How can I, robot, pick up that object with my hand ?

Antonio Morales and **Pedro J. Sanz** and **Angel P. del Pobil** [1]

Abstract. This paper describes a practical approach to the robot grasping problem. An approach that is composed of two different parts. First, a vision-based grasp synthesis system implemented on a humanoid robot able to compute a set of feasible grasps and to execute any of them. This grasping system takes into account gripper kinematics constraints and uses little computational effort.

Second, a learning framework aimed at discovering the visual features that predict a reliable grasp. A grasp characterization scheme based on a set of visual features is developed in order to describe and compare grasps. In addition, a practical measure of grasp reliability is designed and implemented.

Moreover, an algorithm aimed at predicting the performance of an untested grasp using the results observed on previous similar attempts is presented. A second algorithm that actively selects the next grasp to be executed in order to improve the predictive quality of the accumulated experience is introduced, too.

An exhaustive database of experimental data is collected and used to test and validate both algorithms.

1 INTRODUCTION

The ability for manipulating and using objects are some of the most relevant skills that robots have to master in order to interact with its environment and constitute a key component for many robotic applications. Robotic manipulation can be studied at many levels, from the mechanical and physical interactions between different objects, through the proper design of mechanical robot hands, to the purposeful use of different objects. Traditionally, roboticist has focused on the former aspects, and for a good reason. Usually, complex manipulations , from the point of view of a robot, require a precise knowledge of the complex physics involved and the use of carefully designed hands. As a consequence, little attention has been paid on the, high-level, cognitive activities related with the purpose of manipulation and the nature of the manipulated objects.

This paper is the summary of a large project that has been focused on the improvement of the grasping capabilities of a robot in order to be able to grasp objects within unstructured environments. This unstructuredness is derived from the uncertain conditions of the objects to be grasped, and the little practical knowledge of the conditions that make a grasp stable.

We focus on the grasping problem, consisting of determining the kind grasp necessary to carry out certain manipulation tasks on an object. A grasp is defined both by the contacts on the objects surface and the hand and arm configuration necessary to reach them. Moreover, we focus on the pick up task. That is, we grasp the object in order to lift and transport it.

[1] Intelligent Robotics Lab., Univertitat Jaume I, Castellon, Spain. e-mail: {morales, sanzp}@icc.uji.es, pobil@ieee.org

Extensive research on this field during the last two decades has established a strong theoretical framework[15, 13, 2]. However, most of this research has been based on perfect models or ideal operational conditions. These assumptions often become unrealistic in real world applications.

Briefly, the principles of our approach are two: first, the use of sensorial, mainly visual, information to reduce the uncertainty in the environment; second, the development of a learning framework to apprehend the features of the environment that predict the outcome of the actions of the robot.

The development of this project yields two clearly separated parts: the development of a practical grasping system, and the design and implementation of a complete learning scheme.

The main features of the grasping in system (described in sec. 2), is that it makes use of sensorial inputs, mainly vision, to acquire relevant information for the grasping task, in particular the shape and location of the objects to grasp. In addition to this, we also develop a couple of grasp synthesis algorithm able to compute two and three finger grips from this information, using a small computational time, and meeting theoretical stability conditions. Finally, an algorithm to adapt the computed grips to the particular features of the gripper used is necessary, too.

Once this system is developed we face the problem of grip selection. Given a object, many different feasible grips can be performed on it, and it is thus critical to characterize the quality of candidate grips in order to execute the most reliable ones.

In this paper we introduce an ambitious approach that tries to use experience of real grasping actions to tune the behavior and the reliability assessment capabilities of the grasping system. More specifically we follow an active learning approach. According to this paradigm, the agent is allowed to interact with its environment. More specifically, it can execute actions which have an impact on the generation of training data. *Exploration* refers to the process of selecting actions in active learning. In the framework of our problem, the possible actions are the different candidate grips, at a given moment. Actions are selected by the agent in an ”intelligent” way in order to minimize the cost and duration of the learning process.

To reach this goal we develop a learning scheme that is composed of four main parts:

- A grasp characterization scheme that provides a unique description of any grasp (sec. 3). This characterization scheme is based on nine high-level vision-based descriptors. In this way, we represent each grip as a point in a multidimensional space.
- An experimental test (sec. 3.1) by means of which the robot can determine the reliability of a given grasp. This is achieved by executing the grasp and applying on it a set of practical tests to estimate the degree of stability.
- A set of techniques for predicting the reliability of a grasp from its similarity to other grasps (sec. 4). These techniques use the

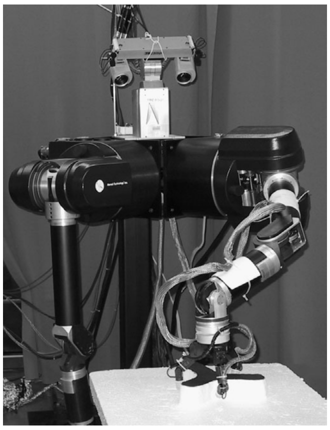

Figure 1. The UMASS Torso. A humanoid robotic system developed at the Laboratory for Perceptual Robotics in the University of Massachusetts[17].

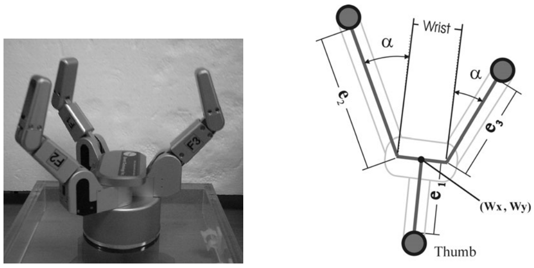

(a) Photo (b) Kinematics

Figure 2. Barrett Hand, `http://www.barretttechnology.com`

characterization schema described in previous point, and are based on pattern classification and recognition techniques.

- An exploration algorithm (sec. 5) that makes use of the problem representation previously built to decide the next action, the grasp to be executed, in order to obtain a better knowledge of the environment with a lower cost, that is, with a minimum number of executions.

Finally, we carry out an experimental validation of these methods using real data from repeated grasping actions of the robot. We collect an extensive set of samples from real grasping executions (sec. 3.2), and use them to tune, test and validate our methods (secs. 4.3 and 5.1).

2 A PRACTICAL GRASPING SYSTEM

We have implemented a robotic grasping system on the UMass humanoid torso, at the Laboratory for Perceptual Robotics in the University of Massachusetts[17]. This humanoid robot consists of two Whole Arm Manipulators from Barrett Technologies, two Barrett hands with tactile sensors at the fingertips and a BiSight stereo head.

The stereo vision system estimates the two-dimensional location of the target object on the table, and provides a monocular image for surface curvature analysis (see [12] for more details). Once a grip is selected (consisting of contact locations and a hand posture), the hand is preshaped and positioned above the object. It moves down, closes the fingers so that the object is grasped, lifted and transported to a designated location.

The main modules/steps of the functioning of this robotic grasping system are the following:

1 **Image processing:** analyzes an image of an unknown planar object, extract its contour and identify triplets of grasping regions.
2 **Grip synthesis:** determines a number of feasible grasps selecting the grasping points for each region triplet; after that, generates finger configurations that could actually be applied to the object in order to perform a grip action.
3 **Grasp selection:** perform an 'intelligent' selection of the grip to execute.
4 **Execution:** execute the grip with support of visual and tactile feedback.

Details about the first, second, and fourth sections of a system of this kind, concerned with the generation of candidate grasping configurations, are fully described in [10, 11, 12], though in the next subsections we introduce the basic concepts.

2.1 Grasp synthesis

We define a *grasp* as the set of three contact points on an object contour, and the corresponding force directions, perpendicular to the contour, which meet in the grasp force focus. We call *hand configuration* each possible grip obtained applying the kinematics constraints of a robot hand to a grasp as defined above.

To avoid misunderstandings, in all this text when referring to grasps and configurations together, the term grip is used.

We assume a real-time system acting in an unstructured environment, which detects unknown objects and, through analysis of visual data, selects and executes a stable grip of such objects.

Fast computation is necessary in order to achieve a real-time interaction with the external world. The ability to cope with uncertainties, in terms of knowledge of friction coefficients or visual and positioning errors, is a must in an uncontrolled environment.

2.2 Configurations

With a perfectly homogeneous three-finger hand, for which the fingers are all the same, the three possible ways of combining fingers with contact points in a grasp are not distinguishable. This is not the case for the Barrett Hand, for which the kinematics of the thumb is different from that of the other two fingers. A photo of the hand is reproduced in Fig. 2(a). Its kinematics are depicted in Fig. 2(b). The hand has four degrees of freedom: the three finger extensions e_1, e_2, e_3 and the spread angle θ.

For each grasp there are three possible positions of the thumb. After deciding where to place the thumb, there are still potentially infinite ways of making the hand touch the object at three contact points. However, when the action line of the thumb is fixed as well, only one solution is possible. A one-dimensional search along all possible thumb force directions gives the best Barrett Hand configuration for a grasp after the thumb position has been defined . Thus, every grasp ideally generates three different configurations, one for each thumb position. When no solutions are found for a thumb position within a grasp, due to the constraints deriving from the hand geometry and kinematics, no corresponding configurations are produced.

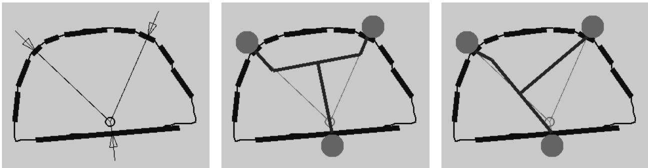

(a) Grasp (b) Config. 1 (c) Config. 2

Figure 3. Generating configurations from a grasp

Typically, dozens of configurations can be generated for an object, mostly depending on the number of regions found. In Fig. 3(b) and 3(c) two configurations generated from the grasp of Fig. 3(a) are depicted.

2.3 Two-finger grips

A particular kind of three-finger grasp is obtained as an extension of two-finger grasps. To generate a two-finger grasp, only two regions are needed, and they must be nearly parallel and facing each other (with friction, regions that are not perfectly parallel can also be used for two-finger grips).

Starting from a real two-finger grasp, if one of the regions is large enough to carry two Barrett Hand fingers, then a *virtual two-finger grasp* is generated. So, there is a special group of three-finger grasps that are computed in a completely different way, and thus have different properties and characteristics. From now on we will refer to them as *two-finger grasps*, meaning that two of the fingers are positioned on the same grasping region.

Each two-finger grasp can generate only one configuration, that is a *two-finger configuration*, as the thumb must be the finger opposed to the other two. An example of a two-finger grasp and its configuration are shown in Fig. 4 (a) and (b).

2.4 Implementation and results

The modules described in the previous sections have been implemented and tested. In a first stage they have been tested isolated, using as inputs images of different objects [10, 11]. These tests show that our implementation obtains the same results as do other classical works [5, 14] employing a few milliseconds on a common PC computer.

On a second stage they have been embedded on the control system of the UMass humanoids torso for building a complete grasping system[12]. Nearly 70 real grasp executions have been performed using this system. These experiments have consisted in placing an object in front of the robot and grasping it by executing one of the hand configurations computed for the object. The selection of the configuration to execute have been done by a human operator.[2]

These experiments show the usefulness and validity of the developed algorithms. However, they also shown the limitations of the grasping system. The first main problem is that the grasp synthesis algorithms produce a large number of possible grips, and there is no clear rule for preferring one to the others. Regarding to this problem, we propose a set quality criteria [3] that gives a value for each grasp. However, this method is not satisfactory enough since it is purely a priory, with no feedback from reality.

A second main problem, is the unexpected bad performance of some *a priory* stable grasp. Though this can be caused by the inaccuracy of the sensor inputs and the execution controllers, it also strongly affected by risks not anticipated during the stability study used to design the grasp synthesis algorithms.

These limitations have motivated the development of the learning framework that uses experience for determining the features of grasps that asses its stability and reliability.

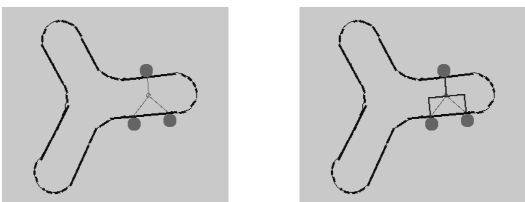

(a) Grasp (b) Configuration

Figure 4. Example of two-finger grip

3 GRASP CHARACTERIZATION SCHEME AND RELIABILITY MEASUREMENT

A characterization scheme to provide a way to describe grasps so that they can be used by the learning procedures has been developed. We have opted for a scheme that measures a set of properties of each grasp. In this way a grasp will be represented by n measurements becoming a point in an *n-dimensional* space. This scheme consists of nine of these high-level features that have been designed in order to meet the next requirements:

Vision-based computation. The features are computed from visually-extracted information.

Hand constraining. Features take into account particular characteristics of the hand.

Location and orientation invariance. Displacements and rotations of the object do not affect the values of the features.

Object independence. Grasps with the same physical properties have the same characterization independently of the object for which they are computed.

Physical meaning. Features are computed to measure physical properties relevant to grasping.

Stability and reliability. Features consider stability and reliability hazards of a grasp.

To summarize, every grip is described by a nine-elements tuple, and therefore, can be abstracted as a point in a nine-dimensions space. This space would contain all the possible grip descriptors.

Due to the limitations of space, we only describe in detail one of the grasp descriptors, as an example of the kind how these requirements are actually applied in the design of the descriptors. For further

[2] In `http://www.robot.uji.es/people/morales/experiments` there is an exhaustive description, including video recordings, of all these experiments.

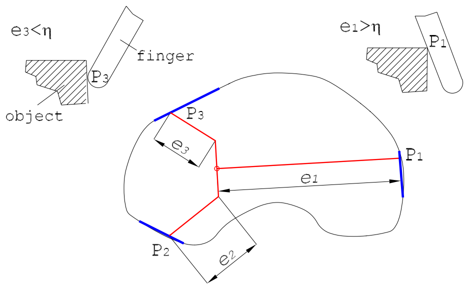

Figure 5. Geometrical representation of the Finger Limit Criterion.

details and a better explanation of all the descriptors the reader is referred to [3].

An example of grasp descriptor: The Finger Limit criterion

When trying to grip large objects, there is a limit in the extension of the fingers . Due to the way the Barrett Hand grips objects, there is a finger extension value that, if overcome, causes the grip to shift from a fingertip grip to a fingerside grip on the part edge, which is more risky and less stable although still possible (see Fig. 5). Therefore, a threshold on the maximum optimal finger extension η has been set in order to avoid marginal contacts: $q_{FG} = \epsilon_1 + \epsilon_2 + \epsilon_3$ where $\epsilon_i = (\frac{e_i - \eta}{\lambda})^2$ if $e_i > \eta$, else 0. The threshold λ is an estimation of the positioning error.

3.1 Experimental measurement of grasp reliability

A key issue in our experimental approach is the definition of a practical measurement of the reliability of a grasp. In order to do this a single object is placed on a table within the robot workspace. Using visual information the robot locates the object and computes a set of feasible grasp configurations. One of the configurations is selected, either manually by a human operator, or automatically by the robot, and executed.

If the robot has been able to lift the object safely, a set of stability tests are applied in sequence. These are aimed at measuring the stability of the current grasp. They consist of three consecutive shaking movements of the hand which are executed with an increasing acceleration. After each movement the tactile sensors are used to check whether the object has been dropped off.

This protocol provides us with a qualitative measure of the success of a grasp. Thus, an experiment may result in five different reliability classes: *E* indicates that the system was not able of lifting the object at all; *D*, *C*, *B* indicate that the object was dropped, respectively, during the first, second, or third series of shaking movements; finally *A* means the object did not fall and was returned successfully to its initial position on the table. Hence, we define $\Omega = \{A, B, C, D, E\}$ as the set of reliability classes.

3.2 Experimental sample dataset

To acquire a sample database large enough to validate the proposed methods, a series of exhaustive experiments have been carried out.

Table 1. SAMPLE DATASETS

	E	D	C	B	A	Total
LIGHT	102	84	33	27	18	**264**
LOW	38.6%	31.8%	12.5%	10.2%	6.8%	**(22)**
LIGHT	51	97	56	38	118	**360**
HIGH	14.2%	26.9%	15.6%	10.6%	32.8%	**(34)**
HEAVY	95	92	29	2	2	**220**
HIGH	43.1%	41.8%	13.2%	0.9%	0.9%	**(23)**

Sample distributions among classes for the different data sets. The figures in brackets in the “Total” column indicates the number of different grip configurations really tested.

Four real objects has been built for this experiment: two with simple shapes and two with more complex shapes. In order to build the sample database the four objects are presented to the grasping system, and a sufficiently large number of grips are executed. The reliability of these grips is obtained applying the test described in section 3.1.

A particular execution of a grip configuration can be influenced by many unpredictable factors. To avoid this problem, each grip is executed a sufficiently large number of times, by varying the location and orientation in the presentation of the object.

The number of feasible grips that are computed for each single object is usually large, varying from several dozens to more than one hundred. The repetition above mentioned could lead to a non practical number of executions, so for each object only a few configuration grips are selected to be executed. This selection consists of the most representative configurations of each object. Each configuration grip is executed 12 times, 4 times for three different orientations of the object.

Since we are also interested in studying the grasping performances in different circumstances, several characteristics of the environment are tested. These are the weight of the objects and the friction coefficient. Two qualitative categories for each of both conditions are distinguished: heavy and light objects, and high and low friction. The different weight is obtained by making two different sets of objects similar in appearance, but made of different material. Different contact friction is achieved by using a latex fingertip to envelope the fingers.

A series of experiments where done following this experimental protocol. Three different combinations of physical properties were tested: light objects and low friction (light/low), heavy objects and high friction (heavy/high); and light objects and high friction (light/high). More than eight hundred samples were obtained from this exhaustive experimentation. Table 1 shows the number of different grips executed and the percentages of grips that resulted in each class of Ω.

4 GRASP RELIABILITY PREDICTION

The learning methodology that we propose is composed of two main algorithmic components. First, a prediction scheme that computes the most likely reliability class of an untested grip, using previous experience as reference. This component assumes the existence of a set of previously executed grips having the values of the descriptors and their reliability class known.

The second component, that will be referred as exploration function, is responsible of building such set of previous attempts by successive selection of the most appropriate grip candidates. In this section we focus on the first component.

In theoretical terms a data set of previous experience is composed of N executed triplets. Each grip $g_i, i = 1 \dots N$ is described by the nine visual features $q_1, \dots q_9$ introduced in subsection 3. The 9-dimensional space G_S is formed by the ranges of the values of the features. Moreover, we have also recorded the performance of the grip and have assigned it to a class $\omega_i \in \Omega$ for each g_i.

A prediction function tries to assess the most likely reliability class for a candidate grasp $g_q \in G_S$ using as reference the previous experience. There exists a wide bibliography on the building of such functions based on the Bayesian decision theory and other non-statistical approaches. In this work we have studied three different approaches for the implementation of the prediction function.

4.1 Density estimation

The first one is a statistical parametric method[4]. It assumes that the samples that belong to every reliability category are distributed in the feature-space according to a particular density function. In our implementation this is a multivariate normal density. We use the existent datasets to estimate the parameters of this density functions, in our case, the *mean* μ_{ω_i} and the *covariance matrix* Γ_{ω_i} where $\omega_i \in \Omega$. For our purposes we are interested of the posterior probability $p(\omega_i|g_q)$.

$$p(\omega_i|g_q) \approx \exp\left(-\frac{1}{2}(x-\mu_{\omega_i})^T\Gamma_{\omega_i}^{-1}(x-\mu_{\omega_i}) - \frac{1}{2}\log\det\Gamma_{\omega_i} + \log p(\omega_i)\right) \quad (1)$$

The most likely class is, then, the one with a higher conditional probability.

4.2 Voting KNN classification rule

A prediction function has the form $F(g) = \bar{\omega}$ where $g \in G_S$ and $\bar{\omega} \in \Omega$. There exists a wide bibliography on the building of such functions based on the Bayesian decision theory [4]. In this paper we have chosen the approach of the non-parametric techniques, in particular the *voting k-nearest neighbor (KNN) rule* [6, 4], for modeling this function. The non-parametric techniques do not assume any density distribution of the features and the classes. To predict the class of a *query* point g_q, the KNN rule counts the K-nearest neighbors and chooses the class that most often appears, the most voted.

In our implementation we have introduced some modifications to the basic schema. First we use the euclidean metric for measuring the distance between the points in the G_S. We weighted the contribution of each of the KNN points according to its distance to the query point. This gives more importance to the closer points. The kernel function used is $K(d) = \frac{1}{1+(d/T)}$, where T is an adjustable parameter, and d is the distance.

We define $KNN(g_q) = \{(g_i, \omega_i), i = 1 \dots k, g_i \in G_S, \omega_i in \Omega\}$ as the k closest points to g_q and d_i their corresponding distances from g_q. The probability corresponding to a class $\bar{\omega}$ are computed using this expression:

$$p(\bar{\omega}, g_q) = \sum_{\substack{g_i \in KNN(g_q) \\ \omega_i = \bar{\omega}}} \frac{K(d_i)}{\sum_{g_j \in KNN(g_q)} K(d_j)} \quad (2)$$

Function P is also an expression of the posterior probability [6]. Our predictor would be defined as $F(g_q) = argmax_{\omega \in \Omega}\{p(\omega, g_q)\}$. That is, the class predicted ω is the one with the largest probability $p(\omega, g_q)$.

Table 2. COMPARISON USING THE LIGHT/HIGH SAMPLE DATASET

	0	**1**	**2**	**3**	**4**	$\bar{e}$
random	23.5%	26.2%	20.3%	20.7%	9.3%	**0.415**
density est.	35.0%	20.3%	15.6%	17.2%	11.9%	**0.365**
knn	51.1%	21.7%	13.3%	11.1%	2.8%	**0.223**

Percentages of misclassifications depending on the error distance. Distance 0 indicates successful classifications.

4.3 Validation and comparison of the methods

Three basic questions need to be answered about the prediction capabilities of the rules described in this section: first, are they able to predict anything at all?; second, are they able to generalize across different objects?; and third, did we have enough data to properly construct a risk function? To answer these questions we have developed a cross-validation method named *leave-one-grasp-out validation* similar to the well known *leave-one-out validation* and *n-fold cross-validation* [4]. This consists of the following steps: 1) given the whole data set, remove all the points of a particular grasp configuration and use this subset as validation set; 2) use the remaining samples for predicting the outcomes of the validation set and compute the mean error; 3) repeat steps 1) and 2) for all configurations. The validation error will be the mean error of the iterations of step 2). The goal of removing all the points of a configuration from the data set is to eliminate points similar to the query grasp in the experience dataset, thus testing generalization properties.

The error metric is based on the concept of *misclassification error distance*. The distance between two consecutive classes is defined as 1, that between A and C as 2, etc. In this way define the error distance $e(g_q) = \{0, \dots, 4\}$ for the prediction of a given query grip. Given a set of predictions $G = \{g_i, i = 1 \dots n\}$, we define the average error metric $\bar{e}(G) = \sum e(g_i)/4$.

Moreover, we compare these prediction methods against the theoretical results that would be obtained by a prediction method that would have chosen *randomly* the predicted class.

The performance of this methods is obtained using the validation procedure described above. Table tab:fullsize shows the results obtained for one the sample datasets (light objects and high friction). The results in the other two cases were similar. The figures obtained indicate the the KNN prediction function improves clearly the other prediction functions, moreover it obtains better results that the naive random prediction.

This results show its validity of KNN function for prediction within this problem. Finally, we also measure the evolution of the performance of the KNN prediction method with different sizes of the sample dataset (fig. 6) and we conclude that the performance improves when the available experience dataset is larger[8, 9].

5 ACTIVE LEARNING FOR EXPERIENCE ACQUISITION

The results of the analysis of prediction methods indicate that it is possible to predict reasonably well the reliability class of a grasp if enough previous experience is available. In this part of the project we question if it is possible to reach a similar degree of performance with less experience. In particular we aim at designing an exploration procedure that guides the continuous execution of grasps with the goal of acquiring the maximum performance possible with the minimum number of executed trials.

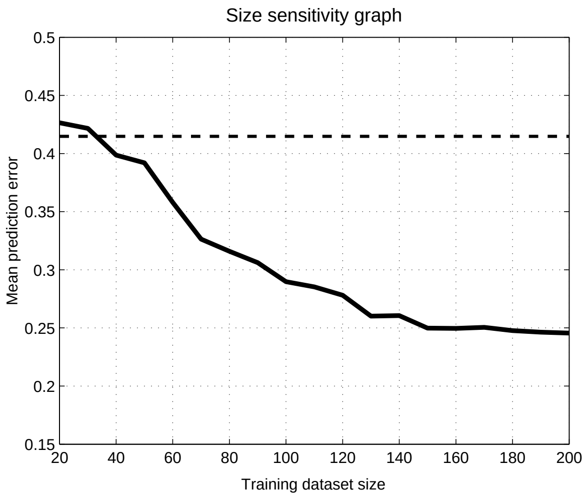

Figure 6. Evolution of the error when the size of the available data set varies. The Solid black line represents the errors obtained by the KNN prediction method, while the dashed line is the threshold of the random error.

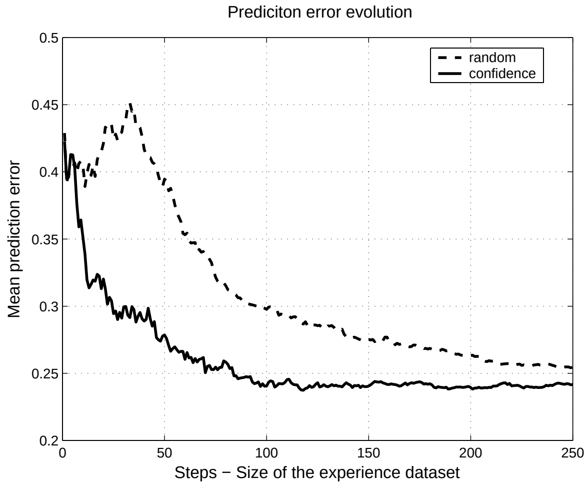

Figure 7. Evolution of the prediction error using the Light/High sample dataset.

In practice, the task of such exploration procedure is to select the next grasp to execute among a set of candidates. This selection must be done in order to improve the predictive capabilities of the stored experience, i.e., the set of already executed grasps.

The algorithm we propose assumes that at any point during the training of the grasping system a set of candidate grips $g_i \in G_S$ is proposed and the algorithm has to select the next grasp to be executed. To accomplish this task, it can take into account the results of previous experiments.

The approach we propose for the selection is inspired in the idea hinted by Thrun [16], "queries are favored that have the least predictable outcome". That is, those candidates which category is less predictable are preferred. This idea is based on the intuition that such candidates are located in areas where the implicit model represented by the experience data set is less clear.

We implement this idea by defining the term *prediction confidence*. For every grip candidate g_i, a class $\omega_i \in \Omega$ is computed using the KNN prediction scheme defined in the previous section. The confidence of that prediction is simply $p(\omega_i, g_i)$. In formal terms the prediction confidence for a grip g_q is defined as $F_{conf}(g_q) = \max\{p(\omega|g_q)\}, \omega \in \Omega$. We use only the KNN prediction function since it proved to obtain the better results in the analyses described in previous section.

Once defined the notion of confidence, it is easy to describe the exploration function. It chooses the candidate with a minimum confidence value. Given a set of m grasp candidates $G_q = \{g_1, \ldots, g_m\} \subset G_S$, the exploration function is defined as,

$$F_{exp}(G_q) = \operatorname*{argmin}_{g_i \in G_q} F_{pred}(g_i) \qquad (3)$$

Hereinafter, we will refer to this method as the *minimum confidence exploration*, or simply the risk exploration function.

Summarizing, this procedure procedure predicts a query point based on its similarity to its neighbors. This is a case of *instance-based* also known as *memory-based learning* [1], which is a numeric variant of the more symbolic *case-based reasoning* [18]. These approaches do not construct an explicit representation of the target function when training samples are provided, but simply store them.

5.1 Validation of the exploration procedure

The performance of the exploration/selection procedure is measured by the predictive capability of the set of samples selected/executed, which reliability class is known. This can be easily measured by using this dataset to predict the class of the samples contained in a secondary *validation test*. We have designed a validation framework that follows this principle. In its design we also take inspiration from the running of the robot in the training environment or in a learning experiment. In this situation the robot will execute a sequence of *selection-execution* actions. Each of these actions will follow the next steps:

1. One or more objects appear in the workspace of the robot. The grasps for them are computed. These are the grasp candidates
2. The robot selects one of them by using the exploration function.
3. The grasp is executed and the reliability test is applied.
4. The new grasp and the performance outcome are added to the experience dataset.

For the execution of the validation algorithm, we take the whole sample dataset available and extract a subset, *validation dataset* from it. The remaining is used as a *pool dataset*. In a sequence of selection steps, a small subset of *candidate samples* are extracted randomly from this pool. The exploration function, in our case, the *minimum confidence* rule, is applied to select one of these candidates. The selected candidate is added to the *experience dataset* and the discarded candidates are returned back to the pool. The performance measurement is done by using the samples in the *experience dataset* for predicting the samples in the *validation set*. The sequence is repeated until the pool dataset is emptied or it contains few samples.

This procedure is repeated a sufficiently large number of times varying the contents of the pool and validation datasets and the performance measurements for each size of the *experience dataset* are averaged.

Figure 7 presents the evolution of the prediction error for different sizes of the Light/High sample dataset, that is equivalent to the number of steps of the algorithm described in the above paragraphs. The graph in dashed lines shows the evolution of the prediction error when the sample to execute is selected randomly among the set of candidates. This case would represent the evolution when no specific exploration rule is applied. From this graph, and similar ones obtained using the other sample datasets, we conclude that the proposed exploration procedure clearly improves the random selection function, and is able to reach maximum performance levels with less than a hundred trials.

6 CONCLUSION

This paper is the summary of large project [7] aimed at improving the grasping skills of a robot to work in the face of unknown conditions and uncertainty. We have approached this problem following two different ways.

The goal of the first part is to develop and to implement a grasping system able to use vision for extracting and using relevant information for grasp synthesis. The visual approach allows the system to deal with unknown objects. We have already emphasized the inclusion of the particular kinematics of the robotic hand within the grasp synthesis algorithms. As a result we have developed a couple of algorithms able to compute two and three-finger grasp for unknown objects using vision as only input, and a third algorithm that constrains their results to the hand geometry.

Moreover, these algorithms have resulted to be fast and suitable to use in real-time manipulation activities. Finally, a complete implementation on the UMass Torso has shown the strengths and limitations of the grasping system. This observations have motivated the approached followed in the next part of the project.

In this second part, we have presented the development of a learning framework for assessing robot grasp reliability. This framework is based on two learning algorithms and a representation of the data, built on a grasp characterization scheme composed of nine high level vision-based descriptors.

The first algorithm is aimed at predicting the reliability of an untested grip from its comparison to previous recorded attempts. The second algorithm, based on the idea of active learning, is an exploration rule that has to select among a set of candidate grips the next one to execute, having the goal of improving the predictive performance of the accumulated experience.

An experimental measurement of the reliability of a grasp have been developed and used to gather an exhaustive database of sample grips. Several validation frameworks that make use of this database, have been designed to test and validate the usefulness and properties of the proposed algorithms.

The results have proved that the algorithms proposed in this work are able to carry out the expected tasks with a reasonable level of performance, despite the complex and unpredictable nature of the task space.

Moreover, the experimental and practical approach followed indicates a possible path that service robotic applications willing to be used in every-day human environments could follow. The inclusion of active learning schemes in robot systems is an appropriate way to improve their adaptability to unmodeled or partially unknown environments and, thus, building real intelligent robot systems.

ACKNOWLEDGMENTS

This work could not have been possible without the priceless help and collaboration of Andy Fagg and Eris Chinellato.

This work has been funded in part by the Ministerio de Ciencia y Tecnología under project DPI2001-3801, by the Generalitat Valenciana under projects GRUPOS 03/153, GV01-244, CTIDIA/2002/195, by the Fundació Caixa-Castelló under project P1-1B2001-28. The second author has been supported by the Ministerio de Ciencia y Tecnología under a FPI Program graduate fellowship.

This work is also partly funded by grants CISE/CDA-9703217 and IRI-9704530, DARPA MARS DABT63-99-1-0004, and NASA/RICIS. The authors wish to thank Roderic Grupen, David Wheeler, Robert Platt and Danny Radhakrishnan, who have provided much of the foundation which allowed the robot experiments.

REFERENCES

[1] D.W. Aha, 'Lazy learning', *Artificial Intelligence Review*, **11**, 7–10, (1997).

[2] A. Bicchi and V. Kumar, 'Robotic grasping and contact: A review', in *IEEE Intl. Conf. on Robotics and Automation*, (April 2000).

[3] E. Chinellato, A. Morales, R.B. Fisher, and A.P. del Pobil, 'Visual features for characterizing robot grasp quality', *IEEE Transactions on Systems, Man and Cybernetics (Part C)*, (2004). In Press.

[4] R.O. Duda, P.E. Hart, and D.G. Stork, *Pattern Classification*, John Wiley & Sons, Inc., 2nd edn., 2001.

[5] B. Faverjon and J. Ponce, 'On computing two-finger force-closure grasps of curved 2D objects', in *IEEE Intl. Conf. on Robotics and Automation*, pp. 424–429, (1991).

[6] T. M. Mitchell, *Machine Learning*, McGraw Hill, 1997.

[7] A. Morales, *Learning to predict grasp reliability with a multifinger robot hand by using visual features*, Ph.D. dissertation, Departament of Computer and Engineering Science, Universitat Jaume I, Castellón, Spain, January 2004. `http://www.robot.uji.es/people/morales/thesis`.

[8] A. Morales, Chinellato E, A.H. Fagg, and A.P. del Pobil, 'Experimental prediction of the preformance of grasps tasks from visual fetures', in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pp. 3423–3428, Las Vegas, Nevada, (October 2003).

[9] A. Morales, Chinellato E, A.H. Fagg, and A.P. del Pobil, 'Using experience for assessing grasp reliability', in *International Conference on Humanoid Robots (Humanoids 2003)*, Karlsruhe, Germany, (October 2003). On CD-ROM.

[10] A. Morales, G. Recatalá, P.J. Sanz, and A.P del Pobil, 'Heuristic vision-based computation of planar antipodal grasps on unknown objects', in *IEEE Intl. Conf. on Robotics and Automation*, pp. 583–588, Seoul, Korea, (May 2001).

[11] A. Morales, P.J. Sanz, and A.P. del Pobil, 'Vision-based computation of three-finger grasps on unknown planar objects', in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pp. 1693–1698, Lausanne, Switzerland, (October 2002).

[12] A. Morales, P.J. Sanz, A.P del Pobil, and A.H. Fagg, 'An experiment in constraining vision-based finger contact selection with gripper geometry', in *IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, pp. 1711–1716, Lausanne, Switzerland, (October 2002).

[13] A.M. Okamura, N.S. Smaby, and M.R Cutkosky, 'An overview of dextereous manipulation', in *IEEE Intl. Conf. on Robotics and Automation*, pp. 255–260, San Francisco, California, (April 2000).

[14] J. Ponce and B. Faverjon, 'On computing three-finger force-closure grasps of polygonal objects', *IEEE Transactions on Robotics and Automation*, **11**(6), 868–881, (1995).

[15] K.B. Shimoga, 'Robot grasp synthesis: A survey', *International Journal of Robotics Research*, **3**(15), 230–266, (June 1996).

[16] S. Thrun, 'Exploration in active learning', in *The Handbook of Brain Theory and Neural Networks*, ed., Michael A. Arbib, 381–384, MIT Press, (1995).

[17] UMass humaniod torso. `http://www-robotics.cs.umass.edu`.

[18] D. L. Waltz, 'Memory-based reasoning', in *The Handbook of Brain Theory and Neural Networks*, ed., Michael A. Arbib, 661–662, MIT Press, (1995).