

Virtual Grasp Release Method and Evaluation

Mores Prachyabrued^{a,*} and Christoph W. Borst^a

^aThe Center for Advanced Computer Studies, University of Louisiana at Lafayette, PO Box 44330, Lafayette, LA 70504, USA

*Corresponding Author: Tel +1 337 9624932

E-mail Addresses: mores_p@hotmail.com (M. Prachyabrued), cwborst@gmail.com (C. W. Borst)

Abstract—We address a “sticking object” problem for the release of whole-hand virtual grasps. The problem occurs when grasping techniques require fingers to be moved outside an object’s boundaries after a user’s (real) fingers interpenetrate virtual objects due to a lack of physical motion constraints. This may be especially distracting for grasp techniques that introduce mismatches between tracked and visual hand configurations to visually prevent interpenetration. Our method includes heuristic analysis of finger motion and a transient incremental motion metaphor to manage a virtual hand during grasp release. We integrate the method into a spring model for whole-hand virtual grasping to maintain the physically-based pickup and manipulation behavior of such models. We show that the new spring model improves release speed and accuracy based on pick-and-drop, targeted ball-drop, and cube-alignment experiments. In contrast to a standard spring-based grasping method, measured release quality does not depend notably on object size. Users subjectively prefer the new approach and it can be tuned to avoid potential side effects such as increased drops or visual distractions. We further investigated a convergence speed parameter to find the subjectively good range and to better understand tradeoffs in subjective artifacts on the continuum between pure incremental motion and rubber-band-like convergence behavior.

Keywords—Interaction techniques; virtual reality; virtual grasping; grasp release

1 INTRODUCTION

WE present a method for improved whole-hand virtual grasping, particularly for the release of grasps. Whole-hand virtual grasping is important for applications that benefit from realistic hand-object interactions. For example, Moehring and Froehlich (2011) showed that users preferred whole-hand interaction over conventional controller-based interaction for functionality assessment in a virtual car interior, since abstract character of the conventional interaction led to a loss of realism and impaired users’ judgement. Good grasping-based interfaces may also have a low learning curve if users can interact with virtual environments naturally.

A “sticking object” grasp release problem occurs when a user’s fingers (real, not rendered) can sink into a virtual object, and the effect may be especially unpleasant when there is a mismatch between tracked and visual hand configurations. For example, there is a mismatch in the spring-based grasping model of Borst and Indugula (2006) to prevent visual interpenetration artifacts. Without physical constraints from a real object, users tend to close their (real) fingers into virtual objects. Since the visual model no longer matches the real hand, an object can appear to stick to the hand (exaggerated finger motions are needed to release the object), and a user can not know precisely when a grasp will release. This led Borst and Indugula to suggest a “light touch” with their approach, and its performance hinges on practice for some users. The problem may be reduced by force feedback, considering such feedback has been shown to reduce hand closing (Fabiani et al., 1996). However, it is also important to support grasping in environments without force feedback, for example, in systems where the

hand is optically tracked and worn or complex devices are not desired. In such environments, additional visual and audio feedback may be useful, to some extent, to reduce hand closing (Fabiani et al., 1996).

A recent study (Prachyabrued and Borst, 2012) showed that preventing hand-object interpenetration is subjectively important for the spring-based grasping approach. However, the prevention increased sticking by increasing real hand closure. Users expected fingers to lift immediately from an object with small release movement. A grasp release method that can match these user expectations while preventing interpenetration would address the tradeoffs.

We propose such a release mechanism in a new spring model based on Borst and Indugula’s (2006) spring model (original). Our model addresses the sticking object problem while retaining the characteristics of physically-based grasping, the prevention of visual interpenetration artifacts, and compatibility with the force rendering method of the original approach. The original spring model couples a simulation-controlled articulated hand model (called the virtual hand or spring hand) to tracked (real) hand configuration using a system of linear and torsional virtual spring-dampers. This resembles the rubber band metaphor (Zachmann and Rettig, 2001) to manage a virtual hand during release of grasps. Instead, our enhanced spring model adds a heuristic analysis of finger motion to detect a user’s intent to release the grasped object, and it uses a transient incremental motion metaphor to manage a virtual hand during a release period.

The contributions described in this paper are:

- We present a spring model for whole-hand virtual grasping that includes a method for improved release.
- We present heuristic analysis of finger motions to detect a user's intent to release the grasped object.
- We present experimental evaluation of our method. It shows that our method improves speed, accuracy, and subjective experience during grasp release, without extra accidental drops or substantial visual problems.
- Our experiments also demonstrate that the sticking problem increases with increasing object size, as release performance of a standard (original) grasping approach decreases notably with increasing object size (our new approach mitigates this).
- Finally, we provide experimental investigation of subjective artifacts related to a convergence motion that follows release. Results provide guidelines for subjectively-optimal convergence speed.

Initial results were presented in a previous paper (Prachyabrued and Borst, 2011), which described targeted ball-drop and subjective comparison experiments. We now present a more complete study, including pick-and-drop and cube-alignment experiments. This generalizes results with more grasp release conditions (object type, release precision requirement, object rotation requirement, and gravity) that may affect release motions. Additionally, we include follow-up studies on possible limitations and optimization of the release mechanisms.

2 PREVIOUS WORK

2.1 Physically-Based Grasping

Physically-based grasping models, such as the one we build on, aim to provide realistic interaction by simulating object motion according to laws of physics. Bergamasco et al. (1994) introduced the use of physically-based object response to achieve whole-hand interaction. They defined grids of control points on a virtual hand to detect contacts between a virtual hand and a virtual object and to compute force vectors acting on the object, including normal contact forces, dynamic frictions, and static frictions. Manipulation was limited to objects with simple shapes.

Using a similar idea, Hirota and Hirose (2003) demonstrated dexterous manipulation of objects with complex shapes in a manipulation system. They used a much larger number of points and a fast collision response computation method.

Borst and Indugula (2005, 2006) extended the concept of virtual coupling to the whole hand. A virtual hand model was coupled to the tracked hand using a system of linear and torsional spring-dampers. These created forces necessary to simulate physically-based grasping using a widely-available simulation tool. Their technique prevented hand-object interpenetration not accounted for in the two previous works.

Jacobs and Froehlich (2011) used a soft body in each finger phalanx to more accurately model contact areas and contact forces for improved finger-based interaction. They used rigid links, instead of virtual spring-dampers, for virtual-tracked hand coupling to avoid spring parameter tuning. They also suggest that the rigid coupling allows faster virtual hand interaction. However, rigid coupling may put more constraints on physics simulation and may cause problems during large hand-object interpenetration.

Allard et al. (2007) used images of a real-world object, captured from different viewpoints, to construct a 3D model representation and inject it into a physically-simulated virtual environment in real-time. This made it possible to rapidly capture approximate hand geometry for coarse hand-object interaction. The captured hand was not very detailed and did not explicitly represent joints.

Wilson et al. (2008) presented physically-based grasping on an interactive surface. They modeled surface contacts as rigid bodies that interacted with virtual objects using physical simulation. This was a limited form of whole-hand grasping near a surface, not a general approach for 3D space.

Microsoft's Holodesk (Hilliges et al., 2012) allows hand interaction with virtual objects in a reach-in augmented reality environment. A user's hand (or another real-world object) is represented with many small sphere particles, with each particle coupled to its tracked position using a spring-damper. Collision response with these spheres provides virtual object response. Grasping in this system is limited by an optical line-of-sight problem, and the real hand is seen to penetrate virtual objects.

2.2 Heuristics-Based Grasping

Heuristics-based grasping refers to grasping approaches that use heuristics to determine grasp state and object motion during grasp. We studied these approaches for our heuristic analysis of finger motions. Purely heuristic approaches are not as general as physically-based grasping, but they may perform well for their intended tasks.

Iwata (1990) tested 16 control points on a virtual hand for contact with a virtual object. The object was grasped when it was touched by the thumb and one of the other fingers. A grasped object's coordinate frame was then attached to the hand coordinate frame so that the object moved with the hand. A similar idea using two fingers was presented by Maekawa and Hollerbach (1998).

In their virtual assembly environment, Wan et al. (2004) abstracted mechanical components into simple primitives (cube, sphere, and cylinder). Possible grasping postures were predefined for each pair of primitive type and size. An object was grasped if collision detection indicated that user hand posture matched one of the previously defined grasping patterns for the object. The object was manipulated by considering its coordinate frame as a child node of hand.

Hilliges et al. (2009) allowed pick-up on an interactive surface by detecting a pinch gesture. An object would be under grasp control (with limited rotation) if a ray,

projected downward from the center of mass of a hole formed by the gesture, intersected the object.

Ullmann and Sauer (2000) presented heuristics, based on contact geometry, for establishing one-hand and two-hand grasps. They presented a fine object manipulation method for computing object motion (not just attaching an object's frame to the hand frame) after grasp had been established.

Holz et al. (2008) and Moehring and Froehlich (2010) presented grasping heuristics and object manipulation methods that are more general. They supported multi-user, multi-hand, multi-finger, and multi-object interactions. They both used the concepts of grasping pairs and friction cones.

Pinch (grasping) detection for a tiny virtual object may be difficult due to imperfect finger tracking. Moehring and Froehlich (2011) modified finger tracking hardware to use conductive stripes of metal at each fingertip for improved pinch detection and improved grasp detection heuristic. They consider an object to be grasped if pinch is detected by this hardware and one of the involved virtual fingers touches the object.

While many heuristic approaches used violation of a grasp condition to determine release state, Moehring and Froehlich (2010) presented explicit release heuristics based on distances of involved grasping pairs. In contrast, our release heuristics consider finger motions.

Osawa (2006) previously considered heuristic release detection to help correct release problems, focusing on release precision problems that result from hand movement. Heuristic analysis detected the release instant and a search backward in time found an adjusted release position (original desired position). In contrast, our work integrates readily with a physically-based grasping model and avoids backtracking that produces discrete jumps in object pose. We show it improves release speed and orientation accuracy in addition to position accuracy.

2.3 Virtual Hand Management

A virtual hand that simply follows a tracked hand configuration typically penetrates virtual objects during interaction (due to lack of motion constraints). There are techniques that prevent the visual interpenetration artifacts with resulting discrepancy between virtual and real hands (which complicates the release of grasps, as pointed out in the introduction). Work by Burns et al. (2006), suggesting that users are more sensitive to visual interpenetration than to visual-proprioceptive discrepancy, motivates the prevention of visual interpenetration. Zachmann and Rettig (2001) discussed two metaphors that can be used to manage a virtual hand after the virtual and real hands separate:

1. The rubber band metaphor: the virtual hand maintains its configuration as close as possible to the real hand.
2. The incremental motion metaphor: the virtual hand moves by the same amount as the real hand.

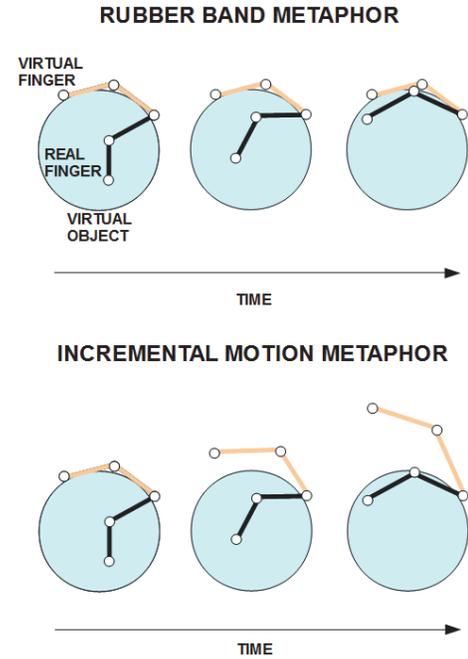


Fig. 1. Hand releasing a virtual object. A rubber band metaphor (top) causes the virtual hand to wait at the object surface, exacerbating sticking. An incremental motion metaphor (bottom) causes the virtual hand to be released from the object more immediately but maintains an offset that can cause grasp problems (Section 8.4.3).

Each metaphor has a drawback. The rubber band metaphor causes the virtual hand to stick to a virtual object's surface upon release (Burns et al., 2006) (Fig. 1 (top)). The phenomenon was similarly observed in other systems using this type of metaphor (Borst and Indugula, 2006; Lindeman et al., 2001). The incremental motion metaphor does not have the sticking problem, but it maintains an offset between the virtual and real hands (Fig. 1 (bottom)). It was reported by Burns et al. (2006) that maintaining an offset between virtual and real hands reduced user performance.

Burns et al. (2007) proposed a third metaphor - MACBETH (Management of Avatar Conflict By Employment of a Technique Hybrid). It involves incremental motion, but it removes position discrepancy by introducing velocity discrepancy that is similarly detectable. Based on their user study comparing MACBETH to the previous metaphors, MACBETH improved user-rated naturalness and user preference while no loss in user performance was detected. However, MACBETH, in its current form, only manages virtual hand base position. Additional work is needed to manage hand orientation and finger joint angles. In contrast, rubber-band and incremental motion metaphors are applicable to both.

A simpler technique to reduce offset between virtual and real hands was used in Immersion's VirtualHand Toolkit (DesRosiers et al., 2001). An offset is gradually reduced to zero when the real hand no longer contacts a virtual object. Full details are not available, but it appears this offset reduction was not designed to improve grasp release, rather just to transition the virtual hand back to the tracked configuration.

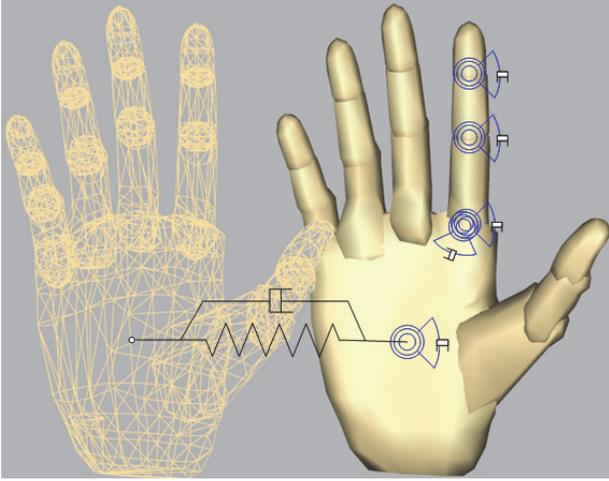


Fig. 2. Borst and Indugula (2006)'s spring model showing tracked hand (left), virtual hand (right), and some of the virtual spring-dampers.

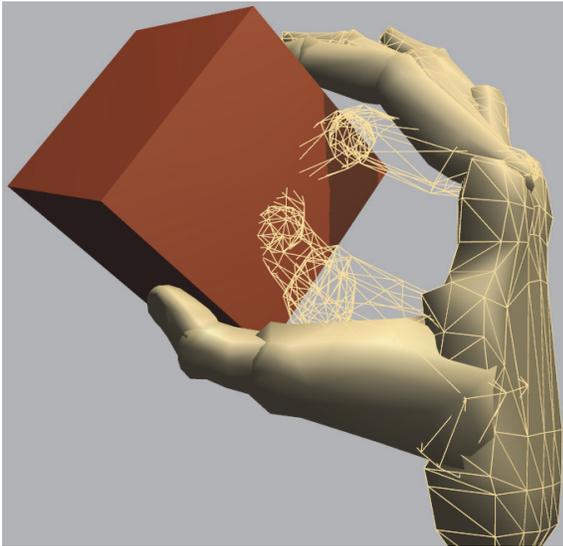


Fig. 3. Grasp showing tracked hand (mesh) that sank into the virtual object and virtual hand (solid) that remained at the object's surface.

We present an approach to virtual whole-hand management that includes the use of incremental motion with offset reduction to manage virtual finger joint angles during grasp release.

3 BORST AND INDUGULA'S SPRING MODEL

3.1 Description of the Spring Model

Borst and Indugula (2006) proposed a physically-based grasping approach that extended the virtual coupling concept to an articulated hand. The approach couples a spring (virtual) hand to a tracked (real) hand using a system of virtual linear and torsional spring-dampers. This produced forces and torques necessary for virtual hand motion using dynamic simulation and for physically-based response of grasped objects via collision response.

Fig. 2 illustrates their spring model. They used 21 torsional and 6 linear virtual spring-dampers. There was one torsional element for each of 20 finger joint degrees of freedom (illustrated only for the index finger), one torsional element and one linear element for the base of the hand (illustrated), and one linear element for each of the five digit tips (not illustrated).

In addition to supporting grasping and manipulation, this spring model addressed the problem of visual interpenetration and included force rendering for force-feedback gloves.

3.2 Grasp Release Problem and the Spring Model

The spring model can be considered a rubber band metaphor to manage a virtual hand: the virtual hand maintains a configuration (palm pose and finger joint angles) pulled toward the tracked hand configuration but subject to constraints. This can cause the virtual hand to stick to a virtual object upon grasp release, as mentioned by Burns et al. (2006) and indicated as a motivation for using a "light touch" by Borst and Indugula (2006). Fig. 3 illustrates the problem. A user closed the fingers further than necessary, and, when the user opens them to release, they may remain inside the object, causing the object to appear stuck (or the hand model to appear unresponsive). The user can exaggerate finger motions to release, but this reduces naturalness and interferes with precision tasks (our experiment will suggest reduced accuracy).

Notably, the problem also occurs to an extent even if the visual hand model is allowed to penetrate objects to match tracked hand configuration. The real fingers still sink into objects due to lack of real motion constraints and small motions may not be sufficient to release grasp. Our visual interpenetration study (Prachyabrued and Borst, 2012) showed that there was slightly less interpenetration and better release performance in this case, but users nonetheless disliked visual interpenetration and believed it increased their hand closure.

4 GRASP RELEASE METHOD AND NEW SPRING MODEL

The two key ideas in our method for improving grasp release are:

1. Heuristic analysis of finger motions (release-heuristic function) to detect a user's intent to release grasp.
2. A transient incremental motion metaphor with subsequent convergence period to manage the virtual hand during grasp release.

4.1 New Spring Model

4.1.1 Three Hand Configurations Concept

Our new spring model behaves similarly to that of Borst and Indugula (2006) except during, and for a short time following, grasp release. To incorporate the incremental motion metaphor for release, the new spring model defines three hand configurations:

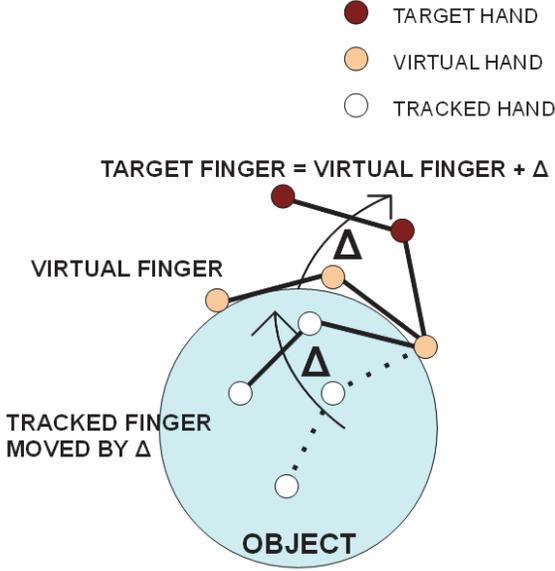


Fig. 4. Target-hand outside the object, causing the virtual hand to open more immediately even when the tracked-hand finger is still inside the object.

1. Tracked hand refers to the real hand configuration as measured by sensing hardware and calibration steps.
2. Spring (virtual, visually-rendered) hand refers to a simulation-controlled virtual hand configuration.
3. Target hand refers to a target configuration for the virtual hand.

The virtual hand is coupled to the target hand (instead of the tracked hand as in (Borst and Indugula, 2006)) using a system of linear and torsional spring-dampers. Fig. 4 illustrates the target hand concept. When a user opens their hand (changing a joint rotation by a “delta” amount) to release a virtual object, we update the target joint configuration to the “current virtual configuration plus delta”. This has the effect of pulling the virtual hand to open by the same delta (resembling the incremental motion metaphor), causing it to release the object more immediately than waiting for fingers of the tracked hand to exit the object’s surface. Subsequently, the target hand is adjusted by a convergence mechanism.

4.1.2 Target Hand Update Algorithm

We update target-hand configuration (palm pose and finger joint angles) for every new tracked-hand configuration. For grasping of unconstrained objects, which is our focus, the grasp release problem comes mostly from finger motions (finger penetrations) and not from palm motions (palm penetration). Therefore, target-hand palm (the base frame for the hand) simply matches tracked-hand palm. For the target-hand finger joint angles, the equations below describe the main update component. We evaluate a release-heuristic function (Section 4.2) prior to the update. For each joint angle in a hand joint model (Section 4.3):

If the release-heuristic function detected release,

$$\theta_{ig1} = \theta_{sp0} + (\theta_{tr1} - \theta_{tr0}) \quad (1)$$

Otherwise,

$$\theta_{ig1} = \theta_{ig0} + (\theta_{tr1} - \theta_{tr0}) \quad (2)$$

where:

θ_{ig1} , θ_{ig0} are next (post-update) and current (pre-update) joint angles of the target hand,
 θ_{tr1} , θ_{tr0} are new and previous joint angles of the tracked hand, and
 θ_{sp0} is the current joint angle of the virtual (spring) hand.

θ_{ig1} is also subject to an additional update mechanism described at the end of this subsection.

Initially, the target and virtual hands are set to the same configuration as the tracked hand. Before release, target-hand finger configuration (finger joint angles) will be equal to tracked-hand finger configuration (they move by the same delta, see (2)). This results in the same virtual-hand behavior (w.r.t. finger motions) as the original spring model. The behavior begins to differ when a release-heuristic function detects release. Target-hand finger configuration will be set to virtual-hand finger configuration plus the change undergone by the tracked-hand fingers, see (1). Later, target-hand finger configuration will be updated using (2) (release-heuristic function no longer detects release). This resembles the incremental motion metaphor to manage virtual-hand fingers. It creates and maintains an offset between the target-hand and tracked-hand finger configurations (also between the virtual and real hands), and the potential exists for the offset to grow with every release of an object.

Maintaining an offset between virtual and real hands reduces user performance (Burns et al., 2006). Therefore, we add a convergence algorithm that gradually adjusts the target-hand finger configuration back to the tracked-hand finger configuration. We define a convergence amount c to be some small angle (see Section 4.3 for example value). At every simulation time step,

1. We compute $\Delta = \theta_{ig1} - \theta_{tr1}$.
2. If $\Delta > c$ then $\theta_{ig1} = \theta_{ig1} - c$. Otherwise,
3. If $\Delta < -c$ then $\theta_{ig1} = \theta_{ig1} + c$. Otherwise,
4. $\theta_{ig1} = \theta_{tr1}$.

This returns the new spring model behavior to the original spring model behavior after some time.

This new spring model preserves the following three important properties of the original spring model:

1. It provides physically-based grasping.
2. It addresses the problem of visual interpenetration.
3. It is compatible with force feedback rendering from (Borst and Indugula, 2006).

4.2 Release-Heuristic Function

Our release-heuristic function analyzes finger motions to detect a user's intent to release a grasped object. The basic idea is to check if the user is releasing the thumb and one of the other fingers from the grasped object. Let:

t = thumb, i = index, m = middle, r = ring, and p = pinky.

$F_t \subseteq \{ \text{thumb joint angles of a hand joint model} \}$ be a set of chosen thumb joint angles used to detect intent to release,

$F_k \subseteq \{ \text{finger joint angles of a hand joint model} \}$ be a set of chosen joint angles of a finger $k \in \{i, m, r, p\}$ used to detect intent to release,

$H(f, cj)$ be a history (L -element cyclic array) of joint angle motions, associated with a chosen joint angle cj of finger $f \in \{t, i, m, r, p\}$,

$th(f, cj)$ be a threshold for joint angle motions, associated with a chosen joint angle cj of finger f ,

$\theta_{tgo}(f, j)$ be a value of a joint angle j of finger f of the current target-hand configuration,

$\theta_{spo}(f, j)$ be a value of a joint angle j of finger f of the current virtual-hand configuration,

$\theta_{tro}(f, j)$ be a value of a joint angle j of finger f of the previous tracked-hand configuration, and

$\theta_{tri}(f, j)$ be a value of a joint angle j of finger f of the new tracked-hand configuration.

We evaluate the release-heuristic function for every new tracked-hand configuration. There are 3 steps:

Step 1: For each finger $f \in \{t, i, m, r, p\}$ and for each chosen joint angle $cj \in F_f$: add $\Delta = \theta_{tri}(f, cj) - \theta_{tro}(f, cj)$ to $H(f, cj)$.

This step adds joint angle motions (Δ) to their corresponding history arrays. Assume, for the remainder of this section, that positive values for Δ indicate opening of the joint angle and negative values indicate closing, then let:

$isOpening(f, cj)$ be a function that returns true if there is at least one element in $H(f, cj)$ that is greater than or equal to the positive threshold $th(f, cj)$ and none of the elements are negative. It returns false otherwise. Basically, this function determines if a chosen joint angle is opening or not by looking at its history. Thresholds are used to prevent false positives.

Step 2: If one of the virtual thumb phalanges contacts the virtual object and there exists a chosen joint angle $cj \in F_t$ for which " $\theta_{tgo}(t, cj) - \theta_{spo}(t, cj) < 0$ " and " $isOpening(t, cj)$ is true", then continue to step 3, otherwise the heuristic function returns false.

This step checks if the user is opening the thumb that is in contact with the virtual object. The condition " $\theta_{tgo}(t, cj) - \theta_{spo}(t, cj) < 0$ " checks if the corresponding joint angle of the virtual hand is active in the current grasp (i.e., it is currently pulled by a torsional spring to grasp the object).

Step 3: If there exists a finger $k \in \{i, m, r, p\}$ with a virtual phalange contacting the virtual object and there is

a chosen joint angle $cj \in F_k$ for which " $\theta_{tgo}(k, cj) - \theta_{spo}(k, cj) < 0$ " and " $isOpening(k, cj)$ is true", then the heuristic function returns true, otherwise it returns false.

This step checks if the user is opening one of the remaining fingers in contact with the virtual object.

So far, the given description is generic. One must specify F_f , L , and $th(f, cj)$ in an implementation (see Section 4.3). The release-heuristic function is customizable, e.g., by setting F_t , F_i , F_m , F_r , F_p , such that the function gives good results for particular grasp types, or by adjusting the thresholds $th(f, cj)$ to account for sensing noise or small unintentional finger movement. Note that the new spring model is independent of the proposed release-heuristic function. We can plug in a different release-heuristic function.

4.3 Implementation Notes

We use a standard hand joint model similar to a CyberGlove joint model (Virtual Technologies Inc., 1994). Each of four fingers has a 2-dof metacarpophalangeal joint (MPJ) for abduction and flexion at the first knuckle and a 1-dof interphalangeal joint (IJ) at each of the remaining two knuckles for flexion (PIJ for the second knuckle and DIJ for the third knuckle). The thumb has a 2-dof trapeziometacarpal joint (TMJ) in the palm for roll and abduction, a 1-dof MPJ for flexion at the first knuckle, and a 1-dof IJ for flexion at the second knuckle.

We use the following values to implement our new spring model and release-heuristic function: $c = 0.035^\circ$, $L = 3$, $F_t = \{\text{TMJ-roll, MPJ-flexion}\}$, $F_i = \{\text{MPJ-flexion, PIJ-flexion}\}$, $F_m = \{\text{MPJ-flexion}\}$, $F_r = \{\text{MPJ-flexion}\}$, $F_p = \{\text{MPJ-flexion}\}$. We set the threshold parameters $th(f, cj)$ to integer multiples of calibrated angular resolutions of finger sensors at the corresponding joint angles (considering calibrated sensor gains). The multipliers for thumb joint angles are 1. The multipliers for the remaining joint angles are 2.

We chose F_t , F_i , F_m , F_r , F_p based on observations of grasp-release motions. We started with the lowest multipliers for the thresholds and increased them to eliminate false positives. We balanced c to produce fast convergence without grasp release difficulty (considered in detail in Section 8). The L value was experimental and needs further investigation: no history ($L=1$) resulted in poorer detection; we did not observe notable effect of increasing L beyond 3 (we tested up to $L=12$).

5 EXPERIMENT

We conducted within-subjects experiments to compare our approach to the standard spring model of Borst and Indugula (2006) (we implement the standard spring model by simply setting target hand to match tracked hand, disabling new mechanisms). The experiments consisted of objective and subjective components. The objective component consisted of a pick-and-drop experiment, a targeted ball-drop experiment, and a cube-alignment experiment, with the following independent variables:

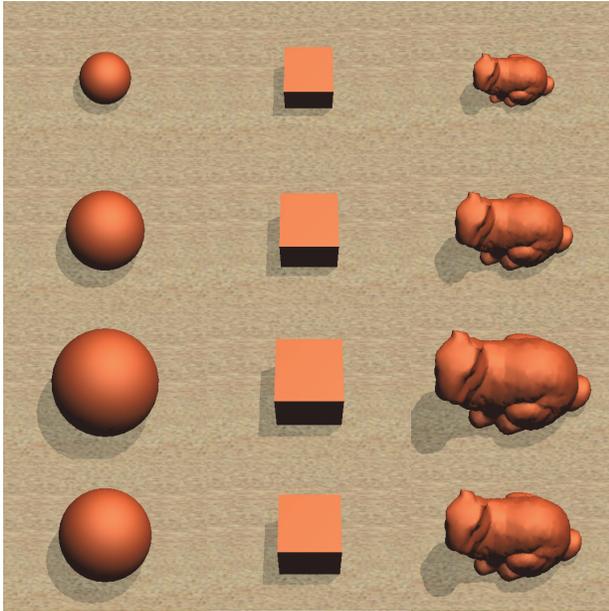


Fig. 5. Object types and object sizes used in the experiments. The top three rows contain small-sized (ball diameter = 6.0 cm, cube size = 5.5 cm), medium-sized (9.0 cm, 6.5 cm), and large-sized objects (12.0 cm, 7.5 cm), used in the objective study, respectively. The fourth row contains objects used in the subjective study (10.5 cm, 7 cm).

1. Grasping Technique - new and old spring models.
2. Object Size - small, medium, and large (see Fig. 5).
3. Object Type (only varied in the pick-and-drop experiment) - ball, cube, and bunny.

The dependent variables (grasp release performance) were:

1. Release Time - amount of time required to release a grasped object.
2. Translation Error (only measured in the targeted ball-drop and the cube-alignment experiments) - translation of an object resulting from grasp release.
3. Rotation Error (only measured in the cube-alignment experiment) - rotation of an object resulting from grasp release.

We included the three experiments to compare the grasp release performance under various grasp release conditions that may affect release motions. This was important to investigate the suitability of the mechanism to different task and grasp types, because, for example, fast and coarse interaction may not benefit from the same mechanism that works with more precise release. The pick-and-drop experiment simply asked subjects to pick an object and drop it into a large pit, requiring only coarse precision for the grasp release. The targeted ball-drop experiment required more precise grasp release by asking subjects to drop a ball at a target position. The cube-alignment experiment also required precise grasp release by asking subjects to align a cube to a floating target cube. However, it included no gravity simulation (which may impact difficulty of release) and thereby resembled a task where a user arranges 3D scene or interface components

using the hand, with objects sticking in place after release. Also, the task required 3D rotation of the cube to align with the target, potentially leading to various hand orientations upon grasp release, compared to the targeted ball-drop experiment that required no rotation alignment. Different object types used in the experiments may also affect grasp types and release motions.

We hypothesized that the new spring model improves speed and accuracy of the grasp release.

The subjective component was a subjective comparison experiment in which a virtual environment contained two objects, using the two different grasp techniques, and users indicated which was easier to release and which was easier to pick up. This allowed us to determine whether or not users could detect quality differences (they were not informed which object used which technique). Object size was not varied for this experiment, but we included three object types: ball, cube, and bunny. The size for each object was a middle size between medium and large sizes from the objective study (see the fourth row of Fig. 5).

We hypothesized that the new technique provides subjectively easier release.

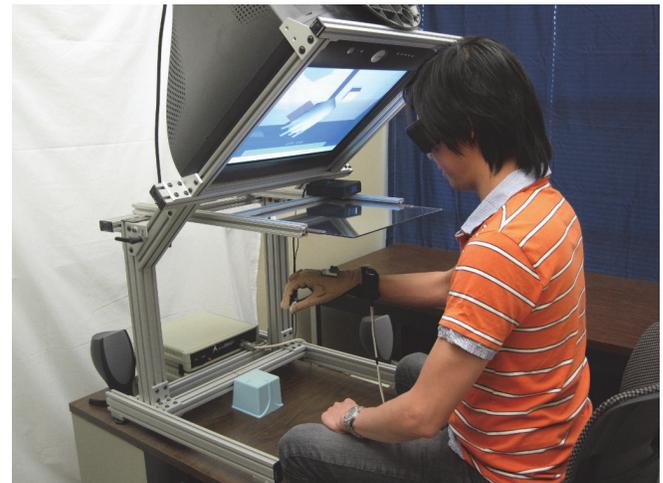


Fig. 6. The grasping system hardware for the experiment.

5.1 Apparatus

Fig. 6 illustrates the grasping system hardware for the experiment. We used a mirror-based “fish tank” VR display (21-inch CRT monitor placed at a 45° angle above a mirror) to co-locate real and virtual workspaces. Monitor resolution was 1024 x 768 and refresh rate was 100 Hz, for time-multiplexed stereoscopic viewing via CrystalEyes LCD shutter glasses. Joint angles were sensed by an 18-sensor right-handed CyberGlove (this glove does not have sensors at distal finger joints, so their angles are computed as two-thirds of the middle knuckle angles). Palm base pose was tracked by an Ascension miniBird 500 system that was synchronized with the monitor refresh to reduce jitter. The head (viewpoint) was

not tracked. Audio output was via ordinary stereo speakers. All software ran on a Dell Precision T5400 with two Intel quad-core Xeon E5450 3.00GHz processors, 8GB RAM, and an NVIDIA QuadroFX 5800 graphics card.

The NVIDIA PhysX SDK (www.nvidia.com) provided physical simulation with collision detection and response. PhysX revolute joints provided torsional springs for finger joint angles. We used equations from (Borst and Indugula, 2006) for the springs at the base of the hand (palm). We omitted the linear fingertip springs from (Borst and Indugula, 2006). Our physical simulation allowed collision shapes to overlap slightly (0.6 cm, set using a parameter in NVIDIA PhysX SDK) for improved contact simulation. To avoid associated visual hand-object interpenetration, hand collision shape was set correspondingly larger than visual hand shape.

Our visual hand model consisted of 16 segments and resembled the model provided with CyberGlove devices. Our OpenGL-based visual rendering system included shadow-mapped shadows. Our application was separated into two main threads: a graphics thread for graphics rendering and an interaction thread for hand data processing and simulation.

5.2 Subjects

28 subjects participated in the experiment: 25 males and 3 females, aged 20 to 33 years (average = 25), 23 right-handed and 5 left-handed. Almost all subjects (27) were students, mostly from computer science and computer engineering programs. Experience levels were mixed: 5 reported previous exposure to virtual grasping (presumably from demos in our lab), 9 others reported exposure to VR systems, and all of the remaining 14 took a graphics class, played video games, or watched 3D movies.

5.3 Design

Considering the four experiments detailed here, subjects performed five total tasks: a learning task, the pick-and-drop experiment, the targeted ball-drop experiment, the cube-alignment experiment, and the subjective-comparison experiment. To reduce possible effects of fatigue and short-term learning, we split experiments into two days, with a different grasping technique presented per subject's day (order randomized per subject such that half of the subjects experienced the new grasping approach on their first day, and half experienced the other approach first). On both days, subjects completed tasks in this order: learning task, pick-and-drop experiment, targeted ball-drop experiment, and cube-alignment experiment. Additionally, subjects completed the subjective-comparison experiment only at the end of the second day, because it involved exposure to both techniques in each of its trials. We calibrated the CyberGlove for each subject before they started per day. Experiment duration was typically 30 to 45 minutes per day.

Within each experiment, there were the following subcomponents:

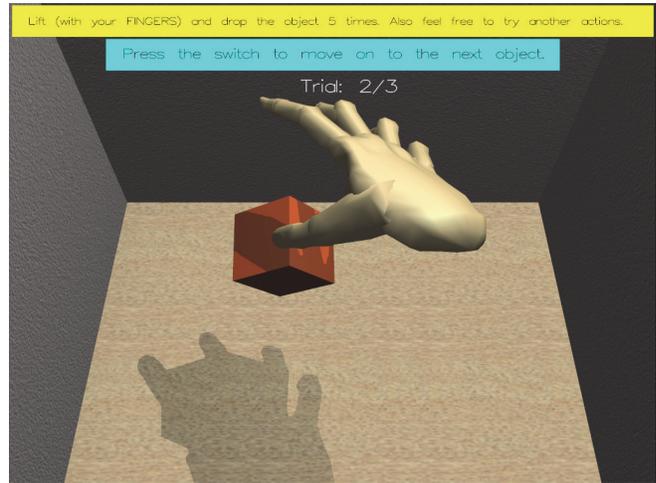


Fig. 7. Learning task that asked subjects to lift and drop the object to practice virtual grasping and releasing.



Fig. 8. Pick-and-drop experiment that asked subjects to drop the object from above the pit to the right side of the scene.

1. A demo session with on-screen instruction to introduce subjects to the task. It demonstrated one trial.
2. A practice session that allowed subjects to practice the task without instruction. It consisted of three trials. As an exception, the subjective-comparison experiment had no practice session.
3. The actual experiment session for measuring performance. It contained no instructions.

5.3.1 Procedure for Learning Task

During the learning task (Fig. 7), subjects practiced virtual grasping in 3 trials, with ball, cube, and bunny objects (one object per trial). They were required to lift and drop an object in each trial at least 5 times to practice grasping and releasing interactions.

5.3.2 Procedure for Pick-and-Drop Experiment

In the pick-and-drop experiment (Fig. 8), subjects picked up an object from the virtual floor at the left side of the scene and dropped it from above the pit at the right side.

The components of the trial are explained by the demo session instructions:

1. Lift the object above a quad. The quad will turn green.
2. Wait for a sound signal (a short beep sound, one second after the quad turns green).
3. Move the object to above the pit (it has to cross beyond the ledge) after the sound signal, using normal speed, and then release it using normal finger motion.

There were 27 trials in the experiment session: 3 object types x 3 object sizes x 3 trials. Condition order was randomized per subject.

The experiment software detected deviations from the intended steps, e.g., moving to the right before the sound signal. The software responded by displaying a warning and restarting the trial. Similar measures were in place for targeted ball-drop and cube-alignment experiments.

5.3.3 Procedure for Targeted Ball-Drop Experiment

In the targeted ball-drop experiment (Fig. 9), subjects picked up a ball from the virtual floor and dropped it from above an X-mark target on the floor. In the demo session, subjects were told that a floating wireframe cube above the target was the best place to drop the ball (the cube center was aligned with the center of the X mark). The components of the trial are explained by the demo session instructions:

1. Pick up the ball and move it inside the cube. The cube will turn green and the (2-second) countdown sound will begin.
2. Wait for the countdown sound to end while holding the hand still.
3. Release the ball immediately at the end of the countdown sound using normal finger motion.

There were 9 trials in the experiment session: three per ball size. Condition order was randomized per subject.

Target placements were chosen during experiment design to include varying positions (and orientations in the next experiment). To encourage precision release, the ball center was required to remain within a predefined threshold distance from the cube center during the countdown sound (or the trial was restarted; this rarely occurred). Similar target placement and precision constraints were in place for the cube-alignment experiment.

5.3.4 Procedure for Cube-Alignment Experiment

In the cube-alignment experiment (Fig. 10), subjects picked up a cube from the virtual floor and aligned it with a floating target wireframe cube. In the demo session, subjects were told that there was no gravity in this experiment. The components of the trial are explained by the demo session instructions:

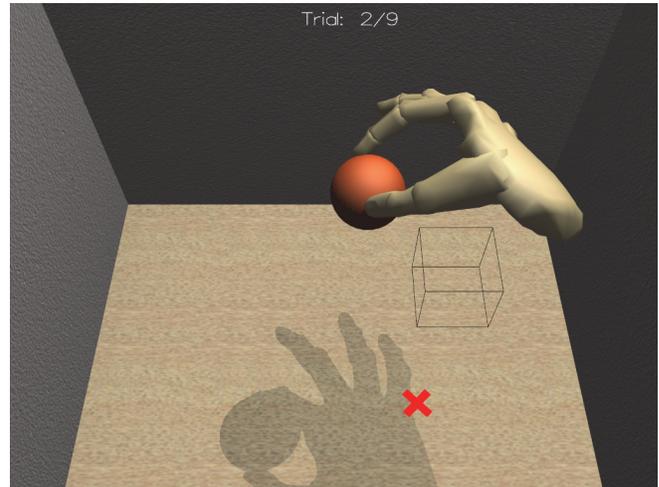


Fig. 9. Targeted ball-drop experiment that asked subjects to drop the ball at the X-mark target on the floor.

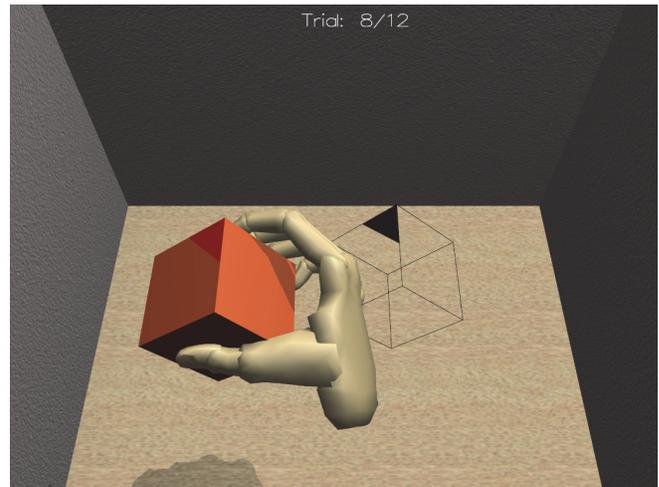


Fig. 10. Cube-alignment experiment that asked subjects to align the cube with the floating wireframe target under no gravity.

1. Pick up the cube and align it with the target. The target will turn green and the (2-second) countdown sound will begin.
2. Wait for the countdown sound to end while holding the hand still.
3. Release the cube immediately at the end of the countdown sound using normal finger motion.

There were 12 trials in the experiment session: four per cube size. Condition order was randomized per subject.

5.3.5 Procedure for Subjective Comparison Experiment

The subjective comparison experiment (Fig. 11) had subjects compare the two grasping techniques directly. In each trial, there were two similar objects at the left and right sides of the scene, separated by an invisible wall at the center (objects could not cross the wall). The left object was manipulated using one grasping technique (randomized per trial) while the right object was manipulated using the other grasping technique. A

question displayed at the top of the scene asked subjects to choose the object that was easier to release. After free exploration, subjects pressed a CyberGlove-mounted switch when they were ready and indicated an object by touching it with the virtual index fingertip for 2 seconds. A second question then asked them to indicate the object that was easier to pick up, and they answered using a similar procedure.

There were 9 trials in the experiment session: three trials each for ball, cube, and bunny. Object type order was randomized per subject.

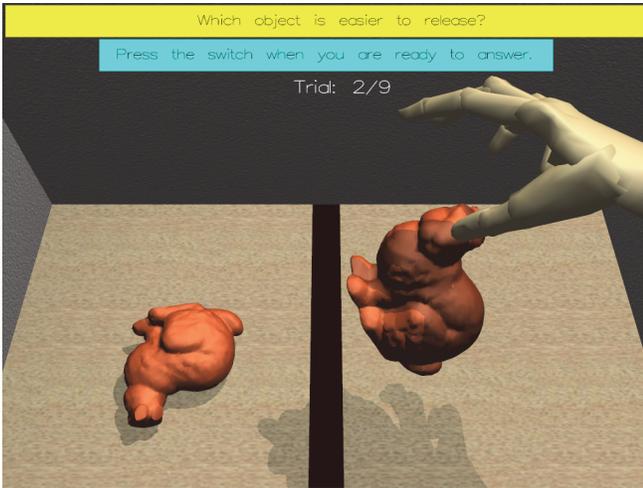


Fig. 11. Subjective comparison experiment that let subjects compare the two grasping techniques directly.

6 RESULTS

6.1 Pick-and-Drop Experiment Results

We computed release time value for the pick-and-drop experiment as follows:

Let:

$t1$ be the time instant when the object crosses beyond the ledge (detected by a boundary plane separating right and left sides).

$t2$ be the time instant when no finger phalanges of the virtual hand touch the object.

Then:

$$\text{Release time} = t2 - t1.$$

Note that this release time consists of a “movement time” component and an “actual grasp release time” component. The movement time is the duration between the instant when the object crosses beyond the ledge ($t1$) and a beginning of grasp release action. The actual grasp release time is the duration between the beginning of grasp release action and successful release of grasp ($t2$).

Fig. 12 summarizes resulting release times. We performed three-way repeated-measures ANOVA on release times. Due to an interaction between technique and size (see Fig. 13), we additionally performed one-way

repeated-measures ANOVA per grasping technique, with object size as the independent variable. Similarly, for an interaction between technique and type (see Fig. 14), we performed one-way repeated-measures ANOVA per grasping technique, with object type as the independent variable. Reported post-hoc test p-values include Bonferroni correction.

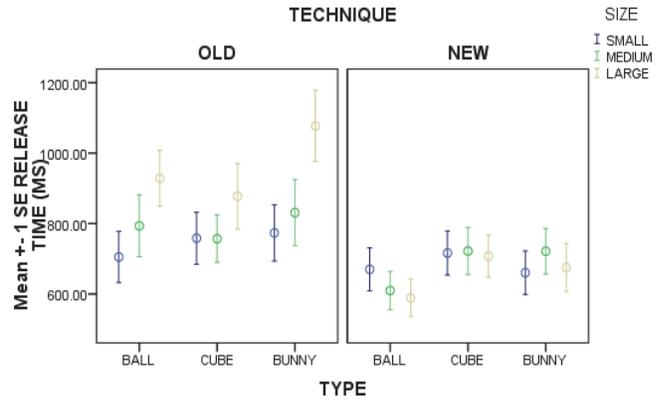


Fig. 12. Release time for the pick-and-drop experiment showing all independent variables (means and standard error bars).

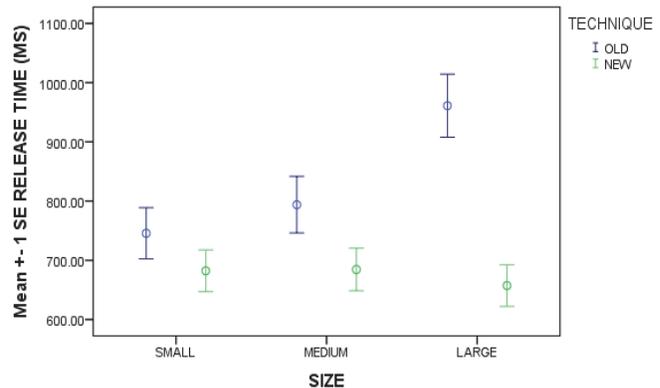


Fig. 13. Release time for the pick-and-drop experiment showing technique-size interaction.

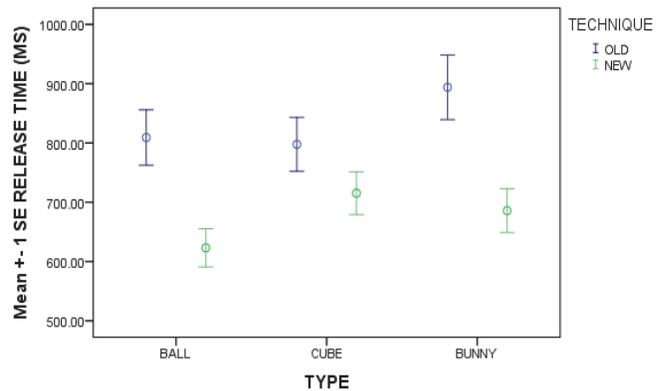


Fig. 14. Release time for the pick-and-drop experiment showing technique-type interaction.

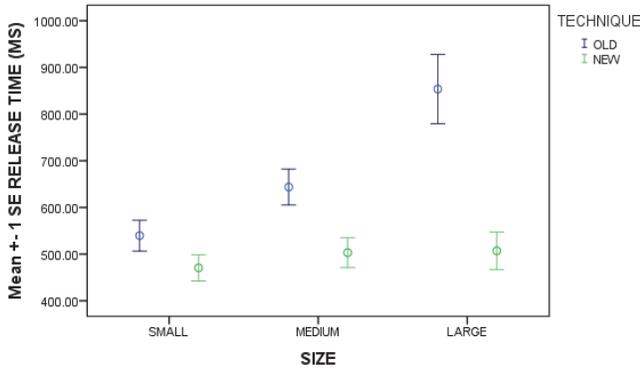


Fig. 15. Release time for the targeted ball-drop experiment.

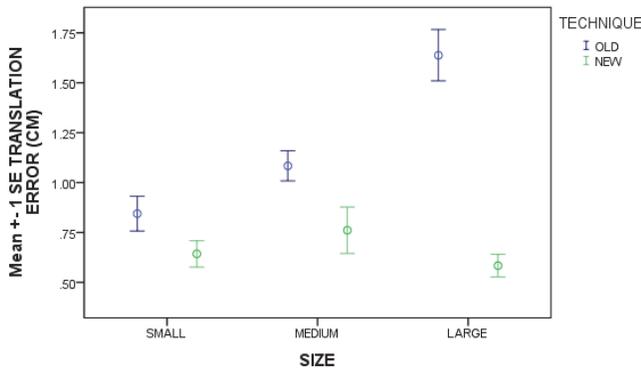


Fig. 16. Translation error for the targeted ball-drop experiment.

For release time:

1. There was a significant effect of grasping technique, $F(1,27) = 19.64$, $p < .001$.
2. There was a significant effect of object size, $F(2,54) = 15.25$, $p < .001$.
3. There was a significant effect of object type, $F(2,54) = 12.22$, $p < .001$.
4. There was a significant technique-size interaction, $F(2,54) = 18.15$, $p < .001$.
5. There was a significant technique-type interaction, $F(2,54) = 6.135$, $p < .005$.

Mean release time with the new spring model was 19% shorter than with the old (standard) spring model on average. Mean release time for the large object was significantly longer than for the medium object ($p < .01$) and for the small object ($p < .001$) by 9% and 13%, respectively. No statistically significant difference was detected in the medium-small pair ($p = .442$). Mean release time for the ball object was significantly shorter than for the cube object ($p < .05$) and for the bunny object ($p < .001$) by 5% and 9%, respectively. No statistically significant difference was detected in the cube-bunny pair ($p = .157$).

The per-technique tests for the technique-size interaction revealed a significant effect of object size for the old spring model ($F(2,54) = 23.45$, $p < .001$) with pairwise comparisons detecting significance in all pairs except the medium-small pair. However, no significant

effect of object size was detected for the new spring model ($F(2,54) = 1.023$, $p = .367$). The per-technique tests for the technique-type interaction revealed a significant effect of object type for both the old spring model ($F(2,54) = 7.049$, $p < .005$, with pairwise comparisons detecting significance in all pairs except the ball-cube pair) and the new spring model ($F(2,54) = 11.176$, $p < .001$, with pairwise comparisons detecting significance in all pairs except the cube-bunny pair).

6.2 Targeted Ball-Drop Experiment Results

We computed release time and translation error values for the targeted ball-drop experiment as follows:

Let:

$t1$ be the time instant when the countdown sound ends.

$t2$ be the time instant when no finger phalanges of the virtual hand touch the ball (the experiment software does not allow multiple grasps in a trial, so this is the end of the single grasp).

$t3$ be the instant when the ball touches the floor.

\mathbf{d} be the projected vector of $(\mathbf{p}_{t3} - \mathbf{p}_{t1})$ on the floor, where \mathbf{p}_{t1} and \mathbf{p}_{t3} are positions of the ball origin at times $t1$ and $t3$, respectively.

Then:

Release time = $t2 - t1$, and

Translation error = $\text{length}(\mathbf{d})$.

Note that this release time consists of a “reaction time” component and the “actual grasp release time” component. The reaction time is the duration between the instant when the countdown sound ends ($t1$) and a beginning of grasp release action. The actual grasp release time is defined as in Section 6.1. Translation error is defined independently of user targeting error. It is a measure of horizontal motion that results from release.

Fig. 15 and Fig. 16 summarize these release times and errors. We performed two-way repeated-measures ANOVA per dependent variable. Due to an interaction, we additionally performed one-way repeated-measures ANOVA per grasping technique, with object size as the independent variable. Reported post-hoc test p-values include Bonferroni correction.

For release time:

1. There was a significant effect of grasping technique, $F(1,27) = 18.02$, $p < .001$.
2. There was a significant effect of object size, $F(2,54) = 18.94$, $p < .001$.
3. There was a significant technique-size interaction, $F(2,54) = 9.99$, $p < .001$.

Mean release time with the new spring model was 27% shorter than with the old (standard) spring model on average. Mean release time for the large ball was significantly longer than for the medium ball ($p < .05$) and for the small ball ($p < .001$) by 19% and 35%, respectively. Mean release time for the medium ball was significantly longer than for the small ball ($p < .001$) by 14%.

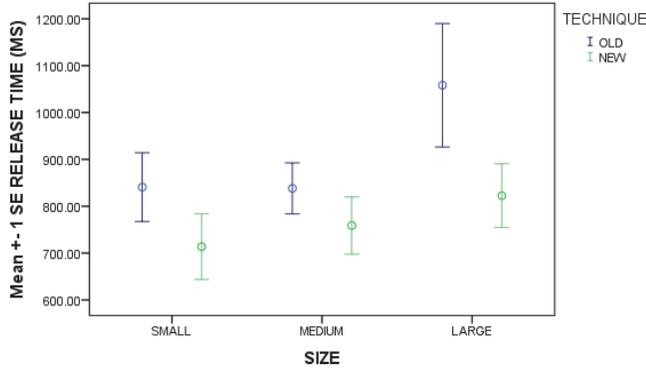


Fig. 17. Release time for the cube-alignment experiment.

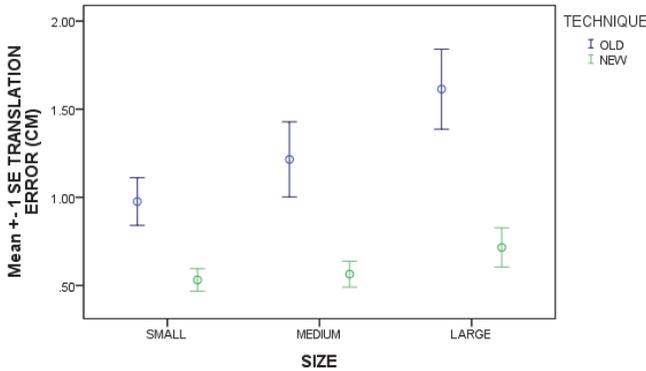


Fig. 18. Translation error for the cube-alignment experiment.

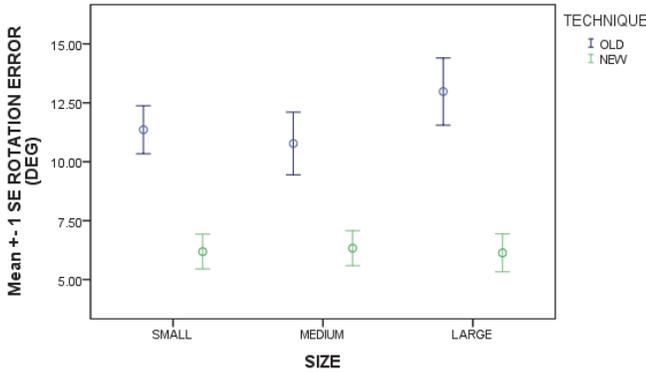


Fig. 19. Rotation error for the cube-alignment experiment.

The per-technique tests revealed a significant effect of object size for the old spring model ($F(2,54) = 17.81$, $p < .001$) with pairwise comparisons detecting significance in all pairs. However, no significant effect of object size was detected for the new spring model ($F(2,54) = .95$, $p = .395$).

For translation error:

1. There was a significant effect of grasping technique, $F(1,27) = 63.32$, $p < .001$.
2. There was a significant effect of object size, $F(2,54) = 11.79$, $p < .001$.
3. There was a significant technique-size interaction, $F(2,54) = 12.34$, $p < .001$.

Mean translation error for the new spring model was 44% smaller than for the old spring model on average. Mean translation error for the large and medium balls was significantly larger than for the small ball by 49% ($p < .005$) and 24% ($p < .01$), respectively. Mean translation error for the large ball was near-significantly larger than for the medium ball ($p = .095$) by 20%.

The per-technique tests revealed a significant effect of object size for the old spring model ($F(2,54) = 18.29$, $p < .001$) with pairwise comparisons detecting significance in all pairs except the medium-small pair (which showed near significance, $p = .065$). However, no significant effect of object size was detected in the new spring model ($F(2,54) = 1.54$, $p = .22$).

6.3 Cube-Alignment Experiment Results

We computed release time, translation error, and rotation error values for the cube-alignment experiment as follows:

Let:

$t1$ be the time instant when the countdown sound ends.
 $t2$ be the time instant when no finger phalanges of the virtual hand touch the cube.

\mathbf{p}_{t1} , \mathbf{p}_{t2} be positions of the cube center at times $t1$ and $t2$, respectively.

\mathbf{q}_{t1} , \mathbf{q}_{t2} be quaternion orientations of the cube at times $t1$ and $t2$, respectively. Then $\mathbf{q}_{t2}(\mathbf{q}_{t1})^*$ describes the cube rotation from $t1$ to $t2$ (* describes quaternion conjugate).

Then:

Release time = $t2 - t1$,

Translation error = $\text{length}(\mathbf{p}_{t2} - \mathbf{p}_{t1})$, and

Rotation error = absolute value of an angle component extracted from the quaternion $\mathbf{q}_{t2}(\mathbf{q}_{t1})^*$ (The angle component was adjusted to fall within $[-\pi, \pi]$ before taking the absolute, as no rotation amount larger than π was observed during the experiment).

Note that this release time consists of the “reaction time” component and the “actual grasp release time” component as described in Section 6.1 and Section 6.2. Translation error and rotation error are defined independently of user targeting error. They are measures of translation and rotation motions that result from release.

Fig. 17, Fig. 18, and Fig. 19 summarize these release times and errors. We performed two-way repeated-measures ANOVA per dependent variable. Even though no significant interaction was detected in this experiment, considering findings from the pick-and-drop and targeted ball-drop experiments (significant effect of object size for old spring model and no detected significant effect of object size for new spring model), we chose to also perform per-technique tests (one-way repeated-measures ANOVA per grasping technique, with object size as the independent variable). Reported post-hoc test p-values include Bonferroni correction.

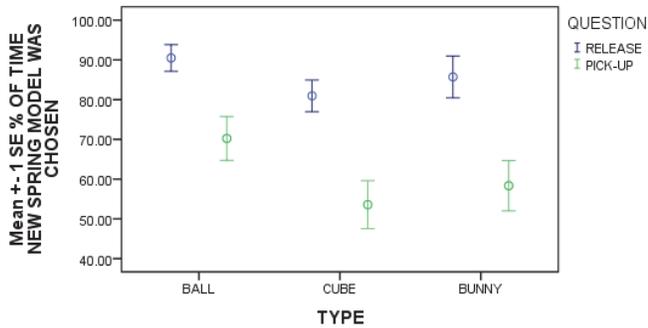


Fig. 20. Percentage of trials for which the new spring model was chosen as easier in subjective comparison experiment (mean and standard error of per-subject scores).

For release time:

1. There was a significant effect of grasping technique, $F(1,27) = 4.398$, $p < .05$.
2. There was a significant effect of object size, $F(2,54) = 6.399$, $p < .005$.
3. No statistically significant technique-size interaction was detected, $F(2,54) = 1.557$, $p = .220$.

Mean release time with the new spring model was 16% shorter than with the old (standard) spring model on average. Mean release time for the large cube was significantly longer than for the medium cube ($p < .05$) and for the small cube ($p < .05$) by 18% and 21%, respectively. No significant difference was detected in the medium-small pair ($p = 1.00$).

The per-technique tests revealed a significant effect of object size for the old spring model ($F(2,54) = 4.472$, $p < .05$) with pairwise comparisons detecting near significance in the large-medium pair ($p = .097$) and the large-small pair ($p = .100$) (these would appear significant without Bonferroni correction, which can be overly conservative). A near-significant effect of object size was detected for the new spring model ($F(2,54) = 3.147$, $p = .051$).

For translation error:

1. There was a significant effect of grasping technique, $F(1,27) = 28.15$, $p < .001$.
2. There was a significant effect of object size, $F(2,54) = 5.369$, $p < .01$.
3. No statistically significant technique-size interaction was detected, $F(2,54) = 1.902$, $p = .159$.

Mean translation error for the new spring model was 52% smaller than for the old spring model on average. Mean translation error for the large cube was significantly larger than for the small cube ($p < .01$) by 55%. No significant difference was detected in the large-medium pair ($p = .244$) and in the medium-small pair ($p = .652$).

Per-technique tests suggest a significant effect of object size for the old spring model ($F(2,54) = 4.056$, $p < .05$) with pairwise comparisons detecting significance in only the large-small pair. However, no significant effect of object size was detected in the new spring model ($F(2,54) = 2.256$, $p = .115$).

For rotation error:

1. There was a significant effect of grasping technique, $F(1,27) = 37.96$, $p < .001$.
2. No statistically significant effect of object size was detected, $F(2,54) = 1.385$, $p = .259$.
3. No statistically significant technique-size interaction was detected, $F(2,54) = 1.343$, $p = .270$.

Mean rotation error for the new spring model was 47% smaller than for the old spring model on average.

Per-technique tests reveal no significant effect of object size for both the old spring model ($F(2,54) = 1.765$, $p = .181$) and the new spring model ($F(2,54) = .046$, $p = .955$).

6.4 Subjective Comparison Experiment Results

For the subjective comparison experiment, we computed a per-subject score as the number of times the subject picked the new spring model over the number of contributing trials (i.e., percentage of trials for which the new technique was chosen as easier). Fig. 20 summarizes the results.

Subjects reported that grasp release was easier for the new spring model than for the old model: overall mean score for the release question was significantly above 0.5 ($t(27) = 12.06$, $p < .001$; all reported tests are two-tailed). Overall, the object manipulated using the new spring model was picked 86% of the time. Furthermore, the result also holds for each object type independently (ball: $t(27) = 12.02$, $p < .001$; cube: $t(27) = 7.75$, $p < .001$; bunny: $t(27) = 6.78$, $p < .001$; p -values were Bonferroni corrected for 3 comparisons).

Subjects reported that pick-up was easier for the new spring model than for the old model: overall mean score for the pick-up question was significantly above 0.5 ($t(27) = 2.25$, $p < .05$). Overall, the object manipulated using the new spring model was picked 61% of the time. Furthermore, the result holds for the ball object ($t(27) = 3.67$, $p < .005$) but not statistically significant for the other objects (cube: $t(27) = .59$, $p = 1.00$; bunny: $t(27) = 1.32$, $p = .597$).

7 DISCUSSION

7.1 Effect of Grasp Technique on Grasp Release

The results from pick-and-drop, targeted ball-drop, and cube-alignment experiments confirm our hypothesis that the new spring model improves speed and accuracy of grasp release. This can be explained by the new spring model requiring less finger extension to release grasped objects due to the use of the incremental motion metaphor during grasp release. Less required finger extension provides faster release and less sticking of grasped objects, which also improves release accuracy.

The subjective comparison results confirm our hypothesis that the new approach provides easier release subjectively, and this is consistent with the objective results discussed previously. Furthermore, the results from the pickup question provide some evidence that the new approach does not induce disturbing pickup problems (reducing possible concerns that the release-

heuristic function could incorrectly trigger during pickup, which could result in the object slipping out of grasp). We expect that there was actually no effect of grasping technique on the pickup action, since new and old (standard) spring models behave similarly during pickup (assuming no side-effect from the use of the release-heuristic function). The results (better subjective pickup with new technique, overall and for the ball object) may reflect overall subject experience with the object during the trial, including release, rather than differences specifically during pickup.

7.2 Effect of Object Size on Grasp Release

The targeted ball-drop results show that it took significantly longer to release larger objects than smaller ones with the old (standard) spring model, with associated reduced (translation) accuracy. This would be explained by larger objects resulting in larger interpenetration (hence more required finger extension). This may simply be due to the larger range of motion available, or to something more complex like tighter grasps learned for larger objects that are expected to be heavier based on real-world experiences. Further support of the reduced performance with increasing object size in the old spring model is found in pick-and-drop and cube-alignment results, where some, but not all relevant performance differences were statistically significant. In cases where statistical significance was not detected, we note the experiment was less sensitive to such differences due to:

1. The movement time component in the release time of the pick-and-drop results (Section 6.1). We observed during the pick-and-drop experiment that distances traveled by objects after crossing beyond the ledge and before grasp release varied (so does movement time). We suspect that this variation of movement time is larger than a variation of reaction time component used in the release time calculation of targeted ball-drop results (Section 6.2). The larger variation of movement time may blur the differences between actual grasp release times of different sized objects.
2. Smaller size variation in the cube object (compared with the other object types, see Fig. 5). Smaller object size difference results in smaller release performance difference in the old spring model. This would help explain the cube-alignment results and the pick-and-drop results.
3. Grasp release might be less natural or more difficult without gravity in the cube-alignment experiment, depending on a subject's approach. We observed that in some trials, subjects took notably longer to release an object, sometimes causing an object to be dragged by a moving hand. These bad trials may affect the cube-alignment results.
4. The subjects used various hand orientations for target rotation alignment in the cube-alignment experiment. We observed that some subjects used different hand orientations for the same target rotation. The choice affected performance as some orientation appeared

less comfortable during release as demonstrated by the subjects. This may affect the cube-alignment results.

Rotation error with the old spring model in the cube-alignment results does not appear to increase notably with increasing object size. This is unexpected and requires further investigation. Possible explanations include that hand-object stickiness does not affect object rotation as strongly as object translation (during release) and that rotation error results could be affected by small cube-size differences, gravity conditions, and choice of hand orientation, as discussed previously.

The new spring model mitigates the problem of increasing sticking with increasing object size, as we detected no statistically significant effect of size in the new spring model and the resulting means and standard errors suggest any present effect would be relatively small. In the new spring model, if the heuristic analysis detects release motion, the virtual hand will open almost immediately independent of the amount of (real) finger penetration. However, release times from the cube-alignment experiment suggest (weakly) some increasing sticking could remain with the new spring model for certain objects, based on a near-significant effect of size. This might be explained by the various hand orientations used in the cube-alignment experiment, where some finger release motions may not be detected well by our heuristic analysis parameters (relating to the specific joints involved). Reduced heuristic detection would make the new spring model behave more like the old spring model.

7.3 Effect of Object Type on Grasp Release

The pick-and-drop results show that it took significantly longer to release the bunny object than the other object types with the old (standard) spring model. This would be explained by relatively large sizes of the bunny resulting in larger interpenetration. We also observed during the pick-and-drop experiment that some subjects occasionally failed to release a bunny when they grasped its neck because the bunny's head got caught at their hand due to concavities (which is why we removed the ears). Subjects could shake their hand to successfully release the bunny, which added to the release time.

The pick-and-drop results also show that it took significantly less time to release the ball object than the other object types with the new spring model. As performance of the new spring model depends on heuristic detection of grasp release, this may be explained by different joints emphasized for different object types during release. Subjects may mainly use flexion at first knuckles and thumb roll to release the ball, which matches the chosen joints in our heuristic implementation, resulting in the best heuristic detection of finger release motions (more frequent uses of incremental motion for release). Subjects may rely increasingly on other joints to release the cube, resulting in reduced heuristic detection and longer release time. For the bunny, longer release may reflect the joints used or the bunny's head getting caught during release.



Fig. 21. A rotating control knob for convergence speed adjustment, at the left side of the display. The knob has no reference points or stops.

The subjective comparison results suggest that subjects were less sensitive to release quality changes with the cube object than with other object types (for the sizes used in the experiment). This may be explained by observing the objective pick-and-drop results in Fig. 12 for medium and large sizes of each object type. The cube object shows the least average performance improvement with the new spring model over the old spring model. This may result from reduced heuristic detection, as discussed above.

7.4 Virtual Hand Management

Our results demonstrate that, for finger joint angles, maintaining pose discrepancy with subsequent convergence can improve user performance and experience. This complements the results of Burns et al. (2007), who introduced discrepancy in hand base position only and showed improved user ratings with no loss in performance for a hand navigation task.

8 FOLLOW-UP EXPERIMENT: HEURISTIC PERFORMANCE AND CONVERGENCE EFFECTS

We conducted follow-up experiments for additional insight into heuristic and convergence behavior of the new technique. The main purpose was to detect and understand possible side-effects of these mechanisms and to provide a starting point for convergence parameter optimization.

8.1 Design

The experiments consisted of a targeted ball-drop experiment, a convergence tuning experiment, and an artifact explanation experiment. The targeted ball-drop experiment studied heuristic trigger accuracy and convergence performance using settings from the earlier experiment. This was to check for potential triggering problems and to estimate the minimum time required between grasps for full convergence. We used the targeted ball-drop approach from the main experiment, with new dependent variables and minor procedure changes (Section 8.3.1). The new dependent variables,

defined further in Section 8.4.1, were:

1. Accidental Drop – number of ball releases outside a release interval.
2. Incorrect Trigger – number of heuristic triggers outside a release interval.
3. Correct Trigger – number of heuristic triggers inside a release interval.
4. Convergence Time – Target hand convergence period after successful release.

The convergence tuning and artifact explanation experiments investigated subjective artifacts of convergence and found subjectively-suitable parameter ranges. The convergence (speed) parameter, c , creates a continuum of virtual hand behaviors. At one end, with zero convergence, the new spring behavior matches the (pure) incremental motion metaphor for release, for which a maintained offset between virtual and real finger configurations can grow with multiple grasps, resulting in unreasonable virtual hand configuration. At the other end, with very fast convergence, the new spring behavior essentially matches the rubber band metaphor (the standard spring behavior), resulting in sticking. The best tradeoff lies somewhere in between, and ideally there is a range of values avoiding both problems.

Convergence tuning found this range by asking subjects to adjust convergence speed to find the lowest, the highest, and the best overall values having good (subjective) performance for each of three ball sizes. Subjects adjusted convergence speed in a range that includes pure-incremental and rubber-band-like behavior at its boundaries.

The artifact explanation experiment provided more understanding of perceivable artifacts and their relative strength by asking subjects to adjust convergence speed freely and explain any unpleasant artifacts they encountered (while also freely switching ball size). We were especially interested to check if subjects would report visual-proprioceptive motion discrepancy (Burns et al., 2006), where virtual and real finger motions disagree, resulting from virtual finger convergence motion.

8.2 Apparatus, Implementation Notes, and Subjects

We used the apparatus from the earlier experiments, with the addition of a Griffin PowerMate knob (Fig. 21). This knob rotates with no stops or reference points that could otherwise bias responses. It varied convergence speed in the range $[0.0, 1.0]$ degrees per simulation step, in 100 increments. Increments were spaced nonlinearly to provide finer control at smaller values (implemented by squaring linearly-spaced values in $[0.0, 1.0]$).

There were 12 participants: all male, aged 21 to 34 years (average = 26), 11 right-handed and 1 left-handed. Most subjects (11) were students, primarily from computer science and computer engineering programs. Experience levels were mixed: 2 previously participated in both the main experiment and another experiment involving only the standard spring model (Prachyabrued

and Borst, 2012), 7 others participated in the other grasping experiment, 1 other reported exposure to a VR system, 1 other played video games and watched 3D movies, and the remaining subject reported minimal experience related to VR.

8.3 Procedure

Subjects performed tasks in the order of presentation in this section. We calibrated the CyberGlove for each subject before they started. Experiment duration was typically 45 to 60 minutes.

A learning task and the targeted ball-drop experiment involved both grasping techniques, with order randomized per subject such that half of the subjects experienced the new spring model first, and half experienced the standard spring model first. The other experiments involved only the new spring model.

8.3.1 Procedures for Learning Task and Targeted Ball-Drop

We used the learning task from the main experiment (Section 5.3.1) except that subjects practiced virtual grasping in three trials with the three ball sizes (one size per trial).

The targeted ball-drop procedure was similar to that in the main experiment (Section 5.3.3), except that the trial would be restarted by the experiment software if an accidental ball drop occurred (defined in Section 8.4.1). Subjects also experienced both grasping techniques in one day.

8.3.2 Procedure for Convergence Tuning

The convergence tuning experiment asked subjects to find lowest, highest, and best parameter values (convergence speed) with “good performance during normal grasp and release”. Subjects were not informed about what the parameter was except that it affected grasp and release.

The first demo trial (Fig. 22) had subjects experience the minimum and maximum values. There were two large balls at the left and right sides of the scene, separated by an invisible wall (balls could not cross the centerline). The left side used the minimum value (0) and the right side used the maximum value (1). An instruction at the top asked subjects to “pick and drop repeatedly using normal release motion” and “notice how the values affect grasp and release”. After free exploration, subjects pressed a CyberGlove-mounted switch to indicate they understood effects and were ready to end the demo.

A second demo trial had subjects freely try other parameter values in a simple ball-drop environment similar to that in Fig. 7. There was one large ball. An instruction asked subjects to “pick and drop repeatedly using normal release motion”, turn the knob to adjust the parameter value, and notice how the value affects grasp and release. Subjects pressed the CyberGlove switch to indicate they understood the behavior and were ready to end the demo.

The first and second demo trials showed parameter value near the top of the scene, but the remaining trials in this experiment did not reveal values, except “MIN” and “MAX”.

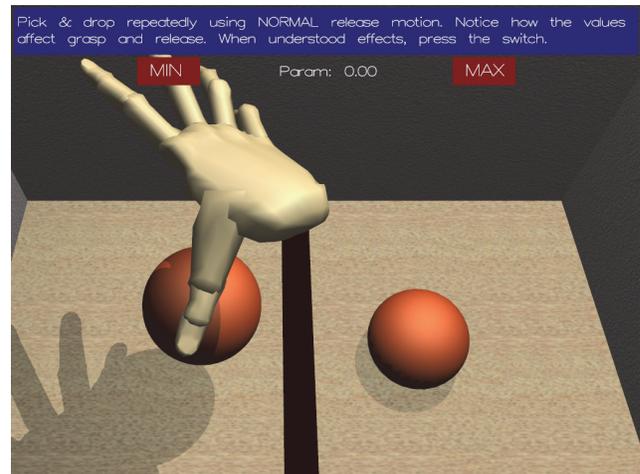


Fig. 22. A demo trial for convergence tuning asked subjects to compare minimum and maximum convergence speeds. The picture also shows an unusual thumb configuration resulting from zero convergence, complicating grasp.

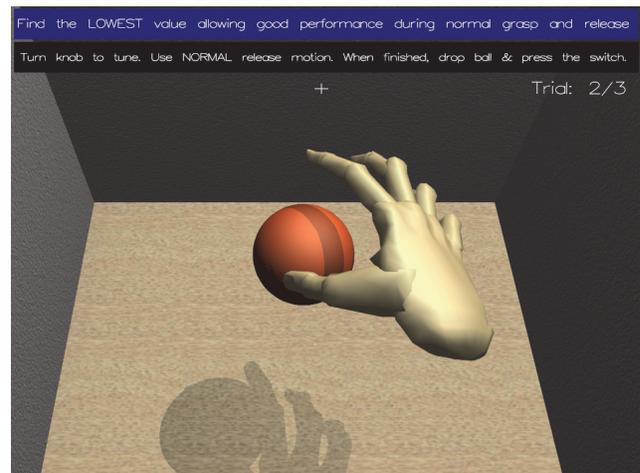


Fig. 23. For convergence tuning, subjects indicated lowest, highest, and best convergence speeds for “good performance during normal grasp and release”. Value was not displayed, but there was a change indicator (+ or -).

A third demo trial simply rehearsed one regular parameter tuning trial (Fig. 23). Per regular trial, subjects picked up and dropped a ball repeatedly while adjusting the parameter in response to three instructions. The initial value was randomized per trial. A value increase or decrease was indicated near the top of the scene with a + or -, respectively. Instructions first stated “Find the LOWEST value allowing good performance during normal grasp and release” (or “HIGHEST”, in randomized order per trial). After adjustment, subjects pressed the CyberGlove switch twice to indicate the current value as the choice (the second press was for confirmation). The second instruction asked for the other extreme value (highest or lowest for good performance), and a final instruction stated “Find the BEST overall value”. There were 3 of the regular 3-part tuning trials (1 per ball size) for a total of 9 tuning questions. Ball size order was randomized per subject.

8.3.3 Procedure for Artifact Explanation

Artifact explanation consisted of asking subjects to explain any unpleasant artifacts while adjusting convergence speed. Subjects picked up and dropped a ball repeatedly in a simple ball-drop environment similar to that in Fig. 7, freely switching between ball sizes by pressing the CyberGlove switch. Initial ball size was large. An instruction stated “Adjust the parameter freely and explain any unpleasant artifacts you find to the experimenter”. The parameter value was shown on screen. Initial value was randomized. Subjects were encouraged by the experimenter to test all ball sizes. To end the experiment, subjects indicated to the experimenter that they had reported all detectable artifacts.

8.4 Results and Discussion

8.4.1 Results and Discussion for Targeted Ball-Drop Follow-up

We computed accidental drop, incorrect trigger, correct trigger, and convergence time for the targeted ball-drop trials as follows:

Let:

- $t1$ be the time when the ball is lifted just above the virtual floor,
- $t2$ be the time when the countdown sound ends,
- $t3$ be the time when no virtual fingers touch the ball (the software does not allow multiple grasps in a trial, so this is the end of the single grasp),
- J be a set of all joint angles of the model described in Section 4.3,
- $\theta_{ig}(j, t)$ be target angle of joint $j \in J$ at time t ,
- $\theta_{tr}(j, t)$ be tracked angle of joint $j \in J$ at time t , and
- c be convergence speed from Section 4.3 (0.035 degrees per simulation time step).

Then:

- Accidental drop = Number of ball drops in $[t1, t2]$. Drops are detected as a lack of phalange-ball contact at $t2$ or a ball hitting the ground in $[t1, t2]$,
- Incorrect trigger = Number of heuristic triggers during $[t1, t2]$,
- Correct trigger = Number of heuristic triggers during $[t2, t3]$, and
- Convergence time = $\max_{j \in J} (|\theta_{ig}(j, t3) - \theta_{tr}(j, t3)|) / c$, (i.e., number of simulation steps for convergence, based on the joint with maximum angle between target and tracked configurations).

Accidental drop: Our release mechanisms did not cause accidental drop problems. During the entire experiment, the new spring model and the standard spring model had only 1 accidental drop each, corresponding to 0.93% of trials. The drop for the new spring model trial was not due to the heuristic, as we detected no associated trigger. The drops were from different subjects. Both drops occurred for small balls and before the countdown sound, when the hand was moving. Based on our observation of

users, we believe a likely cause of drops is quick or careless grasps resulting in undesirable ball position between fingers and subsequent slipping out of grasp.

Incorrect trigger: There was a small number of incorrect triggers (4.6% of trials), although they did not result in ball drops. During the entire experiment, the heuristic triggered for 3 subjects outside of the release interval, in 5 unique trials. One subject had dual incorrect triggers in one trial (medium ball, after the countdown sound began, during which the subjects were instructed to hold the hand still) and a single trigger in another trial (small ball, after the countdown sound began). Another subject had a single trigger in one trial (large ball, before the countdown sound began). The third subject had a single trigger in each of two trials (large ball, before the countdown sound began, and medium ball, after the countdown sound began).

Incorrect triggers may be explained by finger movement (intentional or unintentional), noise in finger sensing, or the subject releasing just a bit too early. A possible reason for an incorrect trigger not resulting in a drop is that finger motion from unintentional movement or noise can be small, such that resulting incremental release motion is not large enough to get the virtual finger outside the object. Our grasping software allowed slight overlap between virtual fingers and the ball for improved contact simulation (see Section 5.1). Another possible reason is that even if the resulting incremental motion was large enough, it might only involve some finger segments while other segments still support the grasp or touch the object (the experiment software showed continual phalange-ball contact during $[t1, t2]$ in trials not involving accidental drops). Resulting release motion could then be corrected by the convergence algorithm before the ball slips out of grasp. This could cause some small shift in grasped object position, but fortunately, the reported percentage of incorrect triggers is small.

Correct trigger: There were correct triggers during release in 81.5% of trials and none in the other 18.5% of trials. Missing triggers (10 trials for small ball, 7 for medium ball, and 4 for large ball) could be explained by finger joints used for release not matching the chosen joints in our heuristic settings, or some early releases (just before the countdown sound ended) classified as incorrect. More missing triggers for smaller objects could be explained by less range of finger release motion resulting in less chance for triggering. Fortunately, the main experiment showed that there is less performance impact of the missing trigger on smaller objects than on larger objects.

10.19% of trials had multiple triggers during release, with one trial having two extra triggers and the others each having a single extra trigger. One possible cause is unintentional movement or noise during the reaction period. Another is that, even when overall release motion is large, triggering motion can be small over a single glove sensing interval at the beginning of release, due to sampling. Sufficiently small motion would not get the

target hand outside of, or far enough away from, the object, due to collision geometry overlap (see Section 5.1). Additionally, some convergence occurs between glove readings due to simulation rate being much higher than the glove update rate, allowing the target hand to sink into the object again when the latest sensed finger is still inside. In such cases, multiple triggers are appropriate.

Convergence time: The new spring model can be expected to support multiple interactions in quick succession if the convergence period is short enough to avoid accumulating excessive joint offsets from multiple interactions. The average convergence times were 186 simulation time steps ($\sigma = 97$), 350 ($\sigma = 164$), and 510 ($\sigma = 244$), for small ball, medium ball, and large ball, respectively, calculated from all trials with nonzero joint offsets at release (t_3). Simulation rate varied between 1252 Hz and 1334 Hz, depending on virtual hand-object contact, with overall average being 1292 Hz. We can use this average rate to estimate convergence time in a more convenient unit (ms): 144 ms, 271 ms, and 395 ms, for small, medium, and large balls, respectively. These are slight overestimates, since the convergence occurred primarily after release, where simulation rate is higher than during contact. Convergence time can be reduced by increasing the convergence speed, with a possible tradeoff of more sticking.

8.4.2 Results and Discussion for Convergence Tuning Experiment

Fig. 24 summarizes subject-tuned lowest, highest, and best parameter values giving “good performance” for each ball size. Results suggest that subjects were not very sensitive to effects of convergence for the small ball, as there is a wide range from lowest to highest values for most subjects. The range decreases with increasing ball size, with a significant reduction in highest good value from small to large sizes ($Z = -2.746$, $p < .01$, Wilcoxon paired signed-rank test). This suggests convergence speed should be optimized with large objects, because good values for large objects are likely to be good for smaller objects. The effect can be explained by larger real finger interpenetration for the large ball resulting in more sticking at high convergence speeds resembling standard rubber band behavior, as suggested by the earlier main experiment (see Section 7.2).

We detected no statistically-significant difference between lowest good values for the large and small balls ($Z = -1.412$, $p = .158$), but we note the sample size is small. At low parameter values, larger real finger interpenetration for larger objects should result in larger virtual-real offsets after release. Very low convergence speed causes offsets to be maintained for a substantial time, causing artifacts investigated in the next section. The tuning results here suggest subjects reached low parameter values before becoming bothered by this, for all tested object sizes.

Subjects seem to prefer slow convergence speed, as the best values are closer to the lowest values than to the highest values (based on medians and distribution

shapes). Subjectively, the best overall convergence speed is likely near the best-value large-ball median, and this is also between most lowest and highest values from the entire tuning study. This median is 0.058 degrees per simulation time step (0.075 degrees per ms), somewhat above the convergence speed from the main experiment, 0.035 (0.045 degrees per ms). The convergence speed from the main experiment still appears good, is above most subjects’ lowest good values, and allowed our release mechanism to have better objective and subjective performance than the old spring model in the main experiment.

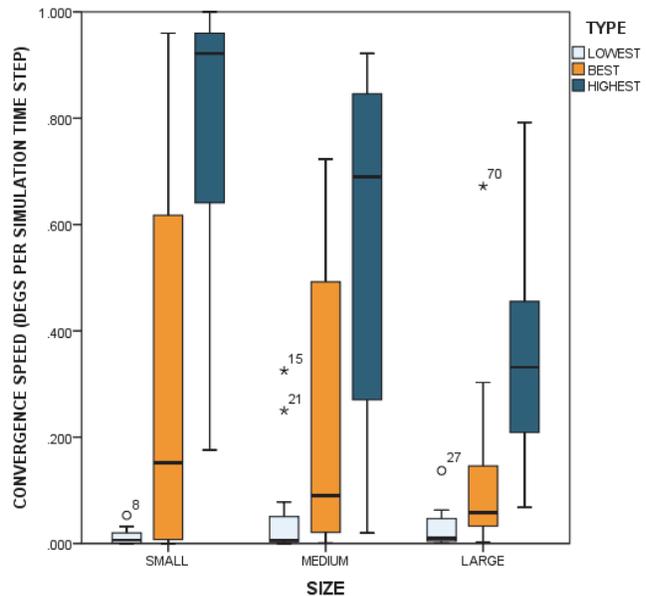


Fig. 24. Lowest, highest, and best convergence speeds giving “good performance”, tuned subjectively, for three ball sizes. Box-and-whiskers plots show each distribution’s median and interquartile boundaries, except for outliers denoted by circles and asterisks. Convergence speed is shown in degrees per simulation step; simulation rate is around 1292 Hz.

8.4.3 Results and Discussion for Artifact Explanation Experiment

Artifact explanations showed more sensitivity to object sticking and to a grasp problem than to visual-proprioceptive finger motion discrepancy. All subjects reported (re)grasping difficulty for near-zero parameter values: 11 subjects reported unrealistic virtual hand configurations preventing a proper grasp, resulting in failed grasp or accidental ball slipping; the other subject only reported accidental drops as the artifact. All subjects also reported sticking during release for high parameter values. All subjects reported reduced artifacts for smaller objects, especially for the small ball compared with the large ball. Overall, statements reflect tradeoffs when going from low to high values: mainly, easier release but more difficult grasping for low parameter values, and easier grasping but more difficult release for high parameter values.

Only one subject reported the visual-proprioceptive motion discrepancy, stating that the virtual thumb motion

did not synchronize with the real thumb motion. This subject stated that the artifact occurred at parameter values between 0.01 and 0.04 (i.e., [0.013, 0.052] degrees per ms) for the large ball and at 0.0625 (0.081 degrees per ms) for the medium ball. We can further consider two types of motion discrepancy (Burns et al., 2006): speed disparity, where the virtual and real fingers move at different speeds; and stationary-self, where the virtual finger moves but the real finger does not. Since we observed that all subjects picked and dropped repeatedly during the experiment, we note that most subjects were insensitive to speed disparity. It is less clear how much subjects are sensitive to the stationary self as the hand was mostly moving during the experiment (although we sometimes observed subjects pausing the hand shortly after release). Based on the results of this experiment, convergence tuning should prioritize the other artifacts over motion discrepancy, as subjects rarely reported distraction from virtual finger motion. Nonetheless, it may be possible to find a good value considering all artifacts: the best value mentioned in the previous section falls outside the range in which the subject reported discrepancy.

As only one subject reported motion discrepancy, we also report our own experience: we believe the convergence motion artifact is most notable at small convergence speeds due to increased duration. We also believe it is mostly noticed at the thumb due to a relatively large interpenetration depth (large offset at release) compared to other digits. It may be possible to reduce detection by making virtual torsional springs for the thumb tighter relative to the fingers, exerting more thumb force to objects to decrease thumb interpenetration while increasing interpenetration at the other fingers. This is relevant especially when grasps involve multiple (non-thumb) fingers, as this increases thumb penetration relative to other digits.

9 CONCLUSION AND FUTURE WORK

We described a grasp release problem that is relevant for any virtual grasping technique that requires fingers to be moved outside an object's boundaries for release. Our enhanced spring model for virtual grasping successfully addresses the sticking object problem, improving grasp release. This was confirmed by experiments showing significantly improved speed, accuracy, and subjective experience, without substantial perceivable side-effects. We demonstrated convergence motion effects after release, but showed convergence parameter ranges for avoiding them.

The convergence parameter study can be extended with objective performance measurement so that convergence speed can be optimized objectively in addition to subjectively. We expect a small convergence speed, such as our default of 0.045 degrees per ms, to be good because it corresponds to aggressive release mechanism behavior and good performance results from our main study. Subjective tuning results approached a low value for large objects, and values that are good for

large objects also appear good for small objects. Small values may raise concerns about user sensitivity to motion discrepancy, although only one subject reported this in our artifact explanation experiment. A more focused experiment can specifically investigate sensitivity to stationary-self motion discrepancy.

It may be possible to further improve convergence with a concept similar to MACBETH (Burns et al., 2007), extended to a fully articulated hand model, to trade off position discrepancy with motion discrepancy in a more intelligent way. For example, convergence speed could vary with object size or penetration depth. Also, different convergence speeds may be best for different release motion speeds, and an enhanced convergence algorithm could account for this.

We want to improve heuristic release detection and understand tradeoffs in our heuristic's parameters. Our main experiment results might reflect that the heuristic analysis worked better for a ball than for a cube, for example. By analyzing finger release motions for different objects and tasks, we may be able to come up with sets F_t , F_i , F_m , F_r , F_p that better support a wide range of objects. Our follow-up experiment showed a small number of missed triggers even for the ball, with a possible relationship to size. In addition to improving chosen joint sets, we may reduce the number of missed triggers by lowering relevant joint angle thresholds for increased sensitivity.

Following results of our work and those of Burns et al. (2007), showing that incremental motion with offset reduction could improve grasp release and hand navigation, respectively, we want to find out if a similar idea could be used to manage a hand base orientation to improve other types of tasks.

REFERENCES

- Allard, J., Menier, C., Raffin, B., Boyer, E., Faure, F., 2007. GrImage: Markerless 3D Interactions. In: *Proceedings of the ACM SIGGRAPH 2007 Emerging Technologies*. San Diego, California, August 5-9, 2007.
- Bergamasco, M., Degl'Innocenti, P., Bucciarelli, D., 1994. A Realistic Approach for Grasping and Moving Virtual Objects. In: *Proceedings of the IEEE/RSJ/GI International Conference on Intelligent Robots and Systems (IROS'94)*. Munich, Germany, pp. 717-724, September 12-16, 1994.
- Borst, C.W., Indugula, A.P., 2005. Realistic Virtual Grasping. In: *Proceedings of the IEEE Virtual Reality Conference (VR'05)*. Bonn, Germany, pp. 91-98, March 12-16, 2005.
- Borst, C.W., Indugula, A.P., 2006. A Spring Model for Whole-Hand Virtual Grasping. *Presence: Teleoperators and Virtual Environments*. 15(1):47-61, February 2006.
- Burns, E., Razaque, S., Panter, A.T., Whitton, M.C., McCallus, M.R., Brooks Jr, F.P., 2006. The Hand Is More Easily Fooled than The Eye: Users Are More Sensitive to Visual Interpenetration than to Visual-Proprioceptive Discrepancy. *Presence: Teleoperators and Virtual Environments*. 15(1):1-15, February 2006.
- Burns, E., Razaque, S., Whitton, M.C., Brooks, F.P., 2007. MACBETH: Management of Avatar Conflict by Employment of a Technique Hybrid. *The International Journal of Virtual Reality*. 6(2):11-20.

- DesRosiers, H., Gomez, D., Tremblay, M., Ullrich, C., 2001. *VirtualHand v.2.5 - Programmer's Guide*. Virtual Technologies Inc., Palo Alto, California.
- Fabiani, L., Burdea, G., Langrana, N., Gomez, D., 1996. Human Interface Using the Rutgers Master II Force Feedback Interface. In: *Proceedings of the IEEE Virtual Reality Annual International Symposium (VRAIS'96)*. Santa Clara, California, pp. 54-59, Mar 30-Apr 3, 1996.
- Hilliges, O., Izadi, S., Wilson, A.D., Hodges, S., Garcia-Mendoza, A., Butz, A., 2009. Interactions in the Air: Adding Further Depth to Interactive Tabletops. In: *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST'09)*. Victoria BC, Canada, pp. 139-148, October 4-7, 2009.
- Hilliges, O., Kim, D., Izadi, S., Weiss, M., Wilson, A.D., 2012. Holodesk: Direct 3D Interactions with a Situated See-Through Display. In: *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI'12)*. Austin, Texas, pp. 2421-2430, May 5-10, 2012.
- Hirota, K., Hirose, M., 2003. Dexterous Object Manipulation Based on Collision Response. In: *Proceedings of the IEEE Virtual Reality Conference (VR'03)*. pp. 232-239, March 22-26, 2003.
- Holz, D., Ullrich, S., Wolter, M., Kuhlen, T., 2008. Multi-Contact Grasp Interaction for Virtual Environments. *Journal of Virtual Reality and Broadcasting*. 5(7), July 2008.
- Iwata, H., 1990. Artificial Reality with Force-Feedback: Development of Desktop Virtual Space with Compact Master Manipulator. *ACM SIGGRAPH Computer Graphics*. 24(4):165-170, August 1990.
- Jacobs, J., Froehlich, B., 2011. A Soft Hand Model for Physically-based Manipulation of Virtual Objects. In: *Proceedings of the IEEE Virtual Reality Conference (VR'11)*. Singapore, pp. 11-18, March 19-23, 2011.
- Lindeman, R.W., Sibert, J.L., Templeman, J.N., 2001. The Effect of 3D Widget Representation and Simulated Surface Constraints on Interaction in Virtual Environments. In: *Proceedings of the IEEE Virtual Reality Conference (VR'01)*. pp. 141-148, 2001.
- Maekawa, H., Hollerbach, J.M., 1998. Haptic Display for Object Grasping and Manipulating in Virtual Environment. In: *Proceedings of the IEEE International Conference on Robotics & Automation*. Leuven, Belgium, pp. 2566-2573, May 16-20, 1998.
- Moehring, M., Froehlich, B., 2010. Enabling Functional Validation of Virtual Cars through Natural Interaction Metaphors. In: *Proceedings of the IEEE Virtual Reality Conference (VR'10)*. Massachusetts, pp. 27-34, March 20-24, 2010.
- Moehring, M., Froehlich, B., 2011. Effective Manipulation of Virtual Objects Within Arm's Reach. In: *Proceedings of the IEEE Virtual Reality Conference (VR'11)*. Singapore, pp. 131-138, March 19-23, 2011.
- Osawa, N., 2006. Automatic Adjustments for Efficient and Precise Positioning and Release of Virtual Objects. In: *Proceedings of the ACM International Conference on Virtual Reality Continuum and Its Applications (VRCIA'06)*. Hong Kong, pp. 121-128, June 14-17, 2006.
- Prachyabrued, M., Borst, C.W., 2012. Visual Interpenetration Tradeoffs in Whole-Hand Virtual Grasping. In: *Proceedings of the IEEE Symposium on 3D User Interfaces (3DUI'12)*. Orange County, California, pp. 39-42, March 4-5, 2012.
- Prachyabrued, M., Borst, C.W., 2011. Dropping the Ball: Releasing a Virtual Grasp. In: *Proceedings of the IEEE Symposium on 3D User Interfaces (3DUI'11)*. Singapore, pp. 59-66, March 19-20, 2011.
- Ullmann, T., Sauer, J., 2000. Intuitive Virtual Grasping for Non Haptic Environments. In: *Proceedings of the IEEE Eight Pacific Conference on Computer Graphics and Applications*. Hong Kong, pp. 373-457, October 3-5, 2000.
- Virtual Technologies Inc., 1994. *CyberGlove™ User's Manual*.
- Wan, H., Gao, S., Peng, Q., 2004. Virtual Grasping for Virtual Assembly Tasks. In: *Proceedings of the IEEE Third International Conference on Image and Graphics (ICIG'04)*. pp. 448-451, December 18-20, 2004.
- Wilson, A.D., Izadi, S., Hilliges, O., Garcia-Mendoza, A., Kirk, D., 2008. Bringing Physics to the Surface. In: *Proceedings of the ACM Symposium on User Interface Software and Technology (UIST'08)*. Monterey, California, pp. 67-76, October 19-22, 2008.
- Zachmann, G., Rettig, A., 2001. Natural and Robust Interaction in Virtual Assembly Simulation. In: *Proceedings of the Eighth ISPE International Conference on Concurrent Engineering: Research and Applications (ISPE/CE2001)*. California, July 2001.