

# Robotic Grasp Detection By Learning Representation in a Vector Quantized Manifold

Mridul Mahajan, Tryambak Bhattacharjee, Arya Krishnan, Priya Shukla, G C Nandi

*Center of Intelligent Robotics*

*Indian Institute of Information Technology, Allahabad*

Prayagraj, India - 211015

mridulmahajan44@gmail.com, priyashuklalko@gmail.com, gcnandi@iiita.ac.in

**Abstract**—For a robot to perform complex manipulation tasks, it is necessary for it to have a good grasping ability. However, vision based robotic grasp detection is hindered by the unavailability of sufficient labelled data. Furthermore, the application of semi-supervised learning techniques to grasp detection is under-explored. In this paper, a semi-supervised learning based grasp detection approach has been presented, which models a discrete latent space using a Vector Quantized Variational AutoEncoder (VQ-VAE). To the best of our knowledge, this is the first time a Variational AutoEncoder (VAE) has been applied in the domain of robotic grasp detection. The VAE helps the model in generalizing beyond the Cornell Grasping Dataset (CGD) despite having a limited amount of labelled data by also utilizing the unlabelled data. This claim has been validated by testing the model on images, which are not available in the CGD. Along with this, we augment the Generative Grasping Convolutional Neural Network (GGCNN) architecture with the decoder structure used in the VQ-VAE model with the intuition that it should help to regress in the vector-quantized latent space. Subsequently, the model performs significantly better than the existing approaches which do not make use of unlabelled images to improve the grasp.

**Index Terms**—Grasp Rectangle, Vector Quantized Variational AutoEncoder (VQ-VAE), Generative Grasp Convolutional Neural Network (GGCNN)

## I. INTRODUCTION

The advancements in the field of automation has led to an explosive expansion in the use of intelligent machines in various applications. But even with such advancements, robots have not yet become a general-purpose utility as a whole. The reason being the ever-changing environment which calls for the need of tremendous adaptive ability and a near-perfect sense of objectivity. A process such as grasping an object, which may easily come to human beings, is a rather complex process when applied to machines. That being said, the ultimate problem boils down to an intact sense of object detection and grasping.

Earlier, tasks such as this were done using analytical approaches, which involved hard-coding the instructions involving the robot's parameters and its world coordinates. These algorithms, called control algorithms, involved defining control over the robot's joints [1] and were designed using the knowledge of human experts. These manually planned approaches achieve efficiency but are restricted by the pro-

grammer's predictions of the robot's environment and also by dynamic environments [2]. The more dynamic the actuator of the robot is intended to be, the more impossible the task becomes of physically planning it. Hence, manual teaching is efficient but exhaustive [3]. Recently however, deep learning has remarkably advanced computer vision in fields such as classification, localisation and detection. It has also been observed that the application of computer vision to the problem of object grasping is analogous to object detection. [4], [5]. Hence, in most previous studies, grasp detection has been presented as a computer vision problem. Owing to an abundance of unlabelled data and an unavailability of sufficient labelled data, we want a model to utilize the unlabelled data too, hence avoiding the expense of a large labelled dataset. Also, the neural network should have the ability to generalize, so that it understands the data and the semantics behind it instead of learning only a mapping from the input space to the output space.

The major contributions of our work are listed as:

- We present a semi-supervised learning based model for vision based robotic grasp detection. The model uses VQ-VAE for regressing in a vector quantized latent space. We show that by utilizing the unlabelled data, the model gives good results despite being trained on limited amount of labelled data. Along with this, we validate the generalization ability of the model as well.
- We augment the robotic grasp detection model's architecture with the decoder structure for better utilization of the vector-quantized latents.

Rest of the paper is organized as follows: Section II analyses previous research in vision based robotic grasp detection along with highlighting the shortcomings in existing approaches. Section III gives a primer on VAE, VQ-VAE and GGCNN, which are important components in our proposed approach, and Section IV presents the results with their analyses. Section V concludes the paper along with highlighting the future research directions.

## II. ANALYSIS OF PREVIOUS RESEARCH

Grasp prediction techniques can be primarily classified into two types: Analytical and Empirical. The analytical techniques involve complex models of geometry, dynamics and kinematics to determine the grasps. [6] reviews such approaches in

detail. However, these approaches are not always preferred due to the underlying complexity as well as the difficulty in modelling them in the real world. On the other hand, empirical techniques involve estimation models and experience-based heuristic approaches, which are discussed below.

Substantial research with the use of deep CNN has been investigated in order to predict grasps on objects. From a broad perspective, grasp detection techniques in literature using deep learning can be divided into two main categories [7]:

- 1) Approaches that involve designing an application-specific model [3], [5], [8].
- 2) Techniques which use a pre-existing model and apply it to grasp detection through transfer learning approaches [1], [9], [10].

Most of them includes a two-stage pipeline [5], [11], [12]: firstly, several grasp candidates are sampled from the image, which are then fed as inputs to a CNN network to figure out the best among the sampled candidates. This leads to a substantial execution time causing the grasps to be executed in open loop, meaning once a grasp has been determined, the robot executes it in a fixed way without taking any feedback from the environment, such as any possible changes in the location or orientation of the object after grasp determination.

[5] uses a two-stage cascaded approach, along with the sliding window concept, to sample a number of grasp candidates. Unlike the two-staged approach, [9] and [1] use adapted versions of AlexNet [13] and ResNet [14] respectively, and are popular object detection models to make grasp predictions. Both of the previously proposed models use a single deep CNN network to regress the final grasp rectangle directly. Directly regressing a *single* grasp rectangle for the entire scene might average all possible grasp rectangles that exist for an object, which itself might not be a good grasp candidate. This is because of the fact that simple averaging does not consider the spatial features of the object to be grasped.

However, [3] possibly addresses both problems of averaging and execution time. The model, named Generative Grasp CNN or simply GGCNN, predicts grasps along with their quality for every single pixel of the image. Unlike [8], GGCNN does not impose a constraint on the objects' shape and is fast enough to be used in real-time applications.

To the best of our knowledge, all the existing models do not consider the fact that the labelled training data for grasp detection is limited. Due to unavailability of sufficient labelled data to train a neural network for vision based grasp detection, the need to explore semi-supervised learning domain becomes evident. Also, the neural network should have the ability to generalize. Hence, we explore these ideas in detail in the next sections.

### III. METHODOLOGY

#### A. Preliminaries

1) *Variational Autoencoder (VAE) [15]*: Explicit modelling of a distribution over the unlabelled training data can uncover latent representations which may help the supervised learning




Fig. 1. The VQ-VAE model. [16]

task of grasp detection. Variational Autoencoders (VAEs) can be used for approximate density estimation of the training data. Essentially, we first define a density function for the training data over the latent variables  $\mathbf{z}$ , which is intractable to be computed for every  $\mathbf{z}$ .

$$\mathbf{p}_\theta(\mathbf{x}) = \int_{\mathbf{z}} \mathbf{p}_\theta(\mathbf{z}) \mathbf{p}_\theta(\mathbf{x}|\mathbf{z}) d\mathbf{z} \quad (1)$$

Assuming that the training data has been generated from the unobserved latent  $\mathbf{z}$ , for the generation task, we sample  $\mathbf{z}$  from the true prior  $\mathbf{p}(\mathbf{z})$  and then sample  $\mathbf{x}$  from the true conditional distribution over the latent  $\mathbf{z}$ ,  $\mathbf{p}_\theta(\mathbf{x}|\mathbf{z})$ . A Gaussian prior is generally used for its simplicity.  $\mathbf{p}_\theta(\mathbf{x}|\mathbf{z})$  is modelled by a neural network whose optimal parameters are obtained by maximizing the likelihood of the training data. However, this data likelihood is intractable as has been discussed before. As a consequence, the posterior distribution becomes intractable as well. Hence, we define  $\mathbf{q}_\phi(\mathbf{z}|\mathbf{x})$  as an approximation to the true posterior. This helps in deriving the lower bound for the data likelihood.

The encoder network here is used for inference, and the decoder network is used for generation. Both of these networks are probabilistic, and therefore, output the mean and the diagonal covariance of the density function they model. The Evidence Lower Bound (ELBO) is a tractable lower bound of the likelihood of the training data whose gradient can be computed by the reparameterization trick to optimize it.

$$ELBO = \mathbf{E}_{\mathbf{z}} [\log \mathbf{p}_\theta(\mathbf{x}|\mathbf{z})] - \mathbf{D}_{\text{KL}}(\mathbf{q}_\phi(\mathbf{z}|\mathbf{x}) || \mathbf{p}_\theta(\mathbf{z})) \quad (2)$$

The first term increases the likelihood of the training data being generated. The second term brings the approximate posterior close to the true prior, which has a closed form solution in the case of a vanilla VAE wherein both these distributions are gaussian.

Therefore, the key highlight of VAE is that the inference network (the encoder) allows the inference of  $\mathbf{q}_\phi(\mathbf{z}|\mathbf{x})$  which can be used for representation learning.

2) *Vector Quantized Variational Autoencoder [16]*: [17] shows that images are modelled better using discrete symbols. Most of the VAE models having a powerful decoder ignore the latent vectors. This situation is known as the posterior collapse. Keeping this in mind, [16] introduced VQ-VAE, wherein the latents are discrete instead of continuous. Further, they show that the posterior collapse issue is solved.

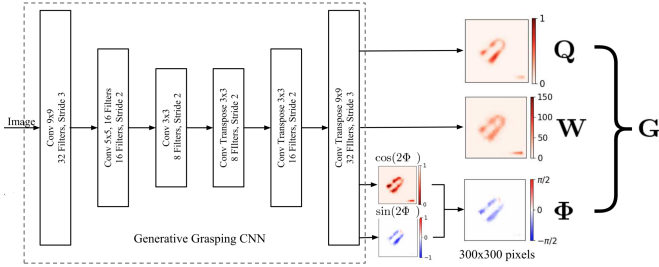


Fig. 2. GGCNN network architecture. [3]

VQ-VAE adds a latent embedding space (also known as the dictionary) to the general VAE framework. Each embedding  $e \in \mathbb{R}^D$ , and  $K$  is the number of such embeddings in the latent embedding space. Here, the encoder is not used to model a gaussian distribution. Instead, its output,  $z_e(x)$ , is used to perform a nearest neighbour lookup on the vectors in the embedding space to obtain  $z_q(x)$ . Hence, the continuous output vectors from the encoder are being quantized. This is known as vector quantization. Next, the decoder uses  $z_q(x)$  for the reconstruction task. Unlike a vanilla VAE, the posterior is a one-hot distribution over the embedding space:

$$q(z = k|x) = \begin{cases} 1 & \text{for } k = \operatorname{argmin}_j \|z_e(x) - e_j\|_2, \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

Along with this, the loss function is the sum of ELBO loss, a dictionary learning term, and a commitment loss. The dictionary learning term moves the latent embeddings closer to the output of the encoder, and the commitment loss makes sure the encoder commits to the embedding by moving the output of the encoder to be closer to the chosen embedding from the dictionary. In the following equation, the stop gradient operation is defined as  $\operatorname{sg}$ .

$$L = \log p(x|z_q(x)) + \|\operatorname{sg}[z_e(x)] - e\|_2^2 + \beta \|z_e(x) - \operatorname{sg}[e]\|_2^2 \quad (4)$$

During training, the prior is chosen to be a uniform distribution over the embedding space. As a result,  $D_{KL}(q(z|x)||p(z))$  becomes a constant equal to  $\log(K)$ . This indeed is independent of the network weights and hence, it is not included in the objective function,  $K$  turns into a hyper-parameter instead. Note that the vector quantization layer does not allow the use of backpropagation due to the arg-min operation. Instead, the gradient is directly copied from the decoder to the encoder, like the straight-through estimator. Fig. 1 gives a visual description of the VQ-VAE model.

3) *Generative Grasp Convolutional Neural Network (GGCNN)* [3]: The Generative Grasp - CNN or the GGCNN model eliminates the need of a two staged pipeline to predict the grasp, thus, causing a huge reduction in both the number of parameters in the model, and execution time. This makes closed loop grasping feasible, causing the robot to visually detect changes in the environment as it reaches to grasp the object, and change the trajectory of the gripper accordingly.

Instead of sampling grasp candidates first and then following the pipeline, the GGCNN network directly predicts grasps on each and every pixel of the image. Given a 2D image  $\mathbf{I}$ , the grasp  $\tilde{\mathbf{g}}$  in the image can be represented as  $\tilde{\mathbf{g}} = (\tilde{p}, \tilde{w}, \tilde{\phi}, q)$ , where  $\tilde{p}$  denotes a pixel position  $(i, j)$  which is the centre of the grasp rectangle in the image,  $\tilde{w}$  and  $\tilde{\phi}$  denote the gripper opening width and the angle of rotation of the gripper respectively. Let the image  $\mathbf{I}$  be of dimension  $H \times W$ . The network architecture approximates a function  $\mathbf{M}$  such that :

$$\mathbf{G}^{3 \times H \times W} = M(\mathbf{I}^{H \times W}) \quad (5)$$

Here,  $\mathbf{G}$  is the *grasp map*, i.e. the set of all grasps over the image space and is denoted by  $\mathbf{G} = (\mathbf{W}, \Phi, \mathbf{Q})^{3 \times H \times W}$ , where each of  $\mathbf{W}$ ,  $\Phi$  and  $\mathbf{Q}$  are of dimension  $H \times W$ , containing the values of  $\tilde{w}$ ,  $\tilde{\phi}$  and  $\tilde{q}$ , for each pixel  $\tilde{p} = (i, j)$  of the image, where  $0 \leq i < H$  and  $0 \leq j < W$ . The optimal grasp rectangle is the one for which  $\tilde{q}$  is maximum.

The network architecture of the GGCNN model consists of multiple convolution layers stack against each other with varying kernels and strides as shown in Fig. 2.

### B. Proposed Approach

Given a dataset with  $n_1$  labelled training pairs of the form  $(x, y)$ , where  $x$  is an RGB image and  $y$  is the grasp vector, and  $n_2$  unlabelled training vectors of the form  $(x)$  with unknown  $y$ , we train a VAE on  $n_1 + n_2$  images in the dataset. The trained VAE should capture important details which should help us in the supervised learning task that follows. We tried a vanilla VAE with normalizing flows, and a VQ-VAE to model the distribution of our training data. By observing the reconstruction quality, we came to the conclusion that VQ-VAE worked very well for our task. This decision was also based on the fact that VQ-VAE does not suffer from posterior collapse. The encoder and the quantization layers are then used to obtain  $n_1$  latent vectors. These latent vectors are then clubbed with their corresponding  $y$  values to obtain training pairs of the form  $(z, y)$  for the supervised learning task, where  $z$  is obtained by passing the corresponding  $x$  through the encoder and the quantizer. These  $(z, y)$  pairs are then used to train our modified GGCNN network. We altered the original GGCNN network by using the decoder architecture used in the VQ-VAE as the initial structure of the modified GGCNN. Intuitively, since the decoder architecture worked well in efficiently using the latent embeddings, it should benefit this task as well. Fig. 3 gives a schemata of the proposed approach.

## IV. RESULTS AND ANALYSIS

We use the Cornell Grasping Dataset (CGD) [18] to train our networks. The dataset contains 885 RGB-D images. Though the dataset is small, D Morrison et al. base the decision to use this dataset for the GGCNN network on the fact that it contains 5110 positive grasp labels, which benefits the prediction of the grasp map. Nevertheless, the dataset is augmented by performing transformations like rotation. As a preprocessing step, we divide every grasp rectangle in the Cornell Grasping Dataset into three parts and the centre

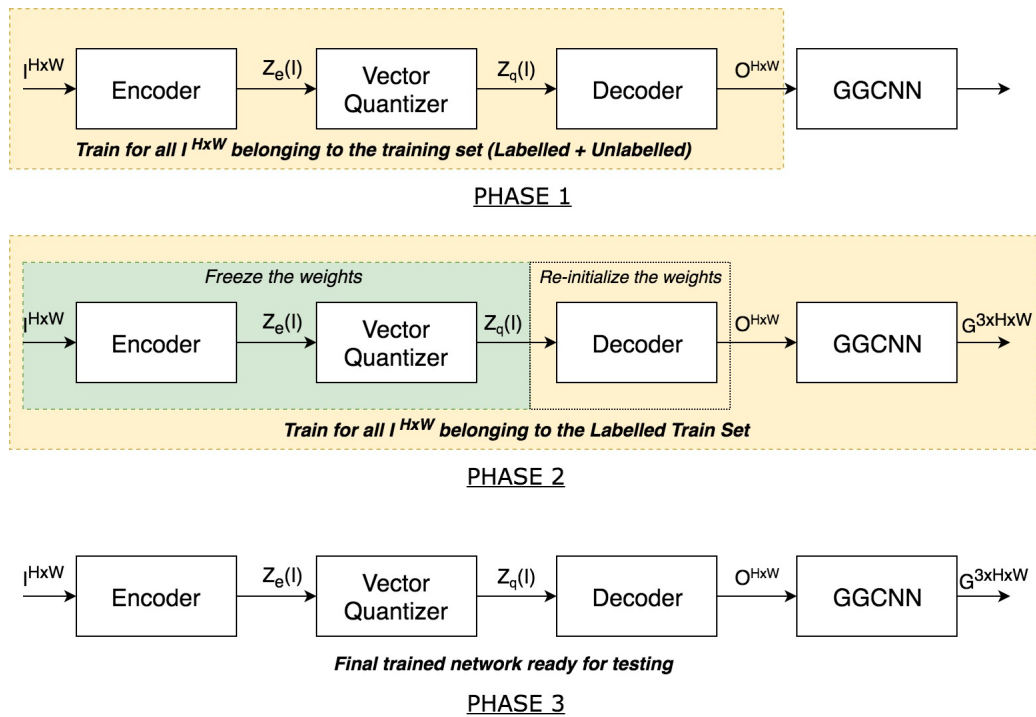


Fig. 3. Proposed Approach.

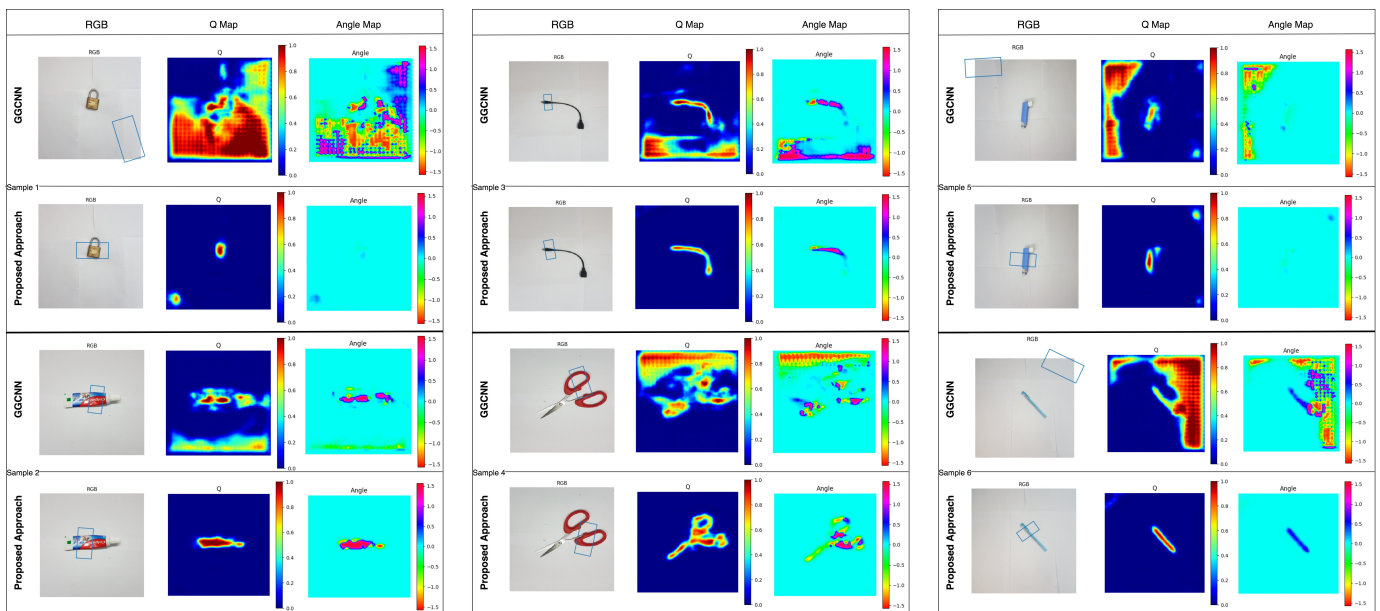


Fig. 4. A comparison between the grasp representations obtained from a GGCNN network and our proposed approach on different unseen objects (a lock, a tooth-paste, a cable, a scissor, an inhaler and a pen)




Fig. 5. A comparison between the GGCNN network and the proposed approach for varying amount of labelled data.

rectangle is considered to be the position of the gripper's centre. Although the Cornell Grasping Dataset contains negative labels as well, we follow the steps of the original GGCNN training procedure and consider any area other than the positive grasp rectangle to be an invalid grasp. 90% of the dataset (10% labelled data and 80% unlabelled data) forms the train set, while the remaining 10% forms the test set. We keep the labelled data percentage low to simulate the unavailability of sufficient labelled data for training. For evaluating the predicted grasps, we use the Jaccard index as the evaluation metric. The predicted grasp is considered to be correct if the Jaccard index of the predicted grasp rectangle and the human-labelled grasp rectangle is greater than 25%.

The first experiment involves training the GGCNN network on only 10% labelled images from the train set and evaluating on the test set. The second experiment is based on our proposed approach. Firstly, all the images in the train set are used as unlabelled images to train the VQ-VAE. Next, the weights of the encoder network and the vector quantization layers are frozen. Thereafter, the decoder network and the cascaded GGCNN network are trained on the labelled images in the train set. In the first experiment, we observe a test accuracy of **76.404%**. In the second experiment, however, the test accuracy is **85.3933%**, which shows the effectiveness of the proposed method.

Fig. 4 shows the Q Map and the Angle Map for multiple object samples, which are not present in the CGD. Color gradients have been used in place of numerical ranges. It is observed that both the networks produce good grasps when tested on images in the test set. However, the results are drastically different when tested on images outside the CGD. Though at times the GGCNN network is able to determine a good quality grasp rectangle (as shown in Fig. 4), the Q Map is far worse in comparison to the one produced by our model in all the cases. This proves that our model is able to generalize well. Fig. 5 plots the accuracy of the two approaches with varying amount of labelled data. Clearly, when the amount of labelled training data is low, our model performs much better than a neural network which utilizes only the labelled data in the training process for grasp detection.

## V. CONCLUSION AND FUTURE WORK

In this work, we have presented a new approach for vision based robotic grasp detection. We have shown that semi-supervised learning approaches can be utilized to make use of limited labelled data available in the grasp detection domain. Further, we have shown that using a Vector Quantized Variational Autoencoder can help in extracting useful features for determining the grasp vector. Our experiments demonstrate that our new approach outperforms current approaches by a significant margin. Also, we have shown that our model was able to generalize as is evident by the results wherein it performed very well on images not present in the Cornell dataset, despite being trained on a small amount of labelled data. The results obtained from the GGCNN network, however, were not good in comparison.

In our approach, the labels had no influence on the estimation of the posterior distribution. Kingma et al. [19] focus on the classification problem and construct the posterior as  $q(\mathbf{z}, \mathbf{y}|\mathbf{x})$  instead and assume it to have a fully factorized form. Here,  $q(\mathbf{y}|\mathbf{x})$  is modelled by the classifier. In future, our approach can be improved by constructing a better posterior which incorporates the GGCNN network and benefits from the labels as well. As a side-benefit, the whole training process would become end-to-end too.

## REFERENCES

- [1] S. Kumra and C. Kanan, "Robotic grasp detection using deep convolutional neural networks," pp. 769–776, 2017.
- [2] J. Kober and J. Peters, "Imitation and reinforcement learning," *Robotics & Automation Magazine, IEEE*, vol. 17, pp. 55 – 62, 07 2010.
- [3] D. Morrison, P. Corke, and J. Leitner, "Closing the loop for robotic grasping: A real-time, generative grasp synthesis approach," *CoRR*, vol. abs/1804.05172, 2018. [Online]. Available: <http://arxiv.org/abs/1804.05172>
- [4] M. Jiang and Saxena, "Efficient grasping from RGBD images: Learning using a new rectangle representation," *IEEE International Conference on Robotics and Automation*, May 2011.
- [5] 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.
- [6] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," in *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065)*, vol. 1. IEEE, 2000, pp. 348–353.
- [7] S. Caldera, A. Rassau, and D. Chai, "Review of deep learning methods in robotic grasp detection," *Multimodal Technologies and Interaction*, vol. 2, no. 3, p. 57, 2018.
- [8] L. Y. Ku, E. G. Learned-Miller, and R. A. Grupen, "Associating grasping with convolutional neural network features," vol. abs/1609.03947, 2016. [Online]. Available: <http://arxiv.org/abs/1609.03947>
- [9] J. Redmon and A. Angelova, "Real-time grasp detection using convolutional neural networks," pp. 1316–1322, May 2015.
- [10] J. Watson, J. Hughes, and F. Iida, "Real-world, real-time robotic grasping with convolutional neural networks," 07 2017, pp. 617–626.
- [11] J. Wei, H. Liu, G. Yan, and F. Sun, "Robotic grasping recognition using multi-modal deep extreme learning machine," *Multidimensional Systems and Signal Processing*, vol. 28, no. 3, pp. 817–833, 2017.
- [12] J. Wang, Q. Hu, and D. Jiang, "A lagrangian network for kinematic control of redundant robot manipulators," *IEEE Transactions on Neural Networks*, vol. 10, no. 5, pp. 1123–1132, 1999.
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," pp. 1097–1105, 2012.
- [14] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," pp. 770–778, 2016.
- [15] D. P. Kingma and M. Welling, "An introduction to variational autoencoders," *CoRR*, vol. abs/1906.02691, 2019. [Online]. Available: <http://arxiv.org/abs/1906.02691>
- [16] A. van den Oord, O. Vinyals, and K. Kavukcuoglu, "Neural discrete representation learning," vol. abs/1711.00937, 2017. [Online]. Available: <http://arxiv.org/abs/1711.00937>
- [17] O. Vinyals, A. Toshev, S. Bengio, and D. Erhan, "Show and Tell: A Neural Image Caption Generator," *arXiv e-prints*, p. arXiv:1411.4555, Nov 2014.
- [18] C. University, "Robot learning lab: Learning to grasp," Available online: [http://pr.cs.cornell.edu/grasping/rect\\_data/data.php](http://pr.cs.cornell.edu/grasping/rect_data/data.php).
- [19] D. P. Kingma, D. J. Rezende, S. Mohamed, and M. Welling, "Semi-Supervised Learning with Deep Generative Models," *arXiv e-prints*, p. arXiv:1406.5298, Jun 2014.