Deep Learning Approaches to Grasp Synthesis: A Review

Rhys Newbury^{1,2}, Morris Gu¹, Lachlan Chumbley¹, Arsalan Mousavian³, Clemens Eppner³, Jürgen Leitner^{1,4}, Jeannette Bohg⁵, Antonio Morales⁶, Tamim Asfour⁷, Danica Kragic⁸, Dieter Fox^{3,9}, Akansel Cosgun¹⁰

Abstract—Grasping is the process of picking up an object by applying forces and torques at a set of contacts. Recent advances in deep-learning methods have allowed rapid progress in robotic object grasping. In this systematic review, we surveyed the publications over the last decade, with a particular interest in grasping an object using all 6 degrees of freedom of the end-effector pose. Our review found four common methodologies for robotic grasping: sampling-based approaches, direct regression, reinforcement learning, and exemplar approaches. Additionally, we found two 'supporting methods' around grasping that use deep-learning to support the grasping process, shape approximation, and affordances. We have distilled the publications found in this systematic review (85 papers) into ten key takeaways we consider crucial for future robotic grasping and manipulation research. An online version of the survey is available at https://rhys-newbury.github.io/projects/6dof/

I. INTRODUCTION

The manipulation of objects is an essential skill for robot in the real-world, with grasping being an integral part of such tasks. Grasping is the process of controlling an object's motion in a desired way by applying forces and torques at a set of contacts. It is an essential ability required for the majority of object manipulation tasks. Synthesizing a good grasp proposal for a specific gripper and an object or a scene made of objects is a high-dimensional search or optimization problem as there are many relative gripper-object poses, joint configurations, and contact conditions. An example of multiple grasp proposals for a single object can be seen in Fig. 1. The quality of each of these grasp hypotheses can be evaluated under a variety of criteria such as grasp stability, which depends on factors such as object or scene geometry, gripper geometry, and kinematics, as well as suitability for a specific manipulation task. Reflecting on more than four decades of research in robotic grasping, we see a change in how grasping is formulated and studied.

Early work on robotic grasping developed a theoretical framework that forms the basis of analytical approaches to grasping [1]. At the core of this framework are contact models, which are typically based on point contacts that define what components of contact forces and torques (i.e. wrenches) can be transmitted at a specific contact and act on the object. In this framework, a *grasp* is defined as the set of wrenches that can be achieved on an object. The goal of grasp synthesis is then often framed [2] as finding a grasp that keeps the object in equilibrium in the presence of disturbances (i.e. *fixturing*) or moves it in a specific way (i.e. *dexterous manipulation*).

¹ Monash University, Australia

- ³ NVIDIA Corporation, USA
- ⁴ LYRO Robotics Pty Ltd, Australia

- ⁶ Jaume I University, Spain
- ⁷ Karlsruhe Institute of Technology, Germany
- 8 KTH Royal Institute of Technology, Sweden
- 9 University of Washington, USA
- ¹⁰ Deakin University, Australia



Fig. 1: Grasp Synthesis is the problem of creating a grasp pose, or a set of grasp poses. Robotic manipulators can grasp objects from multiple angles. This systematic survey reviews deep learning approaches to grasp synthesis that generate grasps utilizing all 6DoF.

Bicchi and Kumar [2] also mention enveloping grasps that wrap the fingers and the palm around the object, achieving more restraining grasps (i.e. *power grasps*). However, the limitation of these analytical approaches is that they assume full knowledge of object shape and geometry, material properties, and dynamics parameters. But in reality, this information is rarely directly observable but can only be inferred from partial, noisy sensory data.

The increased application of data-driven approaches to computer vision has successfully transferred to robotic grasping [3], primarily addressing the complexity and uncertainty in visual perception. To better fit these computer vision techniques, the focus of robotic grasping shifted from concepts around multi-fingered, contact-based representation to pose-based ones. Commonly abstracting the robot's end-effector as an ideal two-fingered gripper approaching the object "top-down". That way, a grasp is parameterized by the position and orientation of the coordinate frame attached to the gripper or the robot wrist. Before that, the degrees of freedom of a grasp were attributed to the robot hand, its kinematic structure, and the ability to control finger movements. This simplification of grasp parameterization is further supported by the increased availability of robust and simple end-effectors - e.g. parallel jaw grippers, suction cups [4] - and under-actuated or soft hands [5, 6] that simplify the control, compared to a dexterous hand.

For the purpose of this review, we reuse terms such as four and six degrees of freedom (DoF) grasps, where the former encodes position in 3D and single rotation of a coordinate system attached to the gripper, relevant for defining "top-down" grasps. Six degrees of freedom grasps relate to approaches that consider the full pose of that coordinate system. Fig. 2 shows examples of both 6-DoF and 4-DoF grasps, for both parallel jaw grippers and dexterous hands. While 6 DoF grasp synthesis is particularly useful in applications that require precise and dexterous manipulation of objects, such as for pick-and-place tasks. However, it can also be more challenging than 4-DoF grasp synthesis due to due to the increased complexity

² The Australian National University, Australia

⁵ Stanford University, USA



Fig. 2: This review focuses on the synthesis of 6 degrees-of-freedom (DoF) grasp pose hypotheses, where the DoF refers to the generated hand/wrist poses defined by the 3D position and orientation of a gripper specific coordinate system. This includes grasp synthesis with parallel grippers and dexterous hands alike.

of the grasping problem.

The dominant framework in the field of computer vision is applying deep learning to images. The robotics community often benefits from advancements in computer vision as the robots are often equipped with RGB and/or depth cameras. Deep learning, particularly when combined with large-scale data, has enabled new robotic capabilities that traditional machine learning or analytical methods cannot match. These include end-to-end manipulation learning [7], ambidextrous grasping [8], in-hand dexterous manipulation [9]), pick-and-place in clutter [10], as well as more recent works such as task-agnostic generalization to new tasks [11] and translating natural language commands to robot actions [12].

Deep learning methods are also crucial for 6-DoF grasp synthesis because they can offer increased accuracy, flexibility, reduced manual engineering, and robustness to variability compared to traditional machine learning or analytical methods. They can learn complex and flexible models that adapt to a wide range of grasping scenarios, without requiring manual feature engineering. Additionally, deep learning methods can account for variations in object shape, pose, and other factors, resulting in more robust grasping performance. Typical stages for grasping an object are illustrated in Fig. 3, with grasp synthesis being the first stage before any action is taken. In this review, we focus on deep-learning approaches applied to 6-DoF grasping, in particular on the grasp synthesis stage of the grasping process.

There are recently published review papers [13–15] similar to ours. Kleeberger et al. [13] covers a broad range of methods for grasping, including model-based, analytical, optimization-based, and learning-based methods. While both papers cover similar topics, our paper aims to provide a more detailed and comprehensive review on the state-of-the-art for 6-DoF grasping. Kroemer et al. [14] provides a general review of machine learning techniques for robot manipulation, including grasping, reaching, and other manipulation tasks. The paper provides a high-level view of the field and is less focused on specific grasping methods. Differently from these two papers, ours has a specialized focus on the specific topic of deep learning for grasp synthesis while offering a comprehensive discussion around future research directions on the topic. To our knowledge, Duan et al. [15] is the closest to our work in terms of taxonomy, focusing on point cloud-based inputs and depth of discussion. The discussion in Duan et al. [15] is mainly around sampling-based methods, which was the dominant method before 2020, their cut-off year for the reviewed papers. In our paper, in addition to Sampling methods (Section III-A), we capture the increasing popularity of end-to-end methods (called Direct Regression in our review - see Section III-B). We also identify supporting methods, such as Affordances (Section IV-B) and Shape Approximation (Section IV-A), that are commonly used together with main grasp synthesis methods. Differently from these review papers, ours provide an analysis on the benchmarking (Section VI) and datasets (Section V) for grasp synthesis. To our knowledge, our review is the only one that focus on 6-DoF grasp synthesis. We believe it is a crucial area of research for future progress in robotic manipulation.

Our systematic review is based on 85 research publications that employ deep-learning methods for grasp synthesis, clustering the work along common methodologies, data sets, and object-sets used. From the methods' viewpoint, we devise a taxonomy along: Grasp Sampling methods, Direct Regression methods, methods incorporating Reinforcement Learning (RL), and Exemplar methods. Furthermore, we identify methods commonly employed in grasping, such as Shape Completion and Affordances. Fig. 5 depicts a visual representation of the structure of the survey paper.

The contributions of this survey are:

- A systematic review of 85 papers, focusing on deep-learning based 6-DoF grasping.
- The synthesis of the papers into 10 key takeaways (discussed in Section VII) which we consider crucial for future research in robotics and manipulation.

II. NOTATIONS AND ANALYTICAL GRASPING

To provide a nuanced discussion regarding the contributions of deep-learning based approaches to grasp synthesis, we briefly overview the necessary fundamentals of grasping and review notations and definitions used to define it.

- A **Grasp Pose** defines the position and orientation of a grasp. Our survey found many different formulations of grasping but most aim to learn the final pose of the robot to generate a successful grasp.
- A **4-DoF grasp** defines a grasp where hand poses are generated and defined by a 3D position and hand orientation about an approach vector that is commonly aligned with the direction of gravity and is therefore often referred to as "top-down grasping". It is often denoted by *x*,*y*,*z*,*yaw*.
- A 6-DoF grasp defines a grasp where hand poses are generated and defined by a 3D position and orientation, thus 6-DOF in total. The major difference to 4DoF grasps is a non-fixed approach vector, providing extra flexibility but increased complexity.

3



Fig. 3: Typical stages for grasping an object. Our review focuses on grasp synthesis, the first stage in the grasping process.

• Affordances refers to the different tasks which can be achieved with an object [16]. This definition has been adopted in previous grasping works (e.g [17]).

An alternative way to frame a grasp, introduced by [18], is using the following three terms:

- An **approach vector** defines the line along which the gripper or robotic hand approach the target object.
- The **grasp center point** is the point in space somewhere along the approach vector where the coordinate frame fixed to the gripper must be positioned before starting to close the fingers.
- The hand orientation defines how the robot hand is oriented around the approach vector when placed on the grasp center point.

Some important terms used throughout the paper while describing approaches to grasping.

- Each point force applied at a contact point on the object surface also generates a torque on the object. A **wrench** summarizes this pair of force and torque applied to the object through a contact in a six-dimensional vector.
- A grasp is said to be in **Force Closure** if the forces that can be applied at the set of *frictional* contacts are sufficient to compensate for *any* external wrench applied to the object [19].
- Hand posture describes the configuration of the gripper or hand fingers when the grasp is started or all the contacts are made
- A **Power Grasp** is a grasp where there are multiple points of contact between the object and the fingers and palm. It maximizes the load carrying ability of a grasp and is highly stable as the enveloping nature of the grasp provides form closure [20].
- Antipodal points are pairs of points on the object surface whose normal vectors are collinear and pointing in the opposite direction [21]. With appropriate finger contact conditions, antipodal point grasps guarantee force closure.
- An **Antipodal grasp** is defined for two-fingered grippers that makes contact with the object at antipodal points [21].

For consistency and a thorough discussion, we provide a short insight on the relevance of analytical approaches.

The majority of methods up to the year 2000 modeled grasping analytically [2, 3]. The focus was on modeling and estimating physical conditions of grasps, such as, for example, grasp stability. A forceclosure grasp was often equated with a stable grasp, although force-

closure is a necessary but insufficient condition for a stable grasp [2, 22]. Physical conditions were usually simplified through approximations such as point contact models, Coulomb friction, and rigid body dynamics [3, 22, 23]. Analytic approaches have been attributed to being complex and not applicable in real-time applications. However, analytic approaches address properties of grasps, while most of the recent methods for grasp synthesis focus on positioning the hand. The advantages of analytical methods are mathematical guarantees on grasp properties, such as force-closure. This makes it easier to assess the conditions of grasps when objects are manipulated after a grasp has been applied. For example, what forces or torques can be exerted on the object before slippage occurs or how an object can be moved using in-hand manipulation. Within the first decade of the 21st century, there has been a rise of data-driven approaches to grasp synthesis [3], thanks to the development of grasping simulators such as GraspIt! [24]. Early approaches often used hand-designed features that corresponded to parts of objects that could be grasped [25-27]. Kamon et al. [28] presented one of the earliest works in 4-DoF grasping that use machine learning (ML) for grasping objects. The authors hand designed a low-dimensional feature space to estimate the quality of grasps. Since then, many works have employed traditional learning methods with a larger feature space [29–33]. One key takeaway from this early research was that while grasping can be done using RGB data only [31], depth information (RGB-D) improved grasping success [29, 32]. We refer to [3] for a more in-depth review of earlier methods for data-driven grasp synthesis. In this review, we focus on deep learning techniques in particular.

III. DEEP LEARNING METHODS IN 6-DOF GRASPING

The number of publications investigating deep learning approaches for 6-DoF grasping have grown significantly in the last few years, as highlighted in Fig. 4. From the systematic review, we identify four main algorithmic methodologies for grasp synthesis using deep-learning based on the reviewed literature: grasp pose sampling, regressing grasp pose directly, reinforcement learning and exemplar methods. These procedures relate to how the grasps are generated at test time.

Sampling approaches consider one or many grasp samples and have learned a function to estimate the quality of a sampled grasp. An essential characteristic of sampling approaches is that each sample is evaluated *individually*. Alternatively, direct



Fig. 4: The number of publications on IEEExplore that includes the keyword "Grasping" in metadata and "6DoF" in the full text is increasing year-by-year. We consider works published after Jan 1, 2012 – when AlexNet [34] was published – in our review.

regression considers the data *globally* and learns a function to predict high-quality grasps. *RL* includes methods that involve the maximization of a cumulative reward function based on a robot's actions or demonstrations. Exemplar methods aim to use a similarity metric between grasps to retrieve high-quality grasps from an existing database which are most similar. These categories are often mutually exclusive, though Aktaş et al. [35] employed both direct regression and sampling while Mahler et al. [36] employed both sampling and exemplar methods.

A. Sampling

TABLE I: Publications with deep-learning focused sampling methods. We cluster the papers based on the space the sample through and how the samples are evaluated. Some approaches further consider an optional refinement stage.

Year	Paper	Sample Space	Evaluation	Refinement
2015	Varley et al. [37]	Priors	Metric	
2015	Kappler et al. [38]	Not Specified	Binary	
2016	Gualtieri et al. [39]	Euclidian	Binary	
2017	ten Pas et al. [40]	Euclidian	Binary	
2017	Zhou and Hauser [41]	Euclidian	Binary	~
2018	Yan et al. [42]	Not Specified	Binary	~
2019	Mousavian et al. [43]	Latent	Binary	~
2019	Liang et al. [44]	Priors	Metric	
2019	Lu et al. [45]	Euclidian	Binary	~
2019	Ottenhaus et al. [46]	Priors	Metric	
2019	Gonçalves and Lima [47]	Euclidian	Binary	
2019	Aktaş et al. [35]	Priors	Binary	~
2020	Riedlinger et al. [48]	Euclidian	Binary	
2020	Murali et al. [49]	Latent	Metric	~
2020	Van der Merwe et al. [50]	Hand Posture	Binary	~
2020	Lundell et al. [51]	Multiple Views	Binary	
2020	Lou et al. [52]	Euclidian	Metric	
2020	Choi et al. [53]	Multiple Views	Binary	
2020	Lu et al. [54]	Hand Posture	Binary	~
2020	Schaub and Schöttl [55]	Multiple Views	Binary	
2020	Murali et al. [56]	Priors	Metric	
2020	Kokic et al. [57]	Euclidian	Metric	
2021	Lundell et al. [58]	Euclidian	Metric	
2021	Lou et al. [59]	Euclidian	Metric	
2021	Lundell et al. [60]	Euclidian	Metric	
2021	Peng et al. [61]	Not Specified	Metric	
2021	Jiang et al. [62]	Euclidian	Binary	
2021	Kasaei and Kasaei [63]	Multiple Views	Binary	
2021	Wang et al. [64]	Priors	Binary	
2021	Munoz [65]	Multiple Views	Binary	
2021	Corsaro et al. [66]	Priors	Binary	
2021	Ren et al. [67]	Not Specified	Metric	
2021	Wen et al. [68]	Not Specified	Metric	

We define sampling methods as any approach that considers each sample *individually* and use information encoded about the sample to make decisions about the grasp. The samples may be sourced from any discrete or continuous n-dimensional space. All reviewed sampling works are shown in Table I.

Deep-learning approaches employing sampling implement the following steps: sample information; evaluate the sample according to a quality estimation function, parameterized by a deep neural network; and optionally refine the sample using an optimization-based approach to achieve a higher quality grasp. We adopt the term 'quality', however, there is no current consensus on the definition. Throughout this paper, we use this term to refer to the confidence of grasp success. Table I presents an overview of all the *Sampling* papers. A popular deep-learning based approach for 4-DOF grasping generates a series of antipodal grasps through sampling while using a neural network to predict grasp quality [69]. This aims to predict the probability of a successful grasp to generate a series of antipodal grasps.

1) Generating Samples

Authors can sample a subset of grasp parameters in an n-dimensional space. This is commonly done using one of two approaches. Random sampling occurs when samples are taken from an arbitrary random distribution. The most common distribution is a uniform distribution, however, some authors sample from other distributions such as Gaussian [50, 70]. The second method is where samples are taken at equispaced intervals within the space.

The subset of grasp parameters can be sampled from: Euclidean space, priors, configuration space, latent space, or multiple views. When sampling in Euclidean space, heuristic-based rules are often used to remove irrelevant grasp candidates.

a) Euclidean Space Sampling

One of the most common approaches is to sample a 3D vector representing the translation part of the grasp pose. This is often achieved by sampling points from a point cloud [39, 52, 59]. The approach vector can be estimated by either using normal information estimated from the sensor data or sampling angles. ten Pas et al. [40] and Riedlinger et al. [48] use the opposite direction of the surface normal to find the grasp approach vector. Riedlinger et al. [48] uses a series of local augmentations on the initial grasp approach vector to generate a sample set of candidate approach vectors. Gualtieri et al. [39] generate candidate grasp approach vectors using the surface normal and curvature axis to generate equispaced orientations orthogonal to the curvature axis.

Some approaches sample angles independently of the surface normals. Lou et al. [52, 59] chose N points from a point cloud and sample the wrist angles randomly for the grasp. However, they restrict the approaching vector of the grasp to be above the table. Kokic et al. [57] randomly sample grasp and roll angles, and offset distances for each point in the point cloud. Both Lu et al. [45], Lundell et al. [58, 60] sample grasp candidates around the center of an object with a random orientation.

Sampling can also be performed with regularly spaced points through euclidean space. Jiang et al. [62] sample regularly spaced points for the position of the grasp pose. Similarly, Gonçalves and Lima [47] sample equispaced points throughout the region of interest.

Alternatively, instead of sampling a xyz vector, angles can be used. Zhou and Hauser [41] chose a random hand orientation and approach vector after which they translate the hand along the approach vector of each sample until a grasp is found that does not collide with the gripper when it is open but collides when closed.



Fig. 5: A visual representation of the core topics covered in this systematic survey.

In Euclidean space, sampled grasp are commonly pruned to remove infeasible grasps based on rules such as:

- The robot hand is in collision with the point cloud when fingers are open [40, 41].
- The closing region of the fingers does not contain at least one point from the point cloud [39–41].

b) Sampling Priors

More complex algorithms can also be used to find a set of feasible grasp candidates. These are often used when creating grasp samples which will then be evaluated for affordances [56] or different types of grasps (e.g power, pinch grasps) [66]. The most commonly [44, 56, 64, 66] used grasping algorithm is an SVM-based approach proposed by Pas and Platt [71]. Mahler et al. [36] use a modification of the grasping algorithm proposed by Smith et al. [72] to generate a series of antipodal grasps. They frame the problem as a multi-armed bandit, to be solved using Thompson sampling [73].

While grasping simulators are often used for generating training data, they have also been used at test time to sample grasps. A common prerequisite for grasping simulators is a full approximation of the shape model. The common simulators for this purpose were GraspIt! [24] or Simox [74]. Ottenhaus et al. [46] generated grasp samples on a reconstructed object using the grasp planner by Simox [74]. Alternatively, Varley et al. [37] sampled grasps using the Simulated Annealing planner [75] for partially visible objects.

A deep-learning approach can also be used to synthesize grasp samples. Aktaş et al. [35] used the direct regression based approach from Kopicki et al. [76] to generate multiple grasps, to be used as grasp samples.

c) Latent Space Sampling

Mousavian et al. [43] train a Variational Auto Encoder [77] and uniformly sample through latent space to generate grasp poses. An example of the grasps generated by this approach is shown in Fig. 6. This sampling approach is also adopted by Murali et al. [49].

d) Hand Posture Space Sampling

Some authors sample grasp configurations from a prior distribution fit to the training set to create an initial hand configuration for their approach [50, 54, 70]. Van der Merwe et al. [50] and [70] fitted a Gaussian Mixture Model to represent a grasp prior function trained on grasp configurations seen during training. Similarly, Lu et al. [54] trained a Mixture Density Network [78] over all grasp attempts from the training set.

e) Multiple Views

A number of papers employ multiple viewpoints of the scene, where these can either be from virtual or real cameras. A camera viewpoint is then sampled, and a grasping approach is employed on the sampled viewpoint. Schaub and Schöttl [55] combine multiple viewpoints around the scene to create a 3D representation of an object. Using the 3D representation, they generate depth images for a series of virtual cameras. They then use a 4-DoF grasping algorithm [79] for each real and virtual camera. This is extended by Schaub et al. [80] who fused depth images around the scene to provide more detailed depth maps. Choi et al. [53] extended a previous 4-DoF approach [81] to 6-DoF by using an iterative improvement



Fig. 6: Mousavian et al. [43] sample through the latent space of a trained model to generate a series of grasp candidates (©2019 IEEE)

algorithm approach to choose an approach direction. Munoz [65] and Kasaei and Kasaei [63] both generate multiple views of the object from virtual cameras using a captured point cloud from a single viewpoint. They proposed a method to select a view according to an entropy-based measure. Lundell et al. [51] use the algorithm proposed by Satish et al. [81] on the depth map from multiple viewpoints. The robot then executes the best grasp from all viewpoints.

2) Sample Evaluation

Once samples are generated, these approaches use a function, commonly parameterized by a neural network, to estimate a numerical value representing a grasp metric. This is commonly used an estimation of the 'quality' of the grasp. The grasp quality can be approximated through an analytical metric or the likelihood of a grasp being successful. While some approaches execute the highest quality sample directly, others refine the samples using optimization-based techniques based on the learned quality function.

a) Binary classification

Predicting grasp success can be treated as a binary classification problem, where the output represents the confidence of a grasp being successful or not. A CNN can be trained for the binary classification of grasp success, with the network either outputting 0 or 1 for a successful grasp [35, 38–43, 45, 47, 48, 50, 54, 64, 66]. During the collection of ground-truth training data to determine if a grasp is successful, the grasp is most commonly executed using a robotics simulation environment [35, 38, 41–43, 45, 48, 50, 54, 62, 66], where a grasp will be labeled successful if the object remains within the gripper after being lifted. Grasp success can also be based on an analytical metric, such as whether the grasp would form a force-closure grasp [40, 47] or if the grasp is antipodal [39].

The binary classification of a sampled grasps can also be consider as one of the outputs of a sampling approach. For example. Jiang et al. [62] designed an approach which learns implicit functions that predicts grasp parameters (quality, rotation and width) from the feature space representation of a randomly sampled query point. Van der Merwe et al. [50] trained a network to predict the success of a grasp, as well as a signed distance function that represents a distance between a query point and the surface of the object.

b) Learning a metric

Learning metrics associated with grasps, rather than a pure binary label, have been proposed to better represent the quality of a grasp. These metrics are often continuous numbers (rather then the previously described binary label) that provide additional information to rank grasp quality or help "guide" deep learning algorithms using them as a fitness score.

Varley et al. [37] learn a series of heatmaps, where each pixel represents the location's efficacy as a fingertip or palm location

for common grasp types found in their training set. Liang et al. [44] designed an approach which learns a grasp quality metric based on the force-closure metric and wrench space analysis [82]. Similarly, Ottenhaus et al. [46] train a CNN to estimate the forceclosure probability of a grasp under a small random perturbation. Lundell et al. [58] and Lundell et al. [60] trained a grasp classifier using a Generative Adversarial Networks (GANs), and use the discriminator loss to help produce realistic-looking grasps. Wen et al. [68] compute a continuous score for each grasp by looking at the stability of randomly sampled grasps in the proximity of the selected grasp. They argue that "grasp stability should be continuous over its 6D neighborhood", therefore, this should allow for more stable grasps. Furthermore, this can allow for imperfections in the robotic grasp position whilst still being likely to execute a stable grasp. Peng et al. [61] and Ren et al. [67] designed an approach to learn the smallest co-efficient of friction which will satisfy a forceclosure grasp. Learned metrics can also represent the quality of the grasp with respect to other aspects of the task, for example, metrics based on the relevance of the grasp to a given affordance [56, 57].

Other approaches consider workspace constraints. Lou et al. [52] designed an approach to learn the probability of grasp success and use an additional network to learn the probability of a grasp being reachable. These probabilities can be multiplied to find the likelihood of success for the entire grasping action. This work is extended by Lou et al. [59] to allow the robot to grasp in constrained environments, such as in boxes, where walls may limit the success of grasping an object.

3) Optimization-based Grasp Refinement

Gradient-based optimization through the trained network can be used to find high-quality grasps candidates. A sampled grasp is taken as the initial seed, which is then refined based on the derivatives of the quality estimation network. This attempts to maximize estimated grasp quality. Zhou and Hauser [41] train a CNN to predict grasp quality given a depth image and a sampled end-effector pose. The end-effector pose is locally optimized using the quasi-newton method on the gradient of the learned quality function. Similarly, some authors take the derivative of the grasp quality with respect to the grasp pose and then use gradient ascent to refine the grasp candidates [42, 43, 45, 49]. An example of the refinement process is shown in Fig. 7, where the initial grasp is a potentially bad sample (blue) and is refined to a higher quality grasp (yellow). Both Van der Merwe et al. [50] and Lu et al. [54] treat finding a grasp configuration as an optimization problem aiming to maximize the probability of grasp success, seeding the process with samples from a prior distribution. Lu et al. [70] extended this using active learning during training, improving their results. In contrast to other works, Aktaş et al. [35] found that optimization of the sample did not improve the grasp success rate in simulation. The authors explore the use of gradientbased optimization and simulated annealing for their optimization.

B. Direct Regression

Direct regression approaches can simultaneously process the entire sample space using a single network to predict either a single or multiple grasps, along with specific grasp properties such as parameters and quality, from visual information.

Direct regression approaches, which can be considered as end-toend methods, utilize a single network to process the entire input to regress an output. These methods attempt to reduce computational



Fig. 7: An example of a grasp refinement process from [43], where the initial grasp is shown in dark blue and the final grasp pose is shown in yellow. (©2019 IEEE)

TABLE II: Approaches to direct regression of 6-DoF grasping using Deep Learning, either generating a single or multiple grasp poses as an output. We found three main approaches: regress to a pose directly (Pose), employing a multi-stage approach, or perform dimensionality reduction (DimRed).

Year	Paper	Grasps	Direct Pose	Multi-stage	DimRed
2017	Veres et al. [83]	Single	~		
2018	Schmidt et al. [84]	Single	~		
2018	Choi et al. [85]	Single			~
2019	Liu et al. [<mark>86</mark>]	Single	~		
2020	Qin et al. [87]	Multiple			~
2020	Yang et al. [88]	Single	~		
2020	Jeng et al. [89]	Multiple			~
2020	Wang and Lin [90]	Single			
2020	Fang et al. [91]	Multiple		~	
2020	Wu et al. [92]	Multiple			~
2020	Ni et al. [93]	Multiple			~
2020	Breyer et al. [94]	Multiple			~
2020	Liu et al. [95]	Single	~		
2020	Shao et al. [96]	Single		~	
2020	Ni et al. [97]	Multiple		~	
2021	Wang et al. [98]	Multiple		~	
2021	Sundermeyer et al. [99]	Multiple			~
2021	Zhao et al. [100]	Multiple		~	
2021	Gou et al. [101]	Multiple			~
2021	Zhu et al. [102]	Multiple		~	
2021	Wei et al. [103]	Multiple		~	
2021	Li et al. [104]	Multiple			~
2021	Li et al. [105]	Single		~	

cost compared to sampling methods by processing data globally through the network in a single pass. Early works in direct regression based on deep learning are inspired by object detection work from the Computer Vision community [106, 107]. Authors would treat finding top-down grasps similar to detecting objects in an image and would use the depth image to recover the 3D position of the grasp [108–111]. Morrison et al. [79] designed a heat-map based approach for 4-DoF grasping that was designed for fast inference, enabling closed-loop grasping and grasping of moving objects. However, as the dimensionality of the output increased, the problem difficulty also increased [43, 112]. This presents scaling issues when directly regressing to high-dimensional outputs such as a 6-DoF grasp pose or a high-DoF dexterous robotic hand configuration. To overcome this, researchers often reduce the DoF of the output from a single network. This can be accomplished through multi-stage approaches, where each stage has a specific task, or by reducing the DoF by conditioning the output on an input. The direct regression approaches found in our systematic review are shown in Table II.

a) Directly Regressing a Pose

Schmidt et al. [84] presented one of the earliest CNN-based works to directly estimate a single 6-DoF grasp pose of an endeffector from input depth images. Similarly, Yang et al. [88] trained a network to estimate the transformation matrix needed to be applied



Fig. 8: Fang et al. [91] uses a direct regression approach to generate a series of grasps from real-world data. (©2020 IEEE)

to the end-effector to produce a successful grasp. However, both approaches assume that there is only a single, most optimal ground truth grasp for each input. This would introduce ambiguities as there may potentially be numerous successful grasps that can be executed for any input. To account for this, Liu et al. [86] designed a loss function that accounts for multiple ground truth grasps in the training data by calculating the loss between the current output and the closest ground truth grasp. This is extended in another work [95] that includes differentiable terms for a grasp metric and self-collisions in the loss function. This allows their approach to work both in an unsupervised manner or using a smaller supervised dataset. Veres et al. [83] create a generative model based on a *Conditional Variational Autoencoder* (CVAE) [113]. They use the CVAE to generate contact positions and contact normals for a multi-fingered robotic hand.

b) Reducing DoF

Due to the difficulty of regressing all 6 DoF of a grasp, some DoF can be reduced by analytically determining some DoF conditioned on the regressed DoF. This is seen in some 4-DoF approaches, for example, the depth of the grasp is not directly regressed but instead uses the depth image to recover the grasp depth conditioned on the grasp position and 2-finger gripper wrist rotation [79].

Sundermeyer et al. [99] reduced the 6-DoF grasping to a 4-DoF representation by ensuring one of the contact points for a twofinger parallel gripper was taken from the point cloud. The 3-DoF rotation and gripper width was then estimated. Brever et al. [94] directly output the predicted grasp quality, orientation, and opening width for each voxel in a queried 3D volume. The 3D position is recovered from the center of a voxel. Jeng et al. [89] propose a coarse-to-fine representation, where the orientation is initially coarsely discretized as a grid with a given confidence value. A refinement step is also used for each grasp pose to allow further flexibility from the discretized coarse representation. Gou et al. [101] estimate an SO(3) orientation and confidence of every pixel directly from an RGB image. An analytical method based on the depth image is then used to find the gripper width and position for each pixel. Ni et al. [93] directly regressed grasps from sparse point clouds. They predict the grasp quality, approaching direction, and opening direction of the gripper for every point in the original point cloud. Likewise, the approach by Li et al. [104] estimates the rotation, width, depth, and quality of a grasp for each point in the point cloud. They combine this with another branch which predicts whether the grasp would be in collision. Alternatively, Choi et al. [85] predicts the most likely grasping direction and wrist orientation from a set of discretized directions and orientations. The translational part of the pose is found by determining the centered contacting voxel within the grasping direction. Wang and Lin [90] assume

they are grasping the centroid of the object and attempt to directly regress the quaternion for the grasp. Wu et al. [92] create a network with three branches, trained to predict the likelihood of the grasp consisting of antipodal points, grasp offset for each of the sampled SE(3) points, and grasp confidence (trained on whether grasp was executed in simulation). Qin et al. [87] trains a neural network based on Qi et al. [114] to predict both grasp point and quality for every point in a point cloud.

c) Multi-Stage Approach

These approaches use multiple stages, where a stage refers to a component of the end-to-end model with a loss function and at least one specific task. This aims to simplify learning by breaking the problem down into smaller parts. Most multi-stage approaches consider three-stage approaches [91, 98, 100, 102, 105], while, Wei et al. [103] propose a two-stage approach. Stages are commonly done in series (one after another), however, the last two stages in the work by Fang et al. [91] are in parallel (simultaneously).

The first stages of the network are commonly used for at least one of the following: predicting grasp quality for subsampled points [91, 98, 100], estimating a grasp for each point [97, 103], estimating a subset of the DoF [91] or employing contrastive loss functions [102]. The middle stages can be used to create grasp proposals (if they were not generated in a previous stage) [98, 100] or further estimate a subset of the DoF [102].

The last stage can act as a refinement stage to improve the regressed grasps [97, 100, 103]. Fang et al. [91] used two final parallel stages to generate grasp proposals and predict the ability of the grasp pose to tolerate larger errors, aiming to improve the robustness to imperfect sensing or control. The final stage can also be used to predict the remaining DoF [102]. Example grasps regressed in the work by [91] from real-world data are shown in Fig. 8.

An alternative multi-stage approach is demonstrated by Shao et al. [96] and Li et al. [105] who propose an approach that uses each stage of the network to predict a single contact point. The subsequently regressed contact point will be conditioned on the previous point. The final grasp is then recovered from the regressed contact points. Furthermore, Shao et al. [96] shows that this approach is generalizable between different robotic hands.

C. Reinforcement Learning

Reinforcement Learning (RL) approaches aim to learn a policy to maximize the cumulative reward commonly over a multi-step task. We only review deep RL, where the policy is parameterized by a deep neural network. We sub-divided the reviewed works into two main approaches: On- and Off-Policy Learning. For a more comprehensive review on deep RL, see [115] and for a deeper exploration of RL in grasping and its open challenges, see [116].

Some seminal works in reinforcement learning for grasping include the work by [117–119]. Levine et al. [117] collected 800k grasp attempts over multiple months. They present a self-supervised approach, and the authors claimed their approach is analogous to an RL formulation. Kalashnikov et al. [118] and Quillen et al. [119] presented works that formulate grasping as a reinforcement learning problem. Kalashnikov et al. [118] focus on real-world data collection, collecting 580k real-world grasping attempt, while Quillen et al. [119] train on purely simulated data. Table III shows and summarizes all reinforcement learning based works found during our review.

TABLE III: Our systematic survey found 10 publications employing *Off Policy* or *On Policy* reinforcement learning (RL). The LfD column indicates approaches which learn from expert demonstration.

Year	Paper	Learning	Algorithm	LfD
2018	Gualtieri and Platt [120]	On Policy	DQN[121]	
2019	Wu et al. [122]	On Policy	PPO[123]	
2019	Merzic et al. [124]	On Policy	TRPO [125]	
2020	Mandikal and Grauman [126]	On Policy	PPO[123]	
2020	Song et al. [127]	Off Policy	Q Learning[128]	~
2020	Wu et al. [112]	On Policy	PPO[123]	
2021	Berscheid et al. [129]	On Policy	Single Step MDP	
2021	Kawakami et al. [130]	On Policy	PPO[123]	~
2021	Tang et al. [131]	Off Policy	Q Learning[128]	
2021	Wang et al. [132]	Off-Policy	DDPG[133]	~

1) On-Policy

We found that out of the reviewed work employing RL, On-Policy methods were more common [112, 120, 122, 126, 129, 130, 132, 134]. In On-Policy RL, training a policy is done using experiences that are collected from the most recent policy. In the work from Kawakami et al. [130], the grasping task is divided into consecutive stages: orienting the end-effector, approaching the target, and closing the gripper. A different RL model is trained for each stage, and curriculum learning is employed that adjusts the reward function based on the success rate of each task. Gualtieri and Platt [120] used a Deep Q-Network [121] with Monte Carlo updates to learn how to grasp an object and place it into a desired configuration. Working with a 24-DoF hand, Mandikal and Grauman [126] trained an actor-critic model. They proposed a two-step architecture: Initially, a CNN, which is trained on ContactDB [135], estimates the pixel regions that belong to a "use" affordance. An RL policy then takes this affordance mask, RGB-D image, and gripper configuration as input and outputs the wrist pose and the 24-DoF robot hand configuration. A sparse reward is awarded when the object is lifted from the ground, and a dense reward is awarded according to a distance metric from the affordance region. Chen et al. [134] used an advantage actor-critic policy gradient to train a policy that will optimize the viewpoint for grasping. They then apply the grasping algorithm developed by Pas and Platt [71] to calculate grasps for the optimized viewpoint. Wu et al. [122] used a single depth image and introduces a novel attention mechanism that learns to focus on sub-regions of the depth image in order to grasp better in cluttered environments. They formulate the problem using a policy gradient method based on PPO [123]. Wu et al. [112] extended this framework to robotic hands with arbitrary degrees of freedom. Merzic et al. [124] propose the use of TRPO to learn control policies that take contact feedback as input. They show that this policy significant improved the robustness of the grasp under both object pose uncertainty and shape complexity.

2) Off-Policy

Off-Policy RL methods use the data collected throughout training to train a new policy. Employing human demonstrations, Song et al. [127] used Q-learning to estimate the optimal Q-function. They simulate future states by giving an action for the robot to complete, allowing the algorithm to forward simulate possible future states conditioned on the current state-action pair. Tang et al. [131] demonstrated collaborative pushing actions to facilitate grasping. Their approach uses Q-learning to learn a deterministic policy for pushing and grasping. No reward is assigned to pushing actions - the agent is only rewarded when the robot successfully grasps the object. Berscheid et al. [129] formulated the problem as a Markov Decision Process with a single action step. They train a fully convolutional neural network to learn a 4-DoF planar grasping system. A model-based controller decides the other two degrees of freedom by avoiding collisions and maximizing grasp quality. Wang et al. [132] trained a grasping policy from demonstrations based on the Deep Deterministic Policy Gradient algorithm [133]. The demonstrations are obtained using a motion and grasp planner, which is assumed to be an 'expert' in their formulation.

D. Exemplar Methods

Exemplar methods attempt to transfer grasps from previous examples. Patten et al. [136] designed an approach to grasp novel objects by learning from experience. This is achieved using metric learning to encode objects with similar geometries nearby in feature space. Finding a successful grasp is framed as a nearest neighbor search through feature space, searching for a previously successful grasp. The approach by Mahler et al. [36] compared a given grasp candidate against a database of successful grasps. They compare grasp candidates using three feature maps: grasp parameters, depth match gradient at local patches around the object, and similarity of the object model assessed by a deep-learning based network. The feature maps form a prior belief distribution on the similarity to all grasps in the database. The grasp with the maximum lower confidence bound of the belief distribution is executed.

IV. SUPPORTING METHODS BASED ON DEEP LEARNING

Deep-learning can be applied to methods throughout the grasping pipeline that aim to improve the success of a grasping task. This task does not have to focus solely on picking up an object. More complex manipulation tasks may require grasping that affords a particular subsequent action. For example, when handing over a full mug, the robot may need to grasp the mug handle. We found two clusters of supporting methods in the reviewed literature: shape approximation techniques and affordances based methodologies.

A. Shape Approximation

The most common form to approximate the shape of an object from partial information found in the literature is *Shape Completion*, which aims to estimate the full object model from a partial input shape (e.g. point cloud of an object from one camera view). We define *Shape approximation* more generally to include any method which approximates shape from an input. This includes the approximation of the actual shape of an object by simple(r) shapes and the fusion of multimodal data to approximate a shape.

1) Shape Completion

Varley et al. [137] trained a 3D CNN to employ shape completion on a single view voxel grid, outputting a voxel grid with shape completed object. Gao and Tedrake [138] instead use the method from Zhang et al. [139] for shape completion, predicting a 3D voxel grid directly from RGB-D images. Kiatos et al. [140] use a variational autoencoder [141] to predict the occluded surface points and associated normals of a partial 3D point cloud. Chavan-Dafle et al. [142] predicted the depth image that estimates the 'back' side of an object from a masked depth image. The front and back sides can then be stitched together quickly to form an object mesh.

The uncertainty around the output of the shape completion can be useful. Gualtieri and Platt [143] incorporate uncertainty in



Fig. 9: Lundell et al. [144] uses drop out layers to generate 20 shape completed samples, where the average of these samples is shown in (d) compared to [137](c). (©2019 IEEE)

their shape completion network, where it represents the estimated probability that each predicted point is accurate. Another approach including uncertainty is Lundell et al. [144] which incorporates a Monte-Carlo drop-out procedure [145] to generate a series of shape completed objects. GraspIt! [24] is then used to plan grasps over the mean object shape. The most suitable grasp over the series of shapes is chosen as the grasp point. An example of this procedure is shown in Fig. 9. Interestingly, uncertainty is ignored in follow-on work from the same authors [51, 58, 60].

2) Auxiliary Tasks

In deep learning, auxiliary tasks can be added into a deep-learning model, which has been shown to boost the performance of the model in some domains [146]. Jiang et al. [62] asserted that 3D reconstruction and grasping are closely related, where both rely on knowledge of an object's local geometries. Authors have proposed learn object reconstruction as an auxiliary task to grasping [42, 62, 88]. Jiang et al. [62] used a self-supervised approach to reconstruct an object and calculate a grasp. Yang et al. [88] simultaneously regressed a grasp pose and reconstructed the point cloud of an object. The grasp pose is then refined by projecting it onto the surface of the point cloud reconstruction. Yan et al. [42] employed two networks, one for shape completion and another for grasp outcome prediction. They demonstrated a performance improvement when the grasping network uses the feature space representation produced by the shape generation network.

3) Other

Avigal et al. [147] do not complete explicit shape prediction, however, they used a network that takes RGB images from multiple viewpoints and generates the corresponding depth maps for the shape. The depth maps are then fed into a 4-DoF grasping algorithm [148]. Ottenhaus et al. [46] used Gaussian Process Implicit Surfaces [149] to fuse visual and tactile sensor inputs. After capturing a point cloud of the object, the robot gathers information about unseen sides of objects using tactile information. Researchers have shown that many everyday objects can be modeled as simple shapes [150, 151]. Using this observation, Torii and Hashimoto [152] approximated objects as a series of 3D primitive shapes (hexahedron, cylinder, sphere). They use a neural network to predict the likelihood of each primitive shape before using a pre-computed database of rules to perform the grasping.

B. Affordances

This section reviews how affordances have been used in the domain of 6-DoF grasping. See [153] for a review of affordances in the more general robotics domain. In addition to considering the success of grasping, these approaches have additional considerations for what kind of task it is used for. For example, when a *handover* task is to be completed for a scissors object, the robot should grasp



Fig. 10: Murali et al. [56] studies task-oriented grasps on unknown objects. The top right visualization shows the grasp with the highest quality considering a given affordance. The bottom right shows all the stable grasp candidates, colored by the relevance to the affordance (green is high)

the scissors by the blade and then pass it over to the human who can grasp it by the handle. The robot's understanding of how an object is used in a particular task or how humans use those objects can lead to higher-level reasoning about the grasping task. This lends itself towards 6-DoF grasping, as different parts of the object need to be grasped in different ways, depending on location of the affordance.

Some researchers use deep neural networks to segment objects for different affordances [126, 154–157]. They then use analytical methods to find grasps within the segmented object portion with an appropriate affordance. Alternatively, the approach by Murali et al. [56] estimated the quality of a sampled grasp given an affordance label. This is shown in Fig. 10, where the grasps shown in green are relevant to the given affordance. Ardón et al. [17] created a knowledge base graph representation using Markov Logic Networks to obtain a probability distribution of grasp affordances. Additionally, both Manuelli et al. [158] and Gao and Tedrake [138] presented an approach based on detecting a fixed number of keypoints for a category of objects.

V. DATASET DESIGN

A. Objects Sets

The objects used for training and testing grasping algorithms are crucial for the reported grasp success and allowing the community to reproduce the results. Researchers commonly use subsets of existing object sets when investigating grasping. However, there is no standard procedure for selecting this subset. This can lead to inconsistencies in objects between different works which are using the same object set. This increases the difficulty of comparing grasping performance between works. The most commonly used object set in the reviewed works is Yale-CMU-Berkeley(YCB) [159], being used almost twice as often than the next most adopted object set (Table IV).

YCB [159] (shown in Fig. 11), BigBIRD [160], KIT [161] and Cornell [109] consist of mostly household items such as food, toys and tools. These object sets are appropriate for service robotics, but may not test the robustness of grasping algorithms on complex objects. ShapeNet [162], 3DNet [163], Grasp [38], PSB [164], ModelNet [165], ObjectNet3D [166] and ContactDB [135] include object model repositories, containing a large number of virtual object models, mainly used for training and testing in simulation.

EGAD! [167] and Procedural [168] consist of procedurally generated object models. EGAD! [167] proposes a set of 3D printable objects that vary in terms of grasping difficulty [148] and object complexity [169]. Procedural [168] generate a simulated object set by attaching rectangular prisms in random orientations and locations.



Fig. 11: A commonly used object set in robotic grasping is the YCB object set [159]. It consists of a set of daily household objects, with a subset of the objects from the 'Food' category shown here. (©2015 IEEE)

TABLE IV: Number of times objects sets are used in reviewed papers.

Object Set/Database	# of times used	Sim//Real
YCB [159]	29	Real
3DNET [163]	15	Sim
BigBIRD [160]	12	Real
KIT [<mark>161</mark>]	12	Real
ShapeNet [162]	11	Sim
Grasp [38]	10	Sim
EGAD! [167]	2	Sim/Real
Cornell [109]	2	Real
Dex-Net [148]	2	Real
PSB [164]	1	Sim
ModelNet [165]	1	Sim
ObjectNet3D [166]	1	Sim
ContactDB [135]	1	Sim
Procedural [168]	1	Sim
Custom	11	Real/Sim

Most real-world object sets do not provide a standardized method to acquire the physical objects consistently (an exception to this is YCB¹). One solution to this is 3D printed datasets such as EGAD!², however, these objects lack semantic meaning.

B. Procedurally Generated Datasets

Even though a majority of the reviewed works use benchmark object sets, they commonly opt to create their own custom datasets with those objects. A number of works used datasets collected using real robots [63, 64, 85, 101, 129] and some combined simulation and real-world data [64, 91, 101]. However, the large majority opted to use purely simulated datasets when training their networks.

Some authors have released their datasets, with reviewed works employing public datasets such as GraspNet-1Billion [91], a hybrid 6-DoF grasping dataset that captures real RGB-D camera data and combines this with simulated grasp poses, Shape Completion Grasping [137], a database of voxel grid pairs for shape completion and ACRONYM [170], a simulation-based dataset for 6-DoF grasping.

C. Expert Datasets

Researchers have proposed various datasets which consist of expert demonstrations, either from a human or an algorithm. Yan et al. [42] generates a dataset of around 1.6k human grasping demonstrations within Virtual Reality (VR) from 5 people. Also in VR, Kawakami et al. [130] creates a system to collect grasping demonstrations. The operator demonstrates the position of the

¹https://www.ycbbenchmarks.com/

²https://dougsm.github.io/egad/



Fig. 12: Song et al. [127] design a low-cost handheld gripper to generate human annotations for grasping. (@2020 IEEE)

TABLE V: Number of times each input format has been used

Input Format	Number of times used
Point Cloud	34
Depth Image	18
RGB-D Image	15
Voxel Grid	12
Segmentation Mask	9
Other	10

arm using a controller with tracked position and pose. Using a handheld gripper (see Fig. 12), Song et al. [127] generated a dataset from human demonstration. The dataset contains 12 hours of gripper-centric RGB-D videos, with each picking attempt separated into short clips to correspond to grasps. Similarly, Diaz Cortes et al. [171] generated 300 demonstrations by kinesthetically teaching the robot trajectories. Osorio et al. [172] created a dataset of human grasps where the human is controlling a robotic gripper with a joystick. Alternatively, the approach by Wang et al. [132] used demonstration from Optimization-based Motion and Grasp Planner (OMG planner) [173]. Taheri et al. [174] presented a dataset of whole-body grasps generated by ten subjects interacting with 51 everyday objects of varying shape and size. The object meshes are annotated with the contacts created by the human hand.

D. Data Representation

There are four major sensor inputs used for the deep learning methods: Point Cloud; Voxel Grid; RGB-D Image; and Depth Image. These four representations are all interchangeable for spatial data, assuming that intrinsic parameters for the camera are known. Table V shows the popularity of each of the input formats and Table VI shows the popularity of different network architectures.

Point Cloud: Point clouds are the most popular data format, especially with the advancements in networks that can learn directly from point clouds, such as PointNet [175]. In addition, the point cloud representation allows researchers to easily fuse data from multiple viewpoints if the relative camera poses are known.

For point cloud based data, researchers commonly used PointNet [175] and PointNet++ [114]. This network backbone type has been used in different methods including direct regression [87, 89, 91–93, 99, 100], RL [132], sampling [43, 44, 49, 56, 61], and shape completion [143].

Point cloud for direct regression tends to subsample points and learn grasps for them [57, 64, 87, 89–91, 93, 97, 100, 103–105]. When a point cloud is used for sampling, one of the following procedures is usually followed. The points inside the gripper are fed to the network [44, 66, 67], the points will be transformed such that the origin aligns with the grasp frame [66, 68], or a representation of the gripper will be rendered, and all of the points will be fed into

TABLE VI: Number of times each network backbone has been used. Some papers use multiple backbones, depending on their network architecture.

Network Backbone	Number of times used
PointNet/PointNet++ [114]	21
ResNet [176]	9
LeNet [179]	3
Other	20
Custom	31

the network [43]. Minimal surveyed work has investigated the use of point clouds in RL apart from [132].

Images: The next most popular formats are depth and RGB-D images. Learning from images has been highly studied in both computer vision and robotics, allowing researchers to train pre-existing architectures for robotic grasp synthesis.

For images, the most common approach is to employ ResNet [176] as the backbone [58, 60, 88, 95, 101, 102, 127, 136, 177]. In addition to ResNet, various other architectures have been used, including VGG [178] LeNet [179], DenseNet [180] and U-Net [181].

Direct regression approaches commonly use the whole image as input [84, 101, 102]. For sampling, some approaches generated a depth image slice representing points within the gripper [39, 40] while others concatenated grasp features after processing images through a CNN [35, 45, 58, 60]. RL approaches input the current camera view to generate an action [112, 120, 122, 127, 129, 131, 171].

Voxel Grids: A voxel represents a value in a regular 3D grid. Voxel Grids are analogous to images in 3D space, where a voxel is similar to a pixel over a 3D grid instead of a 2D image.

For voxel based inputs, VoxNet [182] is most commonly used. VoxNet integrates a volumetric occupancy grid representation of the data with a 3D-CNN. Voxel grids have been used for shape completion [137, 140], sampling [52, 54, 70] and direct regression [85, 86, 94]. For sampling, Voxel grids have been used similarly to point clouds, transforming the origin to align with grasp frame coordinates [52] or adding grasp configuration features after the convolutional layers [46, 54, 70]. Voxel grids have also been used to directly regress a single pose using a subset of the voxel grid corresponding to the object [85, 86] as well as directly regressing a grasp for each voxel [94].

VI. BENCHMARKING

A. Experimental Evaluation

Most approaches implemented their grasping method in the realworld, either as an evaluation of the performance with respect to the metrics described in Section Section VI-C or presenting a demonstration that acts as a proof-of-concept of their method. Demonstrations focus on showing that their system works in the real-world without systematically evaluating their method. However, some approaches only consider evaluating their system using a simulator.

Evaluating in the real-world naturally carries more weight since the goal for robotic grasping is to be ultimately applied to real world. It should be noted that even though most works train their models using purely simulated data, they evaluate their approach in real world experiments. Commonly, the grasping system was directly transferred from simulation to real-world (e.g [39, 40, 43, 91, 92, 137]). TABLE VII: The popularity of various types of hardware systems used throughout the surveyed papers. The most popular category of robotic arms is further analyzed in Table VIII.

Category	Popularity
Robot Arm	66
Humanoid	6
Mobile Arm	3
Not Used	7

TABLE VIII: Popularity of various robotic arm manufacturers throughout the surveyed papers.

Manfacturer	Robot	Popularity	DoF
Franka Emika	Panda	15	7
Universal Robotics	UR5	9	6
Universal Robotics	UR10	4	6
Universal Robotics	UR5e	1	6
Kuka	LBR4	6	7
Kuka	IIWA LBR	3	7
Kinova	Gen3	2	7
Kinova	Mico2	2	6
Kinova	Jaco2	2	6
Rethink Robotics	Baxter	6	7
Staubli	TX60	4	6
ABB	Yumi	2	7
	Custom	2	
	Other	22	
	Not Used	7	

TABLE IX: The popularity of various grippers in the surveyed works, grouped by manufacturer.

Manufacturer	Model	Popularity	DoF
Robotiq	2F	10	1
Robotiq	3F	5	1
Franka Emika	Panda Gripper	12	1
Barret	BarretHand	7	5
Kinova	2F	3	1
Kinova	3F	1	2
Wonik Robotics	Allegro Hand	4	16
Rethink Robotics	Baxter Gripper	4	1
Shadow Robot Company	Shadow Hand	2	20
	Custom	2	
	Other	32	

Techniques such as domain adaptation [183] or domain randomization [184, 185] have also been used to transfer grasping approaches between simulation and real-world. Zhu et al. [102] used contrastive learning [186] to extract invariant features when images are augmented, aiming to improve the model under image sensor noise.

B. Hardware

We found that most works that study 6-DoF grasping use a robotic arm as shown in Table VII. However, a minority of works use other platforms, such as mobile platforms or humanoid robots. The robotic arms were most commonly table-mounted, and the robot would perform table-top grasping. Humanoid robots were commonly used when studying papers relating to affordances [84, 156, 157]. Surveyed works also completed research using mobile robots for grasping, however, none of the works made use of the extra DoF unique to the mobile aspect of the robot.

Table VIII details the frequency each robotic arm is used. Most researchers used either industrial robots or robots designed for human-robot interaction. The most common robotic platform is the Franka Emika Panda. This robot has a redundant DoF which allows more freedom in joint angles when achieving a specific gripper pose.

The most commonly used gripper was a two-finger parallel jaw (51 times) with researchers using the gripper that came with

TABLE X: The frequency of common metrics employed in the reviewed works as a percentage. Success rate is the most commonly used metric throughout all reviewed papers.

Metric	Frequency Used (%)
Success Rate	86
Completion	25
Computation Time	21
Precision	12
Coverage	6
Grasp Prediction Accuracy	8

the robotic arm. Some researchers use grippers with more than two fingers, including the Barret Hand, Allegro Hand, Shadow Hand and, Kinova 3-Finger gripper. Some papers only make use of two-fingers from the Barret hand[112] or switch between two- and three-fingers[90]. Multi-fingered high-DoF hands present a different set of challenges compared to a two-finger gripper. Research on multi-fingered grippers focus on how to generate a grasp pose that considers the large amount of DoF (e.g [54, 66, 86, 95]), or grasping with affordances (e.g [126, 156, 157]). Soft grippers are another potentially interesting line of work with only a few reviewed works making use of them [47, 85, 171].

C. Performance Metrics

A diverse set of performance metrics is used among all these works. The common definitions of performance metrics related to grasping from the reviewed work are listed below and Table X shows the frequency each metric is used. However, the exact definition of each performance metric can vary slightly between different works.

- Grasp Success Rate: The percentage of successful grasps (No. of Successful Grasps / Total No. of Attempted Grasps). The post-grasp steps prior to considering a grasp 'successful' was not consistent across reviewed works.
- Completion / Clearance Rate: The percentage of objects that are removed from the clutter (No. of Objects Grasped / Total No. of Objects in Clutter).
- Computation Time: Time required to compute grasp hypothesis generation.
- 4) **Precision**: The percentage of true positive grasp predictions (No. of True Positive Grasp Predictions / (No. of Selected Positive Grasps)).
- 5) **Coverage**: The percentage of sampled ground truth grasps that are within a threshold distance of any of the generated grasps.
- 6) **Grasp Prediction Accuracy**: The percentage of grasps outcomes correctly predicted (No. of Successful Grasp Predictions / No. of Predictions).

D. Object Configurations

Object configurations are how the objects are arranged in the scene during training or testing. We cluster object configurations into three types.

- 1) **Singulated**: A single object in the scene.
- Piled Clutter: Objects are packed together tightly. Objects are commonly arranged as a pile, for example Fig. 13 (Left).
- 3) **Structured Clutter**: Multiple objects spread out in a scene such that they are not touching, for example Fig. 13 (Right).

Piled clutter is frequently encountered in bin picking applications where a variety of items are compactly arranged in a bin. However, this survey only encompassed a limited number of papers which address bin picking scenarios [39, 40, 59, 103, 127, 129]. In



Fig. 13: Examples of cluttered scenes. We differentiate Piled Clutter (left) from Structured clutter (right). Left image is from [40] (©2017 SAGE) and right image is from [49]. (©2020 IEEE)

bin-picking scenarios, additional constraints are needed for how to approach objects to avoid the bin's walls, making it a more challenging task than table-top scenarios.

Most works surveyed in this study which focus on singulated objects tend to be grasping with a high DOF hand [35, 37, 46, 50, 54, 57, 58, 70, 84, 86, 95, 105, 137, 144]. These papers tend to focus on solving a specific grasping task, rather then aiming for generalized grasping. Other works on singulated objects focus on data representation and the learning process [41, 92, 132, 136]. Singulated object grasping is also used in other contexts such as affordances [17, 56, 126] or manipulation [138]. Researchers generally do not distinguish between structured or piled clutter scenes, however, we consider these distinct scenarios, which may have different solutions.

VII. DISCUSSION AND FUTURE DIRECTIONS

This section discusses the state of the field of deep learning-based grasp synthesis and highlights recommendations for future research directions. We note the key takeaways of each subsection indicated as **Key Takeaways**.

A. Deep Learning Methods

There is currently no consensus on when to use direct regressionbased methods versus sampling-based methods, as both approaches focus on similar tasks with similar success rates.

Direct regression-based approaches have demonstrated the ability to run in real-time [94, 99], but processing speeds can vary widely depending on the chosen architecture. In contrast, sampling-based methods offer the advantage of adjustable processing speed, as the number of samples or the level of optimization can be easily modified. However, based on our review of the literature, no sampling approach has been found to have real-time capabilities.

Depending on the use case of a sampling system, different sampling methods have different advantages. Euclidean sampling tends to quickly generate a wide array of grasps of varying quality. However, a large number of grasps are needed to try to find a high-quality grasp. In contrast, sampling through priors or latent space [43] tends to sample higher quality grasps, however, can come at a cost of time to create the samples.

RL approaches for grasping typically focus not only on grasp generation but also on the trajectory used by the system. However, RL approaches typically generate only one grasp per scene, making it difficult to incorporate additional constraints during execution.

Exemplar methods, which rely on a knowledge database of objects with similar shapes, were the least commonly used approach in our literature review. This is likely because they are limited in their ability to effectively grasp objects that differ significantly from those in the database. Grasping in complex environments has had recent attention in adjacent fields such as agriculture[187], whereas in pure grasping literature, authors typically focus on grasping objects from tabletop scenes where most of the synthesized grasp poses are kinematically feasible. Studying 6-DoF grasping in enclosed spaces such as shelves [188], around obstacles [52, 59], around humans [64], near reachability limits [52] or in-the-wild (e.g orchards [189]) would pose more constraints on the synthesis of the grasp configurations and require the grasp poses to not only be of high quality but also diverse, to increase the probability of finding a feasible trajectory to reach those grasp poses. As 6-DoF grasping approaches are not often tested in such challenging environments or compared directly to 4-DoF ones, the question of whether to use a 6 DoF or a 4-DoF grasping approach for a given application is not adequately answered by the current state-of-the-art.

Our review also found that few research papers provide a ready-to-use implementation of their work, making it easier for other researchers to benchmark their algorithms against. On the other hand, the ones that offer an implementation are commonly used by others as part of their evaluation. For instance, work by ten Pas et al. [40] is used by many others as a benchmark, likely because it is one of the earliest works that provide an open-source implementation wrapped in a ROS package.

Key Takeaways:

- 6-DoF grasping should be studied in more varied environments rather than just tabletop scenarios.
- Researchers should make their algorithms publicly available, ideally in a ready-to-use format (e.g. as a ROS package) to allow for informative benchmarking even on tabletop scenarios.

C. Performance Metrics

Bohg et al. [3] highlighted in 2014 that the grasping community has not yet embraced a consistent set of performance metrics. This observation is still valid today. This can be partly attributed to the large variety of objects, robots, end-effectors, and scenarios used in grasping research. Moreover, there is a divide between subsystem metrics and task-level metrics [190].

There are, however, some common performance metrics that the community has been using. The most popular metric is the grasp success rate, even though the exact definition varies in the literature. Some deem a grasp successful if the object is still held by the gripper after the robot returns to a configuration that is a certain height above the table [37, 92]. Other works impose additional constraints to the success definition, such as if the object is still in the gripper after taking a sequence of actions intended to test the robustness of the grasp [51].

Although success rate is a useful metric, an underutilized type of performance metric is one that measures the time efficacy of a grasping approach. For instance, if the task is to remove objects from a container efficiently, is a slower robot with a 95% success rate preferred to a faster robot with only 75% success rate? The answer depends on the task, since faster grasping approach might be preferred for a task where dropping an object is not detrimental. One metric that could be used for this purpose is Mean Picks Per Hour (MPPH), which is defined as the average number of successful grasps completed in an hour. Researchers should be reminded, however, that a time-based metric such as MPPH does

not only measure the computational efficiency of grasp synthesis but the system as a whole, and will be affected by the robot hardware, trajectory efficiency as well as the compute resources.

The *combination* of employing both a time-based metric and success rate could provide a more comprehensive evaluation of a pure grasping system. This is due to the likely inverse nature of the two metrics, where a system with a high grasp success rate may be very computationally expensive and slow-moving. However, a very fast system may be able to perform grasps and subsequent motions quickly, however likely decreasing the grasp success, which may be undesirable when grasping more fragile objects. As such, considering both of these metrics should provide a more detailed evaluation of the grasping system.

Key Takeaways:

- In addition to reporting the grasp success rate, researchers should consider reporting a more strict definition of the grasp success that tests the robustness of the grasp, similar to Lundell et al. [51].
- We suggest wider adoption of a time-based performance metric, such as the Mean Picks Per Hour (MPPH).

D. Object Sets and Grasping Datasets

Many papers use custom object sets, consisting of daily objects that the authors could find around the lab, making it difficult to compare different grasping approaches. Standardized object sets are very useful in enabling head-to-head comparisons of grasp synthesis algorithms. The most useful object sets for grasping research are those which cover a variety of objects and are easily accessible by the research community. Furthermore, object sets should have accurate 3D models that enable simulation studies. The most commonly used object set currently by today's grasping community is the YCB object set [159]. However, since this may change in the future, researchers should keep an eye on which object sets the community is using.

Models for grasp synthesis are trained on datasets where each data point contains the sensor data, the grasp pose executed by the robot, along with a ground truth label depending whether the grasp was successful or not. The majority of the reviewed works were trained on simulated datasets which offer large-scale data collection. However, there is a lack of realism due to the physics not being modeled accurately [191]. There is a lack of real-world grasping datasets, which tend to offer higher-quality data, are resource intensive to create compared to simulated datasets. While we believe that a comprehensive real-world dataset could be useful for the community analogous to the commonly used Cornell Grasping dataset [109] for 4-DoF top-down grasping, however, most research groups typically lack the time and resources needed to create large-scale real-world grasping datasets, with notable exceptions [117]. A compromise between real and simulated datasets are the hybrid datasets, for instance, Fang et al. [91] whom provided real-world point cloud data, however the grasp hypotheses are evaluated analytically rather than with actual trial-and-error.

Key Takeaways:

- Grasping researchers should adopt one of the widely used object sets. YCB object set [159] is the most commonly used one today, however, in the future, new object sets might find a wider adoption.
- The grasping community would benefit from the introduction of new object datasets that offer variety or specialization in

different aspects such as object geometry, grasping difficulty [167] or deformation.

• We propose authors release the code used to generate the dataset in addition to the dataset itself. This will allow other authors to make changes to the data generation procedure.

E. Trajectory Planning

Motion planning is widely utilized to find collision-free trajectories to execute a chosen grasp hypothesis. However, planning for a collision-free trajectory can be too conservative for densely cluttered scene, for example in cabinets and shelves where it may not possible to grasp the object of interest without nudging the neighboring objects. It has been shown that humans make deliberate contact with the environment during manipulation rather than carefully avoiding it[192]. This strategy exploits how the environment provides physical constraints on how the object and hand move and therefore provides a funnel for uncertainty due to noise in perception and control. This principle has also been shown to work well for robots [192–199]. For example, fixturing is a widely used practice in industry for various application such as machining, assembly and inspection. This suggests that the objective of having collision-free trajectories to acquire a grasp must be relaxed.

Key Takeaway: We encourage the study of motion planning algorithms for tight spaces where it is not possible to reach the targets without nudging other objects or where contact with the environment can be leveraged for more robust grasping.

F. Sensor Modalities

While most of the reviewed papers rely solely on vision for perceiving the world, other sensing modalities have been shown to be very helpful for both manipulation [200] and grasping [201]. Future research should explore how these techniques can be applied to improve the success of 6-DoF grasping. For example, tactile sensors can be used to predict if the object will remain grasped before it is actually lifted [202–204], detect object slip [205], account for uncertainty in object pose [124] and reconstruction of object geometry [46, 206]. Force/Torque sensing is another common sensor modality utilized in robotics there are also other less common modalities include robotic skins [207] and sound [208, 209].

Key Takeaway: Many of the reviewed works are vision-only sensing and considering only a single point-of-view. More research is needed to complement vision with other senses such as touch or hearing.

G. Grippers, hands and beyond

Most of the works in this survey focus on the use of simple two-fingered or three-fingered grippers, with few consider anthropomorphic hands. Simpler grippers greatly reduce the complexity of computing many-Dofs grasp hypotheses, and, most importantly, affordable commercial grippers allow researchers to setup systems for experimentally validating their results. However, this biases the research towards problems suitable for these simpler grippers and makes research on some other problems like grasping with anthropomorphic hand, dexterous, in-hand manipulation or soft-hand grasping fall behind. These directions may become important for enabling robots to go beyond pick and place tasks.

Key Takeaway: Robots have been employed and productized in manufacturing for decades [210–212], with further improvements in manipulation capabilities these robots will be able to perform more complex tasks in more unstructured environments, like the

assembly of unseen parts or on-site construction. Building robust and dexterous hands will result in automating many jobs currently done by humans. The opportunities are immense: new ways of designing hands, new materials, new sensors and ways of actuation, easy exchangeable fingers, redundancy. But it is not only about the hands - making sure the arm can carry the hand, that the hand is a natural extension of the arm, hand and arms in interaction- going beyond one and compensating some degrees of freedom in the hand by skillful dual-arm interaction are just some of the avenues to follow.

H. Grasping as part of a process

All day long, our hands and fingers, touch, push, pull, and enclose objects. We do this with rigid, articulated or deformable objects. We do this in the air, while objects are standing on rigid or soft surfaces or while they are moving. We do this while they are in water or covered in oil, with or without seeing - like finding a key in a pocket. Despite many decades of research and development in the area, human grasping skills are still quite superior to any of the artificial systems so far demonstrated.

In this discussion section, we have identified several areas where further contributions are needed. However, a broader view is necessary to address some if not all the manipulation scenarios mentioned above and, more importantly, consider grasping as part of an entire manipulation process. Specifically, identifying one specific 6DoF grasp pose for an object may not be sufficient to achieve success in these scenarios. One avenue may be to re-think grasp parameterizations to extend them to trajectories defined in joint and/or contact space. This could include exploiting environmental constraints [192] as well and adapting wrist trajectories based on the morphology of the hand itself [213]. New ideas and techniques for online learning and recovering from failure are required. Only new ideas and techniques in these and more areas will push advances that will make robotics systems successfully perform tasks we have been promising for some time such as folding clothing, preparing food, dressing humans. Pushing for and studying grasping beyond pure grasping, is one of the most important challenge we have as a community.

In summary, we have conducted a systematic review of state-of-the-art works which use deep learning-based towards achieving 6-DoF grasping. From this review, we synthesize 10 key takeaways which we believe could enable further progress in this rapidly progressing field.

VIII. ACKNOWLEDGMENTS

We would like to thank Wesley Chan for the helpful discussion at the ideation stages of the work. We would like to thank Lily Tran for helping with table creating and editing.

REFERENCES

- H. Asada, "Studies on prehension and handling by robot hands with elastic fingers," Kyoto University, Tech. Rep., 1979.
- [2] A. Bicchi and V. Kumar, "Robotic grasping and contact: a review," in *IEEE International Conference on Robotics and Automation*, vol. 1, 2000, pp. 348–353.
- [3] J. Bohg, A. Morales, T. Asfour, and D. Kragic, "Data-driven grasp synthesis—a survey," *IEEE Transactions on Robotics*, vol. 30, no. 2, p. 289–309, Apr 2014.
- [4] C. Eppner, S. Höfer, R. Jonschkowski, R. Martín-Martín, A. Sieverling, V. Wall, and O. Brock, "Lessons from the amazon picking challenge: Four aspects of building robotic systems," in *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, ser. IJCAI'17. AAAI Press, 2017, p. 4831–4835.

- [6] M. G. Catalano, G. Grioli, E. Farnioli, A. Serio, M. Bonilla, M. Garabini, C. Piazza, M. Gabiccini, and A. Bicchi, *From Soft to Adaptive Synergies: The Pisa/IIT SoftHand*. Cham: Springer International Publishing, 2016, pp. 101–125. [Online]. Available: https://doi.org/10.1007/978-3-319-26706-7_8
- [7] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.
- [8] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, "Learning ambidextrous robot grasping policies," *Science Robotics*, vol. 4, no. 26, p. eaau4984, 2019.
- [9] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray *et al.*, "Learning dexterous in-hand manipulation," *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020.
- [10] A. Zeng, S. Song, K.-T. Yu, E. Donlon, F. R. Hogan, M. Bauza, D. Ma, O. Taylor, M. Liu, E. Romo *et al.*, "Robotic pick-and-place of novel objects in clutter with multi-affordance grasping and cross-domain image matching," *The International Journal of Robotics Research*, vol. 41, no. 7, pp. 690–705, 2022.
- [11] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich, "Rt-1: Robotics transformer for real-world control at scale," in *arXiv preprint arXiv:2212.06817*, 2022.
- [12] C. Lynch, A. Wahid, J. Tompson, T. Ding, J. Betker, R. Baruch, T. Armstrong, and P. Florence, "Interactive language: Talking to robots in real time," *arXiv* preprint arXiv:2210.06407, 2022.
- [13] K. Kleeberger, R. Bormann, W. Kraus, and M. F. Huber, "A survey on learning-based robotic grasping," *Current Robotics Reports*, vol. 1, no. 4, pp. 239–249, 2020.
- [14] O. Kroemer, S. Niekum, and G. D. Konidaris, "A review of robot learning for manipulation: Challenges, representations, and algorithms," *J. Mach. Learn. Res.*, vol. 22, pp. 30:1–30:82, 2021.
- [15] H. Duan, P. Wang, Y. Huang, G. Xu, W. Wei, and X. Shen, "Robotics dexterous grasping: The methods based on point cloud and deep learning," *Frontiers in Neurorobotics*, vol. 15, p. 73, 2021.
- [16] J. J. Gibson, "The theory of affordances," in *Perceiving, acting, and knowing: toward an ecological psychology.* Hillsdale, N.J. : Lawrence Erlbaum Associates, 1977, pp. pp.67–82. [Online]. Available: https://hal.archives-ouvertes.fr/hal-00692033
- [17] P. Ardón, É. Pairet, R. Petrick, S. Ramamoorthy, and K. Lohan, "Learning grasp affordance reasoning through semantic relations," *IEEE Robotics and Automation Letters*, vol. PP, pp. 1–1, 08 2019.
- [18] A. Morales, T. Asfour, P. Azad, S. Knoop, and R. R. Dillmann, "Integrated grasp planning and visual object localization for a humanoid robot with five-fingered hands," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Beijing, China: Ieee, 2006, pp. 5663–5668.
- [19] V.-D. Nguyen, "Constructing force-closure grasps in 3d," in *IEEE Interna*tional Conference on Robotics and Automation, vol. 4, 1987, pp. 240–245.
- [20] K. Mirza and D. Grin, "General formulation for force distribution in power grasp," in *IEEE International Conference on Robotics and Automation*, 1994, pp. 880–887 vol.1.
- [21] I.-M. Chen and J. Burdick, "Finding antipodal point grasps on irregularly shaped objects," *IEEE Transactions on Robotics and Automation*, vol. 9, no. 4, pp. 507–512, 1993.
- [22] D. Prattichizzo and J. C. Trinkle, "Grasping," in Springer handbook of robotics. Springer, 2016, pp. 955–988.
- [23] R. Murray, Z. Li, and S. Sastry, A Mathematical Introduction to Robotic Manipulation. Taylor & Francis, 1994.
- [24] A. Miller and P. Allen, "Graspit! a versatile simulator for robotic grasping," *IEEE Robotics Automation Magazine*, vol. 11, no. 4, pp. 110–122, 2004.
- [25] S. El-Khoury, A. Sahbani, and V. Perdereau, "Learning the natural grasping component of an unknown object," in *IEEE/RSJ International Conference* on Intelligent Robots and Systems, 2007, pp. 2957–2962.
- [26] F. Kyota, T. Watabe, S. Saito, and M. Nakajima, "Detection and evaluation of grasping positions for autonomous agents," in *International Conference* on Cyberworlds, 2005.
- [27] C. Michel, V. Perdereau, and M. Drouin, "An approach to extract natural grasping axes with a real 3d vision system," in *IEEE International Symposium* on Industrial Electronics, vol. 4, 2006, pp. 3130–3135.
- [28] I. Kamon, T. Flash, and S. Edelman, "Learning to grasp using visual

information," in Proceedings of IEEE International Conference on Robotics and Automation, vol. 3, 1996.

- [29] Y. Jiang, S. Moseson, and A. Saxena, "Efficient grasping from rgbd images: Learning using a new rectangle representation," in *IEEE International Conference on Robotics and Automation*, 2011.
- [30] A. Saxena, J. Driemeyer, J. Kearns, C. Osondu, and A. Y. Ng, "Learning to grasp novel objects using vision," in *Experimental Robotics*. Springer, 2008, pp. 33–42.
- [31] A. Saxena, J. Driemeyer, and A. Y. Ng, "Robotic grasping of novel objects using vision," *The International Journal of Robotics Research*, vol. 27, no. 2, pp. 157–173, 2008.
- [32] A. Saxena, L. Wong, and A. Ng, "Learning grasp strategies with partial shape information." in *Proceedings of the National Conference on Artificial Intelligence*, vol. 3, 2008, pp. 1491–1494.
- [33] A. Morales, E. Chinellato, P. Sanz, A. P. del Pobil, and A. Fagg, "Learning to predict grasp reliability for a multifinger robot hand by using visual features," 2004.
- [34] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information* processing systems, vol. 25, pp. 1097–1105, 2012.
- [35] U. R. Aktaş, C. Zhao, M. Kopicki, and J. L. Wyatt, "Deep dexterous grasping of novel objects from a single view," *International Journal of Humanoid Robotics*, vol. 19, no. 02, p. 2250011, 2022. [Online]. Available: https://doi.org/10.1142/S0219843622500116
- [36] J. Mahler, F. T. Pokorny, B. Hou, M. Roderick, M. Laskey, M. Aubry, K. Kohlhoff, T. Kröger, J. Kuffner, and K. Goldberg, "Dex-net 1.0: A cloudbased network of 3d objects for robust grasp planning using a multi-armed bandit model with correlated rewards," in 2016 ICRA, 2016, pp. 1957–1964.
- [37] J. Varley, J. Weisz, J. Weiss, and P. Allen, "Generating multi-fingered robotic grasps via deep learning," in *IROS*, 2015, pp. 4415–4420.
- [38] D. Kappler, J. Bohg, and S. Schaal, "Leveraging big data for grasp planning," in *ICRA*, 2015, pp. 4304–4311.
- [39] M. Gualtieri, A. ten Pas, K. Saenko, and R. Platt, "High precision grasp pose detection in dense clutter," in *IROS*, 2016, pp. 598–605.
- [40] A. ten Pas, M. Gualtieri, K. Saenko, and R. Platt, "Grasp pose detection in point clouds," *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1455–1473, 2017.
- [41] Y. Zhou and K. Hauser, "6dof grasp planning by optimizing a deep learning scoring function," in *Robotics: Science and Systems Workshop on Revisiting Contact-Turning a Problem into a Solution*, 2017.
- [42] X. Yan, J. Hsu, M. Khansari, Y. Bai, A. Pathak, A. Gupta, J. Davidson, and H. Lee, "Learning 6-dof grasping interaction via deep geometry-aware 3d representations," in *ICRA*, 2018, pp. 3766–3773.
- [43] A. Mousavian, C. Eppner, and D. Fox, "6-dof graspnet: Variational grasp generation for object manipulation," in *IEEE/CVF International Conference* on Computer Vision, 2019, pp. 2901–2910.
- [44] H. Liang, X. Ma, S. Li, M. Gorner, S. Tang, B. Fang, F. Sun, and J. Zhang, "Pointnetgpd: Detecting grasp configurations from point sets," *ICRA*, 2019.
- [45] Q. Lu, K. Chenna, B. Sundaralingam, and T. Hermans, "Planning multifingered grasps as probabilistic inference in a learned deep network," in *Robotics Research*. Springer, 2019, pp. 455–472.
- [46] S. Ottenhaus, D. Renninghoff, R. Grimm, F. Ferreira, and T. Asfour, "Visuo-haptic grasping of unknown objects based on gaussian process implicit surfaces and deep learning," in *IEEE-RAS International Conference on Humanoid Robots*, 2019, pp. 402–409.
- [47] J. Gonçalves and P. Lima, "Grasp planning with incomplete knowledge about the object to be grasped," in *IEEE International Conference on Autonomous Robot Systems and Competitions*, 2019, pp. 1–6.
- [48] M. A. c. Riedlinger, M. Voelk, K. Kleeberger, M. U. Khalid, and R. Bormann, "Model-free grasp learning framework based on physical simulation," in *International Symposium on Robotics*, 2020, pp. 1–8.
- [49] A. Murali, A. Mousavian, C. Eppner, C. Paxton, and D. Fox, "6-dof grasping for target-driven object manipulation in clutter," in *ICRA*, 2020.
- [50] M. Van der Merwe, Q. Lu, B. Sundaralingam, M. Matak, and T. Hermans, "Learning continuous 3d reconstructions for geometrically aware grasping," in *ICRA*, 2020, pp. 11 516–11 522.
- [51] J. Lundell, F. Verdoja, and V. Kyrki, "Beyond top-grasps through scene completion," in *ICRA*, 2020, pp. 545–551.
- [52] X. Lou, Y. Yang, and C. Choi, "Learning to generate 6-dof grasp poses with reachability awareness," in 2020 ICRA, 2020, pp. 1532–1538.
- [53] Y. Choi, H. Kee, K. Lee, J. Choy, J. Min, S. Lee, and S. Oh, "Hierarchical 6-dof grasping with approaching direction selection," in *ICRA*, 2020, pp. 1553–1559.
- [54] Q. Lu, M. Van der Merwe, B. Sundaralingam, and T. Hermans, "Multifingered grasp planning via inference in deep neural networks: Outperforming sampling by learning differentiable models," *IEEE Robotics & Automation Magazine*, vol. 27, no. 2, pp. 55–65, 2020.
- [55] H. Schaub and A. Schöttl, "6-dof grasp detection for unknown objects," in

International Conference on Advanced Computer Information Technologies (ACIT), 2020, pp. 400–403.

- [56] A. Murali, W. Liu, K. Marino, S. Chernova, and A. Gupta, "Same object, different grasps: Data and semantic knowledge for task-oriented grasping," in *Conference on Robot Learning*, 2020.
- [57] M. Kokic, D. Kragic, and J. Bohg, "Learning task-oriented grasping from human activity datasets," *IEEE Robotics and Automation Letters*, vol. PP, pp. 1–1, 02 2020.
- [58] J. Lundell, E. Corona, T. Nguyen Le, F. Verdoja, P. Weinzaepfel, G. Rogez, F. Moreno-Noguer, and V. Kyrki, "Multi-fingan: Generative coarse-to-fine sampling of multi-finger grasps," in 2021 ICRA, 2021, pp. 4495–4501.
- [59] X. Lou, Y. Yang, and C. Choi, "Collision-aware target-driven object grasping in constrained environments," in 2021 ICRA, 2021, pp. 6364–6370.
- [60] J. Lundell, F. Verdoja, and V. Kyrki, "DDGC: Generative Deep Dexterous Grasping in Clutter," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 6899–6906, Oct. 2021.
- [61] G. Peng, Z. Ren, H. Wang, X. Li, and M. O. Khyam, "A self-supervised learning-based 6-dof grasp planning method for manipulator," *IEEE Transactions on Automation Science and Engineering*, 2021.
- [62] Z. Jiang, Y. Zhu, M. Svetlik, K. Fang, and Y. Zhu, "Synergies between affordance and geometry: 6-dof grasp detection via implicit representations," *Robotics: science and systems*, 2021.
- [63] H. Kasaei and M. Kasaei, "Mvgrasp: Real-time multi-view 3d object grasping in highly cluttered environments," *Robotics and Autonomous Systems*, vol. 160, p. 104313, 2023. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0921889022002020
- [64] X. Wang, S. Nisar, and F. Matsuno, "Robust grasp detection with incomplete point cloud and complex background," *Advanced Robotics*, vol. 0, no. 0, pp. 1–16, 2021.
- [65] M. R. Munoz, "Grasping in 6dof: An orthographic approach to generalized grasp affordance predictions," *https://fse.studenttheses.ub.rug.nl/*, 2021.
- [66] M. Corsaro, S. Tellex, and G. Konidaris, "Learning to detect multi-modal grasps for dexterous grasping in dense clutter," in *IROS*, 2021, pp. 4647–4653.
- [67] Z. Ren, G. Peng, J. Yang, and H. Wang, "6-dof grasp planning of manipulator combined with self-supervised learning," in *Chinese Control and Decision Conference (CCDC)*, 2021, pp. 3026–3032.
- [68] B. Wen, W. Lian, K. Bekris, and S. Schaal, "Catgrasp: Learning category-level task-relevant grasping in clutter from simulation," in 2022 International Conference on Robotics and Automation (ICRA), 2022, pp. 6401–6408.
- [69] J. Mahler, M. Matl, X. Liu, A. Li, D. Gealy, and K. Goldberg, "Dex-net 3.0: Computing robust vacuum suction grasp targets in point clouds using a new analytic model and deep learning," in *ICRA*, 2018, pp. 5620–5627.
- [70] Q. Lu, M. Van der Merwe, and T. Hermans, "Multi-fingered active grasp learning," in *IROS*, 2020, pp. 8415–8422.
- [71] A. Pas and R. Platt, Using Geometry to Detect Grasp Poses in 3D Point Clouds, 01 2018, pp. 307–324.
- [72] G. Smith, E. Lee, K. Goldberg, K. F. Böhringer, and J. J. Craig, "Computing parallel-jaw grips," *ICRA*, vol. 3, pp. 1897–1903 vol.3, 1999.
- [73] S. Agrawal and N. Goyal, "Analysis of thompson sampling for the multi-armed bandit problem," in *Conference on learning theory*. JMLR Workshop and Conference Proceedings, 2012, pp. 39–1.
- [74] N. Vahrenkamp, M. Kröhnert, S. Ulbrich, T. Asfour, G. Metta, R. Dillmann, and G. Sandini, *Simox: A Robotics Toolbox for Simulation, Motion and Grasp Planning*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 585–594.
- [75] M. Ciocarlie and P. Allen, "Hand posture subspaces for dexterous robotic grasping," I. J. Robotic Res., vol. 28, pp. 851–867, 06 2009.
- [76] M. S. Kopicki, D. Belter, and J. L. Wyatt, "Learning better generative models for dexterous, single-view grasping of novel objects," *The International Journal of Robotics Research*, vol. 38, no. 10-11, pp. 1246–1267, 2019. [Online]. Available: https://doi.org/10.1177/0278364919865338
- [77] D. P. Kingma and M. Welling, "Auto-encoding variational bayes," in 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings, Y. Bengio and Y. LeCun, Eds., 2014. [Online]. Available: http://arxiv.org/abs/1312.6114
- [78] C. M. Bishop, "Mixture density networks," Aston University, Tech. Rep., 1994.
- [79] D. Morrison, P. Corke, and J. Leitner, "Closing the Loop for Robotic Grasping: A Real-time, Generative Grasp Synthesis Approach," in *Proc. of Robotics: Science and Systems (RSS)*, 2018.
- [80] H. Schaub, A. Schöttl, and M. Hoh, "6-dof grasp detection for unknown objects using surface reconstruction," in 2021 3rd International Congress on Human-Computer Interaction, Optimization and Robotic Applications (HORA), 2021, pp. 1–6.
- [81] V. Satish, J. Mahler, and K. Goldberg, "On-policy dataset synthesis for learning robot grasping policies using fully convolutional deep networks," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1357–1364, 2019.

- [82] C. Ferrari and J. Canny, "Planning optimal grasps," in *Proceedings 1992 IEEE International Conference on Robotics and Automation*, 1992, pp. 2290–2295 vol.3.
- [83] M. Veres, M. Moussa, and G. W. Taylor, "Modeling grasp motor imagery through deep conditional generative models," *IEEE Robotics and Automation Letters*, vol. 2, no. 2, pp. 757–764, 2017.
- [84] P. Schmidt, N. Vahrenkamp, M. Wächter, and T. Asfour, "Grasping of unknown objects using deep convolutional neural networks based on depth images," in 2018 ICRA, 2018, pp. 6831–6838.
- [85] C. Choi, W. Schwarting, J. DelPreto, and D. Rus, "Learning object grasping for soft robot hands," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2370–2377, 2018.
- [86] M. Liu, Z. Pan, K. Xu, K. Ganguly, and D. Manocha, "Generating grasp poses for a high-dof gripper using neural networks," in *IROS*, 2019, pp. 1518–1525.
- [87] Y. Qin, R. Chen, H. Zhu, M. Song, J. Xu, and H. Su, "S4g: Amodal single-view single-shot se (3) grasp detection in cluttered scenes," in *Conference on Robot learning*. PMLR, 2020, pp. 53–65.
- [88] D. Yang, T. Tosun, B. Eisner, V. Isler, and D. Lee, "Robotic grasping through combined image-based grasp proposal and 3d reconstruction," in 2021 ICRA, 2021, pp. 6350–6356.
- [89] K.-Y. Jeng, Y.-C. Liu, Z. Y. Liu, J.-W. Wang, Y.-L. Chang, H.-T. Su, and W. Hsu, "Gdn: A coarse-to-fine (c2f) representation for end-to-end 6-dof grasp detection," in *Proceedings of the 2020 Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, vol. 155. PMLR, 16–18 Nov 2021, pp. 220–231. [Online]. Available: https://proceedings.mlr.press/v155/jeng21a.html
- [90] C.-H. Wang and P.-C. Lin, "Q-pointnet: Intelligent stacked-objects grasping using a rgbd sensor and a dexterous hand," in *IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)*. IEEE Press, 2020, p. 601–606.
- [91] H.-S. Fang, C. Wang, M. Gou, and C. Lu, "Graspnet-Ibillion: a large-scale benchmark for general object grasping," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 444–11 453.
- [92] C. Wu, J. Chen, Q. Cao, J. Zhang, Y. Tai, L. Sun, and K. Jia, "Grasp proposal networks: An end-to-end solution for visual learning of robotic grasps," in *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [93] P. Ni, W. Zhang, X. Zhu, and Q. Cao, "Pointnet++ grasping: Learning an end-to-end spatial grasp generation algorithm from sparse point clouds," in *ICRA*, 2020, pp. 3619–3625.
- [94] M. Breyer, J. J. Chung, L. Ott, S. Roland, and N. Juan, "Volumetric grasping network: Real-time 6 dof grasp detection in clutter," in *Conference on Robot Learning*, 2020.
- [95] M. Liu, Z. Pan, K. Xu, K. Ganguly, and D. Manocha, "Deep differentiable grasp planner for high-dof grippers," in *Robotics: Science and Systems (RSS)*, 2020.
- [96] L. Shao, F. Ferreira, M. Jorda, V. Nambiar, J. Luo, E. Solowjow, J. A. Ojea, O. Khatib, and J. Bohg, "Unigrasp: Learning a unified model to grasp with multifingered robotic hands," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2286–2293, 2020.
- [97] P. Ni, W. Zhang, X. Zhu, and Q. Cao, "Learning an end-to-end spatial grasp generation and refinement algorithm from simulation," *Machine Vision and Applications*, vol. 32, no. 1, pp. 1–12, 2021.
- [98] C. Wang, H.-S. Fang, M. Gou, H. Fang, J. Gao, and C. Lu, "Graspness discovery in clutters for fast and accurate grasp detection," in *Proceedings* of the IEEE/CVF International Conference on Computer Vision (ICCV), October 2021, pp. 15 964–15 973.
- [99] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox, "Contact-graspnet: Efficient 6-dof grasp generation in cluttered scenes," 2021.
- [100] B. Zhao, H. Zhang, X. Lan, H. Wang, Z. Tian, and N. Zheng, "Regnet: region-based grasp network for single-shot grasp detection in point clouds," *arXiv preprint arXiv:2002.12647*, 2020.
- [101] M. Gou, H.-S. Fang, Z. Zhu, S. Xu, C. Wang, and C. Lu, "RGB matters: Learning 7-dof grasp poses on monocular rgbd images," in 2021 ICRA, 2021, pp. 13 459–13 466.
- [102] X. Zhu, L. Sun, Y. Fan, and M. Tomizuka, "6-dof contrastive grasp proposal network," in 2021 ICRA, 2021, pp. 6371–6377.
- [103] W. Wei, Y. Luo, F. Li, G. Xu, J. Zhong, W. Li, and P. Wang, "Gpr: Grasp pose refinement network for cluttered scenes," in 2021 ICRA, 2021, pp. 4295–4302.
- [104] Y. Li, T. Kong, R. Chu, Y. Li, P. Wang, and L. Li, "Simultaneous semantic and collision learning for 6-dof grasp pose estimation," in 2021 IROS, 2021, pp. 3571–3578.
- [105] S. Li, Z. Li, K. Han, X. Li, Y. Xiong, and Z. Xie, "An end-to-end spatial grasp prediction model for humanoid multi-fingered hand using deep network," in 2021 6th International Conference on Control, Robotics and Cybernetics (CRC), 2021, pp. 130–136.
- [106] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *CVPR*, 06 2016, pp. 779–788.

- [107] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in Neural Information Processing Systems*, vol. 28. Curran Associates, Inc., 2015. [Online]. Available: https://proceedings.neurips.cc/paper/2015/file/ 14bfa6bb14875e45bba028a21ed38046-Paper.pdf
- [108] J. Redmon and A. Angelova, "Real-time grasp detection using convolutional neural networks," in *ICRA*, 2015, pp. 1316–1322.
- [109] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *International Journal of Robotics Research*, vol. 34, 01 2013.
- [110] F.-J. Chu, R. Xu, and P. A. Vela, "Real-world multiobject, multigrasp detection," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3355–3362, 2018.
- [111] S. Kumra and C. Kanan, "Robotic grasp detection using deep convolutional neural networks," in *IROS*, 2017, pp. 769–776.
- [112] B. Wu, I. Akinola, A. Gupta, F. Xu, J. Varley, D. Watkins-Valls, and P. K. Allen, "Generative attention learning: a "general" framework for high-performance multi-fingered grasping in clutter," *Autonomous Robots*, pp. 1–20, 2020.
- [113] K. Sohn, H. Lee, and X. Yan, "Learning structured output representation using deep conditional generative models," in *NIPS*, 2015.
- [114] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, ser. NIPS'17, 2017, p. 5105–5114.
- [115] Y. Li, "Deep reinforcement learning: An overview," 2017. [Online]. Available: https://arxiv.org/abs/1701.07274
- [116] M. Mohammed, L. Kwek, and S. C. Chua, "Review of deep reinforcement learning-based object grasping: Techniques, open challenges, and recommendations," *IEEE Access*, vol. 8, pp. 178 450–178 481, 10 2020.
- [117] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *The International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [118] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke, and S. Levine, "Qt-opt: Scalable deep reinforcement learning for vision-based robotic manipulation," in *Conference on Robot Learning (CoRL)*, 06 2018.
- [119] D. Quillen, E. Jang, O. Nachum, C. Finn, J. Ibarz, and S. Levine, "Deep reinforcement learning for vision-based robotic grasping: A simulated comparative evaluation of off-policy methods," in *ICRA*, 2018, pp. 6284–6291.
- [120] M. Gualtieri and R. Platt, "Learning 6-dof grasping and pick-place using attention focus," in *Conference on Robot Learning*. PMLR, 2018, pp. 477–486.
- [121] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski *et al.*, "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [122] B. Wu, I. Akinola, and P. K. Allen, "Pixel-attentive policy gradient for multi-fingered grasping in cluttered scenes," in *IROS*, 2019, pp. 1789–1796.
- [123] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," 2017.
- [124] H. Merzic, M. Bogdanovic, D. Kappler, L. Righetti, and J. Bohg, "Leveraging contact forces for learning to grasp," in 2019 IEEE International Conference on Robotics and Automation, 2019.
- [125] J. Schulman, S. Levine, P. Moritz, M. I. Jordan, and P. Abbeel, "Trust region policy optimization," 2015. [Online]. Available: https://arxiv.org/abs/1502.05477
- [126] P. Mandikal and K. Grauman, "Learning dexterous grasping with object-centric visual affordances," in *ICRA*, 2021.
- [127] S. Song, A. Zeng, J. Lee, and T. Funkhouser, "Grasping in the wild: Learning 6dof closed-loop grasping from low-cost demonstrations," *IEEE Robotics* and Automation Letters, vol. PP, pp. 1–1, 06 2020.
- [128] C. J. Watkins and P. Dayan, "Q-learning," *Machine learning*, vol. 8, no. 3-4, pp. 279–292, 1992.
- [129] L. Berscheid, C. Friedrich, and T. Kröger, "Robot learning of 6 dof grasping using model-based adaptive primitives," in 2021 ICRA, 2021, pp. 4474–4480.
- [130] D. Kawakami, R. Ishikawa, M. Roxas, Y. Sato, and T. Oishi, "Learning 6dof grasping using reward-consistent demonstration," 2021.
- [131] B. Tang, M. Corsaro, G. Konidaris, S. Nikolaidis, and S. Tellex, "Learning collaborative pushing and grasping policies in dense clutter," 2021.
- [132] L. Wang, Y. Xiang, W. Yang, A. Mousavian, and D. Fox, "Goalauxiliary actor-critic for 6d robotic grasping with point clouds," in *5th Annual Conference on Robot Learning*, 2021. [Online]. Available: https://openreview.net/forum?id=jOSWHddP1fZ
- [133] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," arXiv preprint arXiv:1509.02971, 2015.
- [134] X. Chen, Z. Ye, J. Sun, Y. Fan, F. Hu, C. Wang, and C. Lu, "Transferable active grasping and real embodied dataset," in 2020 ICRA, 2020, pp. 3611–3618.

- [135] S. Brahmbhatt, C. Ham, C. C. Kemp, and J. Hays, "Contactdb: Analyzing and predicting grasp contact via thermal imaging," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8709–8719.
- [136] T. Patten, K. Park, and M. Vincze, "Dgcm-net: Dense geometrical correspondence matching network for incremental experience-based robotic grasping," *Frontiers in Robotics and AI*, vol. 7, p. 120, 2020.
- [137] J. Varley, C. DeChant, A. Richardson, J. Ruales, and P. Allen, "Shape completion enabled robotic grasping," in *IROS*. IEEE, 2017, pp. 2442–2447.
- [138] W. Gao and R. Tedrake, "kpam-sc: Generalizable manipulation planning using keypoint affordance and shape completion," in 2021 ICRA, 2021, pp. 6527–6533.
- [139] X. Zhang, Z. Zhang, C. Zhang, J. Tenenbaum, B. Freeman, and J. Wu, "Learning to reconstruct shapes from unseen classes," in *Advances in Neural Information Processing Systems*, vol. 31. Curran Associates, Inc., 2018. [Online]. Available: https://proceedings.neurips.cc/paper/2018/file/ 208e43f0e45c4c78cafadb83d2888cb6-Paper.pdf
- [140] M. Kiatos, S. Malassiotis, and I. Sarantopoulos, "A geometric approach for grasping unknown objects with multifingered hands," *IEEE Transactions* on *Robotics*, pp. 1–12, 2020.
- [141] D. P. Kingma and M. Welling, "An introduction to variational autoencoders," *Foundations and Trends® in Machine Learning*, vol. 12, no. 4, p. 307–392, 2019.
- [142] N. Chavan-Dafle, S. Popovych, S. Agrawal, D. D. Lee, and V. Isler, "Simultaneous object reconstruction and grasp prediction using a cameracentric object shell representation," in 2022 IROS, 2022, pp. 1396–1403.
- [143] M. Gualtieri and R. Platt, "Robotic pick-and-place with uncertain object instance segmentation and shape completion," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 1753–1760, 2021.
- [144] J. Lundell, F. Verdoja, and V. Kyrki, "Robust grasp planning over uncertain shape completions," in *IROS*, 2019, pp. 1526–1532.
- [145] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *International Conference on Machine Learning*, 2016, pp. 1050–1059.
- [146] P. Vafaeikia, K. Namdar, and F. Khalvati, "A brief review of deep multi-task learning and auxiliary task learning," arXiv preprint arXiv:2007.01126, 2020.
- [147] Y. Avigal, S. Paradis, and H. Zhang, "6-dof grasp planning using fast 3d reconstruction and grasp quality cnn," arXiv preprint arXiv:2009.08618, 2020.
- [148] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg, "Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics," in *Robotics: Science and Systems* (*RSS*), 2017.
- [149] M. Björkman, Y. Bekiroglu, V. Hogman, and D. Kragic, "Enhancing visual perception of shape through tactile glances," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3180–3186, 2013.
- [150] Y. Shiraki, K. Nagata, N. Yamanobe, A. Nakamura, K. Harada, D. Sato, and D. N. Nenchev, "Modeling of everyday objects for semantic grasp," in *IEEE International Symposium on Robot and Human Interactive Communication*, 2014, pp. 750–755.
- [151] A. Miller, S. Knoop, H. Christensen, and P. Allen, "Automatic grasp planning using shape primitives," in *ICRA*, vol. 2, 2003, pp. 1824–1829 vol.2.
- [152] T. Torii and M. Hashimoto, "Model-less estimation method for robot grasping parameters using 3d shape primitive approximation," in *IEEE International Conference on Automation Science and Engineering (CASE)*, 2018, pp. 580–585.
- [153] P. Ardón, É. Pairet, K. S. Lohan, S. Ramamoorthy, and R. Petrick, "Affordances in robotic tasks-a survey," *arXiv preprint arXiv:2004.07400*, 2020.
- [154] Y. Li, L. Schomaker, and S. H. Kasaei, "Learning to grasp 3d objects using deep residual u-nets," in *IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2020, pp. 781–787.
- [155] T.-T. Do, A. Nguyen, and I. Reid, "Affordancenet: An end-to-end deep learning approach for object affordance detection," in *ICRA*, 2018, pp. 5882–5889.
- [156] A. Nguyen, D. Kanoulas, D. G. Caldwell, and N. G. Tsagarakis, "Detecting object affordances with convolutional neural networks," in *IROS*, Oct 2016, pp. 2765–2770.
- [157] A. Nguyen, D. Kanoulas, D. G. Caldwell, and N. G. Tsagarakis, "Objectbased affordances detection with convolutional neural networks and dense conditional random fields," in *IROS*, 2017, pp. 5908–5915.
- [158] L. Manuelli, W. Gao, P. Florence, and R. Tedrake, "Kpam: Keypoint affordances for category-level robotic manipulation," in *Robotics Research*, T. Asfour, E. Yoshida, J. Park, H. Christensen, and O. Khatib, Eds. Cham: Springer International Publishing, 2022, pp. 132–157.
- [159] B. Calli, A. Singh, A. Walsman, S. Srinivasa, P. Abbeel, and A. M. Dollar, "The ycb object and model set: Towards common benchmarks for manipulation research," in *International Conference on Advanced Robotics (ICAR)*, 2015, pp. 510–517.
- [160] A. Singh, J. Sha, K. S. Narayan, T. Achim, and P. Abbeel, "BigBIRD:

A large-scale 3D database of object instances," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2014, pp. 509–516.

- [161] 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," *The International Journal of Robotics Research*, vol. 31, no. 8, pp. 927–934, 2012.
- [162] A. X. Chang, T. Funkhouser, L. Guibas, P. Hanrahan, Q. Huang, Z. Li, S. Savarese, M. Savva, S. Song, H. Su, J. Xiao, L. Yi, and F. Yu, "ShapeNet: An Information-Rich 3D Model Repository," Stanford University — Princeton University — Toyota Technological Institute at Chicago, Tech. Rep. arXiv:1512.03012 [cs.GR], 2015.
- [163] W. Wohlkinger, A. Aldoma Buchaca, R. Rusu, and M. Vincze, "3DNet: Large-Scale Object Class Recognition from CAD Models," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2012.
- [164] P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser, "The princeton shape benchmark," in *Shape Modeling Applications*, 2004, pp. 167–388.
- [165] Zhirong Wu, S. Song, A. Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and J. Xiao, "3D ShapeNets: A deep representation for volumetric shapes," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 1912–1920.
- [166] Y. Xiang, W. Kim, W. Chen, J. Ji, C. Choy, H. Su, R. Mottaghi, L. Guibas, and S. Savarese, "Objectnet3d: A large scale database for 3d object recognition," in *European Conference on Computer Vision (ECCV)*. Springer, 2016, pp. 160–176.
- [167] D. Morrison, P. Corke, and J. Leitner, "Egad! an evolved grasping analysis dataset for diversity and reproducibility in robotic manipulation," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4368–4375, 2020.
- [168] 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.
- [169] D. Wang, D. Tseng, P. Li, Y. Jiang, M. Guo, M. Danielczuk, J. Mahler, J. Ichnowski, and K. Goldberg, "Adversarial grasp objects," in 2019 IEEE 15th International Conference on Automation Science and Engineering (CASE), 2019, pp. 241–248.
- [170] C. Eppner, A. Mousavian, and D. Fox, "Acronym: A large-scale grasp dataset based on simulation," in *ICRA*, 2021, pp. 6222–6227.
- [171] D. S. Diaz Cortes, G. Hwang, and K.-U. Kyung, "Imitation learning based soft robotic grasping control without precise estimation of target posture," in *International Conference on Soft Robotics*, 2021, pp. 149–154.
- [172] V. R. Osorio, R. Iyengar, X. Yao, P. Bhattachan, A. Ragobar, N. Dey, and B. Tripp, "37,000 human-planned robotic grasps with six degrees of freedom," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3346–3351, 2020.
- [173] L. Wang, Y. Xiang, and D. Fox, "Manipulation trajectory optimization with online grasp synthesis and selection," in *Robotics: Science and Systems (RSS)*, 2020.
- [174] O. Taheri, N. Ghorbani, M. J. Black, and D. Tzionas, "Grab: A dataset of whole-body human grasping of objects," in *Computer Vision – ECCV* 2020, vol. LNCS 12355. Cham: Springer International Publishing, Aug. 2020, pp. 581–600. [Online]. Available: https://grab.is.tue.mpg.de
- [175] R. Q. Charles, H. Su, M. Kaichun, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *CVPR*, 2017, pp. 77–85.
- [176] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in CVPR, 2016, pp. 770–778.
- [177] M. Bajracharya et al., "A mobile manipulation system for one-shot teaching of complex tasks in homes," in *ICRA*, 2020, pp. 11 039–11 045.
- [178] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *International Conference on Learning Representations (ICLR)*, 2015.
- [179] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [180] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in CVPR, 2017, pp. 2261–2269.
- [181] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical image computing and computer-assisted intervention*. Springer, 2015, pp. 234–241.
- [182] D. Maturana and S. Scherer, "Voxnet: A 3d convolutional neural network for real-time object recognition," in *IROS*, 2015, pp. 922–928.
- [183] 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 2018 ICRA, 2018, pp. 4243–4250.
- [184] J. Tobin, L. Biewald, R. Duan, M. Andrychowicz, A. Handa, V. Kumar, B. McGrew, A. Ray, J. Schneider, P. Welinder, W. Zaremba, and P. Abbeel,

"Domain randomization and generative models for robotic grasping," in *IROS*, 2018, pp. 3482–3489.

- [185] S. James, P. Wohlhart, M. Kalakrishnan, D. Kalashnikov, A. Irpan, J. Ibarz, S. Levine, R. Hadsell, and K. Bousmalis, "Sim-to-real via sim-to-sim: Dataefficient robotic grasping via randomized-to-canonical adaptation networks," in 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE Computer Society, jun 2019, pp. 12 619–12 629. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/CVPR.2019.01291
- [186] P. Khosla, P. Teterwak, C. Wang, A. Sarna, Y. Tian, P. Isola, A. Maschinot, C. Liu, and D. Krishnan, "Supervised contrastive learning," in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 18 661–18 673. [Online]. Available: https://proceedings.neurips.cc/paper/2020/file/d89a66c7c80a29b1bdbab0f2a1a94af8-Paper.pdf
- [187] B. Zhang, Y. Xie, J. Zhou, K. Wang, and Z. Zhang, "State-of-the-art robotic grippers, grasping and control strategies, as well as their applications in agricultural robots: A review," *Computers and Electronics in Agriculture*, vol. 177, p. 105694, 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0168169920311030
- [188] M. Costanzo, G. De Maria, G. Lettera, and C. Natale, "Can robots refill a supermarket shelf?: Motion planning and grasp control," *IEEE Robotics Automation Magazine*, vol. 28, no. 2, pp. 61–73, 2021.
- [189] X. Wang, H. Kang, H. Zhou, W. Au, and C. Chen, "Geometry-aware fruit grasping estimation for robotic harvesting in apple orchards," *Computers* and Electronics in Agriculture, vol. 193, p. 106716, 2022. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0168169922000333
- [190] J. Leitner *et al.*, "The ACRV picking benchmark: A robotic shelf picking benchmark to foster reproducible research," *ICRA*, pp. 4705–4712, 2017.
- [191] C. Rubert, D. Kappler, J. Bohg, and A. Morales, "Predicting grasp success in the real world - a study of quality metrics and human assessment," *Robotics* and Autonomous Systems, vol. 121, p. 103274, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0921889019300247
- [192] C. Eppner, R. Deimel, J. Álvarez Ruiz, M. Maertens, and O. Brock, "Exploitation of environmental constraints in human and robotic grasping," *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 1021–1038, 2015. [Online]. Available: https://doi.org/10.1177/0278364914559753
- [193] M. Kazemi, J.-S. Valois, J. A. Bagnell, and N. Pollard, "Robust object grasping using force compliant motion primitives," 2021.
- [194] L. Righetti, M. Kalakrishnan, P. Pastor, J. Binney, J. Kelly, R. C. Voorhies, G. S. Sukhatme, and S. Schaal, "An autonomous manipulation system based on force control and optimization," *Autonomous Robots*, vol. 36, no. 1-2, pp. 11–30, 2014.
- [195] N. Hudson, T. Howard, J. Ma, A. Jain, M. Bajracharya, S. Myint, C. Kuo, L. Matthies, P. Backes, P. Hebert *et al.*, "End-to-end dexterous manipulation with deliberate interactive estimation," in 2012 IEEE International Conference on Robotics and Automation. IEEE, 2012, pp. 2371–2378.
- [196] M. Toussaint, N. Ratliff, J. Bohg, L. Righetti, P. Englert, and S. Schaal, "Dual execution of optimized contact interaction trajectories," in 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems. IEEE, 2014, pp. 47–54.
- [197] N. C. Dafle, A. Rodriguez, R. Paolini, B. Tang, S. S. Srinivasa, M. Erdmann, M. T. Mason, I. Lundberg, H. Staab, and T. Fuhlbrigge, "Extrinsic dexterity: In-hand manipulation with external forces," in 2014 ICRA. IEEE, 2014, pp. 1578–1585.
- [198] N. Chavan-Dafle and A. Rodriguez, "Prehensile pushing: In-hand manipulation with push-primitives," in *IROS*. IEEE, 2015, pp. 6215–6222.
- [199] L. Shao, T. Migimatsu, and J. Bohg, "Learning to scaffold the development of robotic manipulation skills," in 2020 ICRA. IEEE, 2020, pp. 5671–5677.
- [200] Z. Xia, Z. Deng, B. Fang, Y. Yang, and F. Sun, "A review on sensory perception for dexterous robotic manipulation," *International Journal of Advanced Robotic Systems*, vol. 19, no. 2, p. 17298806221095974, 2022. [Online]. Available: https://doi.org/10.1177/17298806221095974
- [201] Q. Bai, S. Li, J. Yang, Q. Song, Z. Li, and X. Zhang, "Object detection recognition and robot grasping based on machine learning: A survey," *IEEE Access*, vol. 8, pp. 181 855–181 879, 2020.
- [202] Y. Bekiroglu, J. Laaksonen, J. A. Jorgensen, V. Kyrki, and D. Kragic, "Assessing grasp stability based on learning and haptic data," *IEEE Transactions on Robotics*, vol. 27, no. 3, pp. 616–629, 2011.
- [203] R. Calandra, A. Owens, D. Jayaraman, J. Lin, W. Yuan, J. Malik, E. H. Adelson, and S. Levine, "More than a feeling: Learning to grasp and regrasp using vision and touch," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3300–3307, 2018.
- [204] L. Chumbley, M. Gu, R. Newbury, J. Leitner, and A. Cosgun, "Integrating high-resolution tactile sensing into grasp stability prediction," in *Conference* on Robots and Vision, 2022.
- [205] W. Chen, H. Khamis, I. Birznieks, N. F. Lepora, and S. J. Redmond, "Tactile sensors for friction estimation and incipient slip detection—toward dexterous robotic manipulation: A review," *IEEE Sensors Journal*, vol. 18, no. 22, pp.

9049-9064, 2018.

- [206] C. de Farias, N. Marturi, R. Stolkin, and Y. Bekiroglu, "Simultaneous tactile exploration and grasp refinement for unknown objects," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3349–3356, 2021.
- [207] D. Hughes, J. Lammie, and N. Correll, "A robotic skin for collision avoidance and affective touch recognition," *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1386–1393, 2018.
- [208] G. Zöller, V. Wall, and O. Brock, "Active acoustic contact sensing for soft pneumatic actuators," in 2020 ICRA. IEEE, 2020, pp. 7966–7972.
- [209] D. Gandhi, A. Gupta, and L. Pinto, "Swoosh! rattle! thump!-actions that sound," in *Robotics*, ser. Robotics: Science and Systems. MIT Press Journals, 2020.
- [210] P. Carbines, "Welding robot," U.S. Patent 335510B1, Oct. 1999.
- [211] M. Edwards, "Robots in industry: An overview," Applied Ergonomics, vol. 15, no. 1, pp. 45–53, 1984. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0003687084901212
- [212] H. S. Cho, H. J. Warnecke, and D. G. Gweon, "Robotic assembly: a synthesizing overview," *Robotica*, vol. 5, no. 2, p. 153–165, 1987.
- [213] K. Gilday, J. Hughes, and F. Iida, "Wrist-driven passive grasping: interactionbased trajectory adaption with a compliant anthropomorphic hand," *Bioinspir*. *Biomim.*, vol. 16, no. 2, p. 026024, Mar. 2021.
- [214] M. Shishegar, D. Kerr, and J. Blake, "A systematic review of research into how robotic technology can help older people," *Smart Health*, vol. 7, 2018.

IX. APPENDIX

A. Appendix A: Methodology

We conducted a Systematic Literature Review to assess the state-of-the-art of 6-DoF robotic grasping utilizing deep learning approaches. Our review, methodology adapted from [214], searched through six different scholarly libraries – IEEE, Springer, ScienceDirect, SpringerLink, arXiv and Taylor & Francis – with the following search terms and criteria.

B. Search Terms

The following search terms were used to search through the whole paper (where the database allowed this option). The 'AND' function represents a way to enforce the search results to contain all of the search terms. When possible, the search terms were combined using an 'OR' function, otherwise, each term was searched separately. Furthermore, in combination with the search term 'Grasping' was searched in the metadata of the paper, if metadata search was available.

- 6-DoF Grasping
 - Grasping AND Point Cloud AND Deep Learning
 - Shape Completion AND 6-DoF Grasping AND Learning
- Affordances AND Grasping AND 6 DoF

C. Inclusion/Exclusion Criteria

From the papers found in the databases, only publications that met the following criteria are included in this review:

- (a) Paper considered grasping from a table-top scenario,
- (b) All 6-DoF were used for the grasp pose,
- (c) Deep Learning methods were applied in some aspect of the work,
- (d) Published after Jan 1, 2012, (the year Alexnet [34] was published) and
- (e) Written in English.

All criteria was searched manually by the authors. Our search returned a total of 85 papers that matched these criteria and were therefore included in this review. A table of all the reviewed papers and their deep-learning approaches are included in the online version of the survey.

D. Data Extraction/Analysis

Data was extracted from the included papers and verified by at least two of the authors.



Rhys Newbury is a PhD student at Monash University, Australia. He holds a B. Eng and B. Sc (Computer Science) from Monash University. His research interests focus around manipulation and human-robot interaction. He has a focus on systems view to problems and how to apply robots to real-world issues.



Morris Gu is a PhD student at Monash University. Australia. He holds a B. Eng from Monash University. His research interests are in explainability and augmented reality within human-robot interaction.

Lachlan Chumbley is an undergraduate student at Monash University, Australia. His research interests

include robotic grasping and manipulation, particularly

though the use of tactile information.



Antonio Morales is Associate Professor at the Department of Computer Engineering and Science in the Universitat Jaume I of Castelló, Spain. He received his PhD in Computer Science Engineering from Universitat Jaume I in January 2004. Currently, he is director of the Degree Program on Robotics Intelligence and vice-dean of the School on Technology and Experimental Sciences of the Universitat Jaume I. He is a leading researcher at the Robotic Intelligence Laboratory at Universitat Jaume I and his research interests are focused on reactive robot grasping and manipulation. He has been a Principal

Investigator on several European and national research projects.



Tamim Asfour is Professor at the Institute for Anthropomatics and Robotics at KIT, where he holds the chair of Humanoid Robotics Systems and is head of the High Performance Humanoid Technologies Lab (H²T). His research interest is 24/7 humanoid robotics. Specifically, he studies the mechano-informatics of humanoids as the synergetic integration of informatics, artificial intelligence, and mechatronics into complete humanoid robot systems, which are able to perform versatile tasks real world. He is developer and the leader of the development team of the ARMAR humanoid robot family.

Danica Kragic is a Professor at the School of Computer Science and Communication at KTH in Stockholm. She received MSc in Mechanical Engineering from the Technical University of Rijeka, Croatia in 1995 and PhD in Computer Science from KTH in 2001. Danica received the 2007 IEEE Robotics and Automation Society Early Academic Career Award. She is a member of the Swedish Royal Academy of Sciences and Swedish Young Academy. She has chaired the IEEE RAS Technical Committee on Computer and Robot Vision and from 2009 serves as an IEEE RAS AdCom member.



Dieter Fox (Fellow, IEEE) received the Ph.D. degree in computer science from the University of Bonn, Bonn, Germany, in 1998. He is currently a Professor with the Allen School of Computer Science and Engineering, University of Washington, Seattle, WA, USA, where he heads the UW Robotics and State Estimation Lab. He is also Senior Director of Robotics Research with NVIDIA.He has authored or coauthored more than 200 technical papers. He is a co-author of the textbook entitled Probabilistic Robotics (Cambridge, MA, USA: MIT Press, 2005). His research interests include robotics and artificial

intelligence, with a focus on state estimation and perception applied to problems such as mapping, object detection and tracking, manipulation, and activity recognition.



Akansel Cosgun is currently a Senior Lecturer at Deakin University, Australia. He received his Ph.D. degree in robotics from the Georgia Institute of Technology, Atlanta, GA, USA, in 2016. From 2018 to 2022, he was a Research Fellow with Monash University, Australia. He conducts research in robotics, human-robot interaction, and robot learning. From 2019 to 2020, he was the Team Lead for Vision-Based Manipulation with the Australian Centre for Robotic Vision. He has previously worked with Honda Research, Toyota Infotechnology Center, Microsoft Research, and Savioke, a robotics start-up. His

research interests include mobile robots, robotic arms, and self-driving cars with an emphasis on a systems view to problems.



Arsalan Mousavian received the B.Sc. degree from the Iran University of Science and Technology, Tehran, Iran, in 2010, the M.Sc. degree from the University of Tehran, Tehran, in 2013, both in AI and robotics, and the Ph.D. degree in computer science from George Mason University, Fairfax, VA, USA, in 2018. He is currently a Senior Research Scientist with NVIDIA, Seattle, WA, USA. His research interests include 3-D perception methods that help robots accomplish robot manipulation tasks in the real world.



Clemens Eppner is a Research Scientist in the Seattle Robotics Lab at NVIDIA Research. Before joining NVIDIA, he received his Ph.D. at the Robotics and Biology Lab at Technische Universität Berlin, Germany, and M.Sc. from the University of Freiburg, Germany. He is interested in the problem space of grasping and manipulation, including aspects of planning, control, and perception.



Jürgen Leitner is co-founder at LYRO Robotics, Australia. He has designed, developed, and deployed robots fore more than 15 years. His interest lies in building intelligent robots that can safely and reliably interact with the physical world. He holds a PhD (2014), a MSc (Space Science, 2009), a MSc(Tech) (Space Robotics, 2009) and a BSc (Software Engineering, 2007).



Jeannette Bohg is an Assistant Professor of Computer Science at Stanford University. She was a group leader at the Autonomous Motion Department (AMD) of the MPI for Intelligent Systems until September 2017. Before joining AMD in January 2012, Jeannette Bohg was a PhD student at the Division of Robotics, Perception and Learning (RPL) at KTH in Stockholm. Her research focuses on perception and learning for autonomous robotic manipulation and grasping.





