Low-Dimensional Data-Driven Grasping

Peter K. Allen Matei T. Ciocarlie Corey Goldfeder Hao Dang
Department of Computer Science, Columbia University, New York, USA

1 Grasp Synthesis

In general, automatic grasp synthesis can be thought of the task of finding the combination of hand posture (intrinsic degrees of freedom, or DOF's) and position (extrinsic DOF's) that produces a stable grasp, according to a given grasp quality metric. From this perspective, it can be approached as an optimization problem, seeking to maximize the value of the grasp quality Q expressed as a function over a high-dimensional domain:

$$Q = f(\boldsymbol{p}, \boldsymbol{w}) \tag{1}$$

If d is the number of intrinsic hand DOF's, $\boldsymbol{p} \in \mathcal{R}^d$ represents the hand posture and $\boldsymbol{w} \in \mathcal{R}^6$ contains the position and orientation of the wrist. The traditional form for specifying a hand posture is to set a value for each individual DOF of the hand. For complex hands, such as the human hand modeled in our study with $d = 20$, the high dimensionality of the posture space makes direct searches for good grasps intractable. To make this tractable, we perform the grasp planning task in a subspace of highly reduced dimensionality. Our approach is based on the results of Santello *et al.* [8], who have shown that the range of postures that humans use in everyday grasping exhibits significant clustering in the d-dimensional DOF space. However, while human grasping subspaces can be determined through user studies, defining similar subspaces for non-anthropomorphic robotic hands is an interesting open problem. In previous work [4], we have applied this concept to five hand models, using the results of Santello *et al.* for the anthropomorphic models (such as the DLR and Robonaut hands) and empirically derived subspaces for the non-anthropomorphic ones (such as the Barrett hand).

We consider a hand posture subspace defined by a number of d-dimensional basis vectors called **eigengrasps**; the implication is that these vectors can be linearly combined to closely approximate most common grasping positions. By choosing a basis comprising b eigengrasps, a hand posture $\boldsymbol{p}$ placed in the subspace defined by this basis is uniquely defined by the vector $\boldsymbol{a} \in \mathcal{R}^b$ containing the amplitudes along each subspace axis. In previous work [3], we have discussed the feasibility of finding good grasps for dexterous hands by searching a subspace defined by two eigengrasps. This implies a significant dimensionality reduction of the grasp quality function domain, which can be expressed as

$$Q = f(\boldsymbol{a}, \boldsymbol{w}), \quad \boldsymbol{a} \in \mathcal{R}^2 \tag{2}$$

However, this low-dimensional subspace is only useful as long as it contains the hand postures needed for stable grasps of a large variety of objects. The results presented in [3] show that, in general, postures where the hand conforms *perfectly* to the surface of the target can not be found in eigengrasp space. However, by searching this subspace we can usually find a posture that is *very close* to a desired grasp. The eigengrasp space can therefore be thought of as a pre-grasp, or planning space: the best pre-grasps found in this subspace have a good chance of producing stable grasps by simply closing each finger until motion is stopped by contact with the object. This suggests a two-stage grasp planning algorithm: the first stage searches the low-dimensional eigengrasp subspace, while the second stage tests the resulting pre-grasps and outputs the best solutions. We have applied this idea to online automated grasping where the reduced subspace is searched to find a stable grasp for objects [1].

2 A Database of Grasps

Our low-dimensional grasp planner has allowed us to create the Columbia Grasp Database, a freely available collection of hundreds of thousands of form closure grasps for thousands of 3D models [5]. Our primary interest is in using an object's 3D geometry as an index into the database. Given a new 3D object, we can find geometric neighbors in the database, and the accompanying stable grasps for these similar objects. If the number of objects to be grasped in the database is very large and comprehensive then robotic grasping becomes a pre-computed database lookup. While we have not yet achieved this level of performance it is our directional goal.

The most direct way to construct a grasp database is to collect grasping data from human volunteers. We could gather a large set of example objects, outfit an army of graduate students with grasp-capture devices such as datagloves, and record the results. Unfortunately, this approach is prohibitively time consuming for large scale data acquisition. More importantly, data collection from humans can *only* produce grasps with the human hand. Since many popular robotic hands cannot be easily mapped to the human hand, a useful database should include grasps with multiple hands.

The Columbia Grasp Database was created using *GraspIt!*, a publicly available grasp planning and analysis tool developed by our group [7]. The database is intended to be used in conjunction with *GraspIt!* or a similar simulation tool; as we have shown in previous work [6, 2], planning results obtained in simulation can be successfully applied to real robotic hands performing grasping tasks.

Even for grasp planning algorithms that do not rely on simulation, an environment such as *GraspIt!* is an important tool for evaluation, as grasp quality measures are generally impossible to compute in physical experiments. Part of our motivation in producing a grasp database was to provide a benchmark for robotic grasping tasks. Using a common benchmark will make it possible to directly compare grasp planning algorithms, which is currently difficult to do.

We briefly discuss a new database-backed grasp planning algorithm based on the data we have collected (see fig. 1). Using this algorithm, we illustrate the usefulness of a database for grasping, and highlight some of the lessons learned during its construction. We also provide execution results over the entire set of objects in the database at their primary scale. We believe this to be one of the most comprehensive tests found in the grasp planning literature, demonstrating the use of the database as a benchmarking tool.

2.0.1 Acknowledgments.

This work was funded in part by NIH BRP grant 1RO1 NS 050256-01A2.

Figure 1: Three example models and their grasps from neighbors in the grasp database

References

[1] M. Ciocarlie and P. K. Allen. On-line interactive dexterous grasping. In *Haptics: Perception, Devices and Scenarios*, volume 5024, pages 104–113. Springer Lecture Notes in Computer Science, 2008.

[2] M. Ciocarlie and P. K. Allen. On-line interactive dexterous grasping. In *Eurohaptics*, 2008.

[3] M. Ciocarlie, C. Coldfeder, and P. Allen. Dimensionality reduction for hand-independent dexterous robotic grasping. In *IEEE-RSJ Intl. Conf. on Intelligent Robots and Systems*, 2007.

[4] M. Ciocarlie, C. Goldfeder, and P. Allen. Dexterous grasping via eigengrasps: A low-dimensional approach to a high-complexity problem. In *Robotics: Science and Systems Manipulation Workshop - Sensing and Adapting to the Real World*, 2007.

[5] C. Goldefeder, M. Ciocarlie, H. Dang, and P. K. Allen. The columbia grasp database. In *IEEE Intl. Conf. on Robotics and Automation*, 2009.

[6] D. Kragic, A. Miller, and P. K. Allen. Real-time tracking meets online planning. In *Intl. Conf. on Robotics and Automation*, 2001.

[7] A. Miller, P. K. Allen, V. Santos, and F. Valero-Cuevas. From robot hands to human hands: A visualization and simulation engine for grasping research. *Industrial Robot*, 32(1), 2005.

[8] M. Santello, M. Flanders, and J. F. Soechting. Postural hand synergies for tool use. *Journal of Neuroscience*, 18(23):10105–10115, 1998.