

Learning Probabilistic Discriminative Models of Grasp Affordances under Limited Supervision

Ayşe Naz Erkan^{†*}, Oliver Kroemer[†], Renaud Detry[‡], Yasemin Altun[†], Justus Piater[‡], Jan Peters[†]

Abstract—This paper addresses the problem of learning and efficiently representing discriminative probabilistic models of object-specific grasp affordances particularly in situations where the number of labeled grasps is extremely limited. The proposed method does not require an explicit 3D model but rather learns an implicit manifold on which it defines a probability distribution over grasp affordances. We obtain hypothetical grasp configurations from visual descriptors that are associated with the contours of an object. While these hypothetical configurations are abundant, labeled configurations are very scarce as these are acquired via time-costly experiments carried out by the robot. Kernel logistic regression (KLR) via joint kernel maps is trained to map these hypothesis space of grasps into continuous class conditional probability values indicating their achievability. We propose a soft-supervised extension of KLR and a framework to combine the merits of semi-supervised and active learning approaches to tackle the scarcity of labeled grasps. Experimental evaluation shows that combining active and semi-supervised learning is favorable in the existence of an *oracle*. Furthermore, semi-supervised learning outperforms supervised learning, particularly when the labeled data is very limited.

I. INTRODUCTION

Grasping is a fundamental skill for robots that need to interact with their environment in a flexible manner. A wide spectrum of tasks (e.g., emptying a dishwasher, opening a bottle, or using a hammer) depend on the capability to reliably grasp an object or tool as part of a larger planning framework. It is therefore imperative that the robot learns a task-independent model of an object's grasp affordances in an efficient manner. Given such a flexible model, a planner can be used to grasp and manipulate the object for a wide range of tasks. In this paper, we investigate learning probabilistic models of grasp affordances for an autonomous robot equipped with a 3D vision system (see Figure I). An object's affordances refers to the likelihood of a location on the object being graspable, from a specific orientation, by the robot.

Until this decade, the most predominant approach to grasping has been obtaining a full 3D model of the object and then employing various techniques such as friction cones [1] and form- and force- closures [2]. Given the difficulties of obtaining a 3D model with sufficient accuracy to reliably apply these techniques, designing statistical learning methods

for grasping has become an active research field [3], [4], [5], [6]. These new learning methods often employ efficient representations and vision based models, without requiring full 3D reconstruction, in order to provide a more robust alternative to traditional approaches. Much of the previous work focuses only on learning successful grasps [3], [4]. While such generative approaches can be advantageous in cases of a well-defined data distribution, it is well-known that discriminative learning methods have three main advantages over generative models [7]: Firstly, they model class-conditional probabilities of both successful and unsuccessful grasp configurations, leading to a more descriptive model and higher confidences for unsuccessful grasp regions. Secondly, they can incorporate arbitrary feature representations more flexibly. Thirdly, due to the conditional training, they are not affected from any modeling error of the data distribution.

The investigation of discriminative learning methods for grasp affordances presented in this paper continues on from previous approaches of conditional grasp affordance models, namely [5] and [6]. In [5], the authors propose extracting a set of 2D image features and apply a discriminative supervised learning method to model grasp affordance probabilities given the 2D image. In [6], this approach is extended by combining the classifier of [5] with a probabilistic classifier using a set of arm/finger kinematics features in order to identify physically impossible 2D points for the robot to reach. The strength of their approach is the combination of two important kinds of information, i.e., image and kinematic features, in a probabilistic manner.

We propose using Kernel Logistic Regression (KLR) [8] for training grasp affordance models. The main motivation behind this approach is to have the system learn a mapping from local visual features to probabilities directly, as this yields more general models than a comparison of explicit geometric models to those in an object database. While this approach enjoys the advantages of a probabilistic model, it can also capture the non-linear relations between potential grasps efficiently via kernels. This is an essential merit, since our visual grasp features are extracted from the contours of the objects and the orientation of the robot's hand, which results in the grasps lying on a non-linear manifold.

The KLR method provides a principled way of combining information from the object as well as from the robot hand via *joint kernels* [9]. By training a single classifier using joint kernels, as opposed to training two separate classifiers as was previously done [6], our approach can capture non-linear interactions of the morphology of the robot hand and the surface characteristics of the object implicitly. The system

[†] Max Plank Institute for Biological Cybernetics, Spemannstraße 38, Tübingen Germany

{naz, oliverkro, altun, jan.peters}@tuebingen.mpg.de
[‡] Department of Electrical Engineering and Computer Science Montefiore Institute, Université de Liège 4000 Liège Sart Tilman Belgium

{renaud.detry, justus.piater}@ulg.ac.be

* New York University, Computer Science Department New York, NY

therefore does not have to rely on explicit representations such as closed form geometric descriptions or libraries of feasible grasps.

Executing and labeling grasps of novel objects is a time-consuming process that requires human monitoring and may damage the objects. However, a vast number of hypothetical grasp configurations can be generated by a vision model, such as the Early Cognitive Vision reconstructor. These hypothetical grasps can not be given any confident labels, as they have not been empirically tested, and are therefore effectively unlabeled. We investigate using such *unlabeled data* in our KLR approach to reduce the number of grasps that need to be annotated for the affordance model. In particular, we propose combining a novel semi-supervised KLR method with active learning in the context of robot grasping.

Semi-supervised learning and active learning are sub-fields of machine learning that aim to handle the scarcity of labeled data. Semi-supervised learning methods, e.g., [10] and the references therein, use a large set of unlabeled data in order to improve the classification performance by revealing the underlying geometry of the data. Active learning does not rely on a source of unlabeled data, but rather assumes the existence of an annotator, commonly referred to as the *oracle*, that can provide labels to *queries*. In a robotics context, the annotator corresponds to the robot attempting to perform new grasps. The goal of active learning is to guide the robot to evaluate the most informative grasps so that the classification error is reduced with the fewest queries possible.

This framework enables the robot to learn incrementally by autonomously evaluating grasps. We provide comparisons between supervised, semi-supervised as well as a hybrid of semi-supervised and active learning setups, as minimizing the need for large amounts of labeled data is an essential concern. Experimental evaluations show not only that the proposed active learning and semi-supervised learning methods individually improve the system's performance, but that the amount of necessary annotated data is also significantly reduced when supervised learning is combined with active learning.

This paper is organized as follows, in Section II, we describe the details of the acquisition of the features. Section III gives a detailed explanation of the machine learning techniques evaluated in the context of robot grasping. Section IV overviews relevant work in the literature. In Section V, we introduce the experimental setup, give empirical results and provide a comparison of supervised, semi-supervised and active learning approaches. Finally, Section VI provides a discussion and directions for future work.




Fig. 1. Three-finger Barrett hand equipped with a 3D vision system. A table tennis paddle is used in the experiments.




Fig. 2. Kernel logistic regression algorithm is used to discriminate the successful 2(a) and unsuccessful grasps 2(b) lying on separable nonlinear manifolds. The entire hypothesis space 2(c) of potential grasp configurations extracted from pairs of ECV descriptors contains feasible grasps as well as infeasible configurations.

II. VISUAL FEATURE EXTRACTION FOR GRASPING

The inputs of our learning algorithm are represented as grasp configurations generated from Early Cognitive Vision (ECV) descriptors [11], [12], which represent short edge segments in 3D space, as described in [3]. Accordingly, an ECV reconstruction is performed. Next, pose hypotheses for potential grasps are generated from pairs of co-planar ECV descriptors. The grasp position is set to the location of one of the ECV descriptor pairs whereas the grasp orientation is computed from the normal of the plane on which these descriptors lie. The assumption is that two co-planar segments constitute a potential edge of the object that the robot hand can hold. However, this is quite optimistic as many infeasible edges and orientations will be included in the hypothesis space, see Figure II. Hence, we need a learning algorithm to discriminate between the feasible and infeasible grasps contained in this set.

Each grasp is represented with seven values in the object relative reference frame, three for the position and four for the orientation in unit length quaternions. The object relative reference frame is a coordinate system that is attached to the object such that any rigid body transformation applied to the object will also be applied to the coordinate system and objects therein.

III. LEARNING GRASP AFFORDANCES

In this section we outline the key concepts of our learning algorithm. First, we describe a kernel used as a distance metric between pairs of grasp configurations. This kernel decomposes into separate distance measures on the position and rotation parameters. We use this kernel in the KLR algorithm. Later, we propose a soft-supervised variation of the KLR algorithm so that it can accommodate unlabeled data via this distance metric. Finally, we describe the uncertainty criterion to select grasps for the queries in the active learning setting.

A. Joint Kernel

Each grasp configuration $x = (s, r)$ consists of seven parameters, i.e., three from the 3D position s of the robot hand in the object's reference frame, and four from the unit quaternions r defining the rotation. These values have different coordinate systems and have to be treated separately in order to obtain a proper distance metric. This distance metric, which indicates the similarity of two configurations, is employed for both the kernel computation and the similarity measure required by semi-supervised learning, see Equation (2). We define the joint kernel as

$$K(x_a, x_b) = \exp\left(-\frac{\|s_a - s_b\|^2}{2\sigma_s^2} - \frac{f(\theta_{ab})^2}{2\sigma_{f(\theta)}^2}\right),$$

where f is the rotational distance, σ_s and $\sigma_{f(\theta)}$ are the standard deviation of the pose and rotation distances of all pairs of samples respectively. In order to cope with the double cover property [13] of quaternions, we compute the rotational distance $f(\theta_{ab})$, as the smaller angle between the two unit length quaternions r_a and r_b . This definition allows us to use a Gaussian distribution on this rotational distance metric. Here, θ_{ab} is the angle of the 3D rotation that moves r_a to r_b , i.e., $\theta_{ab} = \theta(r_a, r_b) = \arccos(r_a^T r_b)$, and

$$f(\theta_{ab}) = \min\{\theta(r_a, r_b), \theta(r_a, -r_b)\}.$$

For further details on distance computations between unit quaternions see [13]. This joint kernel is similar to that in [14] in the way it decomposes into kernels on position and rotation features. However, there the authors employ a Dimroth-Watson distribution to get the rotational kernel as opposed to the Gaussian distribution, which is preferable due to the computational complexity of the former.

B. Kernel Logistic Regression

Our goal is to model the conditional probability distribution of grasp success $y \in \{-1, 1\}$ given a grasp configuration x as defined in Section III-A. Given labeled data $S = \{(x_i, y_i)\}_{i=1}^l$, KLR achieves this goal by maximizing the regularized log-likelihood of the data $R(w; S)$ defined by

$$R(w; S) = \sum_{i=1}^l \log p(y_i | x_i; w) - \epsilon \|w\|^2, \quad (1)$$

$$p(y = 1 | x; w) = 1 / (1 + \exp(-\langle w, f(x) \rangle)),$$

where $f(x)$ refers to an implicit feature representation induced by a kernel k and w is the corresponding weight vector. It has been shown that this optimization problem can be derived from the Maximum Entropy (MaxEnt) framework, where the goal is to find a conditional probability distribution $p(y|x)$ that matches the data (in the sense that the expected values of features with respect to $p(y|x)$ should match their empirical counterparts) while remaining as simple as possible, or equivalently maximizing the class conditional entropy $H = -\sum_y p(y|x) \log p(y|x)$,

$$\max_p E_{x \sim \tilde{p}_m} [H(p(y|x))] \text{ st.}$$

$$\|E_{x \sim \tilde{p}_m} E_{y \sim p(y|x)} [y f(x)] - E_{(x,y) \sim \tilde{p}_j} [y f(x)]\| \leq \epsilon.$$

Here \tilde{p}_j denotes the empirical joint distribution and \tilde{p}_m denotes the empirical marginal distribution over x . Defining $\tilde{p}_m(x_i) = 1/l$ and $\tilde{p}_j(x_i, y_i) = 1/l$ for all $(x_i, y_i) \in S$ and using duality techniques yield (1).

C. Semi-Supervised Kernel Logistic Regression

The duality relation mentioned in Section III-B suggests that the accuracy of KLR depends on accurate estimates of the empirical marginal and joint distributions. Our goal in the semi-supervised KLR (SSKLR) method is to use unlabeled data to reduce the sampling bias of these distributions. This can be achieved by imposing the *smoothness* of the conditional distribution in the sense that two similar grasp configurations have similar success and failure probabilities. To this end, we propose assigning soft-labels to unlabeled grasp configurations $\{x_i\}_{i=l+1}^n$ that are in the vicinity of labeled grasp configurations with respect to the manifold on which the grasp configurations lie. If the similarity metric conveys the true geometry of the grasp configurations and KLR is trained with respect to the soft success/failure assignments for unlabeled grasp configurations as well as the true labels of labeled grasp configurations, the resulting conditional probability distribution is expected to be smooth.

Similarity based soft-label assignment is equivalent to manipulating the joint distribution \tilde{p}_j to include soft labeled data. We define $\tilde{p}_m(x_i) = 1/n$ and \tilde{p}_j as

$$\tilde{p}_j(x_i, y) = \begin{cases} 1/Z_j & \text{if } 1 \leq i \leq l, y = y_i, \\ s_{ik}/Z_j & \text{if } l < i \leq n, 1 \leq k \leq l, x_i \in N_k, y = y_k, \\ 0 & \text{otherwise,} \end{cases} \quad (2)$$

where N_x is the neighborhood of x and Z_j is the normalization factor for \tilde{p}_j to be a proper probability distribution. Equation (2) allows an unlabeled data to be soft-labeled by multiple labeled data with possibly different labels, which is desirable if an unlabeled data point lies close to multiple label regions. Given these definitions and using duality, we derive the SSKLR problem as maximizing

$$R(w; S) = \sum_{i=1}^n \sum_y \tilde{p}_j(x_i, y) \langle w, y f(x_i) \rangle - \sum_{i=1}^n \tilde{p}_m(x_i) \log \sum_y (\exp \langle w, y f(x_i) \rangle) - \epsilon \|w\|^2. \quad (3)$$

The Representer Theorem [15] states that optimal weight vector of Equation (3) admits the form $w^* = \sum_{i=1}^n y \alpha_i f(x_i)$ [15]. When we substitute the solution into Equation (3), we get a convex optimization over α which can be solved using any convex optimization technique. Inference of a new grasp configuration x is given by the sign of $\sum_{i=1}^n \sum_y y \alpha_i k(x_i, x)$.

D. Uncertainty based active learning

We can employ active learning in scenarios where the robot has the means to choose what to learn. For the active selection of grasps, we use uncertainty sampling [16] which

is straightforward for probabilistic models. In this method, the algorithm queries for the grasps on which it is the least confident. Therefore, at each iteration, the algorithm requests the true label for the grasp, x^* that has the highest class conditional entropy among the set of unlabeled grasps, U

$$x^* = \operatorname{argmax}_{x \in U} H(p(y|x)).$$

In turn, the robot carries out the configuration that corresponds to x^* and labels it accordingly.

IV. RELATED WORK

Efficient representation and vision based modeling of grasp configurations is an active research field [3], [5]. We follow the methodology in [3] to obtain grasp pose candidates and orientations as described in Section II. However, the authors learn grasp densities using successful grasps only, whereas in this paper, we model the class conditional probabilities of both successful and unsuccessful grasp configurations in a discriminative manner. Furthermore, we focus on the scarcity of the labeled data points and we evaluate active and semi-supervised learning algorithms with the smallest number of annotated experiences possible.

Granville et al. [4] present a method where the robot learns a mapping from object representations to grasps from human demonstration. They cluster the orientations of grasps and each cluster is associated with a canonical approach orientation. The authors indicate that limiting the encoding to orientations or excluding position knowledge, is due to their underlying assumption that orientation and position are independent.

As labeled data collection is expensive for most robotics tasks, active learning techniques have already been considered. Salganicoff et al. [17] proposed some of the earliest work on uncertainty based active learning for vision-based grasp learning by modifying the ID3, a decision tree algorithm. Montesano and Lopes [18] also propose a method to learn local visual descriptors of good grasping points via self-experimentation. Their method associates the outputs with confidence values.

In machine learning, various methods to combine semi-supervised and active learning have been proposed to exploit the merits of both approaches [19], [20]. We attempt to be the first in the context of robotics. The active learning methodology in [20] is similar to ours, as the authors employ confidence sampling for active learning based on the probabilistic outputs of a logistic regression classifier. Their method differs from ours since they perform semi-supervised learning via self-training, whereas we propose a soft-labeling approach motivated from the maximum entropy framework.

V. EMPIRICAL EVALUATION

We have empirically evaluated the methods described in Section III on a 3-finger Barrett robot with simple objects such as a table tennis paddle. For supervised learning, we have used a Kernel Logistic Regression classifier and the joint kernel defined on position and orientation features. The

labels were collected by a human demonstrator. For the semi-supervised experiments we have used SSKLR loss given in Section III-C. Details on the experimental setup such as data collection, preprocessing, model selection and the results are given below.

A. Experimental Setup

We collected 200 samples, 100 successful (positive labels) and 100 unsuccessful (negative labels) grasps. We preprocess the data by normalizing the position parameters to zero mean and unit variance. The unit quaternions do not require preprocessing.

All experiments are carried out using the following variation of a fourfold cross validation. We have separated the 200 samples into four non-overlapping validation sets of size 50. The model variance in semi-supervised and active learning can be high as the training set is typically very small. In order to compensate for the resulting high variance, we have generated five random training sets from each of the remaining 150 samples with equal numbers of positive and negative samples. For the data set simulations of the active learning scenario, we used the rest of the samples as the active learning pool for each of the 20 training sets, $\text{trn}_1 \dots \text{trn}_{20}$. Model selection is performed over the averages of the models trained on these 20 training sets and their classification performance is assessed on the corresponding validation sets. To summarize, models trained on sets $\text{trn}_1 \dots \text{trn}_5$ are assessed on validation set val_1 , $\text{trn}_6 \dots \text{trn}_{10}$ on validation set val_2 and so on.

Our framework has two hyper-parameters which are to be set during the model selection. The first parameter, K is the size of the neighborhood in the soft-label assignment step in Equation 2. The second parameter, ϵ is the regularization constant of the kernel logistic regression algorithm. We sweep over a grid of values $K = \{10, 20, 30, 50\}$, and $\epsilon = \{10^{-2}, 10^{-3}, 10^{-4}\}$ and report the error for the hyper-parameters with the cross validation error described above. Note that, for active learning model selection is performed only once, at the initial step.

B. Evaluation on collected data sets

We evaluate the supervised and semi-supervised models with increasing sizes of labeled data. When additional data is selected with uncertainty sampling, we assess the active supervised and active semi-supervised performances. In all experiments, we train initial models with 10 randomly selected labeled samples. We perform model selection in this setup and fix the value of the hyper-parameters for the following experiments. The semi-supervised algorithm uses an additional unlabeled set of size 4000. All results are the averages over the models trained over 20 realizations of the training set and the fourfold cross validation.

First, we empirically evaluate the performance of semi-supervised learning versus supervised learning. Figure 3 shows the improvement of classification error as randomly selected samples are added to the training sets one at a time (hence, classification error of KLR and SSKLR with




Fig. 3. Supervised and semi-supervised logistic regression error on the validation sets versus the number of randomly selected labeled samples added to the initial training of size 10. Model selection is carried out at the initial step with 10 samples. 50 samples are added in an incremental manner and all models are retrained at each iteration. SSKLR uses an unlabeled training set of size 4000. K , the neighborhood size for the similarity based augmentation (Equation 2) is set to 30.

respect to increasing labeled data). As expected, when the size of the labeled data is small, semi-supervised learning is advantageous over supervised learning. The difference diminishes as the dataset gets larger.

An alternative evaluation measure is the *perplexity* of the data, $2^{H(p)} = 2^{(\sum_x -p(x) \log_2 p(x))}$ which measures the uncertainty of the predictions of the trained models. This information theoretic measure is commonly used for probabilistic models in fields such as speech recognition and natural language processing [21]. In Figure 5, we plot the perplexity of KLR and SSKLR. This figure shows that the semi-supervised model is more confident (smaller perplexity) of its predictions than the supervised model, and thus yields preferable results. We also note that the variance of perplexity across different validation sets are smaller in the case of SSKLR, when the dataset is small. This renders semi-supervised learning more robust compared to supervised learning in real-life scenarios.

Secondly, we comparatively demonstrate the impact of active learning. Figure 4 illustrates the performance of both KLR and SSKLR when incrementally trained with uncertainty based sampling. The corresponding perplexity plots are shown in Figure 6. The comparison of KLR and SSKLR in the active learning setting shows a similar behaviour to that of random selection, Figure 3 and 5. Figure 7 illustrates the classification error rate for all four scenarios together. For the supervised classifier, the improvement rate is clearly faster with active learning than random selection. A 10% error rate is achieved with 17 samples whereas to get the same error rate 40 samples are required for the random selection case.

C. On-Policy Evaluation

In order to test our approach in a real life setting we have used a second object, the watering can shown in Figure 8(a). For the experiments we have collected a total




Fig. 4. Supervised and semi-supervised classification error on the validation sets as actively selected samples are queried via uncertainty sampling. The error bars indicate one standard deviation of uncertainty over 20 models. The initial 10 labeled samples are randomly selected.




Fig. 5. Perplexity in random sampling.




Fig. 6. Perplexity in active sampling.




Fig. 7. Classification error rate for KLR, SSKLR, active-KLR and active SSKLR.




Fig. 8. The watering can used for the on-policy evaluation is shown in (a). We initiate the incremental algorithm with 20 labeled configurations shown in (b).

of 20 labeled instances of 10 successful and 10 unsuccessful configurations. Figure 8(b) illustrates these initial training set of data samples where green refers to feasible grasps and red refers to infeasible ones. Later, we trained the system incrementally with 15 more samples separately, both with random (RS) and actively sampled (AS) data. After we stopped training we have identified 10 test configurations on which the AS and RS algorithms disagree the most. When we carried out these configurations on the robot, in 10 out of 10 configurations the decision of the AS was correct and RS failed indicating that the AS is stronger in the decision boundaries.

VI. CONCLUSION AND FUTURE WORK

We have presented a probabilistic approach to model the success likelihoods of grasp configurations from a pool of hypothetical configurations extracted from ECV descriptors. The main bottleneck in the learning process is the scarcity of labeled data due to time-consumption of annotating grasps. Therefore, we have used semi-supervised and active learning approaches in the context of robot grasping. We have experimentally evaluated these approaches in two settings, in the former the data is provided only once as a batch whereas in the latter the agent has the means to query new labeled samples incrementally. We provided the results for three-finger Barrett hand and simple objects. Experimental evaluation indicates that combining semi-supervised and active learning approaches is effective in improving the robot's performance with limited supervision. However, it may not

always be possible to incrementally train a system. When that is not possible, semi-supervised learning is advantageous.

The future direction is to learn visual cues that are shared among various objects so that the grasp affordance models are not object-specific but can be generalized to many object categories. We plan to investigate this direction by using the features proposed in [6] within the joint kernel KLR framework.

REFERENCES

- [1] M. T. Mason and J. K. Salisbury, *Manipulator grasping and pushing operations*. MIT Press, 1985.
- [2] A. Bicchi and V. Kumar, "Robotic grasping and contact: a review," in *IROS*, 2000.
- [3] R. Detry, E. Baseski, M. Popovic, Y. Touati, N. Kruger, O. Kroemer, J. Peters, and J. Piater, "Learning object-specific grasp affordance densities," *International Conference on Development and Learning (ICDL'09)*, vol. 0, pp. 1–7, 2009.
- [4] C. de Granville, J. Southerland, and A. H. Fagg, "Learning grasp affordances through human demonstration," in *Proceedings of the International Conference on Development and Learning (ICDL'06)*, 2006.
- [5] A. Saxena, J. Driemeyer, and A. Y. Ng, "Robotic grasping of novel objects using vision," *The International Journal of Robotics Research*, vol. 27, no. 2, pp. 157–173, 2008.
- [6] A. Saxena, L. Wong, and A. Y. Ng, "Learning grasp strategies with partial shape information," in *AAAI*, 2008.
- [7] C. M. Bishop, *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer, 2007.
- [8] J. Zhu and T. Hastie, "Kernel logistic regression and the import vector machine," in *NIPS*. MIT Press, 2001.
- [9] G. Bakir, J. Weston, and B. Schölkopf, "Learning to find pre-images," in *NIPS*. MIT Press, 2003.
- [10] O. Chapelle, B. Schölkopf, and A. Zien, Eds., *Semi-Supervised Learning*. Cambridge, MA: MIT Press, 2006.
- [11] N. Krüger, M. Lappe, and F. Wörgötter, "Biologically motivated multimodal processing of visual primitives," *Interdisciplinary Journal of Artificial Intelligence the Simulation of Behaviour, AISB Journal*, vol. 1(5), pp. 417–427, 2004.
- [12] N. Pugeault, *Early Cognitive Vision: Feedback Mechanisms for the Disambiguation of Early Visual Representation*. Verlag Dr. Muller, ISBN 978-3-639-09357-5, 2008.
- [13] J. J. Kuffner, "Effective sampling and distance metrics for 3D rigid body path planning," in *IEEE International Conference on Robotics and Automation*, 2004, pp. 3993–3998.
- [14] R. Detry, N. Pugeault, and J. H. Piater, "A probabilistic framework for 3D visual object representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 10, pp. 1790–1803, 2009.
- [15] B. Schölkopf and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. The MIT Press, December 2001.
- [16] D. D. Lewis and J. Catlett, "Heterogeneous uncertainty sampling for supervised learning," in *ICML*, W. W. Cohen and H. Hirsh, Eds. New Brunswick, US: Morgan Kaufmann Publishers, San Francisco, US, 1994, pp. 148–156.
- [17] M. Salganicoff, L. H. Ungar, and R. Bajcsy, "Active learning for vision-based robot grasping," *Machine Learning*, vol. 23, no. 2-3, pp. 251–278, 1996.
- [18] L. Montesano and M. Lopes, "Learning object-specific grasp affordance densities," *International Conference on Development and Learning (ICDL'09)*, 2009.
- [19] X. Zhu, J. Lafferty, and Z. Ghahramani, "Combining active learning and semi-supervised learning using gaussian fields and harmonic functions," in *ICML 2003 Workshop on The Continuum from Labeled to Unlabeled Data in Machine Learning and Data Mining*, 2003, pp. 58–65.
- [20] G. Tür, D. H. Tür, and R. Schapire, "Combining active and semi-supervised learning for spoken language understanding," *Speech Communication*, vol. 45(2), pp. 171–186, 2005.
- [21] D. Jurafsky and J. Martin, *SPEECH and LANGUAGE PROCESSING: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Prentice Hall, 2000.