

A CONSTRAINED OPTIMIZATION FRAMEWORK FOR COMPLIANT UNDERACTUATED GRASPING

Matei Ciocarlie
Willow Garage, Inc.
Menlo Park, California 94025, U.S.A.
matei@willowgarage.com

Peter Allen
Computer Science Department
Columbia University
New York, New York 10027, U.S.A.
allen@cs.columbia.edu

ABSTRACT

This study focuses on the design and analysis of underactuated robotic hands that use tendons and compliant joints to enable passive mechanical adaptation during grasping tasks. We use a quasistatic equilibrium formulation to predict the stability of a given grasp. This method is then used as the inner loop of an optimization algorithm that can find a set of actuation mechanism parameters that optimize the stability measure for an entire set of grasps. We discuss two possible approaches to design optimization using this framework, one using exhaustive search over the parameter space, and the other using a simplified gripper construction to cast the problem to a form that is directly solvable using well-established optimization methods. Computations are performed in 3D, allow arbitrary geometry of the grasped objects and take into account frictional constraints.

INTRODUCTION

In this study¹, we present a framework for the analysis and optimization of a class of passively adaptive underactuated robotic hands. In a broad sense, this is the task of replacing elaborate run-time algorithms (often requiring extensive sensor arrays for input and complex actuation mechanisms for execution) with off-line analysis, performed before the hand is even built. We focus on highly underactuated hand models, where the number of joints far exceeds the number of actuators, noting for example that recent studies have shown reliable grasping performance with even a single actuator for multiple fingers and up to

15 joints [3]. In such cases, the only decision available at runtime is the placement of the hand relative to the target object, emphasizing the importance of the design optimization step. As the fingers are closing, the only computation is performed at an implicit level, by the actuation mechanism itself.

Two important tools for achieving reliable performance when using passively adaptive hands are grasp analysis and design optimization. The former traditionally focuses on a single grasp, and attempts to compute a measure of its stability under a given actuation mechanism. The latter aims to compute the parameters of the actuation mechanism itself, so that stability measures are maximized for a broad range of grasps. As the space of hand design parameters is extremely large, especially when taking into account kinematic considerations such as number of joints, number of fingers, finger link and palm shapes, we narrow the scope of this study as follows:

- we start from a given kinematic design, with a potentially large number of joints grouped into multiple fingers;
- assuming that each joint can be controlled independently, we create a set of stable grasps over a given group of objects. We refer to this set as the *optimization pool*;
- we attempt to derive the design parameters of a passively adaptive underactuation mechanism that increase the stability of the grasps in the set.

There are multiple ways of achieving passive adaptation with a robotic hand design. One of the earliest examples, the Soft Gripper introduced by Hirose and Umetani [4], used tendons for both flexion and extension. Ulrich et al. [5] pioneered the use of a breakaway transmission mechanism which is now used in the Barrett hand (Barrett Technologies, Cambridge, MA). Birglen et al. [6] presented a detailed and encompassing optimization study for underactuated hands, focusing mainly on four-bar linkages but with applications to other transmission mech-

¹This work was partially supported by the National Science Foundation under Grant no. 0904514. Individual parts of this work have been published in conference proceedings in [1, 2]. This study presents a unified view of the entire framework, extends the set of experimental results and discusses in more depth the extended implications of our approach.

anisms as well. Four-bar linkages were also used to construct the MARS hand [7], which later evolved into the SARAH family of hands [8]. These studies have led to the construction of remarkably efficient grippers and hands. In the process, they have highlighted the fact that optimization of a highly underactuated hand is a complex problem; in other words, *simple is hard!*

Our framework is based on the implementation of passively adaptive underactuation using the mechanics of a tendon-actuated hand combined with compliant, spring-like joints. This is based on the work of Dollar and Howe [9], who optimized the actuation and compliance forces of a single tendon design. This design was later implemented in the Harvard Hand [10], which we also use here as one of our case studies. We found the relative ease of constructing a prototype using this actuation paradigm particularly appealing, and believe it can lower the barriers for experimenting with new hand designs and disseminating research results. However, other actuation methods have their own merits, which must be considered in future iterations.

PROBLEM FORMULATION

The starting point for our optimization framework is the quasistatic equilibrium relationship that characterizes a stable grasp. We briefly review this formulation here; for more details we refer the reader to the analysis by Prattichizzo and Trinkle [11].

Consider a grasp with p contacts established between the hand and the target object. For any contact i , the total contact wrench \mathbf{c}_i must obey two constraints. First, the normal component must be positive (contacts can only push, not pull). Second, friction constraints must be obeyed. A common method is to linearize these constraints, by expressing \mathbf{c}_i as a linear combination of normal force and possible friction wrenches:

$$\mathbf{c}_i = \mathbf{D}_i \boldsymbol{\beta}_i \quad (1)$$

$$\boldsymbol{\beta}_i, \mathbf{F}_i \boldsymbol{\beta}_i \geq 0 \quad (2)$$

where the matrices \mathbf{D}_i and \mathbf{F}_i depend only on the chosen friction model, such as linearized Coulomb friction. The contact wrench is now completely determined by the vector of friction and normal wrench amplitudes $\boldsymbol{\beta}_i$, which will be computed as part of the grasp analysis algorithm.

In general, a grasp, as a collection of multiple contacts, is in quasistatic equilibrium if the following conditions are satisfied:

- contact forces are balanced by joint forces (hand equilibrium);
- resultant object wrench is null (object equilibrium);
- contact constraints are met for all contacts in the grasp.

We can assemble the complete grasp description as follows:

$$\mathbf{J}_c^T \mathbf{D} \boldsymbol{\beta} = \boldsymbol{\tau} \quad (3)$$

$$\mathbf{G} \boldsymbol{\beta} = 0 \quad (4)$$

$$\boldsymbol{\beta}, \mathbf{F} \boldsymbol{\beta} \geq 0 \quad (5)$$

where $\boldsymbol{\tau}$ is the vector of joint forces, \mathbf{J}_c is the Jacobian of the contact locations and \mathbf{G} is the grasp map matrix which relates individual contact wrenches to the resultant object wrench. The matrices \mathbf{D} and \mathbf{F} bring together the individual contact constraint matrices \mathbf{D}_i and \mathbf{F}_i for $i = 1 \dots p$ in block diagonal form. The column vector $\boldsymbol{\beta}$ contains all contact amplitudes vectors $\boldsymbol{\beta}_i$ in block column form. In order to avoid the trivial solution where all forces in the system are zero, a constraint can be added requiring total actuator forces to sum to a pre-specified level.

In our framework, we use the term “stable” to refer to grasps that are in quasistatic equilibrium (all of the above constraints are met). It is important to note that, in practice, this is not a sufficient condition for achieving form-closure using the given actuation mechanism. However, it is a necessary prerequisite and, as such, we believe that optimizing a hand to achieve equilibrium under many configurations is a valuable step towards enabling a wide range of grasps. A possible future extension would be to also include a direct measure of grasp quality, according to one of the metrics that have been proposed in the literature [11, 12].

So far, this analysis applies to a hand design regardless of its actuation method. To adapt it to the case of underactuated hands, we must look in more detail at the joint force vector $\boldsymbol{\tau}$, which is a result of the actuation mechanism. We use the common tendon-pulley model (as used for example by Kwak et al. [13]), which assumes that a tendon travels through a number of routing points that it can slide through, but which force it to change direction as it follows the kinematic structure. As a result of this change in direction, the routing points are the locations where the tendon applies force to the links of the finger.

We assume that the hand contains a total of d tendons, each with multiple routing points across different links. In this case, joint forces can be expressed as:

$$\boldsymbol{\tau} = \mathbf{J}_d^T \boldsymbol{\delta} + \boldsymbol{\theta} \mathbf{k} \quad (6)$$

where \mathbf{J}_d is the Jacobian of the tendon routing points and $\boldsymbol{\delta} \in \mathcal{R}^d$ is the vector of applied tendon forces. $\boldsymbol{\theta}$ is a diagonal matrix of joint angle values and \mathbf{k} is the vector of joint spring stiffnesses (without loss of generality, we assume 0 is the rest position for all springs). We now have a complete description of the equilibrium state of the grasp, expressed in eqs. (3) through (6).

In practice, one of the relationships comprised in this formulation is used as an optimization objective, rather than a hard constraint, with two important advantages. First, it provides more information for problems where all the constraints are not feasible in their exact form. Second, problems that have a solution in the exact form will often have an infinity of solutions; formulating an optimization objective allows us to choose an “optimal” one. Which of the above constraints is to be used as an optimization objective depends on the nature of the problem; as a result, this formulation is extremely versatile, and can be adapted to a

number of practical problems in underactuated grasp analysis:

- if the unknown variables include only the contact wrench magnitudes β (and implicitly all the individual contact wrenches c_i), we are computing whether a particular set of actuator forces results in a stable grasp;
- if we extend the set of unknowns to also include the vector δ , we are trying to compute the best actuator forces for a grasp characterized by a particular set of contacts;
- we can even extend the set of unknowns to include components of J_d^T or k , in which case we are computing the best hand design parameters for executing a given grasp (or set of grasps).

GRASP ANALYSIS

Using this framework, we can analyze the stability of a given grasp by using object equilibrium as an optimization objective. We are interested in detecting finger slip on the object surface, as well as any unbalanced forces in the hand-object system; this analysis can be carried out both as the fingers are closing and after the grasp is complete. For a given hand posture and set of contacts, the goal is to determine the contact forces β and actuator forces δ that balance the system, or, if exact equilibrium is infeasible, result in the smallest magnitude wrench on the object:

$$\begin{aligned} \text{minimize } \|\mathbf{G}\beta\| &= \beta^T \mathbf{G}^T \mathbf{G} \beta \text{ subject to :} \\ \begin{bmatrix} \mathbf{J}_c^T \mathbf{D} & -\mathbf{J}_d^T \end{bmatrix} \begin{bmatrix} \beta & \delta \end{bmatrix}^T &= \theta \mathbf{k} \\ \delta, \beta, F\beta &\geq 0 \end{aligned}$$

This is a standard Quadratic Program, with linear constraints. The matrix that defines the quadratic component of the objective function is positive semidefinite by definition, as it is the product of the matrix \mathbf{G} and its transpose. Therefore, the optimization problem is convex, so whenever the conditions are feasible, a global minimum can be determined. In this study, we use the Mosek [14] package to solve all the optimization problems of this form. We can obtain one of three possible results:

- the problem is unfeasible; this indicates that contact forces that obey the constraints can not be supported by actuator forces. The fingers will slip on the surface of the object;
- the problem is feasible and a non-zero global optimum is found; the contacts are stable but some level of *unbalanced force* is applied to the object. If this force is not balanced externally (*i.e.* by interactions between the target object and another surface in the environment), the hand will have to reconfigure itself, also causing the object to move;
- the problem is feasible and the global optimum is zero; the contacts are stable and produce a null resulting wrench. The hand-object system is stable in its current configuration.

One traditional application of grasp analysis methods takes place at run-time. If the object to be grasped can be modeled (or

recognized), this method can be applied to find stable grasps for execution. Another approach uses tactile sensors and proprioception to analyze the grasp currently being executed. However, both of these options require extensive sensing arrays and control options that might not be available in a highly underactuated hand. We believe that a very promising alternative is to use grasp analysis as part of an off-line optimization routine, to improve the overall performance of the hand. We explore this option next.

HAND DESIGN OPTIMIZATION

Using our framework to tackle hand design optimization implies a different approach than the previous section. Grasp analysis focuses on *a single grasp* at one time, and aims to compute the optimal contact or actuator forces specific to that grasp. In contrast, the study of hand design parameters normally implies solving an optimization problem over *a set of grasps*.

The first step of the hand optimization method is thus to create a batch of grasps that we expect the hand to be able to perform. We refer to this set as the *optimization pool*. The task of defining the grasp optimization pool, and executing the optimization procedure, is performed in a simulated environment, which seems natural in the context of design analysis performed before the hand is constructed. We use our publicly available *GraspIt!* simulation engine [15]; all the tools presented here have been integrated in, and are now available as part of *GraspIt!*.

The optimal contact and actuator forces specific to each grasp in the pool are still unknown; now they are joined by a set of unknowns representing actuation parameters which are shared by all the grasps in the pool. A key problem is the decision of which design parameters are the focus of the optimization. We have decided to focus on the parameters of the actuation mechanism, such as tendon route and joint spring stiffness. A second problem, intrinsically related to the first one, regards the method used to perform the optimization. The ideal scenario would intuitively be to assemble a *global optimization problem*, allowing the direct computation of the optimal design parameters over the entire grasp pool. However, such a global approach is not always possible to implement.

Consider for example the problem of optimizing the location of the tendon routing points on their respective links. The effects of the tendon route on the equilibrium condition are encapsulated in the Jacobian of the routing points, \mathbf{J}_d . Changing the location of a routing point on a link has a highly non-linear effect on \mathbf{J}_d . Furthermore, even if we had a linear relationship between tendon route parameters and the routing point Jacobian, the result must be multiplied by the unknown vector of actuation forces δ . As a result, computing both actuation forces *and* optimal tendon route parameters results in a higher order equality constraint which can not be handled by the same optimization tools.

The general case therefore enables us to quantify a given hand design (by separately computing the quality of each indi-




Figure 1. MODEL OF THE HARVARD HAND AND OBJECT MODELS USED FOR GRASP POOL GENERATION.

vidual grasp in the optimization pool), but not to directly compute a global optimum for the design parameters. We envision two possible solutions to this problem:

- a “numerical” approach, where, for each possible set of design parameter values, each grasp in the pool is analyzed independently. A global measure of grasp stability is computed by summing (or averaging) the result for each grasp, and the parameter values that yield the best result are chosen.
- a “global optimization” approach, where new constraints are added to the global formulation (containing design parameter values as explicit unknowns) in order to cast it as a solvable optimization problem, such as a Linear or Quadratic Program. Apart from computational efficiency, this method also has the advantage of producing a provable global optimum. Its drawback is that the introduction of additional constraints limits its applicability to a subset of possible hand designs.

Numerical Optimization

We used the numerical approach to investigate how grasping performance can be improved by changing design parameters for the Harvard Hand [10]. Figure 1 shows the *GraspIt!* model of this hand, which uses a single actuator to drive eight joints that articulate four fingers. We focused on two design parameters: the *actuator torque ratio* between the proximal and distal joints of each finger (circled in the image), and the *spring stiffness ratio* between the same joints. These parameters are determinant for the behavior of the hand, as they affect both the posture of the hand before touching an object and the forces transmitted after contact is made. In particular, we investigated all possible combinations ranging from 0.2 to 1.0 (in steps of 0.2) for the torque ratio and from 0.1 to 1.0 (in steps of 0.1) for the stiffness ratio.

The optimization pool consisted of 2000 possible grasps distributed evenly across 5 object models (glass, flask, toy plane, mug and phone receiver). For each object, the set of grasps was created by sampling multiple approach directions on the faces of the object bounding box and aligning the hand with the axes of the box. For each torque and stiffness combination, we used the analysis method presented in the previous section for all the




Figure 2. THE EFFECT OF HAND DESIGN PARAMETERS ON THE LIKELIHOOD OF OBTAINING A STABLE GRASP. A DARKER COLOR MEANS A HIGHER NUMBER OF STABLE GRASPS.

candidate grasps and reported the number of them that are stable throughout their execution. To enable direct comparison across different objects, each set of results was normalized to a scale of 0 to 1 through division by the maximum number of grasps found for that object. Figure 2 shows the results for each of the objects, as well as their average over the entire set.

The contour maps reveal which areas of the optimization range offer the best performance; in particular they suggest a torque ratio of 0.6 and a stiffness ratio of 0.3. The overall resemblance between the patterns suggests that the global optimum of the average profile is a good compromise, likely to work well on all objects. However, the patterns exhibit enough variation to illustrate the importance of performing this analysis over a large set of models, spanning a wide range of shapes and grasping scenarios. We note that our torque ratio is in agreement with the value found in the optimization study [9] carried out before the construction of the Harvard Hand prototype.

Our analysis consisted of a total of 20,000 grasps for each object (400 candidates for each of the 50 combinations of force and torque ratios); the typical time spent per object was 15 minutes. This performance suggests the possibility of scaling up to significantly larger test sets as, unlike run-time analysis, off-line optimization can benefit from a time budget of weeks or months, as well as massively parallel computing architectures. In addi-




Figure 3. OVERVIEW AND DETAILED JOINT DESCRIPTION FOR TWO-FINGERED GRIPPER USED AS OPTIMIZATION CASE STUDY.

tion, it is possible to extend the set of analyzed parameters, including for example link lengths, number of links, etc.

While our results show that this approach is feasible (at least for a comparable optimization domain), more advanced numerical optimization algorithms can also be used. Examples include simulated annealing or gradient ascent using numerical computation of the gradient. While such algorithms also present disadvantages (like the threat of stopping in local optima), they are generally better equipped to handle larger problems. Another option, discussed in detail in the next section, avoids exhaustive search by using a different problem formulation.

Global Optimization

In order to illustrate our global optimization approach to the hand design problem, we will build up a concrete example, using as a test bed a two-finger model (which we will refer to as a gripper, rather than a hand). We will first describe the starting model, then discuss the reasons for choosing this particular design.

The basic gripper model is presented in Fig. 3. A single tendon provides flexion forces for both fingers, which are co-actuated using a pulley mechanism, similar to the one used in the Harvard Hand. Note that the pulley allows one finger to continue flexing even if the other finger is blocked by contact with the object. Extension forces are provided by spring-like joints. In practice, these joints can be constructed using a compliant, rubber-like material; this design enables distal joints to flex even when proximal joints are stopped. We assume that the kinematic behavior is that of ideal revolute joints, with the center of rotation placed halfway between the connected links. The tendon itself follows a route in the flexion-extension plane of the gripper. The essentially two-dimensional design prevents the links




Figure 4. GRIPPER MODEL FOR GRASPI! AND EXAMPLES OF GRASPS FROM THE OPTIMIZATION POOL.

from leaving this plane without the application of external forces. However, the tendon route inside this plane is not specified, and is one of the targets of the optimizations.

Figure 3 also shows in detail the design parameters of the gripper. The tendon route is determined by the location of the entry and exit points for each link; more specifically, the parameter that we use is the distance between the tendon entry or exit point and the connection between the link and the joint. We also make the simplifying assumption that, for a joint i , the exit point from the proximal link and the entry point in the distal link have the same value for this parameter, which we call l_i . The current value of the joint is θ_i . r is the joint radius (shared by all the joints), while the length of the links is denoted by d .

The reason for using this design and formulation is that they yield a compact and, more importantly, linear relationship between the construction parameters and the joint forces applied through the tendon. If we consider the parameter vector $\mathbf{p} = [l_0 \ l_1 \ l_2 \ l_3 \ l_4 \ l_5 \ r \ d]$, we obtain a relationship of the form:

$$\boldsymbol{\tau} = \boldsymbol{\delta}(\mathbf{B}\mathbf{p} + \mathbf{a}) + \boldsymbol{\theta}\mathbf{k} \quad (7)$$

where the matrix $\mathbf{B} \in \mathcal{R}^{8 \times 8}$ and the vector $\mathbf{a} \in \mathcal{R}^8$ depend only on the joint values $\theta_0 \dots \theta_5$. Furthermore, since we are using a single tendon, we can normalize its value without loss of generality to $\delta = 1N$. The joint force relationship, and by extension the grasp equilibrium conditions, are now linear in all of the unknowns.

Having established the general characteristics of the gripper, the next step was to generate a pool of grasps over which to optimize its performance. We created a kinematic model of the gripper for the *GrasPI!* environment, assuming each joint could be controlled independently. Then, using the interaction tools in the simulator, we manually specified a number of grasps over a set of 3D models of common household objects. The set comprised 70 grasps distributed across 15 objects; the process is illustrated in Fig. 4. Each grasp was defined by the set of gripper joint angles,

the location of the contacts on each link, and the contact surface normals, resulting in a purely "geometric" description of a grasp, with no reference to the actuation mechanism.

Most of the grasps in the pool used different postures for the two fingers of the gripper. We added to the set the "transpose" of each grasp, obtained by rotating the gripper by 180° around the wrist roll axis, reversing the roles of the left and right finger. The complete optimization pool thus comprised 140 grasps. The inclusion of the transposed grasps also ensured that the final optimized parameters, presented in the next section, were symmetrical, with identical results for both fingers.

A key restriction during the creation of the optimization pool was that all the grasps therein were required to have form-closure. *GraspIt!* integrates a number of analysis tools for establishing the form-closure property by building the Grasp Wrench Space, as described by [12]. This formulation is equivalent to the ability of a set of contacts to apply a null resulting wrench on the object while satisfying contact friction constraints, but disregarding any kinematic or actuation constraints.

For each grasp in our optimization pool, we can apply the equilibrium formulation using the actuation mechanism model described earlier in this section:

$$(J_c^j)^T D^j \beta^j = B^j p + a^j + \theta^j k \quad (8)$$

$$G^j \beta^j = 0 \quad (9)$$

$$\beta^j, F^j \beta^j \geq 0 \quad (10)$$

where we use the superscript j to denote the index number of the particular grasp from the optimization pool that we are referring to. The unknowns are the grasp contact forces β^j , the hand parameter vector p and the vector of joint spring stiffnesses k . Note that p and k do not have a superscript as they are shared between all the grasps in the pool.

To obtain a global optimization problem, we assemble these relationships in block form over the entire pool of grasps. The matrices for individual grasps $(J_c^j)^T D^j$, B^j , θ^j , G^j and F^j are assembled in block diagonal form in the matrices $\tilde{J}_c^T \tilde{D}$, \tilde{B} , $\tilde{\theta}$, \tilde{G} and \tilde{F} , respectively. The vectors β^j and a^j are assembled in block columns in the vectors $\tilde{\beta}$ and \tilde{a} . Finally, the joint equilibrium condition (8) assembled for all the grasps in the pool becomes the optimization objective:

$$\text{minimize } \left\| \begin{bmatrix} \tilde{J}_c^T \tilde{D} - \tilde{B} - \tilde{\theta} \\ \tilde{G} \end{bmatrix} \begin{bmatrix} \tilde{\beta} \\ p \\ k \end{bmatrix} - \tilde{a} \right\| \text{ subject to:} \quad (11)$$

$$\tilde{G} \tilde{\beta} = 0 \quad (11)$$

$$\tilde{\beta}, \tilde{F} \tilde{\beta} \geq 0 \quad (12)$$

$$p_{min} \leq p \leq p_{max} \quad (13)$$

$$k_{min} \leq k \leq k_{max} \quad (14)$$

The minimum and maximum values for the construction parameters p and k can be set to reflect constraints in the physical construction of the gripper, as we will show in the next section.

We note that the result is again a convex Quadratic Program that, by construction, always accepts a solution: constraints (11) and (12) are equivalent to each individual grasp having form-closure independently of the actuation mechanism, which we ensured by building our grasp pool accordingly. As a result, the problem can be solved and a global optimum can be computed.

Construction of an Optimized Gripper

The final step of using our framework was physical construction of a gripper according to the results of the optimization. This required setting limits for the optimized parameters that could be implemented in practice. In particular, we used a limit of $-5mm \leq l_i \leq 5mm \forall i$ to ensure that the tendon route was inside the physical volume of each link.

Using a fixed grasp pool affects the parameters that can be part of the analysis: if the values of r (joint radius) and d (link length) can change as part of the optimization, the result can be kinematically different than the model used for the generation of the grasp pool. Taking this aspect into account, we performed two optimizations. For the first one (referred to as Optimization 1), we fixed the values to $r = 5mm$ and $d = 20mm$. For the second one (Optimization 2) we set the limits $3mm \leq r \leq 10mm$ and $15mm \leq d \leq 22mm$. We also added the constraint $r + d = 25mm$, to ensure that the overall length of the fingers would not change.

The joint stiffness levels require additional discussion, as there are two cases to consider. First, during the early stages of the grasp, spring forces and tendon forces play equal parts in the process. Once the fingers are closed however, tendon forces can be increased arbitrarily, while spring forces do not change if the grasp is stable and no joint movement occurs. In the limit, tendon forces dominate to the point that spring forces become negligible. An ideal grasp would be stable in both of these phases. For each of the two optimizations, we used the following convention. Tendon route values were derived by solving the optimization problem without joint springs, in an attempt to ensure grasp stability in the limit. The resulting values were then plugged back into the optimization, in order to compute spring stiffness values, using as limits $1.0 Nmm/rad \leq k_i \leq 2.0 Nmm/rad$.

The dual nature of joint stiffness vs. tendon force optimization presents interesting possibilities and requires more detailed exploration than presented here. In particular, we envision a case where joint springs are used to determine the postures that the hand achieves in "free motion", thus affecting pre-grasp behavior, while tendon routes are optimized for stable grasps in the limit. The transition between these two phases will present additional challenges to address. We believe that the optimization framework presented here can be a first step in this direction.

The results of the optimizations are shown in Table 1. We




Figure 5. UNBALANCED JOINT FORCES OVER THE SET OF GRASPS IN THE OPTIMIZATION POOL.

	l_o	l_1	l_2	k_o	k_1	k_2	r	d
Ad-hoc	5.0	5.0	5.0	1.0	1.0	1.0	5.0	20.0
Optimization 1	5.0	5.0	1.72	1.0	1.0	2.0	5.0	20.0
Optimization 2	5.0	4.64	1.02	1.0	1.0	2.0	7.89	17.11

Table 1. RESULTS OF GRIPPER DESIGN OPTIMIZATION.

only show the values for one of the fingers, since, as mentioned before, the results for the other finger are symmetrical. For a quantitative analysis of the computed optimal configurations, we compared them against an ad-hoc parameter set, with $l_i = 5mm$ and $k_i = 1Nmm/rad \forall i$. The comparison criterion was the level of unbalanced joint forces for each grasp in the limit case. The results are shown in Fig. 5. We notice that the optimized configurations provide significantly more stable grasps across the optimization pool. In addition, allowing a change in the link length and joint radius provides additional stability, but the gains are diminished compared to the case where these parameters are fixed. The total time spent formulating and solving each optimization problem was less than a minute, using a commodity desktop computer. This suggests the future possibility of scaling to much larger grasp optimization pools.

We constructed a prototype gripper using the results of Optimization 1. The links were built using a Stratasys FDM rapid prototyping machine, and assembled using elastic joints cut from a sheet of hard rubber. Each link contained a tendon route with the entry and exit points set according to the optimization results. The width of the strip of rubber was varied for each joint to provide the desired stiffness ratios. As this prototype is intended as a proof-of-concept for the kinematic configuration and design parameters, no motor or sensors were installed; instead, actuation was performed manually.

We found that the prototype gripper is capable of a wide




Figure 6. TWO GRASPS (CENTERED AND ASYMMETRICAL) EXECUTED WITH THE PROTOTYPE GRIPPER.

range of grasping tasks and does not require precise positioning relative to the target object. Its passive adaptation ability is exemplified in Fig. 6, which shows the execution of two grasps. The first one starts from a centered position and leads to relatively similar joint values for both fingers. In contrast, the second grasp requires the joints to conform to an asymmetrical, irregular shape. Both grasps were executed successfully. Figure 7 attempts to provide an illustration of the spectrum of grasps that can be carried out with this gripper. All of the presented grasps were executed successfully and the object was securely lifted off the table, with very little time or effort spent positioning the gripper relative to the target.

DISCUSSION AND CONCLUSIONS

A qualitative analysis of the optimization results can start from the observation that the prototype gripper is capable of executing both fingertip grasps (of varying finger spans) and enveloping grasps (of both regular and irregular shapes). Intuitively, fingertip grasps require relatively low torques on the distal joints, so that fingertip forces are in opposition, rather than oriented towards the palm. Conversely, larger torques on the distal joints benefit enveloping grasps; as a result, the optimization process was required to combine two somewhat opposing goals. The results indicate that the solution indeed enables both kinds of grasps, but the distal joint is both stiffer and less powerful than the proximal ones. In fact, our optimization framework achieves this characteristic by “saturating” many of the hand parameters, which take either the minimum or maximum value allowed.

In this sense, the result of the optimization could be interpreted as meaning that the addition of a third link to the gripper provides little benefit. The resulting gripper comes close to a model with two links per finger, a design also confirmed in the optimization studies of Dollar and Howe [9]. We believe that this is the type of analysis that our framework is natively suited for: in future iterations, we can directly compare two- and three-link models, and compute a *numerical measure* of the benefit provided by the additional link. The relatively simple two-fingered

design that we used here allows an intuitive understanding of the design choices (which makes it well suited for initial testing and proof-of-concept implementations). However, for more complex models, empirical analysis becomes unfeasible, and quantitative tools, such as the one presented here, can prove more valuable.

We also believe that future hand design studies will consist of a combination of exhaustive search and optimization problems for which more efficient algorithms are available. The space of possible hand designs, and implicitly the domain of parameters to be optimized, is practically limitless. Virtually any hand design ever proposed involves some compromise of ad-hoc decisions vs. informed, optimized parameter choices. In our case, we have discussed aspects such as tendon routes and joint stiffness. However, by moving up in the scale at which we are analyzing the hand, we can uncover many more design decisions, which we assumed as given: number and configuration of links, kinematic chains, *etc.* Some of these will likely prove impossible to encapsulate in a solvable optimization problem, thus some contribution from numerical approaches will be unavoidable.

An interesting aspect concerns the on-line algorithms that are used to control the hand during grasping tasks. Traditionally, these algorithms have been designed after the hand was constructed, carefully tuned to extract the best performance from a given mechanical design. Off-line hand optimization enables the opposite approach: the hand mechanism is designed to suit a particular algorithm. The same applies to sensor arrays: we can build a hand that is optimized for the types of grasps that we can expect to perform based on data from a certain sensor. In this way, the hand is intrinsically equipped to handle the shortcomings of the input data. Overall, it seems natural to ask: what comes first, the hand or the algorithm?

As robots with the ability to operate in unstructured settings are constantly evolving, we believe that research on adaptive and underactuated designs has the potential to ultimately provide us with inexpensive and easy-to-build, yet effective robotic hands for a variety of applications in human environments.

REFERENCES

- [1] Ciocarlie, M., and Allen, P., 2009. “A design and analysis tool for underactuated compliant hands”. In IEEE-RSJ Intl. Conf. on Intelligent Robots and Systems.
- [2] Ciocarlie, M., and Allen, P., 2010. “Data-driven optimization for underactuated robotic hands”. In IEEE Intl. Conf. on Robotics and Automation.
- [3] Gosselin, C., Pelletier, F., and Laliberte, T., 2008. “An anthropomorphic underactuated robotic hand with 15 Dofs and a single actuator”. In IEEE Intl. Conf. on Robotics and Automation.
- [4] Hirose, S., and Umetani, Y., 1978. “The development of soft gripper for the versatile robot hand”. *Mechanism and Machine Theory*, **13**, pp. 351–358.




Figure 7. GRASPS EXECUTED WITH THE PROTOTYPE GRIPPER.

- [5] Ulrich, N., Paul, R., and Bajcsy, R., 1988. “A medium-complexity compliant end effector”. In IEEE Intl. Conf. on Robotics and Automation, pp. 434–439.
- [6] Birglen, L., Laliberte, T., and Gosselin, C., 2008. *Underactuated Robotic Hands*. Springer Tracts in Adv. Robotics.
- [7] Gosselin, C., Laliberte, T., and Degoulange, T., 1998. “Underactuated robotic hand”. In Video Proceedings of the IEEE Intl. Conf. on Robotics and Automation.
- [8] Laliberte, T., Birglen, L., and Gosselin, C. M., 2002. “Underactuation in robotic grasping hands”. *Machine Intelligence & Robotic Control*, **4**(3), pp. 1–11.
- [9] Dollar, A., and Howe, R., 2006. “Joint coupling design of underactuated grippers”. In 30th Mech. and Robotics Conf.
- [10] Dollar, A., and Howe, R., 2007. “Simple, robust autonomous grasping in unstructured environments”. In IEEE Intl. Conf. on Robotics and Automation.
- [11] Prattichizzo, D., and Trinkle, J. C., 2008. “Grasping”. In *Springer Handbook of Robotics*. Springer, pp. 671–700.
- [12] Ferrari, C., and Canny, J., 1992. “Planning optimal grasps”. In IEEE Intl. Conf. on Robotics and Automation.
- [13] Kwak, S. D., Blankevoort, L., and Ateshian, G., 2000. “A mathematical formulation for 3d quasi-static multibody models of diarthrodial joints”. *Computer Meth. in Biomech. and Biomed. Eng.*, **3**, pp. 41–64.
- [14] Mosek. Mosek ApS Denmark. <http://www.mosek.com>.
- [15] Miller, A., and Allen, P. K., 2004. “GraspIt!: a versatile simulator for robotic grasping”. *IEEE Robotics and Automation Magazine*, **11**(4), pp. 110–122.