

The Stewart Hand

*A Highly Dexterous,
Six-Degrees-of-Freedom
Manipulator Based on the
Stewart–Gough Platform*



©SHUTTERSTOCK.COM/DIFFERR

By Connor M. McCann, Vatsal V. Patel, and Aaron M. Dollar

In this work, the design, modeling, dimensional synthesis, and experimental characterization of a dexterous robotic hand based on the Stewart–Gough parallel manipulator are presented. The mechanism consists of three parallel linkage fingers with six prismatic actuators that control manipulation. An additional actuator maintains stable grasping forces. A computational model to predict the hand's workspace is then utilized to determine the optimal design parameters that maximize the workspace size and manipulability. A physical prototype based on the optimized parameters is built and experimentally characterized to assess its performance (Figure 1).

The hand demonstrates a large range of motion in all spatial degrees of freedom (DoF). Even with a simple open-loop controller, high accuracy/precision were observed in most of the primary motion directions, with only slightly decreased accuracy in the plane parallel to the palm. The potential of closed-loop control was demonstrated in a teleoperation task, where an operator performed a 6-DoF peg-in-hole

insertion task with variously shaped blocks. The hand design will be made freely available through the Yale OpenHand Project to facilitate future research efforts in this area.

Background

For nearly all robotic applications, it is necessary for a robot to physically interact with its environment, altering the landscape through which it navigates. While some constrained tasks (such as those found on manufacturing assembly lines) can be accomplished with single-purpose end effectors, more complicated unstructured tasks typically require a general-purpose robotic hand that can adapt to objects with varied geometries and mechanical properties. During the past century, hundreds of robotic hands have been proposed in the literature for applications such as assistive robotics, prosthetics, teleoperation, and logistics [1].

While many of the hands proposed to date focus on simply grasping an object, it is often beneficial to manipulate an object to reposition it within the hand. This within-hand manipulation strategy has been demonstrated to increase energy efficiency, accuracy, and obstacle avoidance for manipulation tasks [2]. Given the impressive dexterity of the

Digital Object Identifier 10.1109/MRA.2021.3064750

Date of current version: 6 April 2021

human hand, the vast majority of robotic hands intended for within-hand manipulation adopt an anthropomorphic design, wherein a set of three to five fingers consisting of a serial arrangement of revolute joints is utilized [3]. This




Figure 1. The physical prototype of the proposed 6-DoF Stewart Hand presented in this work, performing a teleoperated dexterous manipulation task.

design has a number of benefits given the simple forward kinematics of the fingers and the relative degree of intuition for human operators. That said, there are a number of drawbacks, as well, due to the high actuation and sensing complexity of the architecture [14], [15].

One of the biggest challenges arises from the fact that anthropomorphic hands are kinematically redundant, having more actuators in the mechanism than spatial DoF at the object. This requires advanced control algorithms to synchronize the motions of all the actuators so as to maintain a stable grasp on an object. If this controller has too much latency or if inadequate sensory information is available, the hand can drop the object. This complexity often limits the utility of these hands for general-purpose applications even if they are capable of complex manipulation in controlled settings. Thus, it is desirable to limit or avoid kinematic redundancy in the mechanism design of dexterous robotic hands.

An alternative design ideology can be found in parallel robotic manipulators [12]. Unlike anthropomorphic hands—which are usually fully actuated at all DoF along the fingers—parallel robots typically have only one actuator on each “leg” of the mechanism, while the remaining joints are passive. This results in underactuated legs, with the remaining kinematic constraints coming from the addition of the end effector platform. Such a manipulator has no kinematic redundancies and has only one actuator for each spatial DoF. This has a number of benefits in terms of simplifying the control of the manipulator.

A number of efforts have recognized similarities between robotic hands and parallel manipulators (wherein the grasped object is analogous to the end effector platform, and the fingers to the mechanism legs). Some efforts have sought to exploit this similarity by applying parallel robot analysis techniques to robotic hands [16]. Additionally, a few robotic hands have utilized parallel mechanisms in their finger designs. Many of these contributions are summarized in Table 1. In one instance, an underactuated

Table 1. A comparison of mechanism architectures.

Work	Kinematic Architecture			Kinematic Constraints		
	Palm	Fingers	Hand	Palm	Fingers	Hand
Anthropomorphic hands [3], [4]	—	Serial	Parallel	—	Exactly constrained	Overconstrained
Yale OpenHand Project [5]	—	Serial	Parallel	—	Underactuated	Underactuated
Laliberté and Gosselin, 2000 [6]	—	Parallel	Parallel	—	Underactuated	Underactuated
Tanikawa et al., 2000 [7]	—	Parallel	Parallel	—	Exactly constrained	Overconstrained
Nefzi et al., 2006 [8]	—	Parallel	Parallel	—	Exactly constrained	Overconstrained
Jin et al., 2020 [9]	—	Parallel	Parallel	—	Exactly constrained	Overconstrained
Yuan et al., 2020 [10]	—	Parallel	Parallel	—	Exactly constrained	Overconstrained
Cui and Dai, 2011 [11]	Parallel	Serial	Parallel	Exactly constrained	Exactly constrained	Overconstrained
<i>Parallel Mechanisms</i> [12]	—	Either	Parallel	—	Underactuated	Exactly constrained
McCann and Dollar, 2017 [13]	—	Parallel	Parallel	—	Underactuated	Exactly constrained

mechanism based on four-bar linkages was utilized to achieve underactuation at the hand level; however, within-hand manipulation was not the objective [17]. Otherwise, to the authors' knowledge, all other published examples of parallel mechanism finger designs have been fully actuated [7]–[10]. Such a system is overconstrained, yielding many of the same control complexities found in anthropomorphic designs since the motions of all joints must be strictly coordinated to maintain a stable grasp on an object.

Our work seeks to address this point by presenting a robotic hand intended for within-hand manipulation based on the Stewart–Gough platform parallel mechanism. Unlike other parallel hand designs, this one has underactuated fingers and is exactly constrained, resulting in a system with as many position-controlled actuators as spatial DoF. A very simple open-loop control strategy is capable of achieving 6-DoF manipulation without the risk of losing grasp stability. Preliminary conference papers on the topic presented a proof-of-concept prototype hand [13] and a computational analysis of the architecture that identified an optimal set of design parameters that balanced manipulability with a 6-DoF workspace size [18]. In this current work, we present a final “optimized” design (Figure 1) that implements the results of the computational study, and we experimentally characterize the hand through a number of physical experiments examining spatial dexterity and accuracy.

Mechanism Concept and Dimensional Synthesis

Mechanism Overview

The design presented in this work is schematically depicted in Figure 2, which illustrates the kinematic joint structure and the pulley differential used to close the hand. Both representations have been simplified for illustrative purposes. The design includes three fingers, each consisting of a revolute joint at the base, a planar linkage with two revolute–prismatic–revolute kinematic chains, and a spherical joint that enables a compliant fingertip to reorient during object manipulation. In each finger, the prismatic joints are actuated in position control (to facilitate manipulation), and the revolute joint at the base is actuated in torque control (to apply a constant, inward torque that maintains frictional contact forces on an object). All other joints are passive. While it is necessary to apply this inward torque at the base of each finger for stability, it is not desirable to impose a kinematic constraint at this joint, as this would overconstrain the mechanism and result in control redundancies. To avoid this issue, a floating pulley differential couples these three joints to a single revolute actuator mounted in the palm of the hand. This actuator is then operated in a torque-controlled mode to avoid imposing any additional kinematic constraints.

The advantages of this approach become apparent when considering the control of the system. The single “grasp actuator” can simply be set to a constant torque value when an object is being grasped and then must merely maintain this torque for the duration of the grasp. Thus, all manipulation occurs as a result of the motion of the six prismatic actuators.

This yields an exactly constrained system that is very simple to control with minimal sensing. By utilizing linear servo motors with built-in, low-level position sensing/control, it is




Figure 2. A schematic representation of the hand mechanism, showing (a) the kinematic joint structure, (b) the pulley differential used to apply equal torque to each of the fingers with a single motor, and (c) the finger normal and contact coordinate frames at each fingertip used during the analysis.

possible to realize complex within-hand manipulation tasks with an otherwise open-loop controller. In fact, the only additional sensors included in the design are potentiometers that measure the angles of the revolute joints at the base of the fingers. These values are used only immediately after an object has been grasped to estimate its shape/size between the fingertips. Once this measurement has been taken, the ensuing manipulation can be carried out by commanding position values to the linear actuators, without any additional sensory input. This greatly simplifies the control of the system.

Theoretical Workspace Modeling

To make the hand as versatile as possible for real-world grasping applications, it is desirable to maximize the 6-DoF workspace size. To assess different hand designs according to this criterion, a computational model was developed to predict the stable workspace [18]. To determine whether a given pose lies within the workspace of the hand, it is necessary to consider kinematic and frictional factors. Kinematic limits can be easily considered by computing the inverse kinematics of the hand and ensuring that all joint limits are obeyed [13]. Frictional considerations are included by calculating the vector of nine independent reaction force components at the fingertips, f_c , when an object undergoes a 6-DoF external wrench, F_e . These contact forces must lie within the stable friction cone of the fingertips, as determined by the coefficient of friction. These contact forces can be mapped to the external wrench via a grasp matrix [19]:

$$\mathbf{G}f_c = -F_e. \quad (1)$$

This system is rank-deficient for a three-fingered hand, so it is typically not possible to uniquely solve for the contact forces given an external wrench. In this case, though, the system can be augmented to yield a unique solution by utilizing the fact that the component of the reaction force normal to the plane of each finger must exactly counter the grasping torque, τ_G , exerted by the differential at the base of the finger. To express this constraint for the i th finger, two coordinate frames are defined [Figure 2(c)]: one (C_i) at the contact location on the fingertip, with its z -axis along the inward normal of an object, and the other (N_i) at the spherical joint, with its z -axis normal to the plane of the finger facing inward toward the object. The contact wrench can be expressed in each frame as

$$F_{N_i} = \mathbf{Ad}_{g_{N_i C_i}}^T F_{C_i}, \quad (2)$$

where the mapping is defined by the adjoint transformation matrix [19]. If the current pose of the hand places the spherical joint at a radius of r_i from the base of the finger, the z component of the force in the N_i frame is known and can be written as

$$\tau_G r_i = \mathbf{U} \mathbf{Ad}_{g_{N_i C_i}}^T \mathbf{Q}_i f_c, \quad (3)$$

where \mathbf{U} and \mathbf{Q}_i are indexing matrices used to select appropriate components (for further details, see our previous work

[18]). Here, F_{C_i} has been expressed in terms of the nine independent components of f_c .

Augmenting the original equilibrium expression with these constraints yields a full-rank system:

$$\begin{bmatrix} \mathbf{G} \\ \mathbf{U} \mathbf{Ad}_{g_{N_i C_i}}^T \mathbf{Q}_1 \\ \mathbf{U} \mathbf{Ad}_{g_{N_i C_i}}^T \mathbf{Q}_2 \\ \mathbf{U} \mathbf{Ad}_{g_{N_i C_i}}^T \mathbf{Q}_3 \end{bmatrix} f_c = \begin{bmatrix} -F_e \\ \tau_G r_1 \\ \tau_G r_2 \\ \tau_G r_3 \end{bmatrix}. \quad (4)$$

This system can be inverted to yield the unknown contact forces, f_c . Once these forces are obtained, the friction cone stability criteria can be checked to determine whether a particular pose lies within the frictionally stable workspace of the hand.

Computational Optimization

The model described in the previous section can determine whether a given pose of a certain hand design is feasible or not based on kinematic and frictional constraints. To utilize this model to optimize the hand design, an overall design quality metric was defined to maximize the 6-DoF workspace volume. In addition to the workspace size, it is also important to consider the manipulability of the hand to ensure that singular configurations are avoided. For this purpose, the local transmission index (LTI) was calculated at each pose [20]. The LTI value ranges between zero and one, with zero indicating a fully singular pose and one representing ideal motion transmission. To assess a given hand design, a large regular grid of workspace poses was sampled and assessed using the model for a number of different external wrench directions. A quality metric was defined by summing the LTI across all poses found to lie within the workspace.

Four key design parameters were identified for the hand (where all length scales have been normalized by the fully extended prismatic actuator length, ℓ), as in the following:

- palm radius, r_p/ℓ
- spherical joint normal offset, n/ℓ
- finger base angle (in the plane of the palm), α
- prismatic actuator stroke length, s/ℓ .

A grid search was performed by varying each of these parameters and testing the design for a number of different object sizes and external wrench directions. Not surprisingly, it was found that maximizing the relative actuator stroke, s/ℓ , always improved the design quality. For the remaining three variables, a less trivial result was found, as illustrated in Figure 3. It can be seen that there exists a locally optimal set of design parameters. The optimal design had the following parameters:

- $r_p/\ell = 0.361$.
- $n/\ell = 0.044$.
- $\alpha = 1.222$ (70°).
- $s/\ell = 0.5$.

It is worth noting that this was the maximum tested value of s/ℓ , but longer actuator strokes would always be desirable if possible, given manufacturing constraints. In addition to identifying this optimal design, one can note from Figure 3

that the performance degrades most steeply when the palm radius is increased. Thus, this is an important variable to prioritize when designing a physical prototype of this hand.

Optimized Prototype Design

Building off the optimization results, a prototype hand was designed using the ideal design parameters [Figure 4(a)]. To promote the use of this novel class of manipulators by other researchers, the design has been made available through the Yale OpenHand Project [5]. The hand was designed to be fabricated from commercially available hardware, 3D-printed parts, and standard servo motors (Actuonix L12-50 and Dynamixel XH-430), making it accessible to others interested in the technology. One of the custom components was found to be most robust when it was crafted from machined/water-jet-cut aluminum sheet stock, but it can be laser-cut from various plastics or wood for low-force applications.

One of the key insights from the design optimization was that maximizing the relative stroke length of the linear actuators is critical to achieving as large a workspace as possible. There is a limited number of commercially available linear servo motors at a size appropriate for robotics hands, with most exhibiting a stroke ratio (s/ℓ) of approximately 0.3. To proportionally increase this value, the actuators were mounted in a “submerged” position underneath the base of the finger [this length offset is depicted in Figure 4(b)]. This necessitated shifting the actuators away from the revolute joint axes at the bottom and top of the finger. Moreover, the linear actuators were staggered to minimize unnecessary additional length at the top of the finger. Through these adjustments, the prototype exhibits $s/\ell = 0.62$, which is greater than the maximum value tested in the grid search and thus should yield even better hand performance. The actuators could not be offset lower underneath the finger, as this would result in collisions with the pulley differential under the palm of the hand.

The fingertip design for this hand exhibits a number of improvements over the iteration presented in [13]. The first

version used a magnetic spherical joint design wherein a spherical magnet and a disk magnet attracted each other [Figure 4(c)], simultaneously holding the fingertip on the hand and providing a restoring moment to realign the fingertip when a grasp was completed [13]. Strong magnets were required to hold the fingertip in place; however, this resulted in large rotational moments that resisted the reorientation of the spherical joint. The new iteration of the design—also shown for comparison in Figure 4(c)—decouples these loads by using a single disk magnet to anchor the fingertip to a ferromagnetic sphere. A set of five cast silicone rubber leaf springs serves to reorient the fingertip. By adjusting the cross section of the leaf springs, the restoring moment can be precisely controlled. The magnetic spherical joint enables a much larger range of motion than most commercially available ball-and-socket joints since, traditionally, the socket must cover more than half of the sphere to maintain geometric closure on the ball—a constraint not present in a magnetic implementation.




Figure 3. The simulated performance landscape as a function of the primary design variables: the palm radius, finger base opening angle, and spherical joint normal offset. The three planes intersect at the identified optimal design.




Figure 4. CAD illustrations of the hand prototype developed using the optimized design parameters, highlighting (a) the full hand architecture, (b) finger design considerations, and (c) fingertip design alterations between an earlier version presented by the authors (right) and the improved scheme (left).

As with any physical prototype, a number of design compromises were made during development. Given the importance of the palm radius as identified in the simulation results, this parameter was considered nonnegotiable. While the final hand design successfully achieves all the design parameters




Figure 5. The simulated prototype hand workspace, considering joint limits and self-collisions along with frictional and kinematic constraints.




Figure 6. An experimental characterization test setup with the Stewart Hand facing down. A marker detection camera points up at the hand. The reset mechanism (consisting of a cable to reset the object and a magnetic arm used to steady the object to prevent swaying) repositions the object before each grasp. The yellow board is employed prior to the experiment to calibrate the global coordinate frame of the camera.

found in the optimization, it does have limits on the motion of the revolute joints at the base of the fingers, due to self-collisions. The simulated workspace considering these factors is shown in Figure 5.

Experimental Characterization

The goal of this study was to determine how the performance of the optimized Stewart Hand compares with its simulated performance. To evaluate this difference, the workspace of the physical hand is quantitatively analyzed through single and multi-DoF characterizations, manipulation range-of-motion and qualitative grasping experiments with a variety of representative objects, and teleoperated tests with the hand mounted on a whole-arm manipulation (WAM) robot arm simulating real-world applications.

Methods

For the range of motion experiments described later in the article, the Stewart Hand was supported on a frame, with the palm facing down and a fixed camera pointing up at the hand. This orientation during object manipulation was seen as a more realistic representation of the hand's interactions with objects when mounted on a robot arm. This also makes manipulation tasks more difficult, as gravity can eject objects from the hand instead of allowing the palm to catch them. The fixed camera at the bottom of the frame was used to record the position and orientation of grasped objects.

For single-axis characterization, a 60-mm-diameter sphere object with an ArUco marker on a small, flat face was used to record position and orientation errors under open-loop control [21], [22]. This object also housed a magnet on the flat face to interface with the reset mechanism shown in Figure 6. The reset mechanism consisted of a swing arm that situated the sphere at the home position before each grasp. This arm also prevented any pregrasp motion of the object and ensured the repeatability of the object's starting position and the location of the finger contacts at the initial grasp. The object was tethered by a thin cable originating from the palm of the hand. The tether served to pull the object toward the palm after a manipulation sequence was completed and to position the object on top of the swing arm. This tether did not support the weight of the object and turned slack before the object was grasped so that it did not interfere with manipulation. The reset mechanism enabled faster and more consistent data collection by automating object repositioning, and it eliminated variability across manipulation sequences by standardizing the initial grasp configuration.

The tendon differential mechanism across the three fingers enables us to perform simple open-loop control for both the grasping and manipulation of arbitrary objects. The grasping actuator is commanded to a closing torque, and the differentially driven fingers conform to the shape of the object. During manipulation, the fingers can give and take tendon freely through the differential, thus enabling

the revolute joints at the base of the fingers to rotate akin to a Stewart–Gough platform. As a result, all grasping and manipulation experiments were carried out using a purely open-loop control scheme that commanded linear actuators to positions determined by the inverse kinematics of the hand. The actual position and orientation of the object were recorded using the marker and the fixed camera but not used for any kind of feedback control. More sophisticated control will be carried out in the future using actual object pose and potentiometer values from the fingers.

In addition to tracking the actual position and orientation of the object for workspace characterization, a slip metric was calculated to quantify the amount of slippage occurring at the contact fingertips along the object's surface. The three potentiometers, which were mounted on the axes of the revolute joints at the base of the fingers, were used to read the angle between the palm and each of the fingers. By knowing this angle and the current stroke length of each linear actuator, the actual position of the three fingertips in space could be calculated. The deformation of the triangle formed by the three fingertips during a manipulation sequence could be quantified by the slip metric at each step. The metric is computed as the maximum difference between the corresponding side lengths of the contact triangle at the current pose ($s_{k,\text{current}}$) from that at the initial grasp ($s_{k,\text{initial}}$) normalized by the initial lengths:

$$\text{slip} = \max_{k=1,2,3} \left\{ \frac{|s_{k,\text{current}} - s_{k,\text{initial}}|}{s_{k,\text{initial}}} \right\}. \quad (5)$$

The metric could also be utilized for thresholding and stopping a manipulation sequence before an object was expected to be dropped, providing a reliable alternative to vision-based drop detection.

Single-DoF Axes Characterization and Range of Motion

The following experiments were carried out to quantitatively evaluate the performance of the Stewart Hand for manipulation along a single axis of motion. The first set of tests characterized the accuracy of these motions for each of the six axes, while the second set of experiments with YCB and foam objects served to assess whether this range of motion extended across objects of different shapes and sizes.

Axes Characterization

As noted previously, a 60-mm-diameter sphere object is used to characterize the performance of the hand in the translation (x , y , and z) and rotation (roll, pitch, and yaw) directions. The circular profile of the object ensures that consecutive grasps are equivalent and that the hand can repeatedly grip a repositioned object at the same starting position. Once the object is grasped from the reset mechanism, the hand is commanded to the kinematic end of the range of motion along a particular axis, and the object's actual position and orientation are recorded as the hand carries out the manipulation sequence. The results of three trials on each of the six axes are shown in Figure 7.

The hand performs reorientations of the object with relatively high accuracy and is also able to execute z translations with minimal error. The ranges of motion for rotations and z -axis translation are limited only by the stroke of the linear actuators. Thus, a larger stroke would result in even larger ranges of motion along these axes. The xy translation performance of the hand is quite linear but deviates substantially from the theoretical case. This is attributed primarily to the competing grasping force from the tendon differential and the manipulation force from the linear actuators. To translate the object in the xy -plane, the linear actuators for at least one of the fingers extend, and the rest of the fingers are required to rotate away from the palm. While the return spring on each finger can aid the latter motion, the grasping motor still continues to apply the same force across all the fingers as a result of the differential.

For instance, the positive y -axis range is truncated, as it points directly in the direction of one of the fingers, requiring it to rotate opposite the grasping torque. Friction in the tendon differential opposes this finger motion, and, subsequently, the motion of the object is hindered. However, the open-loop control continues to extend the linear actuators, and, as a result, the fingertips start to slip on the object. The drop in the actual negative x -axis range was observed to stem from these fingers substantially slipping around the object toward the end of the range and ineffectively applying contact forces in the desired direction of motion. The effect of the grasping torque on the poor xy translation range is evaluated in the “Effect of Grasping Motor Torque on Planar Workspace” section and discussed in more detail in the “Discussion” section.

As mentioned, a slip metric was devised using data from the potentiometers at the base of the fingers and the linear actuator positions to determine the contact triangle at the fingertips. The contact triangle at the initial grasp could then be compared to that during a manipulation sequence. Figure 8 shows a top-down view of how this triangle deforms around the surface of the spherical object as it is translated through the x and y ranges of motion (since these axes experienced the most significant slippage). The contacts can also move in the z -axis along the surface of the sphere, and some of the contact locations appear to cross the dotted circle denoting the object boundary at the grasp (this would not have been possible with a cylindrical object). The plots in Figure 8 describe the trend of the slip metric evaluated at each step of the motion. As is evident from the figure, the amount of slippage at the fingertips (and the corresponding slip metric value) increases away from the center through x and y translation sequences.

Range of Motion With YCB and Foam Objects

We validated the real-world performance of the optimized hand by commanding the hand to grasp 10 objects—eight from the YCB Object and Model Set and two identically sized foam cylinders (one rigid and one soft)—and manipulating them to the end of the axes' ranges of motion [Figure 9(a)] [23]. The objects were selected for their wide variety of shapes and sizes, and they tested the differential's ability to achieve and maintain adaptive finger contacts.

The two foam cylinders were evaluated to assess the effect of object deformability on the range of motion. Each object was tagged with an ArUco marker to record its actual position and orientation.

The hand systematically grasps the object placed approximately near its home position and subsequently manipulates the item to the end of the range of motion along a particular axis. The same motion is carried out for three trials. The limits of the workspace are determined by reading the finger positions from the three potentiometer values after the grasp because every grasp configuration results in a slightly different kinematic workspace size. For example, a really large-diameter object would not be able to be translated in the xy -plane as much as a small diameter object. The object is reset at the home position, and the process is repeated for the three translation and three rotation axes. For consistency, all translations and rotations are reported relative to the initial centroid of the grasping triangle.

The fingers were observed to successfully adapt to the shape of all the objects tested, and none of the objects was dropped within the hand's workspace. Fingertip slipping was noticed on some objects, depending on the surface material

and local surface curvature. The tip of the pear has high curvature, and the object is made up of smooth plastic material that causes the contact at that point to slip easily during a manipulation sequence. Similarly, one of the fingers would naturally slide between the handles when grasping the clamp. For such objects, the large range of motion of the magnetic spherical joints permitted sufficient passive adaptability of the fingertip to the surface curvature. The Rubik's cube and the banana had notable slipping along their surfaces due to their smooth plastic materials. Executing a three-finger grasp on a square profile is quite challenging, too, as a minimal amount of slipping can risk the ejection of the cube. While grasping thin objects, such as the spatula handle, the fingertips occasionally collide with one another since they converge at a single point.

Figure 9(b) plots the mean observed range of motion for each of the objects along the six axes. A few trends are worth noting. The z translation range for all the objects was almost identical, as it was determined almost solely by the stroke length of the linear actuators. Objects with a smaller diameter had larger orientation ranges of motion because the actuator stroke acting closer to the object's center generated a larger rotational angle than the same stroke farther from




Figure 7. Axis characterization results from commanding the object through the kinematic ranges of motion along the six axes. The object's actual position and orientation are recorded for three trials on each axis. Gray dots represent experimental data. The red dashed line indicates a theoretical case where the actual pose matches the commanded pose exactly. (a) x . (b) y . (c) z . (d) Roll. (e) Pitch. (f) Yaw.

the center. For example, the golf ball, banana, and spatula can be rotated across a much larger pitch angle compared to a miniature soccer ball with the same linear actuator stroke. Similarly, the xy translation range is capped by the rotation of the finger bases. It is also worth noting that the x and y translation ranges are different for some objects since these two axes are aligned between and along the fingers, respectively. This planar workspace is explored further in the “Effect of Grasping Motor Torque on Planar Workspace” section. Larger objects, such as the soccer ball, require the fingers to be more open at the initial grasp, and the xy range is thus limited by the finger bases running into hard stops. Finally, deformability did not seem to have a significant effect with the objects tested, since the soft foam cylinder has ranges of motion that are similar to the rigid foam cylinder. This might be because the grasp triangle geometry is evaluated once an object has been grasped; once the deformable cylinder has been squeezed by the fingers, it is effectively just a cylinder with a slightly smaller radius. The hand would likely experience a decrease in accuracy if an object deformed significantly during the manipulation, rather than just during the initial grasping phase.

Effect of Grasping Motor Torque on Planar Workspace

The relatively poor xy translation range was noted in the single-axis characterization experiments, and one of the preliminary

root causes was identified to be the competing grasping and manipulation forces discussed previously, resulting in slipping. The grasping torque is applied by the singular Dynamixel grasping motor through the tendon differential, while manipulation is coordinated by the six linear actuators. To evaluate the effect of the grasping motor torque magnitude on the xy manipulation range of motion, we commanded the hand to grasp and translate the sphere object discussed in the “Single-DoF Axes Characterization and Range of Motion” section to the kinematic limits of the xy workspace at $Z = 0$ for five different values of the grasping motor torque. The positions were commanded open-loop to the hand, and the object’s actual position was recorded using the ArUco marker on its face.

Figure 10 presents the envelope of the actual xy manipulation range of motion of the hand at these five grasp motor torque values. For each run, identical positions were commanded, and the linear actuator parameters were left unchanged. As a reference, this test was also conducted with a fixed platform (replacing the fingertips), essentially converting the hand into a traditional Stewart–Gough platform. No grasping motor torque was required for the fixed platform case since the platform was affixed via a rigid spherical joint. For each of the cases tested, the area of the actual xy workspace envelope was calculated and normalized by the commanded workspace area, as detailed in the table in Figure 10.

For a very low grasping motor torque, the object is observed to easily slip from the hand. The low normal contact

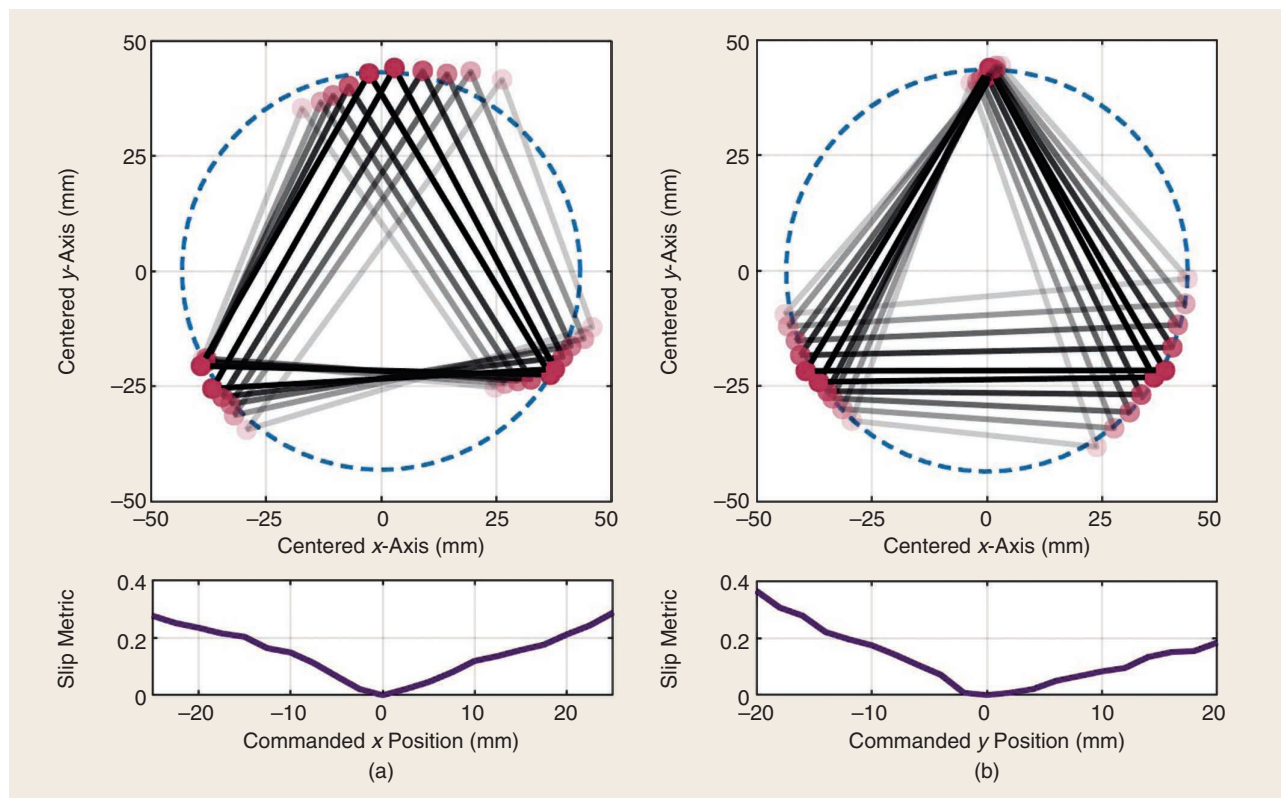


Figure 8. A top-down view of the contact triangle in the object frame deforming around the sphere when commanded through the entire range of (a) the x -axis and (b) the y -axis. The dashed line denotes the object boundary, and the slip metric is shown in plots for the corresponding axes motions.

force leads to insufficient friction between the fingertips and the object, and, consequently, the xy workspace is small. Conversely, a very high grasping motor torque is able to generate a robust grasp. In addition to increased contact forces, however, higher motor torques result in increased tendon differential friction. Thus, the manipulation forces from the linear actuators are too small in comparison, and the fingers are unable to rotate counter to the grasping torque direction, again resulting in a small xy workspace. A grasping motor torque somewhere

between these two limits is ideal, as it balances the requirement of producing sufficient contact forces while allowing the linear actuators to reconfigure the differential during object translation. Note that the grasping motor torque identified as ideal for one object might not be effective for all—the selection of torque is dependent on the surface properties of the object and resulting friction at the fingertips. That is, smoother surfaces might require a high grasping torque to generate any kind of translation motion without the object slipping out of

the hand. Future work might investigate ways to dynamically determine the ideal value by observing the slip metric in real time and adapting the grasping torque.

Multi-DoF Workspace Exploration

The single-axis characterization and range of motion tests are helpful to evaluate the relative performance of the hand across the different DoF. But most real-world manipulation tasks often require coupled multi-DoF motions combining translations and rotations. We tested the multi-DoF performance of the hand by executing up to nine spatial orientations at 100 different points within the Stewart Hand's kinematic workspace. The kinematically feasible workspace of an equivalent Stewart–Gough platform was first simulated and then discretized, and nine orientations (one with no rotations and eight multi-DoF rotations in all three axes) were tested at each point within this theoretical workspace. From this discrete set, 100 points were randomly sampled along with the number of orientations that were kinematically feasible at each point and reachable with a linear trajectory from the home position. Each unique position/orientation pair was independently tested, releasing and re-grasping the object from the home position between each trial, so as to avoid biasing the results with path-dependent slipping effects. The ArUco-tagged sphere object was used for this experiment, and the camera logged the actual position and orientation of the object at the end of each motion.

The slip metric and the translational and rotational errors (the distance and angular difference between the actual and commanded poses, respectively) for each translation point were

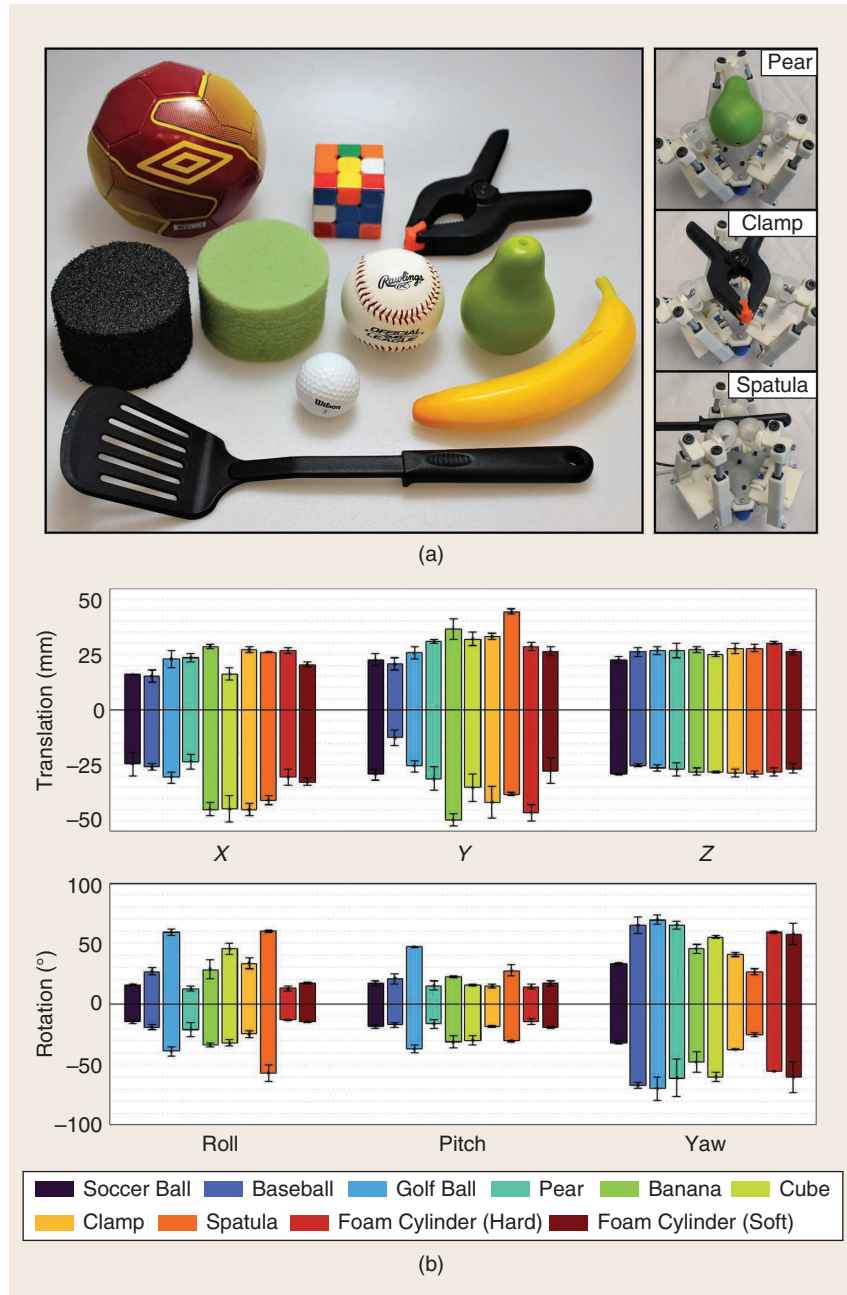


Figure 9. (a) The 10 objects used to validate the real-world performance of the optimized hand. Eight of them are picked from the YCB Object and Model set, and two are identically sized foam cylinders (The black cylinder is rigid and the green cylinder is soft). Sample grasps are also shown for three of the tested objects. (b) The mean range of motion is plotted for the six axes by recording the actual pose of the ArUco marker on each object through three trials along each axis.

averaged across the valid orientations at that point. These averaged metrics appear in Figure 11 at their commanded locations in the workspace. As seen in the characterization experiments, the limited xy manipulation range leads to high translation error at points that are farther from the z -axis. Subsequently, a higher slip metric and rotation error are expected at these points because the contact triangle deforms more for higher x and y values and ineffectually transmits contact forces to the object. In comparison, for points closer to the z -axis, the orientation capability is superior, and the slip metric and translation error diminish. Thus, the most effective manipulation strategy with this hand might be to orient objects immediately after a grasp so as to achieve smaller rotation errors and then carry out any xy translations since this motion is relatively more likely to drop an object and introduce errors.

Teleoperation Experiments on WAM Robot Arm

In addition to commanding the hand open-loop, it is desirable to incorporate real-time feedback to improve control. Since it was not the goal of this work to develop a complex control scheme for the hand but rather to present the mechanical design, teleoperation was used as a proxy for an eventual closed-loop controller, incorporating feedback from a user who could observe the grasped object's actual position and orientation. To provide an intuitive interface, a 3Dconnexion SpaceNavigator 3D mouse was used to control the Stewart Hand mounted on a Barrett WAM robot arm. The mouse outputs a 6-DoF signal that was directly used to command the desired pose of the grasped object. A shape block sorting toy (as detailed in the following) was used to emulate a potential peg-in-hole insertion task—a common class of manipulation problems for industrial assembly tasks.

The cube box was placed at a tilted angle next to the shape block to be inserted. The hand—mounted on the WAM arm—started at a position facing vertically down above the block. Once the hand grasped the block, a preset arm trajectory was executed that roughly positioned the object above the appropriate hole on the cube box. This preset trajectory did not change the orientation of the hand and was determined by visually aligning the palm of the Stewart Hand (while it was not grasping any object) with the center of the target shape hole. Once the preset WAM trajectory was carried out and the grasped block was roughly above the target hole, the teleoperator used the 3D mouse to orient and translate the block until it was partially inserted into the hole. Throughout this step, the teleoperator

could observe the block to obtain visual feedback about its actual pose. Once the shape block was partially inserted, the hand was commanded to release the block, dropping it into the hole. Since the shape blocks were untethered and lightweight, they tended to arbitrarily translate/rotate slightly during the grasping action. It was the responsibility of the operator to counteract this misalignment during the manipulation stage, employing the 6-DoF dexterity of the hand for both reorientation and insertion.

This experiment was carried out with four different shapes (Figure 12), further demonstrating the hand's ability to adapt to irregular object geometries through the differential mechanism. For each of the shapes, the cube box was placed at a different tilt angle to add complexity to the task, on top of the misalignment introduced at the grasping step, as previously described. The holes in the cube box had only a small amount of clearance relative to the corresponding shape blocks. Thus, the task required the precise alignment of the shape block and the hole before insertion was possible. All four shapes were successfully inserted into their respective holes. For more complex object shapes, such as the blue star, the teleoperator iteratively tested the block's alignment by attempting to insert it in the hole, reorienting, and then trying again. While the operator relied on visual feedback during this process, future work might seek to implement this strategy using the fingertip slip metric to ascertain whether the object is coming into contact with the environment, based on how the grasping triangle changes shape as the fingers slide along the object surface. This experiment demonstrated that the performance of the Stewart Hand can be further augmented by closing the feedback loop, as exemplified by simple and effective human teleoperation for a peg insertion task with a variety of shape blocks.

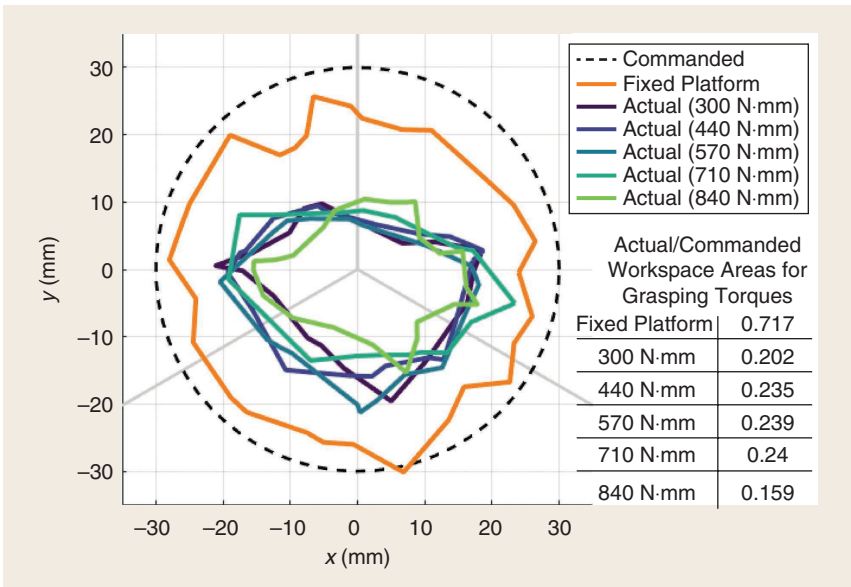


Figure 10. The envelopes for the actual xy ranges of motion at five input grasping torque values, including the envelope of a fixed platform replacing the fingertips. The actual workspace areas normalized by the commanded area of these manipulation envelopes are shown for each of the grasping motor torques and the fixed platform.

This motivates future work to develop automated controllers for this class of manipulator.

Discussion

The experiments on the Stewart Hand to evaluate its manipulation capabilities were all conducted open-loop: the fingers were commanded to close around the object, the potentiometers were used to take a one-time measurement of the object shape, and then the linear actuators were commanded to translate and rotate the object to a desired pose, without any further input regarding the actual pose of the object as it moved or the position of the fingertips. As a result, there was no feedback mechanism to correct for translational and

rotational errors or to adjust for slipping at the contacts. This open-loop control enabled the evaluation of the hand's performance in isolation and highlighted strengths and shortcomings decoupled from those of the control architecture. Despite the absence of true object pose information, the hand performs at par with an equivalent Stewart–Gough platform for rotations and z translations. This is significant because it validates the use of a parallel platform architecture as a robot hand for high-DoF manipulation applications even though a robot hand requires contact friction stability, unlike parallel platforms.

The hand prototype tested in this article was built with optimized design parameters (adjusted for physical

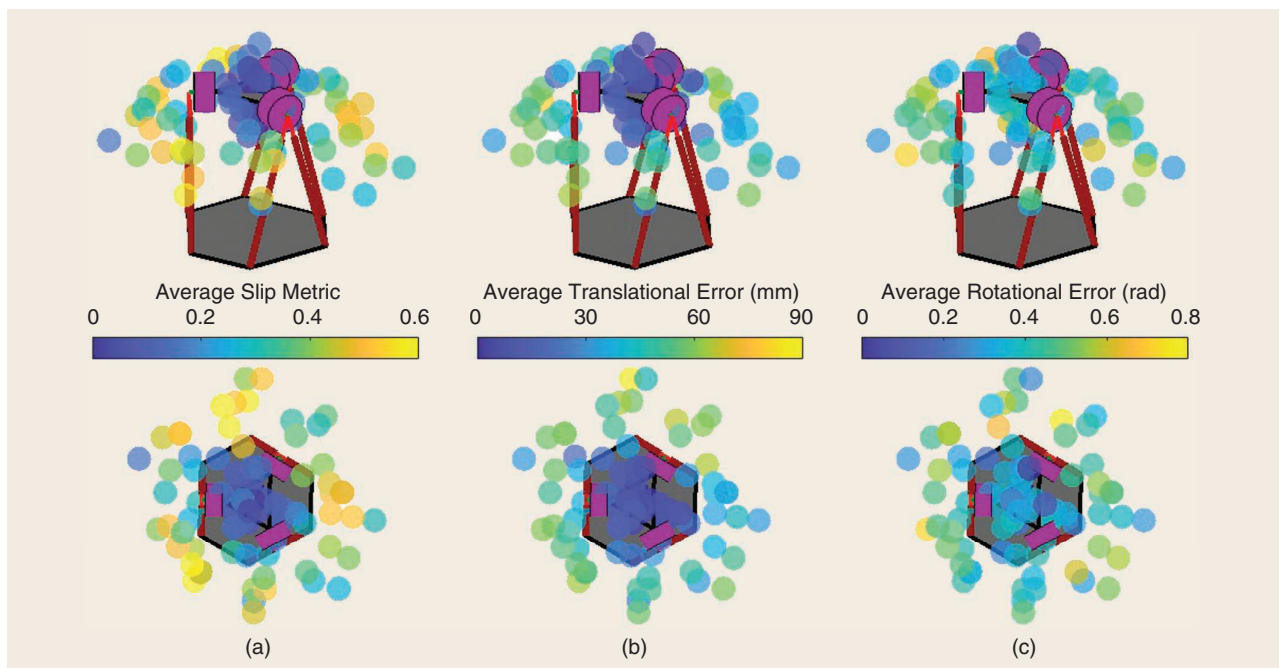


Figure 11. The average of the (a) slip metric, (b) translational error, and (c) rotational error across the orientations attempted at each point in hand's workspace. All data points are plotted at the commanded target locations.

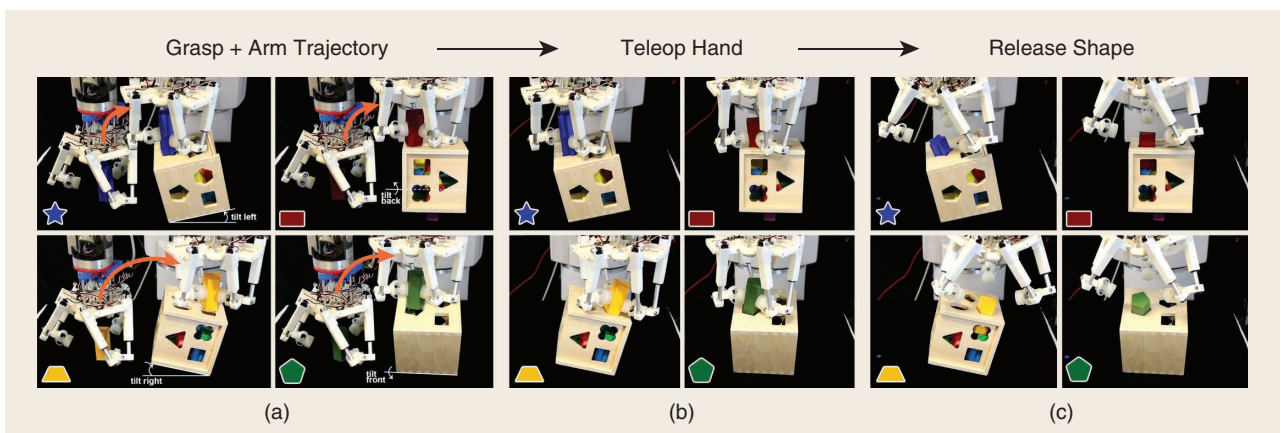


Figure 12. Teleoperating the Stewart Hand to perform the block insertion task with four shapes (a blue star, red rectangle, yellow trapezoid, and green pentagon) on a cube box tilted at different angles. (a) The hand grasps the shape block off the table, and the WAM arm executes a preset trajectory to roughly position the block above the hole. (b) The teleoperator then uses a 3D mouse to translate and orient the shape until the object is partially inserted into the hole. (c) The hand is commanded to release the shape, allowing it to drop through the hole and into the cube box.

implementation constraints) to maximize the dexterous workspace, considering kinematic feasibility, friction stability, and singularities. This optimization predicted a large range of motion in all six DoF, including the xy translation space, where the physical prototype was most challenged. The primary issue with the design optimization was that it assessed static stability at each individual pose but did not consider transitioning from one to another. So, while the grasped object might be kinematically and frictionally stable at poses A and B individually, the optimization does not account for the forces required to push the object between A and B. That is, the grasp would not be able to react the necessary forces required to reconfigure the differential between A and B even if the hand configurations are statically stable at the two poses.

This gap in optimization is more pronounced in the xy workspace because the differential is required to reconfigure the most while executing xy motions. For z translations and pure rotations, the angle between the palm and fingers remains relatively constant, or the change in angle is similar across all fingers, and thus the differential does not need to reconfigure significantly, meaning that the hand can perform manipulations in these DoF. One design choice for applying the necessary manipulation forces might be to actuate each finger with a separate grasp actuator instead of a differential as implemented in this prototype. Unless these actuators are purely torque controlled, such a design would surrender the benefits of underactuated grasping and necessitate a priori knowledge of the object geometry on top of adding to the overall device complexity and cost.

The optimized hand's ranges of motion on YCB and foam objects with the first design iteration in [13] for both the translation and rotation axes are comparable on similarly sized objects. However, the optimized hand is 42% as large as the first design. Normalized by these dimensions, such as the palm radius and the linear actuator stroke lengths, the design parameter optimization yields a significant improvement in the manipulation ranges of motion relative to size. This also highlights the potential scalability of the device. Depending on the application in which the Stewart Hand is deployed, the hand dimensions can be proportioned to grasp and manipulate larger or smaller objects.

The multi-DoF workspace exploration experiment extended the single-axis characterization analysis to motions more akin to real-world tasks. The dome-shaped, kinematically feasible workspace of an equivalent Stewart–Gough platform was simulated to sample points for the exploration. Kinematic stability was chosen over the stricter frictional stability requirements discussed in the “Theoretical Workspace Modeling” section for sampling the workspace to further evaluate the stability calculations employed in the design optimization. Comparing the error and slip metrics of the explored points within the predicted statically stable workspace against those outside it, we can observe that the points deemed frictionally unstable indeed have higher error and slip values. The nonconvex kinematic

workspace was sampled for points approachable only by a linear trajectory from the hand's home position. As a result, a portion of the hand's workspace was not explored—particularly points with significantly negative z coordinates and high absolute xy values—without a more complex path-planning algorithm. Most of these points lie outside the predicted frictionally stable workspace, and, by extension, we can expect that these unexplored points would also have high error and slip metrics. Future work will explore methods for generating nonlinear trajectories with the hand for manipulation sequences to travel through only statically feasible points and more finely explore the entirety of the hand's statically stable workspace.

The teleoperated experiment with the WAM robot demonstrated that the hand is capable of performing dexterous manipulation of unknown objects in a real-world application. The test was human-operated and thus closed-loop, validating the promise of improved hand performance with a closed-loop controller that compensates for slippage and errors between desired and actual poses. Future work on this hand may seek to develop an autonomous controller to accomplish such tasks in addition to continuously computing changes in the contact triangle to monitor slippage.

On the modeling front, in the future, an improved optimization might consider grasp transitions, not just static stability at poses. A physical prototype with individual grasping actuators at each finger could also be developed and characterized for benchmarking the current prototype. This would help to further accredit the benefits of implementing parallel mechanism architectures as real-world robot hands.

Another possible technique to build off this work for industrial and service applications would be to incorporate tactile feedback sensors on the fingertips for grasp and slip detection. This could aid in making the entire manipulation process more autonomous. Such a hand might eliminate the need for expensive high-DoF robot arms by moving the DoF from the arm to the hand. For instance, in manufacturing pick-and-place operations requiring the multi-DoF manipulation of components, a 6-DoF Stewart Hand mounted on a simple SCARA arm could execute more dexterous tasks than a six-axis industrial arm with a parallel-jaw gripper.

Conclusions

In this article, we presented an optimized design for a 6-DoF robotic hand capable of dexterous within-hand manipulation that was based on the Stewart–Gough platform parallel mechanism. The hand design consisted of six linear actuators on three parallel-linkage fingers used for controlling manipulation, while an additional revolute actuator drove the grasping motion of the fingers through an underactuated differential. A computational model facilitated the optimization of the design parameters to maximize workspace size and manipulability. An experimental prototype of the optimal hand design was fabricated based on the optimized parameters, and the design was made available through the Yale OpenHand Project.

This prototype was also characterized through physical experiments that evaluated its single-axis accuracy as well as multi-DoF motions in real-world settings. The hand's actual frictionally stable workspace, where slip metrics and errors were low, closely matched the theoretical computational predictions from the design optimization. Within this workspace, high manipulability and accuracy were observed with an open-loop controller, and a teleoperation task with variously shaped blocks showed real-world performance capabilities of the hand under closed-loop control. The approach of adapting parallel mechanism architectures into dexterous robot hands is validated and warrants further exploration.

Acknowledgment

This work was supported, in part, by the U.S. National Science Foundation, under grants 1928448 and 1900681. Connor McCann and Vatsal Patel contributed equally to this article.

References

- [1] C. Piazza, G. Grioli, M. G. Catalano, and A. Bicchi, "A century of robotic hands," *Annu. Rev. Control. Robot. Auton. Syst.*, vol. 2, no. 1, pp. 1–32, May 2019. doi: 10.1146/annurev-control-060117-105003.
- [2] R. R. Ma and A. M. Dollar, "On dexterity and dexterous manipulation," in *Proc. Int. Conf. Adv. Robot. (ICAR)*, 2011, pp. 1–7. doi: 10.1109/ICAR.2011.6088576.
- [3] E. Nathalia Gama Melo, O. Fernando Avilés Sánchez, and D. Amaya Hurtado, "Anthropomorphic robotic hands: A review," *Rev. Científica Ing. y Desarro*, vol. 32, no. 2, pp. 279–313, 2014. doi: 10.14482/inde.32.2.4715.
- [4] M. Grebenstein et al., "The DLR hand arm system," in *Proc. Int. Conf. Robot. Automat. (ICRA)*, 2011, pp. 3175–3182. doi: 10.1109/ICRA.2011.5980371.
- [5] R. Ma and A. Dollar, "Yale openhand project: Optimizing open-source hand designs for ease of fabrication and adoption," *IEEE Robot. Autom. Mag.*, vol. 24, no. 1, pp. 32–40, Mar. 2017. doi: 10.1109/MRA.2016.2639034.
- [6] T. Laliberté and C. Gosselin, "Actuation system for highly underactuated gripping mechanism," U.S. Patent 6 505 870 B1, 2000.
- [7] T. Tanikawa, T. Arai, and T. Masuda, "Development of micro manipulation system with two-finger micro hand," in *Proc. Int. Conf. Intell. Robots Syst. (IROS)*, 1999, vol. 2, pp. 850–855. doi: 10.1109/IROS.1996.571063.
- [8] M. Nefzi, M. Riedel, and B. Corves, "Development and design of a multi-fingered gripper for dexterous manipulation," *IFAC Proc. Vol.*, vol. 39, no. 16, pp. 133–138, Jan. 2006. doi: 10.3182/20060912-3-DE-2911.00026.
- [9] X. Jin, Y. Fang, D. Zhang, and H. Zhang, "Synthesis of 3-[P][S] parallel mechanism-inspired multimode dexterous hands with parallel finger structure," *J. Mech. Des. Trans. ASME*, vol. 142, no. 8, p. 083301, Aug. 2020. doi: 10.1115/1.4045506.
- [10] S. Yuan, A. D. Epps, J. B. Nowak, and J. K. Salisbury, "Design of a roller-based dexterous hand for object grasping and within-hand manipulation," in *Proc. IEEE Int. Conf. Robot. Automat.*, 2020, pp. 8870–8876. doi: 10.1109/ICRA40945.2020.9197146.
- [11] L. Cui and J. S. Dai, "Posture, workspace, and manipulability of the metamorphic multifingered hand with an articulated palm," *J. Mech. Robot.*, vol. 3, no. 2, p. 021001, 2011. doi: 10.1115/1.4003414.
- [12] J. P. Merlet, *Parallel Mechanisms*, 2nd ed. New York: Springer-Verlag, 2006.
- [13] C. M. McCann and A. M. Dollar, "Design of a Stewart platform-inspired dexterous hand for 6-DOF within-hand manipulation," in *Proc. Int. Conf. Intell. Robots Syst. (IROS)*, 2017, pp. 1158–1163. doi: 10.1109/IROS.2017.8202287.
- [14] A. Bicchi, "Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity," *IEEE Trans. Robot. Autom.*, vol. 16, no. 6, pp. 652–662, 2000. doi: 10.1109/70.897777.
- [15] Z. Xu and E. Todorov, "Design of a highly biomimetic anthropomorphic robotic hand towards artificial limb regeneration," in *Proc. Int. Conf. Robot. Automat. (ICRA)*, 2016, pp. 3485–3492. doi: 10.1109/ICRA.2016.7487528.
- [16] J. Borràs and A. M. Dollar, "Analyzing dexterous hands using a parallel robots framework," *Auton. Robots*, vol. 36, nos. 1–2, pp. 169–180, 2014. doi: 10.1007/s10514-013-9380-x.
- [17] L. Birglen, T. Laliberté, and C. Gosselin, *Underactuated Robotic Hands*, vol. 40. Berlin: Springer-Verlag, 2008.
- [18] C. M. McCann and A. M. Dollar, "Analysis and dimensional synthesis of a robotic hand based on the Stewart-Gough platform," in *Proc. Int. Design Eng. Tech. Conf. (IDETC)*, 2018, p. 8. doi: 10.1115/DETC2018-86149.
- [19] R. M. Murray, Z. Li, and S. S. Sastry, *A Mathematical Introduction to Robotic Manipulation*. Boca Raton, FL: CRC Press, 1994.
- [20] J. Wang, C. Wu, and X.-J. Liu, "Performance evaluation of parallel manipulators: Motion/force transmissibility and its index," *Mech. Mach. Theory*, vol. 45, no. 10, pp. 1462–1476, Oct. 2010. doi: 10.1016/j.mechmachtheory.2010.05.001.
- [21] F. J. Romero-Ramirez, R. Muñoz-Salinas, and R. Medina-Carnicer, "Speeded up detection of squared fiducial markers," *Image Vis. Comput.*, vol. 76, pp. 38–47, Aug. 2018. doi: 10.1016/j.imavis.2018.05.004.
- [22] S. Garrido-Jurado, R. Muñoz-Salinas, F. J. Madrid-Cuevas, and R. Medina-Carnicer, "Generation of fiducial marker dictionaries using mixed integer linear programming," *Pattern Recognit.*, vol. 51, pp. 481–491, Mar. 2016. doi: 10.1016/j.patcog.2015.09.023.
- [23] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar, "Benchmarking in manipulation research: Using the Yale-CMU-Berkeley object and model set," *IEEE Robot. Autom. Mag.*, vol. 22, no. 3, pp. 36–52, Sept. 2015. doi: 10.1109/MRA.2015.2448951.

Connor M. McCann, School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, 02138, USA. Email: cmccann@g.harvard.edu.

Vatsal V. Patel, Department of Mechanical Engineering and Materials Science, Yale University, New Haven, Connecticut, 06511, USA. Email: v.patel@yale.edu.

Aaron M. Dollar, Department of Mechanical Engineering and Materials Science, Yale University, New Haven, Connecticut, 06511, USA. Email: aaron.dollar@yale.edu.

