

Dexterous Grasping via Eigengrasps: A Low-dimensional Approach to a High-complexity Problem

Matei Ciocarlie Corey Goldfeder Peter Allen

Columbia University, New York, USA

E-mail: {cmatei,coreyg,allen}@cs.columbia.edu

Abstract—In this paper, we build upon recent advances in neuroscience research which have shown that control of the human hand during grasping is dominated by movement in a configuration space of highly reduced dimensionality. We extend this concept to robotic hands and show how a similar dimensionality reduction can be defined for a number of different hand models. This framework can be used to derive optimization algorithms that simplify the task of finding stable grasps even for highly complex hand designs. Furthermore, it offers a unified approach for controlling different hands, even if the kinematic structures of the models are significantly different. We illustrate these concepts by building a comprehensive grasp planner that can be used on a large variety of robotic hands under various constraints.

I. INTRODUCTION

One of the hardest problems in robotic grasping is the creation of control algorithms for new hand designs that are beginning to rival the human hand in complexity. Researchers studying robotic grasping have struggled to at least partially replicate human versatility when designing artificial counterparts. This can be seen as a natural consequence of the current demand to push robots out of controlled environments and into the complex and cluttered human surroundings that characterize everyday life.

If we wish to reproduce human-like grasping it would seem natural to draw inspiration not only from the hardware of the human hand, but also from the software; that is, the way the hand is controlled by the brain. This may initially sound like an overly lofty goal: a large part of the human cortex is dedicated to grasping and manipulation, and it would seem reasonable to assume that all of this cognitive machinery is dedicated to finely controlling individual joints and generating highly flexible hand postures. However, recent results in neuroscience research [1] point to the contrary, emphasizing that a majority of the human hand control during common grasping tasks lacks individuation in finger movements.

Attempts to formalize human tendency to simplify the space of possible grasps can be traced back to Napier’s pioneering grasp taxonomy [2], updated later by Cutkosky [3]. While the configuration space of dexterous hands is high-dimensional and very difficult to search directly, these studies show that most useful grasps can be found in the vicinity of a small number of discrete points. These points can be thought of as pre-grasp shapes, or starting positions for finding a good grasp for a new object [4].

In this work we extend this concept by replacing the discrete set of pre-grasp shapes with a continuous subspace derived from analysis of human hand motion during grasping. This subspace can be searched directly and we show that the result of this search is often close enough to a final grasping position that a simple heuristic can be used to derive a force-closure grasp. By performing this search in hand configuration space, we can use forward kinematics to explicitly avoid unfeasible hand positions and collision with obstacles. As a result, our method is well-adapted for operation in cluttered environments.

We note that choosing a good grasp can also be formulated as a problem in the contact space of the object to be grasped, which is almost always lower dimensional than hand configuration space. We refer the reader to [5] for a review of such methods. While the contact space can be discretized and searched completely [6], such approaches usually require inverse kinematics in order to guarantee that the contacts are physically satisfiable by a real robotic hand. Rezzoug and Gorce [7] solve this problem using supervised learning, and produce a hand configuration such that the fingertips satisfy a number of given point contacts (if possible). However, this approach does not guarantee that the hand is not in collision with the object or obstacles at points other than the predefined contacts. An alternative to the use of inverse kinematics is presented by Platt *et al.* [8, 9], starting with the hand in contact with an object and using gradient descent to adjust the contacts.

II. EIGENGRASPS

Any hand posture is fully specified by its joint values, and can therefore be thought of as a point in a high-dimensional joint space. If d is the number of degrees of freedom (DOF) of the hand, than a posture \mathbf{p} can be defined as

$$\mathbf{p} = [\theta_1 \ \theta_2 \ \dots \ \theta_d] \in \mathcal{R}^d \quad (1)$$

where θ_i is the value of i -th degree of freedom.

As we have already mentioned above, previous research suggests that most grasping postures derive from a relatively small set of discrete pregrasp shapes. This would imply that the range of postures used in everyday grasping tasks will exhibit significant clustering in the d -dimensional DOF space. Santello *et al.* [1] verified this hypotheses by collecting a large set of data containing grasping poses from subjects that were asked to shape their hands as if they were grasping a familiar

object. Principal Component Analysis of this data revealed that *the first two principal components account for more than 80% of the variance*, suggesting that a very good characterization of the recorded data can be obtained using a much lower dimensionality approximation of the joint space.

In our work, we will refer to the Principal Components of these postures as *eigengrasps*. The implication is that they form a **low-dimensionality basis for grasp postures**, and can be linearly combined to closely approximate most common grasping positions. Each eigengrasp e_i is a d -dimensional vector and can also be thought of as direction of motion in joint space. Motion along one eigengrasp direction will usually imply motion along all (or most) degrees of freedom of the hand.

$$e_i = [e_{i,1} \ e_{i,2} \ \dots \ e_{i,d}] \quad (2)$$

By choosing a basis comprising b eigengrasps, a hand posture placed in the subspace defined by this basis can be expressed as a function of the amplitudes a_i along each eigengrasp direction:

$$\mathbf{p} = \sum_{i=1}^b a_i e_i \quad (3)$$

and is therefore completely defined by the amplitudes vector $\mathbf{a} = [a_1 \ \dots \ a_b] \in \mathcal{R}^b$.

A. Effective Degrees of Freedom

The first question to consider is how many eigengrasps need to be considered so that the subspace that they define closely approximates the required range of hand postures. Based on the results of Santello *et al.*, we have used the two dominant eigengrasps of the human hand in our work, and will show how they produce good results. It is important to note that our study is primarily concerned with grasp synthesis for common everyday objects and that another choice of eigengrasps might be necessary in a different problem domain such as complex manipulation tasks, or with another dataset, containing unusually shaped or difficult to grasp objects.

An intriguing corollary question is whether the results obtained using such a small set of eigengrasps imply that the other DOF's of the hand are useless. We can provide two arguments to the contrary: as shown in [1], eigengrasps 3 through 6 (in decreasing order of importance), while accounting for less than 20% of the variance in hand posture, do not represent noise and are shown to be related to the object to be grasped. Furthermore, the study presented by Santello *et al.* was performed in the absence of the real object, as subjects reproduced grasps from memory. This suggests that initial grasp planning stages do indeed take place in a low dimensional space, but during the final stages the shape of the object forces the hand to deviate from eigengrasp space in order to conform to the object surface. From this perspective the space defined through eigengrasps can be seen as a pregrasp or grasp planning space, as we shall expand upon later.

B. Application for Robotic Hand Models

Although the work of Santello *et al.* is centered on the study of the human hand, we have found this approach to be extremely useful for robotic hands as well. In our study, we have applied the eigengrasp concept to a total of 4 hand models: the Barrett hand, the DLR hand [10], the Robonaut hand [11] and finally a human hand model. All our hand models, as well as the eigengrasps used in each case, are presented in table I.

For the human hand we have chosen eigengrasp directions based on the published results in [1], taking advantage of the fact that they have been derived through rigorous study over a large number of recorded samples. Since such data is not available for robotic hand models, we have derived eigengrasps attempting to define grasp subspaces similar to the one obtained using human hand eigengrasps. In most cases, such decisions could be made based directly on the similarities with the human hand. For example, the MCP and IP joints can be mapped to the proximal and distal joints of robotic fingers. In the case of the Barrett hand, changes in the spread angle DOF were mapped to human finger abduction. While we found our choices to produce good results, the optimal choice of eigengrasps for non-human hands, as well as the choice of which eigengrasps to use for a particular task, are open questions and interesting directions for future research.

The eigengrasp concept allows us to design flexible control algorithms that operate identically across all the presented hand models. The key to our approach is that the eigengrasps encapsulate the kinematic characteristics of each hand design. This enables control algorithms that operate on eigengrasp amplitudes to ignore low-level operations and concentrate on the high-level task. We believe this method to be similar in spirit to certain aspects of human brain operation, with individual function grouped together in control synergies. Another advantage is the significant dimensionality reduction (by as much as a factor of 10 for complex hands) obtained by operating in the reduced basis eigengrasp space as opposed to the full joint space. In the next section we will derive a grasp planning algorithm that makes use of both these concepts.

III. GRASP PLANNING USING EIGENGRASPS

In essence, the grasp planning task can be thought of as an optimization problem in a high-dimensional space that describes both hand posture (intrinsic DOF's) and position (extrinsic DOF's). Consider the goal of minimizing an energy function of the form:

$$E = f(\mathbf{p}, \mathbf{w}) \quad (4)$$

If d is the number of intrinsic hand DOF's then $\mathbf{p} \in \mathcal{R}^d$ represents the hand posture and $\mathbf{w} \in \mathcal{R}^6$ contains the position and orientation of the wrist.

Intuitively, this energy function has to be related to the quality of the grasp. However, most formulations pose a number of problems. First, it can be very difficult, or even impossible, to compute an analytical gradient. Second, such functions are highly non-linear, as small changes in both finger

















Model	DOFs	Eigengrasp 1			Eigengrasp 2		
		Description	min	max	Description	min	max
Barrett	4	Spread angle opening			Finger flexion		
DLR	12	Prox. joints flexion Finger abduction			Dist. joints flexion Thumb flexion		
Robonaut	14	Thumb flexion MCP flexion Index abduction			Thumb flexion MCP extension PIP flexion		
Human	20	Thumb rotation Thumb flexion MCP flexion Index abduction			Thumb flexion MCP extension PIP flexion		

TABLE I
EIGENGRASPS DEFINED FOR THE ROBOTIC HAND MODELS USED IN THIS PAPER.

posture and wrist position can drastically alter the quality of the resulting grasp. Finally, the legal parameter space is complex, having to satisfy multiple constraints: prevent interpenetration with the object to be grasped as well as potential obstacles, and maintain joint values within their acceptable ranges.

A. Optimization Algorithm

We directly address all of these problems by using **simulated annealing** as the preferred optimization algorithm. We give a brief description of this algorithm here and refer the reader to [12] for an in-depth review.

During each iteration of this algorithm, a neighbor of the current solution is generated by randomly sampling each of the input variables of the energy function. A decision is then made whether to replace the current state with the new one, based on the difference in energy between the two. During early stages, neighboring states are generated by sampling the entire space of the input variables, and the probability of moving to a new state is high even if the jump increases the energy of the system. As the annealing schedule matures, new states sample an increasingly smaller neighborhood around the current solution, and jumps are made only to states that minimize the energy.

The stochastic nature of simulated annealing makes it a particularly good choice for our task. Since new states are generated as random neighbors of the current state, computation of the energy gradient is not necessary, and the algorithm works well on non-linear functions. Furthermore, the possibility of an “uphill move” to a state of higher energy allows it to escape local minima which can trap greedier methods such as gradient descent. However, the random exploration of the input domain




Fig. 1. Desired contact locations for DLR, Robonaut and Human hands

means that high dimensionality of the parameter space will severely affect the computational efficiency of this algorithm.

We therefore propose performing the optimization **in eigengrasp space**, as opposed to DOF space. The energy function takes the form

$$E = f(\mathbf{a}, \mathbf{w}) \quad (5)$$

where $\mathbf{a} \in \mathcal{R}^2$ is the vector of eigengrasp amplitudes. This effectively reduces the parameter space to 8 dimensions (2 eigengrasp amplitudes plus 6 extrinsic DOF’s) from as high as 26 dimensions in the case of the human hand.

The energy function formulation that we propose simply attempts to bring a number of pre-selected contact points on the robotic hand in contact with the object (figure 1). The energy contains two terms: the first one sums the distances between the desired contact points and the object surface while the second one sums the angular differences between the orientation of the surface normals at the contact locations and the closest point on the object. By sampling the palm and all the links of the robotic hand, as in figure 1, we expect the energy function to be minimized when the hand is wrapped around the object generating a large contact area.

In most cases, the resulting hand posture creates an enveloping grasp of the object, especially for complex hand

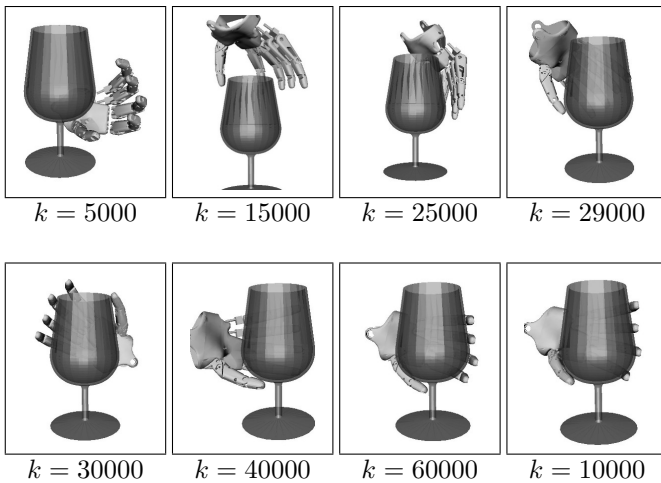


Fig. 2. Simulated annealing example over 100,000 iterations. Each image shows the best state found until iteration k .

models grasping objects similar in size to the hand. However, there exist cases where the desired contact locations are all very close to the object surface without generating a stable grasp. Furthermore, small objects might be impossible to completely wrap the hand around, and an acceptable minimum of the energy function will not exist. We will discuss possible solutions to this problem in the next section.

B. Grasp Planning Results

We have implemented the simulated annealing approach using the publicly available *GraspIt!* simulation engine [13]. For each state generated during the annealing schedule, *GraspIt!* uses forward kinematics to place the robotic hand model in the correct posture and checks for collisions against the object to be grasped as well as other obstacles. If the state is found to be legal, the corresponding energy function is computed and the annealing algorithm proceeds as described above.

We will first analyze the behavior of the simulated annealing algorithm in more detail, using a typical example that involves the Robonaut hand grasping a glass. This optimization, as well as all examples shown in this paper was performed over 100,000 iterations. Figure 2 shows the temporary solution (best state found so far) at various points during the optimization. We can observe what is considered typical behavior for a simulated annealing implementation: at first, the search goes through random states, accepting bad positions as well as good positions. As the annealing schedule progresses, the search space is sampled more often in the vicinity of the good states, while bad states are no longer accepted. Finally, in the later stages, the search is confined in a small neighborhood around the best state, which is progressively refined. The total time required for the optimization was 173 seconds, with the most significant amount of computation used for checking the feasibility of each generated state (*i.e.* checking for collisions and inter-penetrations).

An extensive example set is shown in figure 3: for each hand-object combination the image shows the pre-grasp found by the optimization algorithm. We note that, in most cases,




Fig. 4. Grasp planning taking into account arm and obstacle constraints

planning in the reduced space spanned by only two eigen-grasps does not result in a posture where the robotic hand conforms perfectly to the surface of the object. However, the result is often close enough to such a posture that a stable grasp can be obtained by using a simple heuristic: the pre-grasp is modified by closing each finger until contact with either the object or another finger prevents further motion. This method produces a force-closure grasp in 18 out of the 24 cases shown in figure 3.

So far, we have considered strictly the relationship between a robotic hand and the object to be grasped. Consider however the case of a service robot operating in a human environment: the feasibility of a grasp also depends on the kinematics of the robotic arm that the hand is attached to, as well as any external obstacles. We believe that the dimensionality reduction approach presented here can be successfully used in such situations. For each state that is generated during the simulated annealing search, we have extended the feasibility check to find an appropriate position for the robotic arm using inverse kinematics. If such a position is found and no obstacle collision is detected, then the state is deemed legal. Figure 4 shows pre-grasp results obtained using the eigengrasp planning approach for the Barrett and Robonaut hands attached to a Puma robotic arm.

IV. FUTURE RESEARCH DIRECTIONS

As we have previously mentioned, the energy function formulation used in our search algorithm attempts to wrap the hand around the object and create an enveloping grasp. However, there exist cases where such a grasp might not be feasible, or desirable. For example, obstacle constraints might prevent the hand from wrapping around the object, or a more flexible manipulation-type grasp might be preferable to a power grasp.




Fig. 3. Eigengrasp planner test using 4 hand models to grasp each of 6 objects

The most direct method for obtaining such a grasp is to use only a subset of the desired contact locations shown in figure 1. The subset comprising only fingertip contacts is a natural candidate, but using such a small contact set raises an additional problem: it is generally easy to place all fingertips on the object surface without necessarily obtaining a stable grasp. To address this problem we also propose a modified version of the energy function that includes a built-in notion of grasp quality.

While a number of grasp quality metrics have been presented in the literature, our context is somewhat different: we require a metric that can take into account not only existing contacts between the hand and the object, but also potential contacts that can be realized by small changes in the current state. In this sense, the ideal metric would assess the *potential* of a hand posture, and determine whether the annealing algorithm will search its neighborhood for progressively better states. One possible quality metric that can be modified according to these requirements is the one described by Ferrari and Canny [14]. In its original form, the process involves building the space of wrenches that can be applied by a grasp (the grasp wrench space, or GWS) by taking the convex hull of the wrenches that can be applied through each contact. In our implementation, object contacts are replaced by the desired contact locations exemplified in figure 1. When computing the GWS, we scale the wrenches that can be applied at each




Fig. 5. Grasp planning using only fingertip contacts

desired contact location depending on the distance between the desired contact and the object surface. Thus, if this distance is small, the contact will have a positive contribution to the grasp, and states that bring it closer to the object surface will be rewarded by a higher quality value. If, on the contrary, the desired contact is far from the object, it will not significantly affect the grasp quality measurement.

Once the grasp quality term is computed, it is included in the energy function in negated form, as the annealing algorithm attempts to minimize the energy value. Its contribution biases the search algorithm toward states that not only bring the hand in contact with the object, but also create stable grasping postures. We have found that these formulations work well in practice, as shown in figure 5. We are currently aiming to build upon these results and derive an energy function formulation which would guarantee stable grasps while eliminating the need to pre-specify desired contact locations on the hand.

Finally, the tests presented in this paper have been performed using a grasping simulator and taking advantage of complete knowledge of the object geometry. We believe that this environment is ideal for evaluating the theoretical possibilities of such a method, as it also provides a direct verification of the results using rigorous quality metrics. However, general robotics applications usually have less information available about the object to be grasped, such as a single laser scan or perhaps sparse stereo, with substantial occlusions. There are two general approaches for dealing with such incomplete data: attempt to construct or recover a full object model [15], or use only the observed data without any additional model building [16]. The eigengrasp dimensionality reduction presented here makes no assumptions about the object at all. However, the particular cost functions we use do require a full 3D model. Platt *et al.* [8] showed how a grasp control function can be constructed even with limited knowledge of object geometry, and a promising direction for future research will be designing similar cost functions for use with our method.

V. CONCLUSIONS

In this paper we have built upon recent results in neuroscience research, which show that human hand control for common grasping tasks mostly takes place in a space of much lower dimensionality than the number of degrees of freedom of the human hand. We have extended this concept for a number of robotic hands: for each model, we have defined a low dimensional subspace of the degrees of freedom space, determined by a number of basis vectors which we call *eigengrasps*.

As long as the eigengrasp space provides a good approximation of the hand motion required for a given task, control algorithms can be designed to operate in this space and take advantage of the dimensionality reduction. In the case of grasp planning, data collected from human users has shown that this is indeed the case. In this paper we show that this is also the case for complex robotic hands: after optimizing the pre-grasp hand posture in eigengrasp space, we can use simple heuristics to find stable grasps even for complex hand models that have traditionally been very difficult to plan for.

The eigengrasp framework acts not only to reduce control complexity, but also as an interface between the kinematic structure of the hand and higher-level task planning. Therefore, for a given task, it is possible to use a unified treatment for a number of robotic hand models, even though the kinematic specifications may be significantly different. We have illustrated this concept by using the eigengrasp planner on four robotic hands, with the number of intrinsic DOF's ranging between 4 and 20. The results show that it is indeed possible to apply an identical control algorithm to all of these hand models and obtain consistent results. Furthermore, the planning method we have presented can take into account robotic arm constraints as well as external obstacles.

While this work has been focused on the task of grasping everyday objects, we believe that eigengrasp-like control synergies can be found for many other domains. Since the

published experimental data we draw upon was collected under such assumptions, we found it unjustified to generalize our particular choices of eigengrasps without further analysis. However, the effectiveness of the grasp planning algorithm based on relatively few eigenvectors of hand motion suggests that identifying similar dimensionality reduction strategies for other domains will prove a fruitful area of future research.

REFERENCES

- [1] M. Santello, M. Flanders, and J. F. Soechting, "Postural hand synergies for tool use," *Journal of Neuroscience*, vol. 18, no. 23, pp. 10 105–10 115, 1998.
- [2] J. R. Napier, "The prehensile movements of the human hand," *Journal of Bone and Joint Surgery*, vol. 38, pp. 902–913, 1956.
- [3] M. R. Cutkosky, "On grasp choice, grasp models, and the design of hands for manufacturing tasks," *IEEE Transactions on Robotics and Automation*, vol. 5, pp. 269–279, 1989.
- [4] A. T. Miller, S. Knoop, H. I. Christensen, and P. K. Allen, "Automatic grasp planning using shape primitives," *IEEE Intl. Conf. on Robotics and Automation*, vol. 2, pp. 1824–1829, 2003.
- [5] A. Bicchi and V. Kumar, "Robotic grasping and contact: A review," *IEEE Intl. Conf. on Robotics and Automation*, pp. 348–353, 2000.
- [6] Y.-H. Liu, M.-L. Lam, and D. Ding, "A complete and efficient algorithm for searching 3-d form-closure grasps in the discrete domain," *IEEE Transactions on Robotics*, vol. 20, no. 5, pp. 805–816, 2004.
- [7] N. Rezzoug and P. Gorce, "A biocybernetic method to learn hand grasping posture," *Kybernetes*, vol. 32, no. 4, pp. 478–490, 2003.
- [8] R. Platt, A. H. Fagg, and R. Grupen, "Nullspace composition of control laws for grasping," in *IEEE Intl. Conf. on Robotics and Automation*, Washington, D.C., May 2002, pp. 1717–1723.
- [9] —, "Manipulation gaits: Sequences of grasp control tasks," in *IEEE Intl. Conf. on Robotics and Automation*, New Orleans, LA, April 2004, pp. 801–806.
- [10] C. S. Lovchik and M. A. Diftler, "The robonaut hand: A dextrous robot hand for space," in *IEEE Intl. Conf. on Robotics and Automation*, 1998, pp. 907–912.
- [11] J. Butterfass, G. Hirzinger, S. Knoch, and H. Liu, "Dlr's articulated hand, part i: Hard- and software architecture," in *IEEE Intl. Conf. on Robotics and Automation*, 1998, pp. 2081–2086.
- [12] L. Ingber, "Very fast simulated re-annealing," *J. Mathl. Comput. Modelling*, vol. 12, no. 8, pp. 967–973, December 1989.
- [13] A. Miller and P. K. Allen, "Graspit!: A versatile simulator for robotic grasping," *IEEE Robotics and Automation Magazine*, vol. 11, no. 4, pp. 110–122, December 2004.
- [14] C. Ferrari and J. Canny, "Planning optimal grasps," in *IEEE Intl. Conf. on Robotics and Automation*, 1992, pp. 2290–2295.
- [15] 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 Intl. Conf. on Intelligent Robots and Systems*, Beijing, China, October 2006, pp. 5663–5668.
- [16] A. Saxena, J. Driemeyer, J. Kearns, and A. Ng, "Robotic grasping of novel objects," in *Advances in Neural Information Processing Systems*. Cambridge, MA: MIT Press, 2006.