})^T$, with $n \in \{1, 2\}$, around which the Gaussian distributions are centered correspond to the 2D centers (in Cartesian coordinates) of the gripper plates. The distance of the means $\|\boldsymbol{\mu}^{(1)} - \boldsymbol{\mu}^{(2)}\|_2 = w$ represents the width of the grasp. The Gaussian distributions are elliptical with $\Sigma = \text{diag}(\sigma_x^{(n)}, \sigma_y^{(n)})^2$. The primary axis of the ellipse represents the grasp height h . The orientation of the Gaussian kernels is adjusted by the rotation matrix $R(\theta)$ to make up for the correct grasping pose with respect to the object. The mixture model can be then defined as

$$G(\mathbf{p}) = \sum_{n=1}^N \frac{\exp\left(-\frac{1}{2}(\mathbf{p} - \boldsymbol{\mu}^{(n)})^T R(\theta) \Sigma^{-1} R(\theta)^T (\mathbf{p} - \boldsymbol{\mu}^{(n)})\right)}{2\pi N \sigma_x^{(n)} \sigma_y^{(n)}}, \quad (1)$$

where \mathbf{p} denotes a pixel's location inside the belief map. An illustration of our adapted grasp belief maps is shown in Figure 2.

Grasp belief maps encode the same information as the grasp rectangles, while expressing an encoding of the inherent spatial uncertainty around a grasp location. The proposed representation encourages the encoding of image structures, so that a rich image-dependent spatial model of grasp choices can be learned. Moreover, the ampli-

tude as well as variance of the predicted belief maps can act as a measure of confidence for the exact location and orientation of the grasp. In Figure 3, we show all possible grasp configurations for an item using both the traditional bounding box representation and our adapted continuous approach based on belief maps.

A model equipped with grasp belief maps can express its uncertainty spatially in the map, while direct regression of rectangles makes it harder to model spatial uncertainty. Further, such mixture models can be seamlessly extended to model grasp representations of other, more complex types of grippers, such as hand prostheses.

In practice, we create heat maps by constructing Gaussian kernels according to Equation 1, parametrized by the centers and dimensions of the gripper fingers. The centers of the gripper plates specify the means of the Gaussian kernels, σ_x is proportional to the gripper height and σ_y is a chosen constant value.

3.2 CNN Regression

For the regression of confidence maps, a common design choice among deep learning methods have been fully convolutional networks (FCNs) [23]. For our purpose, we use the fully convolutional residual architecture proposed in [17], which has shown competitive performance for dense prediction tasks, in particular depth estimation, in real time. The encoder is based on ResNet-50 [11], which embeds the input into a low dimensional latent representation. The decoder is built from custom residual up-convolutional blocks, which increase the spatial resolution up to half of the input resolution. The architecture is shown in Figure 4.

Given our problem definition, the network is trained to perform a mapping from a monocular RGB input to a single-channel heatmap comprised of the Gaussian mixture which represents the grasp belief. Since there are typically more than one viable grasp per object, choosing a single ground truth grasp becomes an ambiguous problem. When training in the single-grasp setup, we choose the most stable available grasp as ground truth, that is the one with the maximum grasping area. To this end, the objective function to be minimized is the Euclidean norm between the predicted belief map \tilde{G} and the chosen ground truth map:

$$\mathcal{L}(\tilde{G}, G) = \|\tilde{G} - G\|_2^2 \quad (2)$$

3.3 Multiple Grasp Predictions

Training the model with a single viable grasp is not optimal and could harm generalization, because the model gets penalized for predicting grasps which are potentially valid, but do not exactly match the ground truth. In other words, the samples that the model learns from do not cover the entire grasp distribution. Thus, in the case of known objects, the model would overfit to the single grasp possibility it has seen, while in the case of previously unseen objects the uncertainty which arises would prevent the model from producing a sharp and reliable belief map (Figure 1).

To overcome this limitation we propose a multi-grasp estimation setup, Instead of forcing the model to produce exactly one grasp, we allow the same model to produce multiple simultaneous outputs $\tilde{G} = \{\tilde{G}^{(m)}\}_m$, $m \in \{1, 2, \dots, M\}$. In practice, we




Fig. 4. The architecture of the fully convolutional residual network used in this paper. M refers to the number of grasp map predictions.

replicate the last layer M times. Our goal is to then train the model such that it approximates the entire distribution of viable grasps. This problem can be formulated as an oracle meta-loss \mathcal{M} that acts on top of the problem-specific objective function \mathcal{L} . By denoting the cost value of each grasp output as

$$\mathcal{L}_m = \mathcal{L}(\tilde{G}^{(m)}, G) \quad (3)$$

we can then define the meta-loss through the following minimum formulation:

$$\mathcal{M}(\tilde{G}, G) = (1 - \epsilon) \min_{m=1, \dots, M} \mathcal{L}_m + \frac{\epsilon}{M - 1} \sum_{m' \neq \arg \min_m \mathcal{L}_m} \mathcal{L}_{m'} \quad (4)$$

The proposed algorithm works as follows. At each training step, a grasp belief map is chosen randomly as the ground truth label among all available ground truth possibilities for the given input sample. In this way, the entire grasp distribution for each sample will be seen during training. Since the model cannot know which ground truth belief map will be chosen for a specific image, it will learn to disentangle the possibilities into the M grasping hypotheses. This is achieved by the loss \mathcal{M} in Equation 4. This objective is based on the hindsight loss, which only considers the output $\tilde{G}^{(m)}$ which is closest to the given ground truth G . Here we formulate it in a more intuitive way by using a soft approximation in which the oracle selects the best grasp with weight $1 - \epsilon$ and $\frac{\epsilon}{M - 1}$ for all the other predictions, where $\epsilon = 0.05$. This enables the output branches to be trained equally well, especially if they were not initially selected.

3.4 Grasp Option Ranking

Our previously described model predicts M grasp hypotheses. For this system to be used in practice, we need a method to assess the hypotheses quality and find which one should be selected. Therefore, it is desirable to find a way to rank all candidate grasps and pick one with a high probability of successful grasping. As we train the model to produce two multivariate normal distributions, one way to rank the predicted belief

maps is by fitting a two-component Gaussian mixture model (GMM) to each output map using finite mixture model estimation [25].

The main parameters of a Gaussian mixture model are the mixture component weights ϕ_k and the component means μ_k and variances/covariances σ_k with K being the number of components. The mathematical description of a GMM distribution over all the components is

$$p(\mathbf{x}) = \sum_{k=1}^K \phi_k \mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}_k, \sigma_k), \quad \sum_{k=1}^K \phi_k = 1 \quad (5)$$

where $\mathcal{N}(\mathbf{x} \mid \boldsymbol{\mu}, \sigma)$ represents a normal distribution with mean $\boldsymbol{\mu}$ and variance σ . Mixture models can be estimated via the expectation maximization (EM) algorithm [7], as finding the maximum likelihood analytically is intractable. EM iteratively finds a numerical solution to the maximum likelihood estimation of the GMM. The EM algorithm follows two main steps: (E) computes an expectation of component assignments for each given data point given the current parameters and (M) computes a maximum likelihood estimation and subsequently updates the model parameters. The model iterates over E and M steps until the error is less than a desired threshold.

We fit the same parametric model that was used to create the ground truth belief maps (Equation 1) and use the likelihood of the fit for each of the M predictions for ranking and choose the best fitted prediction as the system’s final output.

4 Experiments and Results

In this section, we evaluate our method experimentally on a public benchmark dataset and compare to the state of the art. Further, we investigate the influence of the number of grasp hypotheses M on the performance of the method.

4.1 Dataset

Cornell dataset We evaluate our approach on the Cornell grasp detection dataset [20], which consists of 885 RGB-D images from 240 graspable objects with a resolution of 640×480 pixels. The annotated ground truth includes several grasping possibilities per object represented by rectangles. The dataset is mainly suited for 2D grippers with parallel plates, but as the grasp size and location are included in the representation, it has the potential to be used also for other types of grippers as it is used in [9] for a 3-finger gripper. There are 2 to 25 grasp options per object of different scales, orientations and locations, however, these annotated labels are not exhaustive and do not contain every possible grasp. Figure 5 shows some cropped samples of the dataset as used in this work. Here we only use the RGB images and disregard the depth maps.

Data splits We follow a cross-validation setup as in previous work [1, 9, 15, 20, 29, 38], using image-wise and object-wise data splits. The former split involves training with all objects, while some views remain unseen to evaluate the intra-object generalization capability of the methods. However, even an over-fitted model could perform well on this



Fig. 5. A representation of a subset of the objects of the Cornell grasp detection dataset [20].

split. The object-wise split involves training on all available views of the same object and testing on new objects and thus is suitable for evaluating inter-object performance. However, the unseen objects are rather similar to ones used in training.

It is worth noting that none of the previous methods studied the potential of generalizing to truly novel *shapes*, as the dataset includes a variety of similar objects. For example, there are several objects with different colors but of the same shape. Therefore, the object-wise split may not be a good measure for generalization to novel shapes. To investigate our framework’s performance on unseen shapes, we have created an additional *shape-wise* split, to encourage larger variation in objects between the train and test sets. We pick the train and test folds such that all the objects of similar shapes, *e.g.* various kinds of hats, are included in one of the test/train folds only and therefore novel when testing. Both image-wise and object-wise splits are validated in five folds. We perform two-fold cross validation for the shape-wise split, where we use the first 20% of objects for testing and the remainder for training. The second fold uses the same split but with reversed order of objects.

4.2 Implementation details

In all our experiments we pre-process the images and annotations as detailed in the following. As the images contain a large margin of background around the objects, we crop them and their corresponding grasp maps to 350×350 pixels and then bilinearly down-sample the image to 256×256 and the grasp map to 128×128 . Prior to cropping we apply data augmentation techniques. We sample a random rotation in $[-60^\circ, 60^\circ]$, center crops with a translation offset of $[-20, 20]$ pixels and scaling between 0.9 and 1.1. Each image is augmented six times. Thus, the final dataset contains 5310 images after augmentations. All the images and labels are normalized to a range of $[0, 255]$.

To train the single grasp prediction model, we choose the largest ground truth grasp rectangle as label since area is a good indicator for probability and stability of the grasp.

This selection may be trivial, but training a single grasp model is not feasible without pre-selection of a fixed ground truth among the available grasps.

On the other hand, our multiple grasp prediction model can deal with a *variable* number of ground truth grasp maps per image. At each training step, we randomly sample one of the available ground truth annotations. We also add hypothesis dropout with rate 0.05 as regularization [33]. We investigate and report the performance of our framework for different numbers M of grasp hypotheses. To rank multiple predicted grasps, we performed EM steps for up to 1000 iterations and calculated the negative log-likelihood for the parameters σ_k and μ_k .

Training was performed on an NVIDIA Titan Xp GPU using MatConvNet [36]. The learning rate was set to 0.0005 in all experiments. For regularization we set weight decay to 0.0005 and add a dropout layer with rate equal to 0.5. The models were trained using stochastic gradient descent with momentum of 0.9 for 50 epochs and a batch size of 5 and 20 for training multiple and single prediction models respectively.

4.3 Grasp Detection Metric

We report quantitative performance using the rectangle metric suggested by [12] for a fair comparison. A grasp is counted as a valid one only when it fulfills two conditions:

- The intersection over union (IoU) score between the ground truth bounding box (B) and the predicted bounding box (B^*) is above 25%, where

$$\text{IoU} = \frac{B \cap B^*}{B \cup B^*} \quad (6)$$

- The grasp orientation of the predicted grasp rectangle is within 30° of that of the ground truth rectangle.

This metric requires a grasp rectangle representation, while our network predicts grasp belief maps. We therefore calculate the modes μ_1 and μ_2 as the centers of each elliptical Gaussian for every predicted belief map. The Euclidean distance between these modes should be equal to the grasp rectangle’s width w (Figure 2). We compute the height h of the grasp rectangle as the major axis of the ellipse (after binarization of the belief map with a threshold of 0.2). We determine the gripper’s orientation θ by calculating the angle of the major axis as $\arctan \frac{d_1}{d_2}$; where d_1 and d_2 are the vertical and horizontal distance between the centers of elliptical Gaussian maps respectively. We can then convert the belief maps to a grasping rectangle representation. Under high uncertainty, *i.e.* when a grasp map is considerably noisy, we discard the hypothesis as a rectangle cannot be extracted. We note that a valid grasp meets the aforementioned conditions with respect to *any* of the ground truth rectangles and compute the percentage of valid grasps as the *Grasp Estimation Accuracy*.

4.4 Evaluation and Comparisons

In the following, we compare our multiple grasp prediction method with the single-grasp baseline and state-of-the-art methods. As there are several ground truth annotations per object, we compare the selected prediction to all the ground truth grasp rectangles to find the closest match. Among the predictions there can be some which are

Table 1. Comparison of the proposed method with the state of the art. *multiple* refers to our multiple prediction models, while *multiple / reg* are the models trained with diversity regularization.

Method	Input	Grasp Estimation Accuracy (%)		
		Image-wise	Object-wise	Shape-wise
Lenz et al. [20]	RGB-D	73.9	75.6	-
Wang et al. [38]	RGB-D	85.3	-	-
Redmon et al. [29]	RGB-D	88.0	87.1	-
Asif et al. [1]	RGB-D	88.2	87.5	-
Kumra et al. [15]	RGB-D	89.2	89.0	-
Guo et al. [9]	RGB-D, tactile	93.2	89.1	-
Kumra et al. [15]	RGB	88.8	87.7	-
<i>single</i> $M = 1$	RGB	83.3	81.0	73.7
<i>multiple</i> $M = 5$	RGB	91.1	90.6	85.3
<i>multiple / reg</i> $M = 5$	RGB	89.1	89.2	82.5
<i>multiple</i> $M = 10$	RGB	91.5	90.1	86.2

not viable, while others are perfect matches. The selected prediction for each image is one with the maximum GMM likelihood.

Quantitative results We report the results in Table 1, where M indicates the number of hypotheses and consequently $M = 1$ refers to the regression of single belief map and can be seen as a baseline in the following experiments. The proposed model with $M = 5$ predicted grasps shows significant improvement in performance over the single-grasp model (the average number of grasps per object in the dataset is also approximately five). This performance boost reveals the potential of modeling ambiguity in robotic grasp detection. To study the effect of the number of grasping options, we also evaluated our approach with $M = 10$.

While it only relies on RGB data as input, our multiple grasp approach outperforms all state-of-the-art methods that use additional depth information, except for Guo *et al.* [9] who also leverage tactile data. Moreover, both single and multiple grasp models have a faster grasp prediction run-time than the state of the art at 56 ms. GMM maximum likelihood estimation for hypothesis selection increases the run-time to 95 ms. Increasing the number of outputs M does not have a negative effect on speed.

It is worth noting that the comparable performance of the models in the image- and object-wise splits (also in prior work) suggests that task difficulty does not change much between the two scenarios. With the more challenging shape-wise scenario that we have proposed, we can better evaluate performance on novel objects. In this case, the accuracy of the single grasp baseline drops significantly. On the contrary, the multiple grasp model is still able to handle the increased difficulty with a large performance boost over the baseline. It can be observed that with an increasing number of grasp hypotheses the performance gap of the multiple-grasp over the single-grasp model is the highest for the shape-wise split, with over 12% increase in accuracy for unseen shapes/objects.

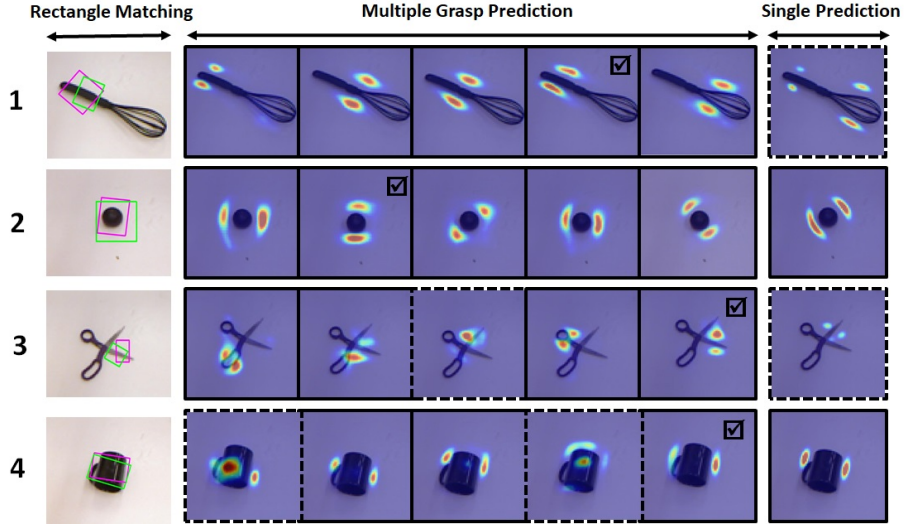


Fig. 6. Five and single grasp map predictions of sample objects in the dataset. A solid frame around an image is an indicator of grasp detection success, while a dashed line shows an incorrect detection. The predictions marked with \checkmark are the top-ranked ones according to the GMM likelihood estimation. These predictions are converted back to grasp rectangles (magenta) and compared with the closest ground truth grasp (green).




Fig. 7. Examples of diversity within predicted grasp maps (converted to rectangles).

Diversity of predictions We also examine the diversity of the predicted hypotheses for each image. We have performed experiments adding a repelling regularizer [32] (weighted by a factor of 0.1) to the hindsight loss to further encourage diverse predictions. The accuracy of this model with $M = 5$ (Table 1) is slightly worse than our multiple prediction model without the regularizer. As a measure of hypothesis similarity, we calculate the average cosine distance among all predictions given an input. The average similarity for the object-wise split decreases only marginally from 0.435 (without regularizer) to 0.427 (with regularizer), suggesting that the multiple prediction framework does not really benefit by explicitly optimizing diversity. Our framework can naturally produce diverse predictions, which we intuitively attribute to the hypothesis dropout regularization used during training.

Qualitative results In Figure 6 we show qualitative examples from our multi-grasp framework (with $M = 5$) and a comparison to the single grasp ($M = 1$) model’s

Table 2. Average grasp estimation accuracy of all hypotheses (lower limit) and average grasp success (upper limit).

Method	Image-wise	Object-wise	Shape-wise
lower limit ($M = 5$)	80.0	77.4	75.0
lower limit ($M = 10$)	76.5	73.3	72.1
upper limit ($M = 5$)	98.0	98.5	96.3
upper limit ($M = 10$)	99.2	98.4	99.1

predictions, noting the advantage of multiple grasp predictions both in terms of accuracy and variability. We observe that for objects that have several distinct grasping options, our multiple prediction framework models the output distribution sufficiently. Object 3 (scissors) is undoubtedly a challenging object with many different grasping poses, which are successfully estimated via multiple predictions. In Figure 7 we further emphasize the diversity among the grasp hypotheses, by showing multiple extracted rectangles for various objects.

4.5 Evaluation over Multiple Grasps

In Table 2 we report the lower and upper detection accuracy limits of the multi-grasp models. Instead of evaluating only the top-ranked grasp hypothesis, we first evaluate *all* predictions provided by our model. This evaluation gives the lower limit of the model’s performance, as it computes the success rate of all hypotheses, including even those with a low probability of being chosen. This result suggests that the estimated belief maps correspond, in most cases, to valid grasps (75% overall accuracy compared to 85.3% for one chosen grasp in shape-wise split, when $M = 5$). This lower bound decreases as M increases, *i.e.* it is more likely to have a (noisy) prediction that does not match any of ground truth grasp rectangles with higher M . However, thresholding the “good” matches based on GMM likelihood can counteract this drop in performance while leaving multiple grasping choices to the robot.

Another observation is that the top-ranked prediction is not necessarily the best one in terms of grasping performance. This can be seen in the upper limit evaluation, in which if there exists at least one matching grasp detection among all hypotheses, it counts overall as successful. For $M = 10$ the upper limit exceeds 98% accuracy for the object-wise split. This implies that there is in almost all cases at least one valid prediction returned by our model, although GMM fitting might not always result in correct ranking. Still, the top-ranked prediction performance in Table 1 is closer to the upper rather than the lower limit.

4.6 Generalization

Finally, to evaluate the performance of the proposed model in a real-world scenario, we test it on several common household objects, such as cutlery, keys and dolls, in an




Fig. 8. The top ranked grasp map selected by the GMM likelihood estimation module for a $M = 5$ model evaluated on common household objects in real-time. Objects 1-5 have similar shapes to the objects in the Cornell grasp dataset. Objects 6-12, however, represent novel shapes and textures compared to the dataset used for training. Despite variations from the training distribution, our method produces reasonable grasp maps for all tested objects.

own setup —and not test images from the same dataset. The differences to the Cornell dataset are not only in the type of objects used, but also in the camera views and illumination conditions. Through this setup we evaluate the generalization ability of the model under different conditions and challenging novel shapes and textures. Figure 8 illustrates the evaluated objects and the estimated grasp that is chosen as the maximum GMM likelihood. Our model is robust against these variations and results in viable and confident grasping options for all tested objects.

5 Conclusion

We have developed an efficient framework for robotic grasp detection. The representation of a grasp is redefined from an oriented rectangle to a 2D Gaussian mixture belief map that can be interpreted as the confidence of a potential grasp position. This allows us to handle the ambiguity stemming from the many possible ways to grasp an object. We employ a fully convolutional network for belief map regression and estimate a variety of viable grasp options per object. This approach embraces the ambiguous nature of the grasping task and provides a better approximation of the grasp distribution. This property manifests itself in the majority of the predicted grasps being viable solutions and the improvement over the single-grasp baseline becoming larger when tackling scenarios with increased difficulty, such as novel objects, shapes and textures. Our ranking approach selects the grasp positions with the highest likelihood, which result in real-time, state-of-the-art performance. Considering the fact that our belief map formulation also contains a measure of size, an interesting future direction could be the application of this method to prosthetic hands.

Acknowledgments This work is supported by UK Engineering and Physical Sciences Research Council (EP/R004242/1). We also gratefully acknowledge the support of NVIDIA Corporation with the donation of a Titan Xp GPU used for the experiments.

References

1. Asif, U., Bennamoun, M., Sohel, F.A.: RGB-D object recognition and grasp detection using hierarchical cascaded forests. *IEEE Transactions on Robotics* **33**(3), 547–564 (2017)
2. Belagiannis, V., Zisserman, A.: Recurrent human pose estimation. In: *International Conference on Automatic Face & Gesture Recognition (FG 2017)* (2017)
3. Bicchi, A., Kumar, V.: Robotic grasping and contact: A review. In: *Proceedings of 2000 IEEE International Conference on Robotics and Automation (ICRA)*. vol. 1, pp. 348–353. IEEE (2000)
4. Bulat, A., Tzimiropoulos, G.: Human pose estimation via convolutional part heatmap regression. In: *European Conference on Computer Vision (ECCV)*. pp. 717–732. Springer (2016)
5. Bulat, A., Tzimiropoulos, G.: Super-fan: Integrated facial landmark localization and super-resolution of real-world low resolution faces in arbitrary poses with gans. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018)
6. Cao, Z., Simon, T., Wei, S.E., Sheikh, Y.: Realtime multi-person 2d pose estimation using part affinity fields. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
7. Dempster, A.P., Laird, N.M., Rubin, D.B.: Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)* (1977)
8. Du, X., Kurmann, T., Chang, P.L., Allan, M., Ourselin, S., Sznitman, R., Kelly, J.D., Stoyanov, D.: Articulated multi-instrument 2d pose estimation using fully convolutional networks. *IEEE Transactions on Medical Imaging* (2018)
9. Guo, D., Sun, F., Liu, H., Kong, T., Fang, B., Xi, N.: A hybrid deep architecture for robotic grasp detection. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE (2017)
10. Guzman-Rivera, A., Kohli, P., Glocker, B., Shotton, J., Sharp, T., Fitzgibbon, A., Izadi, S.: Multi-output learning for camera relocalization. In: *Conference on Computer Vision and Pattern Recognition* (2014)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
12. Jiang, Y., Moseson, S., Saxena, A.: Efficient grasping from RGB-D images: Learning using a new rectangle representation. In: *International Conference on Robotics and Automation (ICRA)*. IEEE (2011)
13. Kehoe, B., Patil, S., Abbeel, P., Goldberg, K.: A survey of research on cloud robotics and automation.
14. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: *Advances in Neural Information Processing Systems* (2012)
15. Kumra, S., Kanan, C.: Robotic grasp detection using deep convolutional neural networks. *arXiv preprint arXiv:1611.08036* (2016)
16. Laina, I., Rieke, N., Rupprecht, C., Vizcaíno, J.P., Eslami, A., Tombari, F., Navab, N.: Concurrent segmentation and localization for tracking of surgical instruments. In: *International conference on medical image computing and computer-assisted intervention*. pp. 664–672. Springer (2017)
17. Laina, I., Rupprecht, C., Belagiannis, V., Tombari, F., Navab, N.: Deeper depth prediction with fully convolutional residual networks. In: *2016 Fourth International Conference on 3D Vision (3DV)*. IEEE (2016)
18. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. *Proceedings of the IEEE* (1998)

19. Lee, S., Prakash, S.P.S., Cogswell, M., Ranjan, V., Crandall, D., Batra, D.: Stochastic multiple choice learning for training diverse deep ensembles. In: *Advances in Neural Information Processing Systems* (2016)
20. Lenz, I., Lee, H., Saxena, A.: Deep learning for detecting robotic grasps. *The International Journal of Robotics Research* (2015)
21. Levine, S., Pastor, P., Krizhevsky, A., Quillen, D.: Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *arXiv preprint arXiv:1603.02199* (2016)
22. Li, Z., Chen, Q., Koltun, V.: Interactive image segmentation with latent diversity. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
23. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (2015)
24. 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)
25. McLachlan, G., Peel, D.: *Finite mixture models*. John Wiley & Sons (2004)
26. Merget, D., Rock, M., Rigoll, G.: Robust facial landmark detection via a fully-convolutional local-global context network. In: *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018)
27. Miller, A.T., Allen, P.K.: Graspit! a versatile simulator for robotic grasping. *IEEE Robotics & Automation Magazine* (2004)
28. Papandreou, G., Zhu, T., Kanazawa, N., Toshev, A., Tompson, J., Bregler, C., Murphy, K.: Towards accurate multi-person pose estimation in the wild. *Conference on Computer Vision and Pattern Recognition (CVPR)* (2017)
29. Redmon, J., Angelova, A.: Real-time grasp detection using convolutional neural networks. In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE (2015)
30. Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
31. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: *Advances in Neural Information Processing Systems* (2015)
32. Roohan, M., Ye, L., Wang, Y.: Video summarization using fully convolutional sequence networks. *arXiv preprint arXiv:1805.10538* (2018)
33. Ruppel, C., Laina, I., DiPietro, R., Baust, M., Tombari, F., Navab, N., Hager, G.D.: Learning in an uncertain world: Representing ambiguity through multiple hypotheses. In: *International Conference on Computer Vision (ICCV)* (2017)
34. Saxena, A., Driemeyer, J., Ng, A.Y.: Robotic grasping of novel objects using vision. *The International Journal of Robotics Research* (2008)
35. Varley, J., DeChant, C., Richardson, A., Nair, A., Ruales, J., Allen, P.: Shape completion enabled robotic grasping. *arXiv preprint arXiv:1609.08546* (2016)
36. Vedaldi, A., Lenc, K.: Matconvnet – convolutional neural networks for matlab. In: *Proceeding of the ACM International Conference on Multimedia* (2015)
37. Viereck, U., Pas, A., Saenko, K., Platt, R.: Learning a visuomotor controller for real world robotic grasping using simulated depth images. In: *Conference on Robot Learning* (2017)
38. Wang, Z., Li, Z., Wang, B., Liu, H.: Robot grasp detection using multimodal deep convolutional neural networks. *Advances in Mechanical Engineering* (2016)
39. Wei, S.E., Ramakrishna, V., Kanade, T., Sheikh, Y.: Convolutional pose machines. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2016)
40. Zapata-Impata, B.S.: Using geometry to detect grasping points on 3D unknown point cloud. In: *International Conference on Informatics in Control, Automation and Robotics* (2017)