



Título artículo / Títol article:

**Data-driven grasp
synthesis—a survey**

Autores / Autors:

**Bohg, J. ; Morales, A. ;
Asfour, T. ; Kragic, D.**

Revista:

IEEE Transactions on Robotics

Versión / Versió:

Postprint

**Cita bibliográfica / Cita
bibliogràfica (ISO 690):**

**BOHG, Jeannette, et al. Data-
driven grasp synthesis—a
survey. *Robotics, IEEE
Transactions on*, 2014, vol. 30,
no 2, p. 289-309.**

url Repositori UJI:

[http://hdl.handle.net/10234/
132550](http://hdl.handle.net/10234/132550)

Data-Driven Grasp Synthesis - A Survey

Jeannette Bohg, *Member, IEEE*, Antonio Morales, *Member, IEEE*, Tamim Asfour, *Member, IEEE*,
Danica Kragic *Member, IEEE*

Abstract—We review the work on data-driven grasp synthesis and the methodologies for sampling and ranking candidate grasps. We divide the approaches into three groups based on whether they synthesize grasps for known, familiar or unknown objects. This structure allows us to identify common object representations and perceptual processes that facilitate the employed data-driven grasp synthesis technique. In the case of known objects, we concentrate on the approaches that are based on object recognition and pose estimation. In the case of familiar objects, the techniques use some form of a similarity matching to a set of previously encountered objects. Finally, for the approaches dealing with unknown objects, the core part is the extraction of specific features that are indicative of good grasps. Our survey provides an overview of the different methodologies and discusses open problems in the area of robot grasping. We also draw a parallel to the classical approaches that rely on analytic formulations.

Index Terms—Object grasping and manipulation, grasp synthesis, grasp planning, visual perception, object recognition and classification, visual representations

I. INTRODUCTION

Given an object, *grasp synthesis* refers to the problem of finding a grasp configuration that satisfies a set of criteria relevant for the grasping task. Finding a suitable grasp among the infinite set of candidates is a challenging problem and has been addressed frequently in the robotics community, resulting in an abundance of approaches.

In the recent review of Sahbani et al. [1], the authors divide the methodologies into *analytic* and *empirical*. Following Shimoga [2], analytic refers to methods that construct force-closure grasps with a multi-fingered robotic hand that are *dexterous*, in *equilibrium*, *stable* and exhibit a certain *dynamic behaviour*. Grasp synthesis is then usually formulated as a constrained optimization problem over criteria that measure one or several of these four properties. In this case, a grasp is typically defined by the *grasp map* that transforms the forces exerted at a set of contact points to object wrenches [3]. The criteria are based on geometric, kinematic or dynamic formulations. Analytic formulations towards grasp synthesis have also been reviewed by Bicchi and Kumar [4].

Empirical or *data-driven* approaches rely on sampling grasp candidates for an object and ranking them according to a specific metric. This process is usually based on some existing

grasp experience that can be a heuristic or is generated in simulation or on a real robot. Kamon et al. [5] refer to this as the *comparative* and Shimoga [2] as the *knowledge-based* approach. Here, a grasp is commonly parameterized by [6, 7]:

- the grasping point on the object with which the *tool center point* (TCP) should be aligned,
- the *approach vector* which describes the 3D angle that the robot hand approaches the grasping point with,
- the wrist orientation of the robotic hand and
- an initial finger configuration

Data-driven approaches differ in how the set of grasp candidates is sampled, how the grasp quality is estimated and how good grasps are represented for future use. Some methods measure grasp quality based on analytic formulations, but more commonly they encode e.g. human demonstrations, perceptual information or semantics.

A. Brief Overview of Analytic Approaches

Analytic approaches provide guarantees regarding the criteria that measure the previously mentioned four grasp properties. However, these are usually based on assumptions such as simplified contact models, Coulomb friction and rigid body modeling [3, 8]. Although these assumptions render grasp analysis practical, inconsistencies and ambiguities especially regarding the analysis of grasp dynamics are usually attributed to their approximate nature.

In this context, Bicchi and Kumar [4] identified the problem of finding an accurate and tractable model of contact compliance as particularly relevant. This is needed to analyze statically-indeterminate grasps in which not all internal forces can be controlled. This case arises e.g. for under-actuated hands or grasp synergies where the number of the controlled degrees of freedom is fewer than the number of contact forces. Prattichizzo et al. [9] model such a system by introducing a set of springs at the contacts and joints and show how its dexterity can be analyzed. Rosales et al. [10] adopt the same model of compliance to synthesize feasible and prehensile grasps. In this case, only statically-determinate grasps are considered. The problem of finding a suitable hand configuration is cast as a constrained optimization problem in which compliance is introduced to simultaneously address the constraints of contact reachability, object restraint and force controllability. As is the case with many other analytic approaches towards grasp synthesis, the proposed model is only studied in simulation where accurate models of the hand kinematics, the object and their relative alignment are available.

In practice, systematic and random errors are inherent to a robotic system and are due to noisy sensors and inaccurate object, robot, etc. models. The relative position of object

J. Bohg is with the Autonomous Motion Department at the MPI for Intelligent Systems, Tübingen, Germany, e-mail: jbohlg@tuebingen.mpg.de.

A. Morales is with the Robotic Intelligence Lab. at Universitat Jaume I, Castelló, Spain, e-mail: Antonio.Morales@uji.es.

T. Asfour is with the KIT, Karlsruhe, Germany, e-mail: asfour@kit.edu.

D. Kragic is with the Centre for Autonomous Systems, Computational Vision and Active Perception Lab, Royal Institute of Technology KTH, Stockholm, Sweden, e-mail: dank@kth.se.

This work has been supported by FLEXBOT (FP7-ERC-279933).

and hand can therefore only be known approximately which makes an accurate placement of the fingertips difficult. In 2000, Bicchi and Kumar [4] identified a lack of approaches towards synthesizing grasps that are robust to positioning errors. Since then, this problem has shifted into focus. One line of research follows the approach of *independent contact regions* (ICRs) as defined by Nguyen [11]: a set of regions on the object in which each finger can be independently placed anywhere without the grasp losing the force-closure property. Several examples for computing them are presented by Roa and Suárez [12] or Krug et al. [13]. Another line of research towards robustness against inaccurate end-effector positioning makes use of the caging formulation. Rodriguez et al. [14] found that there are caging configurations of a three-fingered manipulator around a planar object that are specifically suited as a way point to grasping it. Once the manipulator is in such configuration, either opening or closing the fingers is guaranteed to result in an equilibrium grasp without the need for accurate positioning of the fingers. Seo et al. [15] exploited the fact that two-fingered immobilizing grasps of an object are always preceded by a caging configuration. Full body grasps of planar objects are synthesized by first finding a two-contact caging configuration and then using additional contacts to restrain the object. Results have been presented in simulation and demonstrated on a real robot.

Another assumption commonly made in analytic approaches is that precise geometric and physical models of an object are available to the robot which is not always the case. In addition, we may not know the surface properties or friction coefficients, weight, center of mass and weight distribution. Some of these can be retrieved through interaction: Zhang and Trinkle [16] propose to use a particle filter to simultaneously estimate the physical parameters of an object and track it while it is being pushed. The dynamic model of the object is formulated as a mixed nonlinear complementarity problem. The authors show that even when the object is occluded and the state estimate cannot be updated through visual observation, the motion of the object is accurately predicted over time. Although methods like this relax some of the assumptions, they are still limited to simulation [14, 10] or consider 2D objects [14, 15, 16].

B. Development of Data-Driven Methods

Up to the year 2000, the field of robotic grasping¹ was clearly dominated by analytic approaches [11, 4, 17, 2]. Apart from e.g. Kamon et al. [5], data-driven grasp synthesis started to become popular with the availability of GraspIt! [18] in 2004. Many highly cited approaches have been developed, analyzed and evaluated in this or other simulators [19, 20, 21, 22, 23, 24]. These approaches differ in how grasp candidates are sampled from the infinite space of possibilities. For grasp ranking they rely classical metrics based on analytic formulations such as the widely used ϵ -metric proposed in Ferrari and Canny [17]. It constructs the *grasp wrench space* (GWS) by

computing the convex hull over the wrenches at the contact points between the hand and the object. ϵ ranks the quality of a force closure grasp by quantifying the radius of the maximum sphere still fully contained in the GWS.

Developing and evaluating approaches in simulation is attractive because the environment and its attributes can be completely controlled. A large number of experiments can be efficiently performed without having access to expensive robotics hardware that would also add a lot of complexity to the evaluation process. However, it is not clear if the simulated environment resembles the real world well enough to transfer methods easily. Only recently, several articles [39, 40, 24] have analyzed this question and come to the conclusion that the classic metrics are not good predictors for grasp success in the real world. They do not seem to cope well with the challenges arising in unstructured environments. Diankov [24] claims that in practice grasps synthesized using this metric tend to be relatively *fragile*. Balasubramanian et al. [39] systematically tested a number of grasps in the real world that were stable according to classical grasp metrics. Compared to grasps planned by humans and transferred to a robot by kinesthetic teaching on the same objects, they under-performed significantly. A similar study has been conducted by Weisz and Allen [40]. It focuses on the ability of the ϵ -metric to predict grasp stability under object pose error. The authors found that it performs poorly especially when grasping large objects.

As pointed out by Bicchi and Kumar [4] and Prattichizzo and Trinkle [8], grasp closure is often wrongly equated with stability. Closure states the existence of equilibrium which is a necessary but not sufficient condition. Stability can only be defined when considering the grasp as a dynamical system and in the context of its behavior when perturbed from an equilibrium. Seen in this light, the results of the above mentioned studies are not surprising. However, they suggest that there is a large gap between reality and the models for grasping that are currently available and tractable.

For this reason, several researchers [25, 26, 27] proposed to let the robot learn how to grasp by experience that is gathered during grasp execution. Although, collecting examples is extremely time-consuming, the problem of transferring the learned model to the real robot is non-existent. A crucial question is how the object to be grasped is represented and how the experience is generalized to novel objects.

Saxena et al. [28] pushed machine learning approaches for data-driven grasp synthesis even further. A simple logistic regressor was trained on large amounts of synthetic labeled training data to predict good grasping points in a monocular image. The authors demonstrated their method in a household scenario in which a robot emptied a dishwasher. None of the classical principles based on analytic formulations were used. This paper spawned a lot of research [29, 30, 31, 32] in which essentially one question is addressed: What are the object features that are sufficiently discriminative to infer a suitable grasp configuration?

From 2009, there were further developments in the area of 3D sensing. Projected Texture Stereo was proposed by Konolige [41]. This technology is built into the sensor head of the PR2 [42], a robot that is available to comparatively many

¹Citation counts for the most influential articles in the field. Extracted from scholar.google.com in October 2013. [11]: 733. [4]: 490. [17]: 477. [2]: 405. [5]: 77. [18]: 384. [19]: 353. [20]: 100. [21]: 110. [22]: 95. [23]: 96. [24]: 108. [25]: 38. [26]: 156. [27]: 39. [28]: 277. [29]: 75. [30]: 40. [31]: 21. [32]: 43. [33]: 77. [34]: 26. [35]: 191. [36]: 58. [37]: 75. [38]: 39.

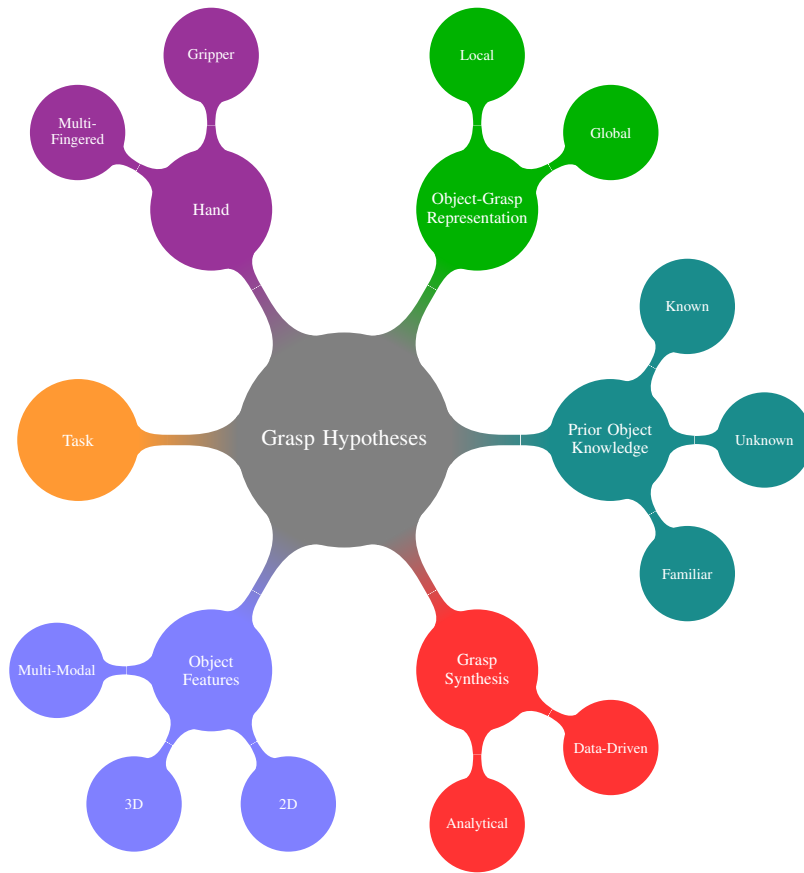


Figure 1: We identified a number of aspects that influence how the final set of grasp hypotheses is generated for an object. The most important one is the assumed **prior object knowledge** as discussed in Section I-D. Numerous different **object-grasp representations** are proposed in the literature that are relying on **features** of different modalities such as 2D or 3D vision or tactile sensors. Either local object parts or the object as a whole are linked to specific grasp configurations. **Grasp synthesis** can either be analytic or data-driven. The latter is further detailed in Fig. 2. Very few approaches explicitly address the **task** or **hand** kinematics of the robot.

robotics research labs and running on the OpenSource middle-ware ROS [43]. In 2010, Microsoft released the Kinect [44], a highly accurate depth sensing device based on the technology developed by PrimeSense [45]. Due to its low price and simple usage, it became a ubiquitous device within the robotics community. Although the importance of 3D data for grasping has been previously recognized, many new approaches were proposed that operate on real world 3D data. They are either heuristics that map structures in this data to grasp configurations directly [33, 34] or they try to detect and recognize objects and estimate their pose [35, 46].

Furthermore, we have recently seen an increasing amount of robots fulfilling very specific tasks such as towel folding [37] or preparing pancakes [38]. In these scenarios, grasping is embedded into a sequence of different manipulation actions.

C. Analytic vs. Data-Driven Approaches

Contrary to analytic approaches, methods following the data-driven paradigm place more weight on the object representation and the perceptual processing, e.g., feature extraction, similarity metrics, object recognition or classification and pose estimation. The resulting data is then used to retrieve grasps from some knowledge base or sample and rank them by comparison to existing grasp experience. The parameterization of the grasp is less specific (e.g. an approach vector instead of fingertip positions) and therefore accommodates for uncertainties in perception and execution. This provides a natural

precursor to reactive grasping [47, 48, 49, 33, 50], which, given a grasp hypothesis, considers the problem of robustly acquiring it under uncertainty. Data-driven methods cannot provide guarantees regarding the aforementioned criteria of dexterity, equilibrium, stability and dynamic behaviour [2]. They can only be verified empirically. However, they form the basis for studying grasp dynamics and further developing analytic models that better resemble reality.

D. Classification of Data-Driven Approaches

Sahbani et al. [1] divide the data-driven methods based on whether they employ object features or observation of humans during grasping. We believe that this falls short of capturing the diversity of these approaches especially in terms of the ability to transfer grasp experience between similar objects and the role of perception in this process. In this survey, we propose to group data-driven grasp synthesis approaches based on what they assume to know *a priori* about the query object:

- **Known Objects:** These approaches assume that the query object has been encountered before and that grasps have already been generated for it. Commonly, the robot has access to a database containing geometric object models that are associated with a number of good grasp. This database is usually built offline and in the following will be referred to as an *experience database*. Once the object has been recognized, the goal is to estimate its pose and retrieve a suitable grasp.

- *Familiar Objects*: Instead of exact identity, the approaches in this group assume that the query object is similar to previously encountered ones. New objects can be *familiar* on different levels. Low-level similarity can be defined in terms of shape, color or texture. High-level similarity can be defined based on object category. These approaches assume that new objects similar to old ones can be grasped in a similar way. The challenge is to find an object representation and a similarity metric that allows to transfer grasp experience.
- *Unknown Objects*: Approaches in this group do not assume to have access to object models or any sort of grasp experience. They focus on identifying structure or features in sensory data for generating and ranking grasp candidates. These are usually based on local or global features of the object as perceived by the sensor.

We find the above classification suitable for surveying the data-driven approaches since the assumed prior object knowledge determines the necessary perceptual processing and associated object representations for generating and ranking grasp candidates. For known objects, the problems of recognition and pose estimation have to be addressed. The object is usually represented by a complete geometric 3D object model. For familiar objects, an object representation has to be found that is suitable for comparing them to already encountered object in terms of graspability. For unknown objects, heuristics have to be developed for directly linking structure in the sensory data to candidate grasps.

Only a minority of the approaches discussed in this survey cannot be clearly classified to belong to one of these three groups. Most of the included papers use sensor data from the scene to perform data-driven grasp synthesis and are part of a real robotic system that can execute grasps.

Finally, this classification is well in line with the research in the field of neuroscience, specifically, from the theory of the dorsal and ventral stream in human visual processing [51]. The *dorsal* pathway processes immediate action-relevant features while the *ventral* pathway extracts context- and scene-relevant information and is related to object recognition. The visual processing in the ventral and dorsal pathways can be related to the grouping of grasp synthesis for familiar/known and unknown objects, respectively. The details of such links are out of the scope of this paper. Extensive and detailed reviews on the neuroscience of grasping are offered in [52, 53, 54].

E. Aspects Influencing the Generation of Grasp Hypotheses

The number of *candidate grasps* that can be applied to an object is infinite. To sample some of these candidates and define a quality metric for selecting a good subset of *grasp hypotheses* is the core subject of the approaches reviewed in this survey. In addition to the prior object knowledge, we identified a number of other factors that characterize these metrics. Thereby, they influence which grasp hypotheses are selected by a method. Fig. 1 shows a mind map that structures these aspects. An important one is how the quality of a candidate grasp depends on the object, i.e., the *object-grasp representation*. Some approaches extract local object




Figure 2: Data-driven Grasp Synthesis can either be based on heuristics or on learning from data. The data can either be provided in the form of offline generated labeled training data, human demonstration or through trial and error.

	Object-Grasp Represen.		Object Features			Grasp Synthesis							
	Local	Global	2D	3D	Multi-Modal	Heuristic	Human Demo	Labeled Data	Trial & Error	Task	Multi-Fingered	Deformable	Real Data
Glover et al. [55]	✓		✓			✓							✓
Goldfeder et al. [21]	✓			✓	✓	✓							✓
Berenson et al. [56]	✓			✓	✓	✓							✓
Miller et al. [19]	✓			✓	✓	✓							✓
Przybylski et al. [57]	✓			✓	✓	✓							✓
Roa et al. [58]	✓			✓	✓	✓							✓
Detry et al. [27]	✓			✓	✓	✓	✓						✓
Detry et al. [59]	✓			✓	✓	✓			✓				✓
Huebner et al. [60]	✓			✓	✓	✓							✓
Faria et al. [61]	✓			✓	✓	✓				✓			✓
Diankov [24]	✓			✓	✓	✓							✓
Balasubramanian et al. [39]	✓	✓		✓	✓	✓							✓
Borst et al. [22]	✓			✓	✓	✓							✓
Brook et al. [62]	✓			✓	✓	✓							✓
Ciocarlie and Allen [23]	✓			✓	✓	✓							✓
Romero et al. [63]	✓			✓	✓	✓							✓
Papazov et al. [64]	✓			✓	✓	✓							✓
Morales et al. [7]	✓			✓	✓	✓							✓
Collet Romea et al. [65]	✓			✓	✓	✓							✓
Kroemer et al. [66]	✓			✓	✓	✓							✓
Ekvall and Kragic [6]	✓			✓	✓	✓							✓
Tegin et al. [67]	✓			✓	✓	✓							✓
Pastor et al. [49]	✓			✓	✓	✓							✓
Stulp et al. [68]	✓			✓	✓	✓							✓

Table I: Data-Driven Approaches for Grasping Known Objects

attributes (e.g. curvature, contact area with the hand) around a candidate grasp. Other approaches take global characteristics (e.g. center of mass, bounding box) and their relation to a grasp configuration into account. Dependent on the sensor device, *object features* can be based on 2D or 3D visual data as well as on other modalities. Furthermore, *grasp synthesis* can be analytic or data-driven. We further categorized the latter in Fig. 2: there are methods for *learning* either from *human demonstrations*, *labeled examples* or *trial and error*. Other methods rely on various *heuristics* to directly link structure in sensory data to candidate grasps. There is relatively little work on task-dependent grasping. Also, the applied robotic hand is usually not in the focus of the discussed approaches. We will therefore not examine these two aspects. However, we will indicate whether an approach takes the task into account and whether an approach is developed for a gripper or for the more complex case of a multi-fingered hand. Table I-III list all the methods in this survey. The table columns follow the structure proposed in Fig. 1 and 2.

II. GRASPING KNOWN OBJECTS

If the object to be grasped is known and there is already a database of grasp hypotheses for it, the problem of finding a feasible grasp reduces to estimating the object pose and then filtering the hypotheses by reachability. Table I summarizes all the approaches discussed in this section.

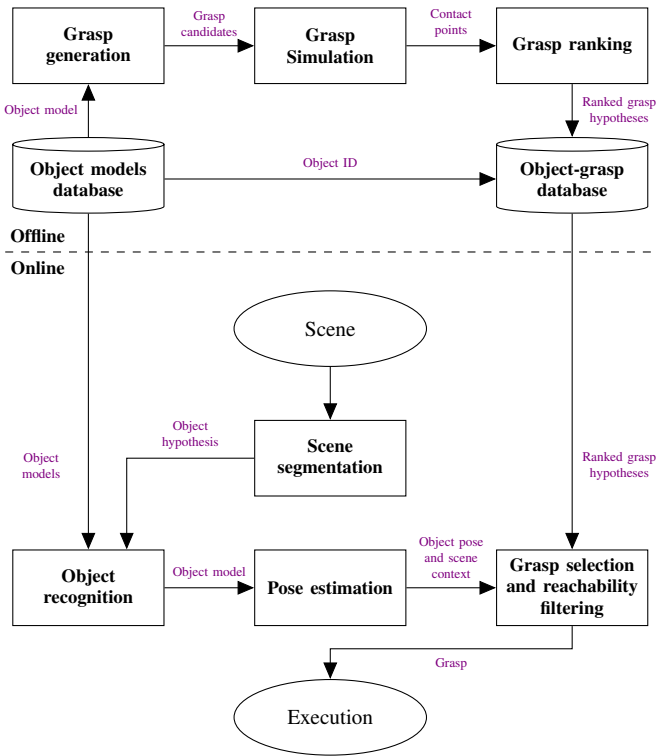


Figure 3: Typical functional flow-chart for a system with offline generation of a grasp database. In the offline phase, every object model is processed to generate grasp candidates. Their quality is evaluated for ranking. Finally, the list of grasp hypotheses is stored with the corresponding object model. In the online phase, the scene is segmented to search and recognize object models. If the process succeeds, the associated grasp hypotheses are retrieved and unreachable ones are discarded. Most of the following approaches can be summarized with this flowchart. Some of them only implement the offline part. [21, 56, 19, 57, 58, 60, 24, 39, 22, 62, 23, 7, 67]

A. Offline Generation of a Grasp Experience Database

First, we look at approaches for generating the experience database. Figs. 3 and 5 summarize the typical functional flow-chart of these type of approaches. Each box represents a processing step. Please note, that these figures are abstractions that summarize the implementations of a number of papers. Most reviewed papers focus on a single module. This is also true for similar figures appearing in Sections IV and III.

1) *3D Mesh Models and Contact-Level Grasping*: Many approaches in this category assume that a 3D mesh of the object is available. The challenge is then to automatically generate a set of good grasp hypotheses. This involves sampling the infinite space of possible hand configurations and ranking the resulting candidate grasps according to some quality metric. The major part of the approaches discussed in the following use force closure grasps and rank them according to the previously discussed ϵ -metric. They differ mostly in the way the grasp candidates are sampled. Fig. 3 shows a flow-chart of which specifically the upper part (Offline) visualizes the data flow for the following approaches.

Some of them approximate the object’s shape with a constellation of primitives such as spheres, cones, cylinders and boxes as in Miller et al. [19], Hübner and Kragic [69] and Przybylski et al. [57] or superquadrics (SQ) as in Goldfeder et al. [21]. These shape primitives are then used to limit the amount of candidate grasps and thus prune the search tree for finding the most stable set of grasp hypotheses. Examples

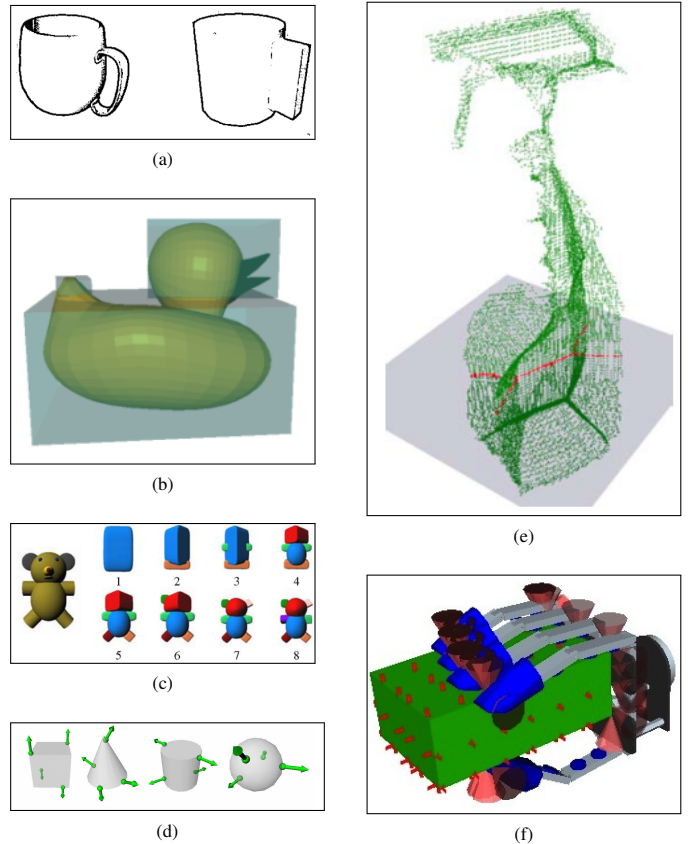


Figure 4: Generation of grasp candidates through object shape approximation with primitives or through sampling. 4a) Primitive Shape Decomposition [19]. 4b) Box Decomposition [69]. 4c) SQ Decomposition [21]. 4d) Randomly sampled grasp hypotheses.[22]. 4e) Green: Centers of a union of spheres. Red: Centers at a slice through the model [70, 57]. 4f) Grasp candidate sampled based on surface normals and bounding box [71].

for these approaches are shown in Fig. 4a-4c and Fig. 4e. Borst et al. [22] reduce the number of candidate grasps by randomly generating a number of them dependent on the object surface and filter them with a simple heuristic. The authors show that this approach works well if the goal is not to find an optimal grasp but instead a fairly good grasp that works well for “everyday tasks”. Diankov [24] proposes to sample grasp candidates dependent on the objects bounding box in conjunction with surface normals. Grasp parameters that are varied are the distance between the palm of the hand and the grasp point as well as the wrist orientation. This scheme is implemented in OpenRave [71]. The authors find that usually a relatively small amount of 30% from all grasp samples is in force closure. Examples for these sampling approaches are shown in Fig. 4d and 4f. Roa et al. [58] present an approach towards synthesizing power grasps that is not based on evaluating the force-closure property. Slices through the object model and perpendicular to the axes of the bounding box are sampled. The ones that best resemble a circle are chosen for synthesizing a grasp.

All these approaches are developed and evaluated in simulation. As claimed by e.g. Diankov [24], the biggest criticism towards ranking grasps based on force closure and the ϵ -metric is that relatively *fragile grasps* might be selected. A common approach to filter these, is to add noise to the grasp parameters and keep only those grasps in which a certain percentage of the neighboring candidates also yield force closure. Weisz and

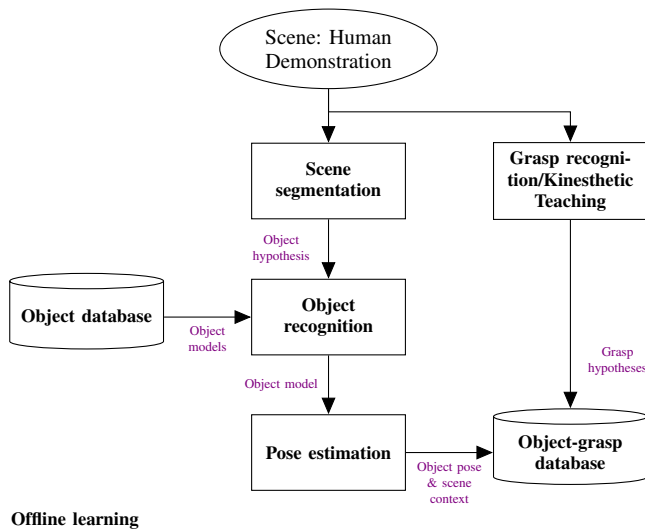


Figure 5: Typical functional flow-chart of a system that learns from human demonstration. The robot observes a human grasping a known object. Two perceptual processes are followed in parallel. On the left, the object is recognized. On the right, the demonstrated grasp configuration is extracted or recognized. Finally, object models and grasps are stored together. This process could replace or complement the offline phase described in Fig. 3. The following approaches follow this approach: [27, 61, 39, 63, 66, 49, 68].

Allen [40] followed a similar approach that focuses in particular on the ability of the ϵ -metric to predict grasp stability under object pose uncertainty. For a set of object models, the authors used GraspIt! [18] to generate a set of grasp candidates in force closure. For each object, pose uncertainty is simulated by perturbing it in three degrees of freedom. Each grasp candidate was then re-evaluated according to the probability of attaining a force closure grasp. The authors found that their proposed metric performs superior especially on large objects. Also Berenson et al. [56] consider the ϵ -metric as only one factor in selecting an appropriate grasp. The second factor in computing the final score of a grasp is the whole scene geometry.

Balasubramanian et al. [39] question classical grasp metrics. The authors systematically tested a number of grasps in the real world that were stable according to classical grasp metrics. These grasps under-performed significantly when compared to grasps planned by humans through kinesthetic teaching on the same objects. The authors found that humans optimize a *skewness* metric, i.e., the divergence of alignment between hand and principal object axes.

2) *Learning from Humans*: A different way to generate grasp hypotheses is to observe how humans grasp an object. This is usually done offline following the flow-chart in Fig. 5. This process produces an experience database that is exploited online in a similar fashion as depicted in Fig. 3.

Ciocarlie and Allen [23] exploit results from neuroscience that showed that human hand control takes place in a space of much lower dimension than the hand's degrees of freedom. This finding was applied to directly reduce the configuration space of a robotic hand to find pre-grasp postures. From these so called *eigengrasps* the system searches for stable grasps.

Detry et al. [27] model the object as a constellation of local multi-modal contour descriptors. Four elementary grasping actions are associated to specific constellations of these features resulting in an abundance of grasp candidates. They are modeled as a non-parametric density function in the space

of 6D gripper poses, referred to as a *bootstrap* density. Human grasp examples are used to build an object specific *empirical* grasp density from which grasp hypotheses can be sampled. This is visualized in Fig. 8f and 8g.

Kroemer et al. [66] represent the object with the same features as used by Detry et al. [27]. How to grasp specific objects is learned through a combination of a high-level reinforcement learner and a low level reactive grasp controller. The learning process is bootstrapped through imitation learning in which a demonstrated reaching trajectory is converted into an initial policy. Similar initialization of an object specific grasping policy is used in Pastor et al. [49] and Stulp et al. [68].

Faria et al. [61] use a multi-modal data acquisition tool to record humans grasping specific objects. Objects are represented as probabilistic volumetric maps that integrate multi-modal information such as contact regions, fixation points and tactile forces. This information is then used to segment these maps into object parts and determine whether they are graspable or not. However, no results are provided on how a robot can use this information to synthesize a grasp. Romero et al. [63] present a system for observing humans visually while they interact with an object. A grasp type and pose is recognized and mapped to different robotic hands in a fixed scheme. For validation of the approach in the simulator, 3D object models are used. This approach has been demonstrated on a humanoid robot by Do et al. [72]. The object is not explicitly modeled. Instead, it is assumed that human and robot act on the same object in the same pose.

In the method presented by Ekvall and Kragic [6], a human demonstrator wearing a magnetic tracking device is observed while manipulating a specific object. The grasp type is recognized and mapped through a fixed schema to a set of robotic hands. Given the grasp type and the hand, the best approach vector is selected from an offline trained experience database. Unlike Detry et al. [27] and Romero et al. [63], the approach vector used by the demonstrator is not adopted. Ekvall and Kragic [6] assume that the object pose is known. Experiments are conducted with a simulated pose error. No physical experiments have been demonstrated. Examples for the above mentioned ways to teach a robot grasping by demonstration are shown in Fig. 6.

3) *Learning through Trial and Error*: Instead of adopting a fixed set of grasp candidates for a known object, the following approaches try to refine them by *trial and error*. In this case, there is no separation between offline learning and online exploitation as can be seen in Fig. 7. Kroemer et al. [66], Stulp et al. [68] apply reinforcement learning to improve an initial human demonstration. Kroemer et al. [66] uses a low-level reactive controller to perform the grasp that informs the high level controller with reward information. Stulp et al. [68] increase the robustness of their non-reactive grasping strategy by learning shape and goal parameters of the motion primitives that are used to model a full grasping action. Through this approach, the robot learns reaching trajectories and grasps that are robust against object pose uncertainties. Also Detry et al. [59], adopt an approach in which the robot learns to grasp specific objects by experience. Instead of human examples as in [27], successful grasping trials are used to build the object-




Figure 6: Robot grasp learning from human demonstration. 6a) Kinesthetic Teaching [73]. 6b) Human-to-robot mapping of grasps using a data glove [6]. 6c) Human-to-robot mapping of grasps using visual grasp recognition [63]

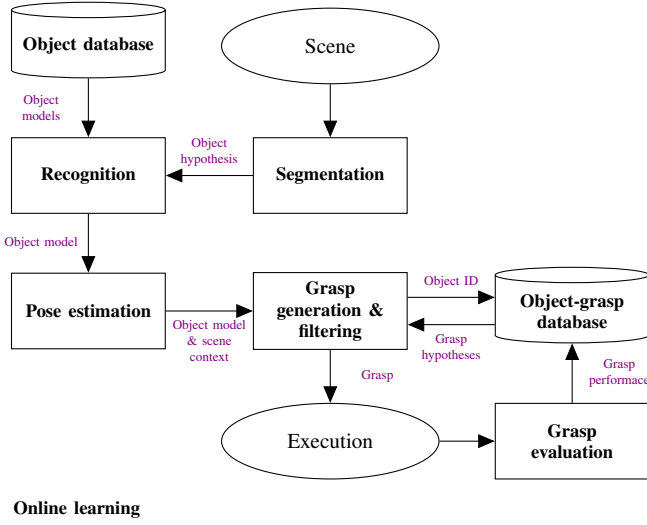


Figure 7: Typical functional flow-chart of a system that learns through trial and error. First, a known object in the scene is segmented and recognized. Past experiences with that object are retrieved and a new grasp hypothesis is generated or selected among the already tested ones. After execution of the selected grasp, the performance is evaluated and the memory of past experiences with the object is updated. The following approaches use trial-and-error learning: [27, 59, 66, 68].

specific empirical grasp density. This non-parametric density can then be used to sample grasp hypotheses.

B. Online Object Pose Estimation

In the previous section, we reviewed different approaches towards grasping known objects regarding their way to generate and rank candidate grasps. During online execution, an object has to be recognized and its pose estimated before the offline trained grasps can be executed. Furthermore, from the set of hypotheses not all grasps might be feasible in the current scene. They have to be filtered by reachability. The lower part of Fig. 3 visualizes the data flow during grasp execution and how the offline generated data is employed.

Several of the aforementioned grasp generation methods [66, 27, 59] use the probabilistic approach towards object representation and pose estimation proposed by Detry et al. [74]. It is visualized in Fig. 8e. Grasps are either selected by sampling from densities [27, 59] or a grasp policy refined from a human demonstration is applied [66]. Morales et al. [7] use the method proposed by Azad et al. [75] to recognize an object and estimate its pose from a monocular image as shown in Fig. 8a. Given this information, an appropriate grasp configuration can be selected from a grasp experience database that has been acquired offline. The whole system is demonstrated on the robotic platform described in Asfour

et al. [76]. Huebner et al. [60] demonstrate grasping of known objects on the same humanoid platform and using the same method for object recognition and pose estimation. The offline selection of grasp hypotheses is based on a decomposition into boxes. Ciocarlie et al. [77] propose a robust grasping pipeline in which known object models are fitted to point cloud clusters using standard ICP [78]. The search space of potential object poses is reduced by assuming a dominant plane and rotationally-symmetric objects that are always standing upright. An example scene is shown in Fig. 8b. Papazov et al. [64] demonstrate their previous approach on 3D object recognition and pose estimation [79] in a grasping scenario. Multiple objects in cluttered scenes can be robustly recognized and their pose estimated. No assumption is made about the geometry of the scene, shape of the objects or their pose.

The aforementioned methods assume a-priori known rigid 3D object model. Glover et al. [55] consider known deformable objects. Probabilistic models of their 2D shape are learned offline. The objects can then be detected in monocular images of cluttered scenes even when they are partially occluded. The visible object part serve as a basis for planning a grasp under consideration of the global object shape. An example for a successful detection is shown in Fig. 8c.

Collet Romea et al. [80] use a combination of 2D and 3D features as an object model. Examples for objects from an earlier version of the system [65] are shown in Fig. 8d. The authors estimate the object's pose in a scene from a single image. The accuracy of their approach is demonstrated through a number of successful grasps.

III. GRASPING FAMILIAR OBJECTS

The idea of addressing the problem of grasping *familiar* objects originates from the observation that many of the objects in the environment can be grouped together into categories with common characteristics. In the computer vision community, objects within one category usually share similar visual properties. These can be, e.g., a common texture [81] or shape [82, 83], the occurrence of specific local features [84, 85] or their specific spatial constellation [86, 87]. These categories are usually referred to as *basic level categories* and emerged from the area of cognitive psychology [88].

For grasping and manipulation of objects, a more natural characteristic may be the functionality that they afford [30]: similar objects are grasped in a similar way or may be used to fulfill the same task (pouring, rolling, etc). The difficulty is to find a representation that encodes these common affordances. Given the representation, a similarity metric has to be

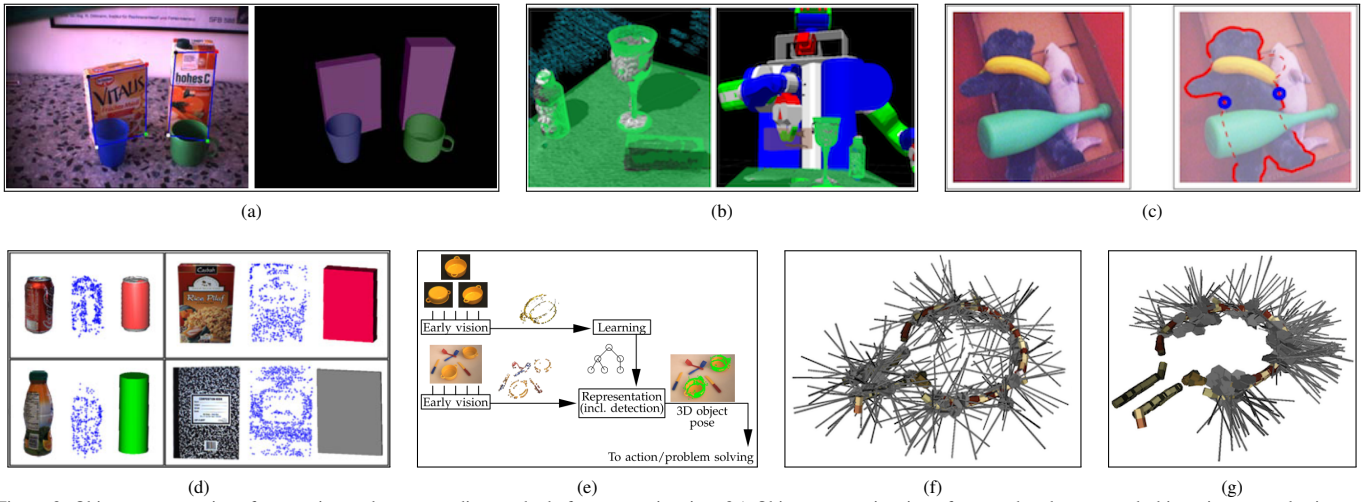


Figure 8: Object representations for grasping and corresponding methods for pose estimation. 8a) Object pose estimation of textured and untextured objects in monocular images [75]. 8b) ICP-based object pose estimation from segmented point clouds [77]. 8c) Deformable object detection and pose estimation in monocular images[55]. 8d) Multi-view object representation composed of 2D and 3D features [65]. 8e) Probabilistic and hierarchical approach towards object pose estimation [74]. 8f) Grasp Candidates linked to groups of local contour descriptors. [27]. 8g) Empirical grasp density built by trial and error. [27]

	Object-Grasp Represen.		Object Features			Grasp Synthesis							
	Local	Global	2D	3D	Multi-Modal	Heuristic	Human Demo	Labeled Data	Trial & Error	Task	Multi-Fingered	Deformable	Real Data
Song et al. [89]	✓	✓	✓	✓	✓		✓	✓		✓	✓	✓	✓
Li and Pollard [90]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
El-Khoury and Sahbani [91]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Hübner and Kragic [69]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Kroemer et al. [92]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Detry et al. [93]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Detry et al. [94]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Herzog et al. [73]	✓	✓	✓	✓	✓		✓	✓		✓	✓	✓	✓
Ramisa et al. [95]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Bouliarais et al. [96]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Montesano and Lopes [97]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Stark et al. [30]	✓	✓	✓	✓	✓		✓	✓		✓	✓	✓	✓
Saxena et al. [28]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Saxena et al. [29]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Fischinger and Vincze [98]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Le et al. [31]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Bergström et al. [99]	✓	✓	✓	✓	✓		✓	✓		✓	✓	✓	✓
Hillenbrand and Rosa [100]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Bohg and Kragic [32]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Bohg et al. [101]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Curtis and Xiao [102]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Goldfeder and Allen [103]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Marton et al. [104]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Rao et al. [105]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Speth et al. [106]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Madry et al. [107]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Kamon et al. [5]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Montesano et al. [26]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Morales et al. [25]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Pelossos et al. [20]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓
Dang and Allen [108]	✓	✓	✓	✓	✓			✓		✓	✓	✓	✓

Table II: Data-Driven Approaches for Grasping Familiar Objects

found under which objects of the same functionality can be considered as alike. The approaches discussed in this survey are summarized in Table II. All of them employ learning mechanisms and showed that they can generalize the grasp experience on training data to new but familiar objects.

A. Discriminative Approaches

First, there are approaches that learn a discriminative function to distinguish between good and bad grasp configurations. They mainly differ in what object features are used and thereby in the space over which objects are considered similar. Furthermore, they parameterize grasp candidates differently. Many of them only consider whether a specific part of the object is graspable or not. Others also learn multiple contact points or full grasp configurations. A flow-chart for the approaches discussed in the following is presented in Fig. 9.

1) Based on 3D Data: El-Khoury and Sahbani [91] distinguishes between graspable and non-graspable parts of an

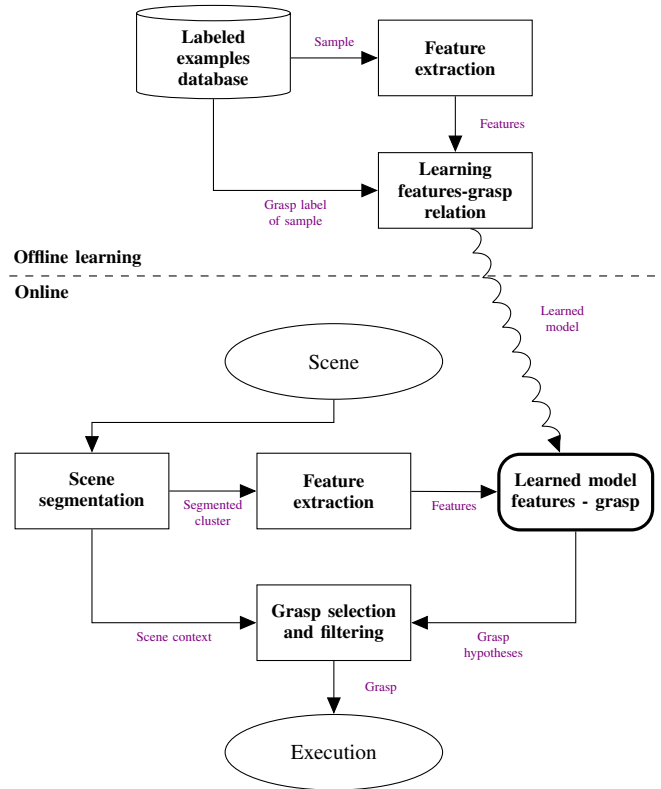


Figure 9: Typical functional flow-chart of a system that learns from labeled examples. In the offline learning phase a database is available consisting of a set of objects labeled with grasp configurations and their quality. Database entries are analyzed to extract relations between specific features and the grasps. The result is a learned model that given some features can predict grasp qualities. In the online phase, the scene is segmented and features are extracted from the scene. Given this, the model outputs a ranked set of promising grasp hypotheses. Unreachable grasps are filtered out and the best is executed. The following approaches use labeled training examples: [89, 90, 91, 69, 93, 94, 95, 96, 30, 28, 29, 31, 99, 100, 32, 101, 102, 103, 104, 107, 20, 108]

object. A point cloud of an object is segmented into parts and each part then approximated by a *superquadric* (SQ). An artificial neural network (ANN) is used to classify whether or not the part is prehensile. The ANN is trained offline on human-labeled SQs. If one of the object parts is classified as prehensile, an n-fingered force-closure grasp is synthesized on this object part. Grasp experience is therefore only used to decide where to apply a grasp, not how the grasp should be




Figure 10: a) Object model. b) Part segmentation. c) SQ approximation. d) Graspable part and contact points [91].

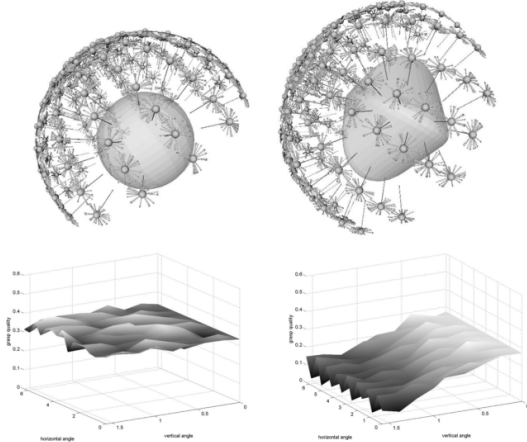


Figure 11: Top) Grasp candidates performed on SQ. Bottom) Grasp quality for each candidate [20].

configured. These steps are shown for two objects in Fig. 10.

Pelossof et al. [20] approximate an object with a single SQ. Given this, their goal is to find a suitable grasp configuration for a Barrett hand consisting of the approach vector, wrist orientation and finger spread. A *Support Vector Machine* (SVM) is trained on data consisting of feature vectors containing the SQ parameters and a grasp configuration. They are labeled with a scalar estimating the grasp quality. This training data is shown in Fig. 11. When feeding the SVM only with the shape parameters of the SQ, their algorithm searches efficiently through the grasp configuration space for parameters that maximize the grasp quality.

Both of the aforementioned approaches are evaluated in simulation where the central assumption is that accurate and detailed 3D object models are available: an assumption not always valid. An SQ is an attractive 3D representation due to its low number of parameters and high shape variability. However, it remains unclear whether an SQ could equally well approximate object shape when given real-world sensory data that is noisy and incomplete.

Hübner and Kragic [69] decompose a point cloud into a constellation of boxes. The simple geometry of a box reduces the number of potential grasps significantly. To decide which of the sides of the boxes provides a good grasp, an ANN is trained offline on synthetic data. The projection of the point cloud inside a box to its sides provides the input to the ANN. The training data consists of a set of these projections from different objects labeled with the grasp quality metrics. These are computed by GraspIt! while performing the according grasps.

Boularias et al. [96] model an object as a *Markov Random Field* (MRF) in which the nodes are points in a point cloud and edges are spanned between the six nearest neighbors of a point. The features of a node describe the local point distribution




Figure 12: Labeled training data. 12a) One example for each of the eight object classes in training data in [28] along with their grasp labels (in yellow). 12b) Positive (red) and negative examples (blue) for grasping points [96].

around that node. A node in the MRF can carry either one of two labels: a good or a bad grasp location. The goal of the approach is to find the maximum a-posteriori labeling of point clouds for new objects. Very little training data is used which is shown in Fig. 12b. A handle serves as a positive example. The experiments show that this leads to a robust labeling of 3D object parts that are very similar to a handle.

Although both approaches [69, 96] also rely on 3D models for learning, the authors show examples for real sensor data. It remains unclear how well the classifiers would generalize to a larger set of object categories and real sensor data.

Fischinger and Vincze [98] propose a *height-accumulated* feature that is similar to Haar basis functions as successfully applied by e.g. Viola and Jones [109] for face detection. The values of the feature are computed based on the height of objects above e.g. the table plane. Positive and negative examples are used to train an SVM that distinguishes between good and bad grasping points. The authors demonstrate their approach for cleaning cluttered scenes. No object segmentation is required for the approach.

2) *Based on 2D Data:* There are number of experience-based approaches that avoid the complexity of 3D data and mainly rely on 2D data to learn to discriminate between good and bad grasp locations. Saxena et al. [28] propose a system that infers a point at where to grasp an object directly as a function of its image. The authors apply machine learning to train a grasping point model on labeled synthetic images of a number of different objects. The classification is based on a feature vector containing local appearance cues regarding color, texture and edges of an image patch in several scales and of its neighboring patches. Samples from the labeled training data are shown in Fig. 12a. The system was used successfully to pick up objects from a dishwasher after it has been additionally trained for this scenario.

Instead of assuming the availability of a labeled data set, Montesano and Lopes [97] allow the robot to autonomously explore which features encode graspability. Similar to [28], simple 2D filters are used that can be rapidly convolved with an image. Given features from a region, the robot can compute the posterior probability that a grasp applied to this location will be successful. It is modeled as a Beta distribution and estimated from the grasping trials executed by the robot and their outcome. Furthermore, the variance of the posterior can be used to guide exploration to regions that are predicted to have a high success rate but are still uncertain.

Another example of a system involving 2D data and grasp experience is presented by Stark et al. [30]. Here, an object is represented by a composition of prehensile parts. These so called *affordance cues* are obtained by observing the




Figure 13: Scene segmentation with (right) and without (left) depth information. [105]




Figure 14: Three grasp candidates for a cup represented by two local patches and their major gradient as well as their connecting line. [31]

interaction of a person with a specific object. Grasp hypotheses for new stimuli are inferred by matching features of that object against a codebook of learned *affordance cues* that are stored along with relative object position and scale. How to grasp the detected parts is not solved since hand orientation and finger configuration are not inferred from the affordance cues. Similar to Boularias et al. [96], especially locally very discriminative structures like handles are well detected.

3) *Integrating 2D and 3D Data*: Although the above approaches have been demonstrated to work well in specific manipulation scenarios, inferring a full grasp configuration from 2D data alone is a highly under-constrained problem. Regions in the image may have very similar visual features but afford completely different grasps. The following approaches integrate multiple complementary modalities, 2D and 3D visual data and their local or global characteristics, to learn a function that can take more parameters of a grasp into account.

Saxena et al. [29] extend their previous work on inferring 2D grasping points by taking the 3d point distribution within a sphere centered around a grasp candidate into account. This enhances the prediction of a stable grasp and also allows for the inference of grasp parameters like approach vector and finger spread. In earlier work [28], only downward or outward grasp with a fixed pinch grasp configuration were possible.

Rao et al. [105] distinguish between graspable and non-graspable object hypotheses in a scene. Using a combination of 2D and 3D features, an SVM is trained on labeled data of segmented objects. Among those features are for example the variance in depth and height as well as variance of the three channels in the Lab color space. These are some kind of *meta* features that are used instead of the values of e.g. the color channels directly. Rao et al. [105] achieve good classification rates on object hypotheses formed by segmentation on color and depth cues. Fig. 13 compares the segmentation results with and without taking depth into account. Le et al. [31] model grasp hypotheses as consisting of two contact points. They apply a learning approach to rank a sampled set of fingertip positions according to graspability. The feature vector consists of a combination of 2D and 3D cues such as gradient angle or depth variation along the line connecting the two grasping points. Example grasp candidates are shown in Fig. 14.




Figure 15: Example shape contexts descriptor for the image of a pencil. 15a) Input image. 15b) Canny edges. 15c) Top) All vectors from one point to all other sample points. Bottom) Sampled points of the contour with gradients. 15d) Histogram with four angle and five log-radius bins comprising the vectors in 15c Bottom) [32].

Bohg and Kragic [32] propose an approach that instead of using local features, encodes global 2D object shape. It is represented relative to a potential grasping point by shape contexts as introduced by Belongie et al. [83]. Fig. 15 shows a potential grasping point and the associated feature.

Similar to Saxena et al. [29], Bergström et al. [99] see the result of the 2D based grasp selection as a way to search in a 3D object representation for a full grasp configuration. The authors extend their previous approach [32] to work on a sparse edge-based object representation. They show that integrating 3D and 2D based methods for grasp hypotheses generation results in a sparser set of grasps with a good quality.

Different from the above approaches, Ramisa et al. [95] consider the problem of manipulating deformable objects, specifically folding shirts. They aim at detecting the shirt collars that exhibit deformability but also have distinct features. The authors show that a combination of local 2D and 3D descriptors works well for this task. Results are presented in terms of how reliable collars can be detected when only a single shirt or several shirts are present in the scene.

B. Grasp Synthesis by Comparison

The aforementioned approaches study what kind of features encode similarity of objects in terms of graspability and learn a discriminative function in the associated space. The methods we review next take an exemplar-based approach in which grasp hypotheses for a specific object are synthesized by finding the most similar object or object part in a database to which good grasps are already associated.

1) *Synthetic Exemplars*: Li and Pollard [90] treat the problem of finding a suitable grasp as a shape matching problem between the human hand and the object. The approach starts off with a database of human grasp examples. From this database, a suitable grasp is retrieved when queried with a new object. Shape features of this object are matched against the shape of the inside of the available hand postures. An example is shown in Fig. 16.

Curtis and Xiao [102] build upon a knowledge base of 3D object types. These are represented by Gaussian distributions over very basic shape features, e.g., the aspect ratio of the object's bounding box, but also over physical features, e.g. material and weight. Furthermore, they are annotated with a set of representative pre-grasps. To infer a good grasp for a new object, its features are used to look up the most similar




Figure 16: Matching contact points between human hand and object [90].

object type in the knowledge base. If a successful grasp has been synthesized in this way and it is similar enough to the object type, the mean and standard deviation of the object features are updated. Otherwise a new object type is formed in the knowledge base.

While these two aforementioned approaches use low-level shape features to encode similarity between objects, Dang and Allen [108] present an approach towards *semantic* grasp planning. In this case, *semantic* refers to both, the object category and the task of a grasp, e.g. pour water, answer a call or hold and drill. The authors define a *semantic affordance map* on a prototypical object that links object features to an approach vector and to semantic grasp features (task label, joint angles and tactile sensor readings). For planning a task-specific grasp on a novel object of the same category, the object features are used to retrieve the optimal approach direction and associated grasp features. The approach vector serves as a seed for synthesizing a grasp with the Eigengrasp planner [23]. The grasp features are used as a reference to which the synthesized grasp should be similar.

Hillenbrand and Roa [100] frame the problem of transferring functional grasps between objects of the same category as pose alignment and shape warping. They assume that there is a source object given on which a set of functional grasps is defined. Pose clustering is used to align another object of the same category with it. This is followed by establishing correspondences between these two objects and subsequent transfer of the fingertip contact points from the source object to the target object. A feasible force-closed grasp is then re-planned. The authors present experimental results in which 80% of grasps on the source object were also stable and reachable on the target object. These results are however limited to the category of cups containing six instances.

All four approaches [90, 102, 108, 100] compute object features that rely on the availability of 3D object meshes. The question remains how these ideas could be transferred to the case where only partial sensory data is available to compute object features and similarity to already known objects. One idea would be to estimate full object shape from partial or multiple observations as proposed by the approaches in Sec. IV-A and use the resulting potentially noisy and uncertain meshes to transfer grasps. The above methods are also suitable to create experience databases offline that require only little labeling. In the case of category-based grasp transfer [108, 100] only one object per category would need to be associated with grasp hypotheses and all the other objects would only need a category label. No expensive grasp simulations for many grasp candidates would need to be executed as for the approaches in Section II-A1. Dang and Allen [108] followed this idea and demonstrated a few grasp trials on a real robot. 3D models




Figure 17: Lower dimensional space in which similar pairs of grasps and object parts are close to each other [93].

of the query objects are available in an experience database that was built with the proposed category-based grasp transfer. The algorithm by Papazov and Burschka [79] was used to recognize objects in the scene and estimate their pose.

Also Goldfeder and Allen [103] built their knowledge base only from synthetic data on which grasps are generated using the previously discussed Eigengrasp planner [23]. Different from the above approaches, observations made with real sensors from new objects are used to look up the most similar object and its pose in the knowledge base. This is done in two steps. In the first one, descriptors of the object views are matched against the synthetically generated ones from the database. Given the most similar object model, viewpoint constraints are exploited to find its pose and scale. Once this is found, the associated grasp hypotheses can be executed on the real object. Although experiments on a real platform are provided, it is not entirely clear how many trials have been performed on each object and how much the pose of an object was varied. As discussed earlier, the study conducted by Balasubramanian et al. [39] suggests that the employed grasp planner is not the optimal choice for synthesizing grasps that also work well in the real world. However, the methodology of retrieving similarly shaped objects from a grasp experience database is promising.

Detry et al. [93] aim at generalizing grasps to novel objects by identifying parts to which a grasp has already been successfully applied. This look-up is rendered efficient by creating a lower-dimensional space in which object parts that are similarly shaped relative to the hand reference frame are close to each other. This space is shown in Fig. 17. The authors show that similar grasp to object part configurations can be clustered in this space and form prototypical grasp-inducing parts. An extension of this approach is presented by Detry et al. [94] where the authors demonstrate how this approach can be used to synthesize grasps on novel objects by matching these prototypical parts to real sensor data.

2) *Sensor-based Exemplars*: The above mentioned approaches present promising ideas towards generalizing prior grasp experience to new objects. However, they are using 3D object models to construct the experience database. In this

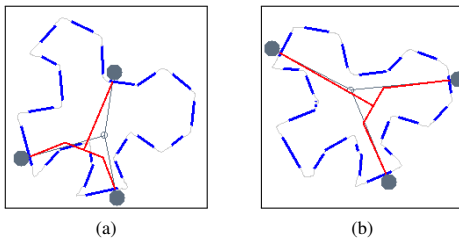


Figure 18: 18a) Successful grasp configuration for this object. 18b) Unsuccessful grasp configuration for the same object [25].

section, we review methods that generate a knowledge base by linking object representations from real sensor data to grasps that were executed on a robotic platform. Fig. 19 visualizes the flow of data that these approaches follow.

Kamon et al. [5] propose one of the first approaches towards generalizing grasp experience to novel objects. Their aim is to learn a function $f : Q \rightarrow G$ that maps object- and grasp-candidate-dependent quality parameters Q to a grade G of the grasp. An object is represented by its 2D silhouette, its center of mass and main axis. The grasp is represented by two parameters f_1 and f_2 from which in combination with the object features the fingertip positions can be computed. Learning is bootstrapped by the offline generation of a knowledge database containing grasp parameters along with their grade. This knowledge database is then updated while the robot gathers experience by grasping new objects. The system is restricted to planar grasps and visual processing of top-down views on objects. It is therefore questionable how robust this approach is to more cluttered environments and strong pose variations of the object.

Morales et al. [25] use visual feedback to infer successful grasp configurations for a three-fingered hand. The authors take the hand kinematics into account when selecting a number of planar grasp hypotheses directly from 2D object contours. To predict which of these grasps is the most stable one, a k -nearest neighbour (KNN) approach is applied in connection with a grasp experience database. The experience database is built during a trial-and-error phase executed in the real world. Grasp hypotheses are ranked dependent on their outcome. Fig. 18 shows a successful and unsuccessful grasp configuration for one object. The approach is restricted to planar objects. Speth et al. [106] showed that their earlier 2D based approach [25] is also applicable when considering 3D objects. The camera is used to explore the object and retrieve crucial information like height, 3D position and pose. However, all this additional information is not applied in the inference and final selection of a suitable grasp configuration.

The approaches presented by Herzog et al. [73] and Kroemer et al. [92] are also maintaining a database of grasp examples. They combine learning by trial and error on real world data with a part based representation of the object. There is no restriction of object shape. Each of them bootstrap the learning by providing the robot with a set of positive example grasps. However, their part representation and matching is very different. Herzog et al. [73] store a set of local templates of the parts of the object that have been in contact with the object during the human demonstration. Given a segmented object point cloud, its 3D convex hull is constructed. A template is a height map that is aligned with one polygon of this

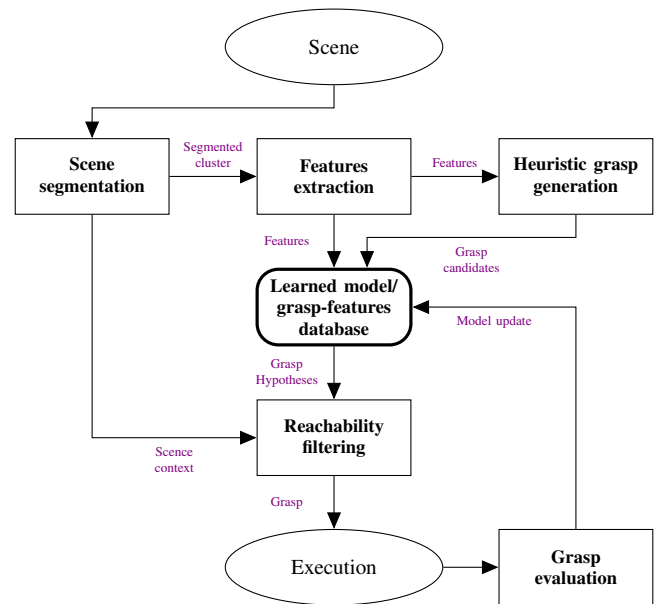


Figure 19: Typical functional flow-chart of a system that learns from trial and error. No prior knowledge about objects is assumed. The scene is segmented to obtain object clusters and relevant features are extracted. A heuristic module produces grasp candidates from these features. These candidates are ranked using a previously learned model or based on comparison to previous examples. The resulting grasp hypotheses are filtered and one of them is finally executed. The performance of the execution is evaluated and the model or memory is updated with this new experience. The following approaches can be summarized by this flow chart: [92, 73, 97, 106, 5, 26, 25]




Figure 20: Example query and matching templates [73].

hull. Together with a grasp hypotheses, they serve as positive examples. If a local part of an object is similar to a template in the database, the associated grasp hypothesis is executed. Fig. 20 shows example query templates and the matched template from the database. In case of failure, the object part is added as a negative example to the old template. In this way, the similarity metric can weight in similarity to positive examples as well as dissimilarity to negative examples. The proposed approach is evaluated on a large set of different objects and with different robots.

Kroemer et al. [92] use a pouring task to demonstrate the generalization capabilities of the proposed approach to similar objects. An object part is represented as a set of points weighted according to an isotropic 3D Gaussian with a given standard deviation. Its mean is manually set to define a part that is relevant to the specific action. When shown a new object, the goal of the approach is to find the sub-part that is most likely to afford the demonstrated action. This probability is computed by kernel logistic regression whose result depends on the weighted similarity between the considered sub-part and the example sub-parts in the database. The weight vector

is learned given the current set of examples. This set can be extended with new parts after each time an action has been executed and either succeeded or not. In this way, the weight vector can be updated to take the new experience into account. Herzog et al. [73] and Kroemer et al. [92] both do not adapt the similarity metric itself under which a new object part is compared to previously encountered examples. Instead the probability of success is estimated taken all the examples from the continuously growing knowledge base into account.

C. Generative Models for Grasp Synthesis

Very little work has been done on learning generative models of the whole grasp process. These kind of approaches identify common structures from a number of examples instead of finding a decision boundary in some feature space or directly comparing to previous examples under some similarity metric. Montesano et al. [26] provide one example in which affordances are encoded in terms of an action that is executed on an object and produces a specific effect. The problem of learning a joint distribution over a set of variables is posed as structure learning in a Bayesian network framework. Nodes in this network are formed by object, action and effect features that the robot can observe during execution. Given 300 trials, the robot learns the structure of the Bayesian network. Its validity is demonstrated in an imitation game where the robot observes a human executing one of the known actions on an object and is asked to reproduce the same observed effect when given a new object. Effectively, the robot has to perform inference in the learned network to determine the action with the highest probability to succeed.

Song et al. [89] approach the problem of inferring a full grasp configuration for an object given a specific task. As in [26], the joint distribution over the set of variables influencing this choice is modeled as a Bayesian network. Additional variables like task, object category and task constraints are introduced. The structure of this model is learned given a large number of grasp examples generated in GraspIt! and annotated with grasp quality metrics as well as suitability for a specific task. The authors exploit non-linear dimensionality reduction techniques to find a discrete representation of continuous variables for efficient and more accurate structure learning. The effectiveness of the method is demonstrated on the synthetic data for different inference tasks. The learned quality of grasps on specific objects given a task is visualized in Fig. 21.

D. Category-based Grasp Synthesis

Most of the previously discussed approaches link low-level information of the object to a grasp. Given that a novel object is similar in shape or appearance to a previously encountered one, then it is assumed that they can also be grasped in a similar way. However, objects might be similar on a different level. Objects in a household environment that share the same functional category might have a vastly different shape or appearance. However they still can be grasped in the same way. In Section III-B1, we have already mentioned the work by Dang and Allen [108], Hillenbrand and Roa [100] in which task-specific grasps are synthesized for objects of the same




Figure 21: Ranking of approach vectors for different objects given a specific task. The brighter an area the higher the rank. The darker an area, the lower the rank [89].

category. The authors assume that the category is known a-priori. In the following, we review methods that generalize grasps to familiar objects by first determining their category.

Marton et al. [110] use different 3D sensors and a thermo camera for performing object categorization. Features of the segmented point cloud and the segmented image region are extracted to train a Bayesian Logic Network for classifying object hypotheses as either boxes, plates, bottles, glasses, mugs or silverware. In [110] no grasping results are shown. A modified approach is presented in [104]. A layered 3D object descriptor is used for categorization and an approach based on the *Scale-invariant feature transform* (SIFT) [111] is applied for view based object recognition. To increase robustness of the categorization, the examination methods are run iteratively on the object hypotheses. A list of potential matching objects are kept and reused for verification in the next iteration. Objects for which no matching model can be found in the database are labeled as novel. Given that an object has been recognized, associated grasp hypotheses can be reused. These have been generated using the technique presented in [112].

Song et al. [89] treat object category as one variable in the Bayesian network. Madry et al. [107] demonstrate how the category of an object can be robustly detected given multi-modal visual descriptors of an object hypothesis. This information is fed into the Bayesian Network together with the desired task. A full hand configuration can then be inferred that obeys the task constraints. Bohg et al. [101] demonstrate this approach on the humanoid robot ARMAR III [76]. For robust object categorization the approach by Madry et al. [107] is integrated with the 3D-based categorization system by Wohlkinger and Vincze [113]. The pose of the categorized object is estimated with the approach presented by Aldoma and Vincze [114]. Given this, the inferred grasp configuration can be checked for reachability and executed by the robot.

Recently, we have seen an increasing amount of new approaches towards pure 3D descriptors of objects for categorization. Although, the following methods look promising, it has not been shown yet that they provide a suitable base for generalizing grasps over an object category. Rusu et al. [115, 116] provide extensions of [35] for either recognizing or categorizing objects and estimating their pose relative to the

	Object-Grasp Represen.		Object Features			Grasp Synthesis							
	Local	Global	2D	3D	Multi-Modal	Heuristic	Human Demo	Labeled Data	Trail & Error	Task	Multi-Fingered	Deformable	Real Data
Kraft et al. [118]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Popović et al. [119]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Bone et al. [120]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Richtsfield and Vincze [121]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Maitin-Shepard et al. [37]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Hsiao et al. [33]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Brook et al. [62]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Bohg et al. [122]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Stückler et al. [123]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Klingbeil et al. [34]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Maldonado et al. [124]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Marton et al. [112]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Lippiello et al. [125]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Dunes et al. [126]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Kehoe et al. [127]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Morales et al. [128]	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table III: Data-Driven Approaches for Grasping Unknown Objects

viewpoint. While in [116] quantitative results on real data are presented, [115] uses simulated object point clouds only. Lai et al. [36] perform object category and instance recognition. The authors learn an instance distance using the database presented in [46]. A combination of 3D and 2D features is used. However, the focus of this work is not grasping. Gonzalez-Aguirre et al. [117] present a shape-based object categorization system. A point cloud of an object is reconstructed by fusing partial views. Different descriptors (capturing global and local object shape) in combination with standard machine learning techniques are studied. Their performance is evaluated on real data.

IV. GRASPING UNKNOWN OBJECTS

If a robot has to grasp a previously unseen object, we refer to it as *unknown*. Approaches towards grasping known objects are obviously not applicable since they rely on the assumption that an object model is available. The approaches in this group also do not assume to have access to other kinds of grasp experiences. Instead, they propose and analyze heuristics that directly link structure in the sensory data to candidate grasps.

There are various ways to deal with sparse, incomplete and noisy data from real sensors such as stereo cameras: we divided the approaches into methods that i) approximate the full shape of an object, ii) methods that generate grasps based on low-level features and a set of heuristics, and iii) methods that rely mostly on the global shape of the partially observed object hypothesis. The reviewed approaches are summarized in Table III. A flow chart that visualizes the data flow in the following approaches is shown in Fig. 22.

A. Approximating Unknown Object Shape

One approach towards generating grasp hypotheses for unknown objects is to approximate objects with shape primitives. Dunes et al. [126] approximate an object with a quadric whose minor axis is used to infer the wrist orientation. The object centroid serves as the approach target and the rough object size helps to determine the hand pre-shape. The quadric is estimated from multi-view measurements of the global object shape in monocular images. Marton et al. [112] show how grasp selection can be performed exploiting symmetry by fitting a curve to a cross section of the point cloud of an object. For grasp planning, the reconstructed object is imported to

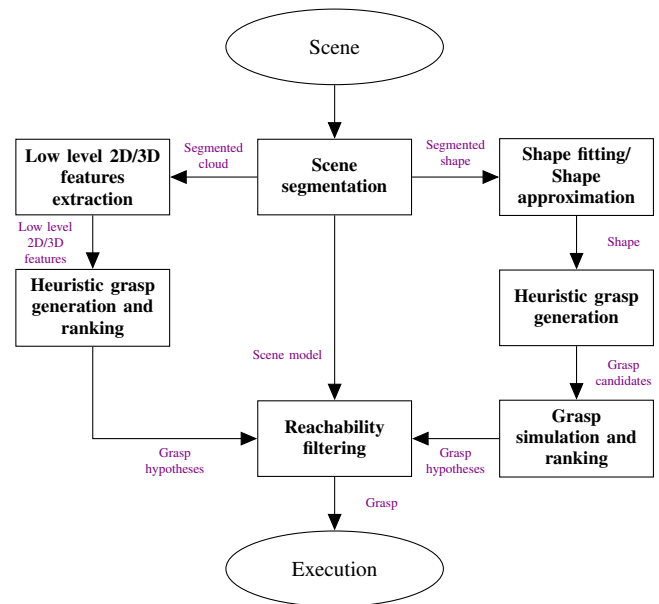


Figure 22: Typical functional flow-chart of a grasping system for unknown objects. The scene is perceived and segmented to obtain object hypotheses and relevant perceptual features. Then the system follows either the right or left pathway. On the left, low level features are used to generate heuristically a set of grasp hypotheses. On the right, a mesh model approximating the global object shape is generated from the perceived features. Grasp candidates are then sampled and executed in a simulator. Classical analytic grasp metric are used to rank the grasp candidates. Finally non reachable grasp hypotheses are filtered out, and the best ranked grasp hypothesis is executed. The following approaches use the left pathway: [118, 119, 121, 37, 33, 62, 123, 34, 124, 128]. The following approaches estimate a full object model: [120, 122, 112, 125, 126, 127]




Figure 23: Estimated full object shape by assuming symmetry. 23a) Ground Truth Mesh. 23b) Original Point Cloud. 23c) Mirrored Cloud with Original Points in Blue and Additional Points in Red. 23d) Reconstructed Mesh [122].

a simulator. Grasp candidates are generated through randomization of grasp parameters on which then the force-closure criteria is evaluated. Rao et al. [105] sample grasp points from the surface of a segmented object. The normal of the local surface at this point serves as a search direction for a second contact point. This is chosen to be at the intersection between the extended normal and the opposite side of the object. By assuming symmetry, this second contact point is assumed to have a contact normal in the direction opposite to the normal of the first contact point. Bohg et al. [122] propose a related approach that reconstructs full object shape assuming planar symmetry which subsumes all other kinds of symmetries. It takes the complete point cloud into account and not only a local patch. Two simple methods to generate grasp candidate on the resulting completed object models are proposed and evaluated. An example for an object whose full object shape is approximated with this approach is shown in Fig. 23.

Opposed to the above mentioned techniques, Bone et al. [120] make no prior assumption about the shape of the object. They apply shape carving for the purpose of grasping with




Figure 24: Unknown object shape estimated by shape carving. 24a Left) Object Image. Right) Point cloud. 24b Left) Model from silhouettes. Right) Model merged with point cloud data [120].

a parallel-jaw gripper. After obtaining a model of the object, they search for a pair of reasonably flat and parallel surfaces that are best suited for this kind of manipulator. An object reconstructed with this method is shown in Fig. 24.

Lippiello et al. [125] present a related approach for grasping an unknown object with a multi-fingered hand. The authors first record a number of views from around the object. Based on the object bounding box in each view, a polyhedron is defined that overestimates the visual object hull and is then approximated by a quadric. A pre-grasp shape is defined in which the fingertip contacts on the quadric are aligned with its two minor axes. This grasp is then refined given the local surface shape close to the contact point. This process is alternating with the refinement of the object shape through an elastic surface model. The quality of the grasps is evaluated by classic metrics. As previously discussed, it is not clear how well these metrics predict the outcome of a grasp. It remains to be shown whether grasps synthesized in this way are more successful than e.g. the just closing the fingers in the pre-grasp configuration.

B. From Low-Level Features to Grasp Hypotheses

A common approach is to map low-level 2D or 3D visual features to a predefined set of grasp postures and then rank them dependent on a set criteria. Kraft et al. [118] use a stereo camera to extract a representation of the scene. Instead of a raw point cloud, they process it further to obtain a sparser model consisting of local multi-modal contour descriptors. Four elementary grasping actions are associated to specific constellations of these features. With the help of heuristics, the large number of resulting grasp hypotheses is reduced. Popović et al. [119] present an extension of this system that uses local surfaces and their interrelations to propose and filter two and three-fingered grasp hypotheses. The feasibility of the approach is evaluated in a mixed real-world and simulated environment. The object representation and the evaluation in simulation is visualized in Fig. 25a.

Hsiao et al. [33] employ several heuristics for generating grasp hypotheses dependent on the shape of the segmented point cloud. These can be grasps from the top, from the side or applied to high points of the objects. The generated hypotheses are then ranked using a weighted list of features such as for example number of points within the gripper or distance between the fingertip and the center of the segment. Some examples for grasp hypotheses generated in this way are shown in Fig. 25b. This method is integrated into the grasping pipeline proposed in [77] for the segmented point cloud clusters that did not get recognized as a specific object.

The main idea presented by Klingbeil et al. [34] is to search for a pattern in the scene that is similar to the 2D cross section




Figure 26: PR2 gripper and associated grasp pattern [34].




Figure 27: Mapping global object shape to grasps. 27a) Simplified hand model and grasp parameters to be optimized [124]. 27b) Planar object shape uncertainty model Left) Vertices and center of mass with Gaussian position uncertainty ($\sigma = 1$). Right) 100 samples of perturbed object models [127].

of the robotic gripper interior. This is visualized in Fig. 26. The idea is similar to the work by Li and Pollard [90] as shown in Fig. 16. However, in this work the authors do not rely on the availability of a complete 3D object model. A depth image serves as the input to the method and is sampled to find a set of grasp hypotheses. These are ranked according to an objective function that takes pairs of these grasp hypotheses and their local structure into account.

Maitin-Shepard et al. [37] propose a method for grasping and folding towels that can vary in size and are arranged in unpredictable configurations. Different from the approaches discussed above, the objects are deformable. The authors propose a border detection methods that relies on depth discontinuities and then fit corners to border points. These then serve as grasping points. Examples for grasping a towel are shown in Fig. 25c. Although this approach is applicable to a family of deformable objects, it does not detect grasping points by comparing to previously encountered grasping points. Instead it directly links local structure to a grasp. For this reason, we consider it as an approach towards grasping unknown objects.

C. From Global Shape to Grasp Hypothesis

Other approaches use the global shape of an object to infer one good grasp hypothesis. Morales et al. [128] extracted the 2D silhouette of an unknown object from an image and computed two and three-fingered grasps taking into account the kinematics constraints of the hand. Richtsfeld and Vincze [121] use a segmented point cloud from a stereo camera. They search for a suitable grasp with a simple gripper based on the shift of the top plane of an object into its center of mass. A set of heuristics is used for selecting promising fingertip positions. Maldonado et al. [124] model the object as a 3D Gaussian. For choosing a grasp configuration, it optimizes a criterion in which the distance between palm and object is minimized while the distance between fingertips and the object is maximized. The simplified model of the hand and optimization variables are shown in Fig. 27a.

Stückler et al. [123] generate grasp hypotheses based on eigenvectors of the object's *footprints* on the table. Footprints




Figure 25: Generating and ranking grasp hypotheses from local object features. 25a) Generation of grasp candidates from local surface features and evaluation in simulation [119]. 25b) Generated grasp hypotheses on point cloud clusters and execution results [33]. 25c) Top) Grasping a towel from the table. Bottom) Re-grasping a towel for unfolding [37].

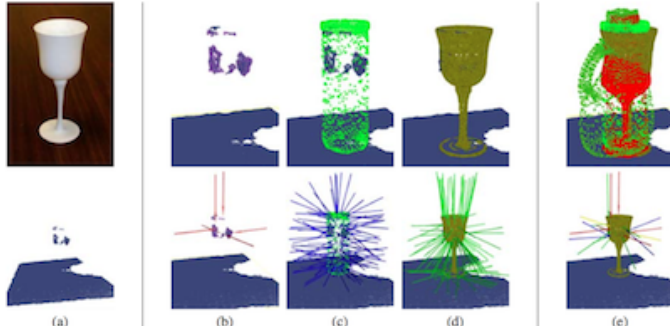


Figure 28: a) Object and point cloud. b,c,d) Object representation and grasp hypotheses. e) Overlaid representations and list of consistent grasp hypotheses [62, 130].

refer to the 3D object point cloud projected onto the supporting surface. A similar approach has been taken by [129].

Kehoe et al. [127] assume an overhead view of the object and approximate its shape with an extruded polygon. The goal is to synthesize a zero-slip push grasp with a parallel jaw gripper given uncertainty about the precise object shape and the position of its center of mass. For this purpose, perturbations of the initial shape and position of the centroid are sampled. For an example of this, see Fig. 27b. For each of these samples, the same grasp candidate is evaluated. Its quality depends on how often it resulted in force closure under the assumed model of object shape uncertainty.

V. HYBRID APPROACHES

There are a few data-driven grasp synthesis methods that cannot clearly be classified as using only one kind of prior knowledge. One of these approaches has been proposed in Brook et al. [62] with an extension in [130]. Different grasp planners provide grasp hypotheses which are integrated to reach a consensus on how to grasp a segmented point cloud. The authors show results using the planner presented in [33] for unknown objects in combination with grasp hypotheses generated through fitting known objects to point cloud clusters as described in [77]. Fig. 28 shows the grasp hypotheses for a segmented point cloud based on the input from these different planners. Another example for a hybrid approach is the work by Marton et al. [104]. A set of very simple shape primitives like boxes, cylinders and more general rotational objects are considered. They are reconstructed from segmented point clouds by analysis of their footprints. Parameters such as circle radius and the side lengths of rectangles are varied; curve parameters are estimated to reconstruct more complex rotationally symmetric objects. Given these reconstructions, a

look-up is made in a database of already encountered objects for re-using successful grasp hypotheses. In case no similar object is found, new grasp hypotheses are generated using the technique presented in [112]. For object hypotheses that cannot be represented by the simple shape primitives mentioned above, a surface is reconstructed through triangulation. Grasp hypotheses are generated using the planner presented in [33].

VI. DISCUSSION AND CONCLUSION

We have identified four major areas that form open problems in the area of robotic grasping:

a) *Object Segmentation*: Many of the approaches that are mentioned in this survey usually assume that the object to be grasped is already segmented from the background. Since segmentation is a very hard problem in itself, many methods make the simplifying assumption that objects are standing on a planar surface. Detecting this surface in a 3D point cloud and performing Euclidean clustering results in a set of segmented point clouds that serve as object hypotheses [116]. Although the dominant surface assumption is viable in certain scenarios and to shortcut the problem of segmentation, we believe that we need a more general approach to solve this.

First of all, some objects might usually occur in a specific spatial context. This can be on a planar surface, but it might also be on a shelf or in the fridge. Aydemir and Jensfelt [131] propose to learn this context for each known object to guide the search for them. One could also imagine that this context could help segmenting foreground from background. Furthermore, there are model-based object detection methods [75, 55, 74, 80, 64] that can segment a scene as a by-product of detection and without making strong assumptions about the environment. In case of unknown objects, some methods have been proposed that employ the interaction capabilities of a robot, e.g. visual fixation or pushing movements with the robot hand, to segment the scene [132, 118, 133, 134, 135]. A general solution towards object segmentation might be a combination of these two methods. The robot first interacts with objects to acquire a model. Once it has an object model, it can be used for detecting and thereby segmenting it from the background. Once again, how the object is represented plays a crucial role in the segmentation problem.

b) *Learning to Grasp*: Let us consider the goal of having a robotic companion helping us in our household. In this scenario, we cannot expect that the programmer has foreseen all the different situations that this robot will be confronted

with. Therefore, the ideal household robot should have the ability to continuously learn about new objects and how to manipulate them while it is operating in the environment. In the future, we will also not be able to rely on having 3D models readily available of all objects the robot could possibly encounter. This requires the ability to learn a model that could generalize from previous experience to new situations. Many open questions arise: How is the experience regarding one object and grasp represented in memory? How can success and failure be autonomously quantified? How can a model be learned from this experience that would generalize to new situations? Should it be a discriminative, a generative or exemplar-based model? What are the features that encode object affordances? Can these be autonomously learned? In which space are we comparing new objects to already encountered ones? Can we bootstrap learning by using simulation or by human demonstration? The methods that we have discussed in Section III about grasping familiar objects approach these questions. However, we are still far from a method that answers all of them in a satisfying way.

c) Autonomous Manipulation Planning: Recently, more complex scenarios than just grasping from a table top have been approached by a number of research labs. How a robot can autonomously sequence a set of actions to perform such a task is still an open problem. Towards this end, Tenorth et al. [136] propose a cloud robotics infrastructure under which robots can share their experience such as action recipes and manipulation strategies. An inference engine is provided for checking whether all requirements are fulfilled for performing a full manipulation strategy. It would be interesting to study how the uncertainty in perception and execution can be dealt with in conjunction with such a symbolic reasoning engine.

When considering a complex action, grasp synthesis cannot be considered as an isolated problem. On the contrary, higher-level tasks influence what the best grasp in a specific scenario might be, e.g. when grasping a specific tool. Task constraints have not yet been considered extensively in the community. Current approaches, e.g. [89, 108], achieve impressive results. As an open question stands how to scale to life-long learning.

d) Robust Execution: It has been noticed by many researchers that inferring a grasp for a given object is necessary but not sufficient. Only if execution is robust to uncertainties in sensing and actuation, a grasp can succeed with high probability. There are a number of approaches that use constant tactile or visual feedback during grasp execution to adapt to unforeseen situations [47, 33, 49, 137, 138, 139, 129]. Tactile feedback can be from haptic or force-torque sensors. Visual feedback can be the result from tracking the hand and object simultaneously. Also in this area, there are a number of open questions. How can tactile feedback be interpreted to choose an appropriate corrective action independent of the object, the task and environment? How can visual and tactile information be fused in the controller?

A. Final Notes

In this survey, we reviewed work on data-driven grasp synthesis and propose a categorization of the published work.

We focus on the type and level of prior knowledge used in the proposed approaches. We looked at what assumptions are commonly made about the objects being manipulated as well as about the complexity of scenes in which grasping is demonstrated and evaluated. We identified recent trends in the field and provided a discussion about the remaining challenges.

An important issue is the current lack of general benchmarks and performance metrics suitable for comparing the different approaches. Although various object-grasp databases are already available e.g. the Columbia Grasp database [140], the VisGraB data set [141] or the playpen data set [142] they are not commonly used for comparison. We acknowledge that one of the reasons is that grasping in itself is highly dependent on the employed sensing and manipulation hardware. There have also been robotic challenges organized such as the DARPA Arm project [143] or RoboCup@Home [144] and a framework for benchmarking has been proposed by Ulbrich et al. [145]. However, none of these successfully integrate all the subproblems relevant for grasping.

Given that data-driven grasp synthesis is an active field of research and lots of work has been reported in the area, we set up a web page that contains all the references in this survey at www.robotic-grasping.com. They are structured according to the proposed classification and tagged with the mentioned aspect. The web page will be constantly updated with the most recent approaches.

REFERENCES

- [1] A. Sahbani, S. El-Khoury, and P. Bidaud, "An overview of 3D object grasp synthesis algorithms," *Robotics and Autonomous Systems*, Aug. 2011.
- [2] K. Shimoga, "Robot Grasp Synthesis Algorithms: A Survey," *Int. Jour. of Robotic Research*, vol. 15, no. 3, pp. 230–266, 1996.
- [3] R. N. Murray, Z. Li, and S. Sastry, *A Mathematical Introduction to Robotics Manipulation*. CRC Press, 1994.
- [4] A. Bicchi and V. Kumar, "Robotic grasping and contact," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, San Francisco, Apr. 2000, invited paper.
- [5] I. Kamon, T. Flash, and S. Edelman, "Learning to grasp using visual information," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 1994, pp. 2470–2476.
- [6] S. Ekvall and D. Kragic, "Learning and Evaluation of the Approach Vector for Automatic Grasp Generation and Planning," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2007, pp. 4715–4720.
- [7] A. Morales, T. Asfour, P. Azad, S. Knoop, and R. Dillmann, "Integrated Grasp Planning and Visual Object Localization For a Humanoid Robot with Five-Fingered Hands," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Beijing, China, Oct. 2006, pp. 5663–5668.
- [8] D. Prattichizzo and J. Trinkle, *Handbook of Robotics*. Berlin, Heidelberg: Springer, 2008, ch. 28. Grasping, pp. 671–700.
- [9] D. Prattichizzo, M. Malvezzi, M. Gabiccini, and A. Bicchi, "On the manipulability ellipsoids of underactuated robotic hands with compliance," *Robotics and Autonomous Systems*, vol. 60, no. 3, pp. 337 – 346, 2012.
- [10] C. Rosales, R. Suarez, M. Gabiccini, and A. Bicchi, "On the Synthesis of Feasible and Prehensile Robotic Grasps," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012.
- [11] V.-D. Nguyen, "Constructing force-closure grasps," *Int. Jour. Robotic Res.*, vol. 7, no. 3, pp. 3–16, 1988.
- [12] M. A. Roa and R. Suárez, "Computation of independent contact regions for grasping 3-d objects," *IEEE Trans. on Robotics*, vol. 25, no. 4, pp. 839–850, 2009.
- [13] R. Krug, D. N. Dimitrov, K. A. Charusta, and B. Iliev, "On the efficient computation of independent contact regions for force closure grasps," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, 2010, pp. 586–591.
- [14] A. Rodriguez, M. T. Mason, and S. Ferry, "From Caging to Grasping," in *Robotics: Science and Systems (RSS)*, Apr. 2011.

- [15] J. Seo, S. Kim, and V. Kumar, "Planar , Bimanual , Whole-Arm Grasping," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012, pp. 3271–3277.
- [16] L. E. Zhang and J. C. Trinkle, "The Application of Particle Filtering to Grasping Acquisition with Visual Occlusion and Tactile Sensing," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012.
- [17] C. Ferrari and J. Canny, "Planning optimal grasps," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, vol. 3, may 1992, pp. 2290–2295.
- [18] A. T. Miller and P. K. Allen, "Graspit! a versatile simulator for robotic grasping," *Robotics & Automation Magazine, IEEE*, vol. 11, no. 4, pp. 110–122, 2004.
- [19] A. T. Miller, S. Knoop, H. I. Christensen, and P. K. Allen, "Automatic Grasp Planning Using Shape Primitives," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2003, pp. 1824–1829.
- [20] R. Pelossof, A. Miller, P. Allen, and T. Jebera, "An SVM learning approach to robotic grasping," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2004, pp. 3512–3518.
- [21] C. Goldfeder, P. K. Allen, C. Lackner, and R. Pelossof, "Grasp Planning Via Decomposition Trees," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2007, pp. 4679–4684.
- [22] C. Borst, M. Fischer, and G. Hirzinger, "Grasping the Dice by Dicing the Grasp," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2003, pp. 3692–3697.
- [23] M. Ciocarlie and P. Allen, "Hand posture subspaces for dexterous robotic grasping," *The Int. Jour. of Robotics Research (IJRR)*, vol. 28, pp. 851–867, July 2009.
- [24] R. Diankov, "Automated construction of robotic manipulation programs," Ph.D. dissertation, Carnegie Mellon University, Robotics Institute, Aug 2010.
- [25] A. Morales, E. Chinellato, A. Fagg, and A. del Pobil, "Using Experience for Assessing Grasp Reliability," *Int. Jour. of Humanoid Robotics*, vol. 1, no. 4, pp. 671–691, 2004.
- [26] L. Montesano, M. Lopes, A. Bernardino, and J. Santos-Victor, "Learning object affordances: From sensory–motor coordination to imitation," *IEEE Trans. on Robotics*, vol. 24, no. 1, pp. 15–26, Feb. 2008.
- [27] R. Detry, E. Bašeski, N. Krüger, M. Popović, Y. Touati, O. Kroemer, J. Peters, and J. Piater, "Learning object-specific grasp affordance densities," in *IEEE Int. Conf. on Development and Learning*, 2009, pp. 1–7.
- [28] A. Saxena, J. Driemeyer, and A. Y. Ng, "Robotic grasping of novel objects using vision," *The Int. Jour. of Robotics Research (IJRR)*, vol. 27, no. 2, pp. 157–173, Feb. 2008.
- [29] A. Saxena, L. Wong, and A. Y. Ng, "Learning Grasp Strategies with Partial Shape Information," in *AAAI Conf. on Artificial Intelligence*, 2008, pp. 1491–1494.
- [30] M. Stark, P. Lies, M. Zillich, J. Wyatt, and B. Schiele, "Functional Object Class Detection Based on Learned Affordance Cues," in *Int. Conf. on Computer Vision Systems (ICVS)*, ser. LNAI, vol. 5008. Springer-Verlag, 2008, pp. 435–444.
- [31] Q. V. Le, D. Kamm, A. F. Kara, and A. Y. Ng, "Learning to grasp objects with multiple contact points," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2010, pp. 5062–5069.
- [32] J. Bohg and D. Kragic, "Learning grasping points with shape context," *Robotics and Autonomous Systems*, vol. 58, no. 4, pp. 362–377, 2010.
- [33] K. Hsiao, S. Chitta, M. Ciocarlie, and E. G. Jones, "Contact-reactive grasping of objects with partial shape information," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Taipei, Taiwan, Oct 2010, pp. 1228–1235.
- [34] E. Klingbeil, D. Rao, B. Carpenter, V. Ganapathi, A. Y. Ng, and O. Khatib, "Grasping with application to an autonomous checkout robot," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2011, pp. 2837–2844.
- [35] R. B. Rusu, N. Blodow, and M. Beetz, "Fast point feature histograms (fpfh) for 3d registration," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*. Piscataway, NJ, USA: IEEE Press, 2009, pp. 1848–1853.
- [36] K. Lai, L. Bo, X. Ren, and D. Fox, "Sparse distance learning for object recognition combining rgb and depth information," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, Shanghai, China, May 2011, pp. 4007–4013.
- [37] J. Maitin-Shepard, M. Cusumano-Towner, J. Lei, and P. Abbeel, "Cloth grasp point detection based on Multiple-View geometric cues with application to robotic towel folding," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2010, pp. 2308–2315.
- [38] M. Beetz, U. Klank, I. Kresse, A. Maldonado, L. Mösenlechner, D. Pangercic, T. Rühr, and M. Tenorth, "Robotic Roommates Making Pancakes," in *IEEE/RAS Int. Conf. on Humanoid Robots (Humanoids)*, Bled, Slovenia, Oct. 2011.
- [39] R. Balasubramanian, L. Xu, P. D. Brook, J. R. Smith, and Y. Matsuoka, "Physical human interactive guidance: Identifying grasping principles from human-planned grasps," *IEEE Trans. on Robotics*, vol. 28, no. 4, pp. 899–910, Aug 2012.
- [40] J. Weisz and P. K. Allen, "Pose Error Robust Grasping from Contact Wrench Space Metrics," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012, pp. 557–562.
- [41] K. Konolige, "Projected texture stereo," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, May 2010, pp. 148–155.
- [42] Willow Garage, "PR2," www.willowgarage.com/pages/pr2/overview.
- [43] "ROS (Robot Operating System)," www.ros.org.
- [44] Microsoft, "Kinect- Xbox.com," www.xbox.com/en-US/KINECT.
- [45] PrimeSense, www.primesense.com.
- [46] K. Lai, L. Bo, X. Ren, and D. Fox, "A Large-Scale Hierarchical Multi-View RGB-D Object Dataset," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2011, pp. 1817–1824.
- [47] J. Felip and A. Morales, "Robust sensor-based grasp primitive for a three-finger robot hand," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. Piscataway, NJ, USA: IEEE Press, 2009, pp. 1811–1816.
- [48] K. Hsiao, P. Nangeroni, M. Huber, A. Saxena, and A. Y. Ng, "Reactive grasping using optical proximity sensors," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2009, pp. 2098–2105.
- [49] P. Pastor, L. Righetti, M. Kalakrishnan, and S. Schaal, "Online Movement Adaptation based on Previous Sensor Experiences," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, San Francisco, USA, Sep 2011, pp. 365–371.
- [50] J. Romano, K. Hsiao, G. Niemeyer, S. Chitta, and K. Kuchenbecker, "Human-inspired robotic grasp control with tactile sensing," *IEEE Trans. on Robotics*, vol. 27, no. 6, pp. 1067–1079, Dec 2011.
- [51] M. A. Goodale, "Separate Visual Pathways for Perception and Action," *Trends in Neurosciences*, vol. 15, no. 1, pp. 20–25, 1992.
- [52] U. Castiello, "The neuroscience of grasping," *Nature Reviews Neuroscience*, vol. 6, no. 9, pp. 726–736, 2005.
- [53] J. C. Culham, C. Cavina-Pratesi, and A. Singhal, "The role of parietal cortex in visuomotor control: what have we learned from neuroimaging?" *Neuropsychologia*, vol. 44, no. 13, pp. 2668–2684, 2006.
- [54] E. Chinellato and A. P. Del Pobil, "The neuroscience of vision-based grasping: a functional review for computational modeling and bio-inspired robotics," *Jour. of Integrative Neuroscience*, vol. 8, no. 2, pp. 223–254, 2009.
- [55] J. Glover, D. Rus, and N. Roy, "Probabilistic models of object geometry for grasp planning," in *Proceedings of Robotics: Science and Systems IV*. Zurich, Switzerland, June 2008.
- [56] D. Berenson, R. Diankov, K. Nishiwaki, S. Kagami, and J. Kuffner, "Grasp planning in complex scenes," in *IEEE/RAS Int. Conf. on Humanoid Robots (Humanoids)*, Dec. 2007, pp. 42–48.
- [57] M. Przybylski, T. Asfour, and R. Dillmann, "Planning grasps for robotic hands using a novel object representation based on the medial axis transform," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. IEEE, sept. 2011, pp. 1781–1788.
- [58] M. A. Roa, M. J. Argus, D. Leidner, C. Borst, and G. Hirzinger, "Power Grasp Planning for Anthropomorphic Robot Hands," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012.
- [59] R. Detry, D. Kraft, A. G. Buch, N. Krüger, and J. Piater, "Refining grasp affordance models by experience," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2010, pp. 2287–2293.
- [60] K. Huebner, K. Welke, M. Przybylski, N. Vahrenkamp, T. Asfour, D. Kragic, and R. Dillmann, "Grasping known objects with humanoid robots: A box-based approach," in *Int. Conf. on Advanced Robotics (ICAR)*, 2009, pp. 1–6.
- [61] D. R. Faria, R. Martins, J. Lobo, and J. Dias, "Extracting data from human manipulation of objects towards improving autonomous robotic grasping," *Robotics and Autonomous Systems*, Aug. 2011.
- [62] P. Brook, M. Ciocarlie, and K. Hsiao, "Collaborative grasp planning with multiple object representations," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2011, pp. 2851–2858.
- [63] J. Romero, H. Kjellström, and D. Kragic, "Modeling and evaluation of human-to-robot mapping of grasps," in *Int. Conf. on Advanced Robotics (ICAR)*. IEEE, 2009, pp. 228–233.
- [64] C. Papazov, S. Haddadin, S. Parusel, K. Krieger, and D. Burschka, "Rigid 3d geometry matching for grasping of known objects in cluttered scenes," *Int. Jour. of Robotics Research (IJRR)*, vol. 31, no. 4, pp. 538–553, Apr. 2012.
- [65] A. Collet Romea, D. Berenson, S. Srinivasa, and D. Ferguson, "Object recognition and full pose registration from a single image for robotic manipulation," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*,

- May 2009, pp. 48 – 55.
- [66] O. B. Kroemer, R. Detry, J. Piater, and J. Peters, “Combining active learning and reactive control for robot grasping,” *Robotics and Autonomous Systems*, vol. 58, pp. 1105–1116, Sep 2010.
- [67] J. Tegin, S. Ekvall, D. Kragic, B. Iliev, and J. Wikander, “Demonstration based Learning and Control for Automatic Grasping,” *Jour. of Intelligent Service Robotics*, vol. 2, no. 1, pp. 23–30, Aug 2008.
- [68] F. Stulp, E. Theodorou, M. Kalakrishnan, P. Pastor, L. Righetti, and S. Schaal, “Learning motion primitive goals for robust manipulation,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, San Francisco, USA, Sep 2011, pp. 325 – 331.
- [69] K. Hübner and D. Kragic, “Selection of Robot Pre-Grasps using Box-Based Shape Approximation,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2008, pp. 1765–1770.
- [70] M. Przybylski and T. Asfour, “Unions of balls for shape approximation in robot grasping,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*. Taipei, Taiwan: IEEE, Oct 2010, pp. 1592–1599.
- [71] R. Diankov and J. Kuffner, “Openrave: A planning architecture for autonomous robotics,” Robotics Institute, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-08-34, July 2008.
- [72] M. Do, J. Romero, H. Kjellström, P. Azad, T. Asfour, D. Kragic, and R. Dillmann, “Grasp Recognition and Mapping on Humanoid Robots,” in *IEEE/RAS Int. Conf. on Humanoid Robots (Humanoids)*, Paris, France, Dec 2009.
- [73] A. Herzog, P. Pastor, M. Kalakrishnan, L. Righetti, T. Asfour, and S. Schaal, “Template-Based Learning of Grasp Selection,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012.
- [74] R. Detry, N. Pugeault, and J. Piater, “A probabilistic framework for 3D visual object representation,” *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 31, no. 10, pp. 1790–1803, 2009.
- [75] P. Azad, T. Asfour, and R. Dillmann, “Stereo-based 6d object localization for grasping with humanoid robot systems,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2007, pp. 919–924.
- [76] T. Asfour, P. Azad, N. Vahrenkamp, K. Regenstern, A. Bierbaum, K. Welke, J. Schröder, and R. Dillmann, “Toward Humanoid Manipulation in Human-Centred Environments,” *Robotics and Autonomous Systems*, vol. 56, pp. 54–65, Jan. 2008.
- [77] M. Ciocarlie, K. Hsiao, E. G. Jones, S. Chitta, R. B. Rusu, and I. A. Sucan, “Towards reliable grasping and manipulation in household environments,” in *Int. Symposium on Experimental Robotics (ISER)*, New Delhi, India, Dec 2010.
- [78] P. J. Besl and N. D. McKay, “A method for registration of 3-d shapes,” *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 14, pp. 239–256, 1992.
- [79] C. Papazov and D. Burschka, “An efficient ransac for 3d object recognition in noisy and occluded scenes,” in *ACCV (1)*, ser. Lecture Notes in Computer Science, R. Kimmel, R. Klette, and A. Sugimoto, Eds., vol. 6492. Springer, 2010, pp. 135–148.
- [80] A. Collet Romea, M. Martinez Torres, and S. Srinivasa, “The moped framework: Object recognition and pose estimation for manipulation,” *Int. Jour. of Robotics Research (IJRR)*, vol. 30, no. 10, pp. 1284 – 1306, Sep 2011.
- [81] J. Shotton, J. Winn, C. Rother, and A. Criminisi, “Texonboost: Joint appearance, shape and context modeling for multi-class object recognition and segmentation,” in *European Conf. Computer Vision (ECCV)*, 2006, pp. 1–15.
- [82] V. Ferrari, L. Fevrier, F. Jurie, and C. Schmid, “Groups of adjacent contour segments for object detection,” *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 30, no. 1, pp. 36–51, 2008.
- [83] S. Belongie, J. Malik, and J. Puzicha, “Shape Matching and Object Recognition Using Shape Contexts,” *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 24, no. 4, pp. 509–522, 2002.
- [84] C. Dance, J. Willamowski, L. Fan, C. Bray, and G. Csurka, “Visual categorization with bags of keypoints,” in *ECCV Int. Workshop on Statistical Learning in Computer Vision*, 2004.
- [85] F.-F. L. and P. Perona, “A bayesian hierarchical model for learning natural scene categories,” in *IEEE Computer Society Conf. on Computer Vision and Pattern Recognition (CVPR)*, vol. 2, 2005, pp. 524–531.
- [86] B. Leibe, A. Leonardis, and B. Schiele, “An implicit shape model for combined object categorization and segmentation,” in *Toward Category-Level Object Recognition*, 2006, pp. 508–524.
- [87] S. Lazebnik, C. Schmid, and J. Ponce, “Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories,” in *IEEE Conf. on Computer Vision and Pattern Recognition*, vol. 2, 2006, pp. 2169–2178.
- [88] E. Rosch, C. B. Mervis, W. D. Gray, D. M. Johnson, and P. Boyes-Braem, “Basic objects in natural categories,” *Cognitive Psychology*, vol. 8, no. 3, pp. 382–439, 1976.
- [89] D. Song, C. H. Ek, K. Hübner, and D. Kragic, “Multivariate discretization for bayesian network structure learning in robot grasping,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, Shanghai, China, May 2011, pp. 1944–1950.
- [90] Y. Li and N. Pollard, “A Shape Matching Algorithm for Synthesizing Humanlike Enveloping Grasps,” in *IEEE/RAS Int. Conf. on Humanoid Robots (Humanoids)*, Dec. 2005, pp. 442–449.
- [91] S. El-Khoury and A. Sahbani, “Handling Objects By Their Handles,” in *IROS-2008 Workshop on Grasp and Task Learning by Imitation*, 2008.
- [92] O. Kroemer, E. Ugur, E. Oztop, and J. Peters, “A Kernel-based Approach to Direct Action Perception,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012.
- [93] R. Detry, C. H. Ek, M. Madry, J. Piater, and D. Kragic, “Generalizing grasps across partly similar objects,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012.
- [94] —, “Learning a Dictionary of Prototypical Grasp-predicting Parts from Grasping Experience,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2013, to appear.
- [95] A. Ramisa, G. Alenyà, F. Moreno-Noguer, and C. Torras, “Perception and Manipulation of deformable objects,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012.
- [96] A. Boularias, O. Kroemer, and J. Peters, “Learning robot grasping from 3-d images with markov random fields,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2011, pp. 1548–1553.
- [97] L. Montesano and M. Lopes, “Active learning of visual descriptors for grasping using non-parametric smoothed beta distributions,” *Robotics and Autonomous Systems*, vol. 60, pp. 452–462, 2012.
- [98] D. Fischinger and M. Vincze, “Empty the basket - a shape based learning approach for grasping piles of unknown objects,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Oct 2012, pp. 2051–2057.
- [99] N. Bergström, J. Bohg, and D. Kragic, “Integration of visual cues for robotic grasping,” in *Computer Vision Systems*, ser. Lecture Notes in Computer Science, vol. 5815. Springer Berlin / Heidelberg, 2009, pp. 245–254.
- [100] U. Hillenbrand and M. A. Roa, “Transferring functional grasps through contact warping and local replanning,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Oct 2012, pp. 2963–2970.
- [101] J. Bohg, K. Welke, B. León, M. Do, D. Song, W. Wohlkinger, M. Madry, A. Aldóma, M. Przybylski, T. Asfour, H. Martí, D. Kragic, A. Morales, and M. Vincze, “Task-based grasp adaptation on a humanoid robot,” in *Int. IFAC Symposium on Robot Control (SYROCO)*, Dubrovnik, Croatia, Sep 2012.
- [102] N. Curtis and J. Xiao, “Efficient and effective grasping of novel objects through learning and adapting a knowledge base,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, sept. 2008, pp. 2252 –2257.
- [103] C. Goldfeder and P. Allen, “Data-driven grasping,” *Autonomous Robots*, vol. 31, pp. 1–20, 2011.
- [104] Z. C. Marton, D. Pangercic, N. Blodow, and M. Beetz, “Combined 2D-3D Categorization and Classification for Multimodal Perception Systems,” *Int. Jour. of Robotics Research (IJRR)*, 2011, accepted.
- [105] D. Rao, Q. V. Le, T. Phoka, M. Quigley, A. Sudsang, and A. Y. Ng, “Grasping novel objects with depth segmentation,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Taipei, Taiwan, Oct 2010, pp. 2578–2585.
- [106] J. Speth, A. Morales, and P. J. Sanz, “Vision-Based Grasp Planning of 3D Objects by Extending 2D Contour Based Algorithms,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2008, pp. 2240–2245.
- [107] M. Madry, D. Song, and D. Kragic, “From object categories to grasp transfer using probabilistic reasoning,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012.
- [108] H. Dang and P. K. Allen, “Semantic grasping: Planning robotic grasps functionally suitable for an object manipulation task,” in *IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*, 2012, pp. 1311–1317.
- [109] P. A. Viola and M. J. Jones, “Rapid object detection using a boosted cascade of simple features,” in *IEEE Computer Society Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2001, pp. 511–518.
- [110] Z. C. Marton, R. B. Rusu, D. Jain, U. Klank, and M. Beetz, “Probabilistic categorization of kitchen objects in table settings with a composite sensor,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, St. Louis, MO, USA, Oct 2009, pp. 4777 – 4784.
- [111] D. G. Lowe, “Object Recognition from Local Scale-Invariant Features,” in *Int. Conf. on Computer Vision*, ser. ICCV ’99, vol. 2. Washington,

- DC, USA: IEEE Computer Society, 1999.
- [112] Z. C. Marton, D. Pangercic, N. Blodow, J. Kleinehellefort, and M. Beetz, "General 3D Modelling of Novel Objects from a Single View," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Taipei, Taiwan, Oct 2010, pp. 3700 – 3705.
- [113] W. Wohlkinger and M. Vincze, "Shape-based depth image to 3d model matching and classification with inter-view similarity," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, San Francisco, USA, Sept 2011, pp. 4865–4870.
- [114] A. Aldoma and M. Vincze, "Pose alignment for 3d models and single view stereo point clouds based on stable planes," *Int. Conf. on 3D Imaging, Modeling, Processing, Visualization and Transmission*, vol. 0, pp. 374–380, 2011.
- [115] R. B. Rusu, G. Bradski, R. Thibaux, and J. Hsu, "Fast 3d recognition and pose using the viewpoint feature histogram," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Taipei, Taiwan, Oct 2010, pp. 2155 – 2162.
- [116] R. B. Rusu, A. Holzbach, G. Bradski, and M. Beetz, "Detecting and segmenting objects for mobile manipulation," in *Proceedings of IEEE Workshop on Search in 3D and Video (S3DV), held in conjunction with the 12th IEEE Int. Conf. on Computer Vision (ICCV)*, Kyoto, Japan, Sep 2009.
- [117] D. I. Gonzalez-Aguirre, J. Hoch, S. Rohl, T. Asfour, E. Bayro-Corrochano, and R. Dillmann, "Towards shape-based visual object categorization for humanoid robots," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, IEEE, 2011, pp. 5226–5232.
- [118] D. Kraft, N. Pugeault, E. Baseski, M. Popovic, D. Kragic, S. Kalkan, F. Wörgötter, and N. Krueger, "Birth of the object: Detection of objectness and extraction of object shape through object action complexes," *Int. Jour. of Humanoid Robotics*, pp. 247–265, 2009.
- [119] M. Popović, G. Kootstra, J. A. Jørgensen, D. Kragic, and N. Krüger, "Grasping unknown objects using an early cognitive vision system for general scene understanding," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, San Francisco, USA, Sep 2011, pp. 987 – 994.
- [120] G. M. Bone, A. Lambert, and M. Edwards, "Automated Modelling and Robotic Grasping of Unknown Three-Dimensional Objects," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, Pasadena, CA, USA, May 2008, pp. 292–298.
- [121] M. Richtsfeld and M. Vincze, "Grasping of Unknown Objects from a Table Top," in *ECCV Workshop on 'Vision in Action: Efficient strategies for cognitive agents in complex environments'*, Marseille, France, Sep 2008.
- [122] J. Bohg, M. Johnson-Roberson, B. León, J. Felip, X. Gratal, N. Bergström, D. Kragic, and A. Morales, "Mind the Gap - Robotic Grasping under Incomplete Observation," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, May 2011.
- [123] J. Stückler, R. Steffens, D. Holz, and S. Behnke, "Real-Time 3D Perception and Efficient Grasp Planning for Everyday Manipulation Tasks," in *European Conf. on Mobile Robots (ECMR)*, Örebro, Sweden, Sep 2011.
- [124] A. Maldonado, U. Klank, and M. Beetz, "Robotic grasping of unmodeled objects using time-of-flight range data and finger torque information," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, Oct. 2010, pp. 2586–2591.
- [125] V. Lippiello, F. Ruggiero, B. Siciliano, and L. Villani, "Visual grasp planning for unknown objects using a multifingered robotic hand," *IEEE/ASME Trans. on Mechatronics*, vol. 18, no. 3, pp. 1050–1059, June.
- [126] C. Dunes, E. Marchand, C. Collwet, and C. Leroux, "Active Rough Shape Estimation of Unknown Objects," in *IEEE Int. Conf. on Intelligent Robots and Systems (IROS)*, 2008, pp. 3622–3627.
- [127] B. Kehoe, D. Berenson, and K. Goldberg, "Toward Cloud-Based Grasping with Uncertainty in Shape: Estimating Lower Bounds on Achieving Force Closure with Zero-Slip Push Grasps," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012.
- [128] A. Morales, P. J. Sanz, A. P. del Pobil, and A. H. Fagg, "Vision-based three-finger grasp synthesis constrained by hand geometry," *Robotics and Autonomous Systems*, vol. 54, no. 6, pp. 496 – 512, 2006.
- [129] X. Gratal, J. Romero, J. Bohg, and D. Kragic, "Visual servoing on unknown objects," *IFAC Mechatronics: The Science of Intelligent Machines*, 2011.
- [130] K. Hsiao, M. Ciocarlie, and P. Brook, "Bayesian grasp planning," in *ICRA 2011 Workshop on Mobile Manipulation: Integrating Perception and Manipulation*, 2011.
- [131] A. Aydemir and P. Jensfelt, "Exploiting and modeling local 3d structure for predicting object locations," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2012.
- [132] G. Metta and P. Fitzpatrick, "Better Vision through Manipulation," *Adaptive Behavior*, vol. 11, no. 2, pp. 109–128, 2003.
- [133] J. Kenney, T. Buckley, and O. Brock, "Interactive segmentation for manipulation in unstructured environments," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, ser. ICRA'09. Piscataway, NJ, USA: IEEE Press, 2009, pp. 1343–1348.
- [134] N. Bergström, C. H. Ek, M. Björkman, and D. Kragic, "Generating object hypotheses in natural scenes through human-robot interaction," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, San Francisco, USA, Sep 2011, pp. 827–833.
- [135] D. Schiebener, A. Ude, J. Morimoto, T. Asfour, and R. Dillmann, "Segmentation and learning of unknown objects through physical interaction," in *IEEE/RAS Int. Conf. on Humanoid Robots (Humanoids)*, oct. 2011, pp. 500 –506.
- [136] M. Tenorth, A. C. Perzylo, R. Lafrenz, and M. Beetz, "The RoboEarth language: Representing and Exchanging Knowledge about Actions, Objects, and Environments," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, St. Paul, MN, USA, May 2012.
- [137] Y. Bekiroglu, R. Detry, and D. Kragic, "Joint observation of object pose and tactile imprints for online grasp stability assessment," in *Manipulation Under Uncertainty (ICRA 2011 Workshop)*, 2011.
- [138] N. Hudson, T. Howard, J. Ma, A. Jain, M. Bajracharya, S. Myint, C. Kuo, L. Matthies, P. Backes, P. Hebert, T. Fuchs, and J. Burdick, "End-to-End Dexterous Manipulation with Deliberate Interactive Estimation," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2012.
- [139] M. Kazemi, J.-S. Valois, J. A. D. Bagnell, and N. Pollard, "Robust object grasping using force compliant motion primitives," in *Robotics: Science and Systems (RSS)*, Sydney, Australia, July 2012.
- [140] C. Goldfeder, M. Ciocarlie, H. Dang, and P. K. Allen, "The columbia grasp database," in *IEEE Int. Conf. on Robotics and Automation (ICRA)*. Piscataway, NJ, USA: IEEE Press, 2009, pp. 3343–3349.
- [141] Mærsk Mc-Kinney Møller Institutet, University of Southern Denmark, "VisGraB-a benchmark for vision-based grasping of unknown objects," www.robwork.dk/visgrab.
- [142] Healthcare Robotics Labs, Georgia Institute of Technology, "Pr2 playpen," ros.org/wiki/pr2_playpen.
- [143] DARPA, "ARM | Autonomous Robotic Manipulation," www.thearmrobot.com.
- [144] "RoboCup@Home," www.ai.rug.nl/robocupathome.
- [145] S. Ulbrich, D. Kappler, T. Asfour, N. Vahrenkamp, A. Bierbaum, M. Przybylski, and R. Dillmann, "The OpenGRASP benchmarking suite: An environment for the comparative analysis of grasping and dexterous manipulation," in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, sept. 2011, pp. 1761 –1767.



Jeannette Bohg is a research scientist at the Autonomous Motion Department, Max-Planck-Institute for Intelligent Systems in Tübingen, Germany. She holds a Diploma in Computer Science from the Technical University Dresden, Germany and a M.Sc. in Applied Information Technology from Chalmers in Göteborg, Sweden. In 2011, she received her PhD from the Royal Institute of Technology (KTH) in Stockholm, Sweden. Her research interest lies at the intersection between robotic grasping and Computer Vision. Specifically, she is interested in

the integration of multiple sensor modalities and information sources for enhanced scene understanding. She demonstrated how this work can be used in an active perception framework and leads to improved grasping and manipulation.



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. 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, and on development of robot simulation. He has been a Principal Investigator of the European Cognitive Systems Integrated Project GRASP and on several national and locally funded research projects.

He has served as Associated Editor for the IEEE International Conference on Robotics and Automation and for the IEEE/RSJ International Conference on Intelligent Robots and Systems. He has also served as reviewer for multiple relevant journals and conferences. He is member of the IEEE-RAS society since 1998.



Tamim Asfour is Professor at the Institute for Anthropomatics, Karlsruhe Institute of Technology (KIT). He received his diploma degree in Electrical Engineering and his PhD in Computer Science from the University of Karlsruhe. In 2003 he was awarded with the Research Center for Information Technology (FZI) prize for his outstanding Ph.D. thesis on sensorimotor control in humanoid robotics and the development of the humanoid robot ARMAR. He is leading the development team of the humanoid robot series ARMAR. He is European Chair of the

IEEE RAS Technical Committee on Humanoid Robots and member the Executive Board of the German Association of Robotics (DGR: Deutsche Gesellschaft für Robotik). His major research interest is humanoid robotics. In particular, his research focuses on learning from human observation, goal-directed imitation learning, dexterous grasping and manipulation, active vision and active touch, whole-body motion planning, system integration, robot software and hardware control architecture, motor control and mechatronics.



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. Her research is in the area of computer vision, object grasping and manipulation and human-robot interaction. Her recent work explores different learning methods for formalizing the models for integrated representation of objects and actions that can be applied on them. This work has demonstrated how robots can achieve scene understanding through active exploration and how full body tracking of humans can be made more efficient.