Effects and Optimization of Visual-Proprioceptive Discrepancy Reduction for Virtual Grasping

Mores Prachyabrued*
University of Louisiana at Lafayette

Christoph W. Borst†
University of Louisiana at Lafayette

Abstract

A virtual hand can intentionally deviate from the real hand configuration to improve interaction and experience. Subsequent discrepancy reduction (convergence) is important to avoid undesirable side effects such as grasping difficulty. Prior work did not study convergence motion effects adequately to understand finger motions and speed effects. We present the first grasping-focused studies on convergence motion speed effects. The results are important for optimizing convergence motion and can contribute to guidelines for related interaction types.

Keywords: Virtual grasping, visual-proprioceptive discrepancy.

Index Terms: H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Artificial, augmented, and virtual realities

1 Introduction

We previously showed benefits of separating a virtual hand from the real hand configuration. Preventing visual hand-object interpenetration improves subjective grasping experience [1], and maintaining resulting joint angle offsets to lift fingers at release improves performance [2]. Other motivations include improved hand navigation [3] and pseudohaptic effects [4].

Subsequent virtual-real discrepancy reduction (convergence that removes the offset) is important to avoid artifacts, including unreasonable virtual hand configuration [2]. Researchers suggest convergence algorithms (Section 2) but did not study convergence motion effects adequately to understand finger motions and speed effects. Very slow speed maintains discrepancy for too long and causes grasp problems [2]. Fast speed can negate discrepancy benefits so that virtual fingers stick to object surfaces during release in case of substantial real finger penetration [2]. Visual-proprioceptive motion discrepancy [3] may affect user experience as converging virtual finger motion differs from real motion.

Understanding convergence effects is important for optimizing interactions, e.g., finding convergence speeds for good grasp and release performance or user-preferred finger motions. We present a grasping-focused study on convergence motion effects and finger convergence speed. It investigates performance effects of speed, a possible relationship to user motion speeds, and a subjective user adjustment of the related motion discrepancy.

2 Related Work: Virtual Hand Management

There are two base metaphors for virtual hand management after separation between virtual and real hands [5]. A rubber band metaphor immediately pulls the virtual hand to the real hand. It suffers from sticking between virtual fingers and objects as noted above, and a similar effect has been seen in non-grasping hand interactions [3][6]. An incremental motion metaphor maintains offsets, moving virtual state by the same amount as real motion. It mitigates sticking, but it can produce unreasonable hand states, grasp problems [2], and reduced pointing performance [3].

We addressed this with a release method [2] using incremental motion with a convergence mechanism for joint-angle offset reduction. It used fixed convergence speed: each angle offset was reduced a constant amount per time step. Several recent grasping methods (e.g., [7][8][9]) use rubber-band-like behavior and may improve with a similar release method. Our new study helps explain the difference and allows such methods to be optimized.

Incremental motion with convergence was previously shown in MACBETH for managing hand base position [3]. MACBETH adapted convergence speed based on position discrepancy magnitude and user motion speed to balance between user detection of position and motion discrepancies. However, without considering performance effects further, it may reach a convergence speed where significant performance drops occur. In contrast, we consider this performance drop, in addition to focusing on grasping effects rather than hand base position.

3 Method: Grasping Model and Convergence

Following a basic spring-based approach [10], we couple a simulation-driven virtual hand to target (real) hand configuration with virtual linear and torsional springs. This drives the virtual hand to follow the real hand, while collision response provides grasped object motion and hand constraints. Upon heuristically-detected release (the user opens the thumb and another digit while virtual counterparts touch the object) [2], the grasping model triggers the release mechanism mentioned above.

The basic incremental component works as follows. Per tracker reading, each virtual hand joint has a target angle updated by the rotation amount sensed for the corresponding real joint. Also, at the instant of heuristically-detected release, the target value is set directly to the virtual hand joint angle.

The convergence component, per simulation step, updates each joint's target angle by adding or subtracting constant convergence amount c to reduce offset between virtual and real joints (clamped to avoid overshoot). Parameter c is convergence speed in degs per simulation step (convert to degs per ms using simulation rate, averaging 1863 Hz). Upon full convergence (no remaining offset), hand behavior returns to basic spring-based grasping behavior until the release heuristic triggers again.

Parameter c gives a continuum of virtual hand behaviors. At value zero, grasp release matches pure incremental motion, for which a maintained offset grows with multiple grasps. For very high (fast) values of c, virtual grasp matches the rubber band approach. Ideally, there would be values avoiding all problems.

Implementation Notes: We use a standard hand model [2] with 4 dof per digit and 16 total hand segments. Nvidia PhysX provides physical simulation with collision response. PhysX revolute joints provide torsional springs for finger rotations. Hand base (palm) springs are implemented as in [10]. Collision shapes can overlap slightly (~0.6 cm) to improve contact simulation. To avoid visual interpenetration, hand collision shape is correspondingly larger than visual shape. Default convergence speed is 0.045° per ms.

*e-mail: mores_p@hotmail.com
†e-mail: cwborst@gmail.com

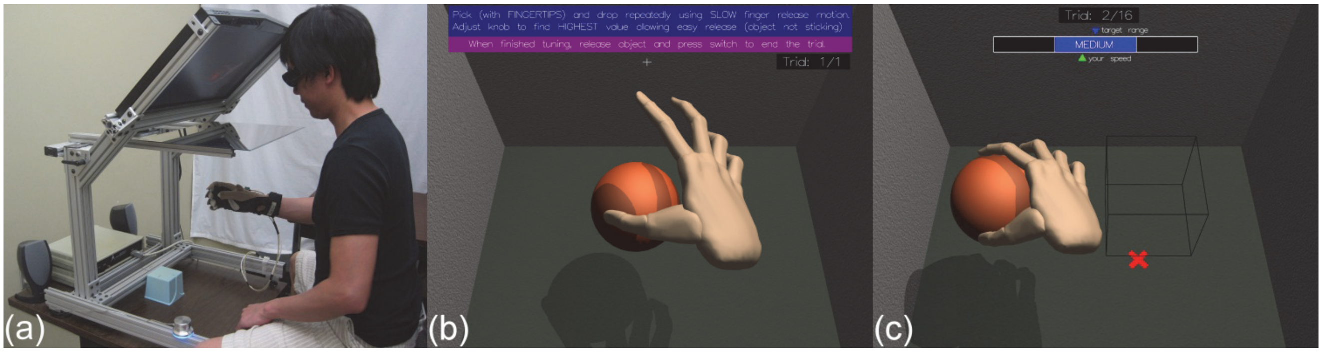

Figure 1: Experiment environment: (a) apparatus, (b) convergence tuning experiment, (c) targeted ball-drop experiment.

4 Experiment

4.1 Initial Study

In an initial study, users gave subjectively-good *c* ranges and reported any artifacts while handling virtual balls. Details are in [2], except we added more subjects later (20 vs. 12). Results suggested tuning *c* should focus on large objects first: small objects had a wider good *c* range. Motion speed discrepancy was only reported by one subject, so other artifacts should be prioritized. These were hand state and (re)grasping problems at very low *c* and sticking during release at high *c*. Hand base motion or selective attention (focus on object) may hide finger motion artifacts. Subjects seemed to favor low *c*. Subjectively-best *c* was around 0.054° per ms.

4.2 Design, Apparatus, and Subjects

The new experiments had within-subjects designs. A targeted ball-drop experiment measured release performance to understand convergence speed effects objectively and to investigate a possible relationship of release motion speeds to convergence effects. The independent variables were:

1. *Convergence Speed* (*c* from Section 3) – very slow (0.005° per ms), slow (0.090°), fast (0.283°), and very fast (0.584°).
2. *Release Speed* – a user finger motion target range: slow (≤ 0.33 mm per ms), medium ([0.28, 0.73]), fast (≥ 0.68).

The *c* values were spaced by a square function in a range taken from tuning results for a large ball from the initial study (25% and 75% quartile boundaries). Release speed ranges were based on 3 pilot subjects. Experiments used the large ball size (6 cm radius) and only used a ball shape because we expect similar effect patterns across object types as in past experiences [2] and we consider release speeds to be more interesting.

A convergence tuning experiment investigated a possible relationship between the release motion speeds and convergence effects subjectively by asking subjects to adjust *c* to the highest value allowing easy release (no sticking). Adjustment range included pure-incremental to rubber-band-like behaviors.

A motion perception experiment considered user responses to visual-proprioceptive motion discrepancy and had subjects adjust a thumb-coupling spring stiffness to balance between thumb and finger motion. Increased stiffness increases thumb force on objects, reducing thumb penetration and convergence motion while increasing these for other digits [11]. This may provide subjective improvement of motion discrepancy. Subjects were asked their preference and expectations. In contrast to the initial study, we focused subjects specifically on convergence motion. This also differs from [11], which focused on overall quality, did not consider convergence speeds, and considered grasp variations.

Apparatus and Software Notes: We used an Acer GD235HZ LCD with 1920x1080 stereo images, Nvidia 3D Vision glasses (no head tracking), a miniBird 500 for palm tracking, an 18-sensor CyberGlove (distal joint angles were set to two thirds of middle knuckle angles), and a PC with speakers. See Figure 1a.

Subjects adjusted *c* and stiffness using a PowerMate knob without stops or reference points. It varied *c* in the range [0.0°, 1.863°] per ms and a stiffness multiplier in the range [0.5, 32.0], in 100 increments, spaced by a square function for finer control at smaller values as in [2]. The stiffness multiplier scaled spring and damping constants for each of the 4 torsion springs of the thumb.

A release speed enforcement monitored changes in thumb-to-index tip distance. Subject release speed was measured based on the maximum change during any 100-ms period in a time interval beginning 1000 ms before release end (the instant when no virtual phalanges touch the object) and ending 100 ms after release end.

Subjects: 20 subjects participated: 17 males and 3 females, aged 22 to 36 years (average 27), all right-handed. All were students: 19 from computer science or computer engineering.

4.3 Procedure

Subjects performed tasks in the presented order below. We calibrated the CyberGlove per subject before starting.

Learning Task: Subjects practiced grasping and release speeds in a ball-drop environment as in Figure 1b. There were 3 trials: one per release speed. A graphical bar (Figure 1c) showed the target range and the latest speed. When speed was out of range, a 2-s warning message stated "TOO SLOW" or "TOO FAST".

Targeted Ball-Drop: This experiment included three trial sets: one per release speed, in random order. In each set, subjects first practiced a task four times: once with on-screen instructions and three times without. The task required subjects to pick up the ball from the virtual floor, move it into a wireframe cube above an X mark on the floor (the cube was centered above the X), and release it on expiration of an audible countdown timer, using release speed according to trial condition (Figure 1c). The cube switched between black and bright green as the ball was centered in it. The color switched at a threshold distance of 1.5 cm. This also triggered the 2-s countdown timer (tick-tick-beep). Ball center was required to stay within 1.75 cm of the cube center during the countdown or the trial restarted. The trial also restarted if the release speed target was missed or release was premature (between pickup and beep sound). A graphical speed bar and warning message were used as in the learning task. Per release speed set, after practice trials, there were 16 real trials, i.e., 4 trials for each of the 4 convergence speeds, in random order.

Convergence Tuning: This experiment included three trial sets: one per release speed, in random order. Per set, subjects first practiced a task once to get used to a release speed target and then performed one real trial. Per trial, subjects picked up and dropped the ball repeatedly in the ball-drop environment (Figure 1b) while adjusting c with the instruction "find the HIGHEST value allowing easy release (object not sticking)". The starting c value was 0. The task did not reveal c's value, except "MIN", "MAX", and increase or decrease indicated by a + or –. Subjects indicated completion with a glove-mounted switch. Warning messages for off-target release speeds were shown, but the speed bar was not shown, to minimize distraction away from hand visuals.

Motion Perception: Here, subjects practiced a task once and then performed 3 real trials, one per convergence speed, in random order. Trials were similar to convergence tuning with the following differences: the tuned parameter was thumb spring stiffness and started at a random value, and instructions were "tune for PREFERRED hand motion after release" and "use NORMAL release motion" (no speed enforcement). Additionally, the experimenter asked subjects to stop the real hand after release and watch virtual finger motions before starting a new grasp. At the end of the experiment, subjects explained to the experimenter how they tuned and what their preferences were.

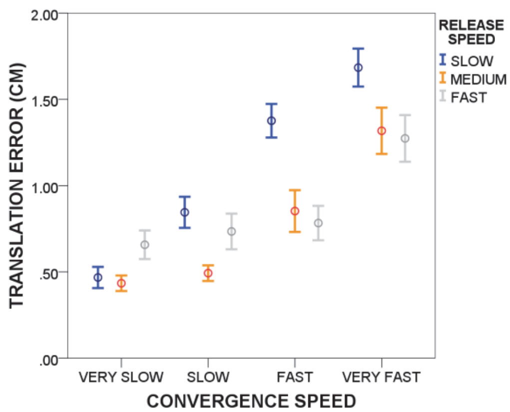

Figure 2: Translation error for the targeted ball-drop experiment (mean and standard error bars).

4.4 Results and Discussion

4.4.1 Results and Discussion for Targeted Ball-Drop

The dependent variables are:

1. *Release time*: time between countdown timer end (release start) and release end (when no virtual fingers touch the ball; note we checked for, and detected no, regrasps).
2. *Translation error*: horizontal magnitude (parallel to floor) of ball translation from release start to floor contact time.
3. *Convergence time*: time required for convergence to real hand configuration, based on c and max offset at release end.

Figures 2 and 3 show results. Reported statistics use 2-way repeated-measures ANOVA with Bonferroni-corrected followups.

Translation error: Increasing convergence speeds reduced release accuracy throughout the tested range ($F(3,57) = 69.81$, $p < .001$; all pairwise comparisons: $p < .001$). This is from reduced release motion of virtual fingers, for increasing convergence speeds, making the virtual hand behave more like the rubber-band metaphor and increasing sticking.

Overall, faster release speed improved release accuracy over slow release ($F(2,38) = 7.50$, $p = .002$; medium vs. slow: $p = .006$; fast vs. slow: $p = .041$; fast vs. medium: not significant $p = .884$). This may be from increased release motion of virtual fingers at faster release speeds, reducing sticking.

There was interaction between release speed and convergence speed effects ($F(6,114) = 6.33$, $p < .001$). Faster release allowed higher convergence speeds before notable accuracy loss. This is because larger and faster finger release motions avoid sticking problems except at higher convergence speeds. However, fast release may cause more palm motion, contributing some error.

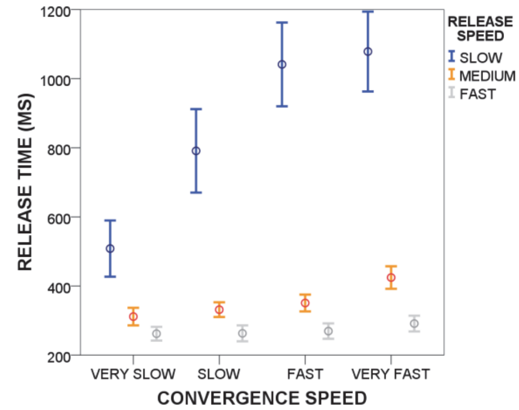

Figure 3: Release time for the targeted ball-drop experiment.

Release time: Increasing convergence speeds increased release time overall ($F(3,57) = 30.87$, $p < .001$; slow vs. very-slow: $p = .008$; very-fast vs. fast: not-significant $p = .404$; all other pairs: $p < .001$). Increasing release speeds reduced release time (as is straightforward). Faster release allowed higher convergence speeds before notable performance drops (significant interaction: $F(6,114) = 19.23$, $p < .001$).

Convergence time: Convergence times for convergence speeds from very-slow to very-fast were 4209 ms, 159 ms, 20 ms, and 2 ms. The large very-slow time can result in excessive offset accumulation, resulting in unreasonable virtual hand state and (re)grasping problems.

Convergence times increased for release speeds from slow to fast: 975 ms, 1128 ms, 1189 ms ($F(2,38) = 10.63$, $p < .001$; medium vs. slow: $p = .006$; fast vs. slow: $p = .006$; fast vs. medium: not significant $p = .373$). This may be from less sticking in faster release resulting in more joint offsets at release end for incremental-motion-like behavior. We note reduced release-speed effects for increasing convergence speeds (significant interaction: $F(6,114) = 7.23$, $p < .001$).

Summary for convergence speed: There is a tradeoff between release performance drops for increasing speeds and problematic convergence times for very low speeds, supporting subjective effects reported in the initial study. Faster release allows higher convergence speeds before substantial performance drops.

4.4.2 Results and Discussion for Convergence Tuning

Figure 4 shows subject-tuned highest convergence speed for "easy release" per release speed. Results echo the relationship between release speed and convergence effects reported above: faster release allowed higher convergence speeds before release problems. Tuned value increased as release speed increased (slow-medium: $Z = -3.88$, $p < .001$; medium-fast: $Z = -3.30$, $p < .001$, Wilcoxon paired signed-rank tests).

4.4.3 Results and Discussion for Motion Perception

Figure 5 shows the subject-tuned thumb stiffness multiplier for balancing convergence motion (and interpenetration) between the thumb and other digits. Higher stiffness decreases thumb convergence motion while increasing other digit motions; the default value is 1.0. Median tuned values for very-slow, slow, and fast conditions were 6.60, 3.93, and 4.25, notably larger than the default. Based on our experiences, 3.93 and 4.25 show moderately more finger motion than thumb motion, depending on number of fingers involved (more fingers give less per-finger motion). 6.60 has noticeably more finger motions than thumb motion. Subject explanations help explain large variations in tuned values in fast and very-slow conditions. With fast convergence, 10 subjects stated that they did not detect much motion difference from tuning, so we expect they tune indiscriminately. In the very-slow condition, 3 subjects tuned randomly: two were confused by the motions, and the other didn't like any motion at all. Overall, 6 other subjects made statements suggesting sensitivity to thumb motion over finger motion, and we believe they tune high to reduce thumb motion. Note that [11] found even higher stiffness values for best overall subjective quality and a value around 2.1 for best overall release accuracy.

Subjects reported various preferences. Some comments reflected what subjects wanted to see without matching actual available behaviors. 10 subjects preferred no motions for all digits. 5 others preferred small motions for all digits. 2 others preferred small thumb motion and no finger motions, while 2 others preferred the opposite (no thumb motion and small finger motions). 1 other stated "not sure". 2 of the 5 that preferred small motions for all digits referenced slight inward motions from real world grasping: digits may close somewhat after real-world release. One of these two stated preferring slightly more finger motion than thumb motion. Comments may add some insight to the initial study results about user insensitivity to convergence. Some users may not regard the motions as undesirable artifacts.

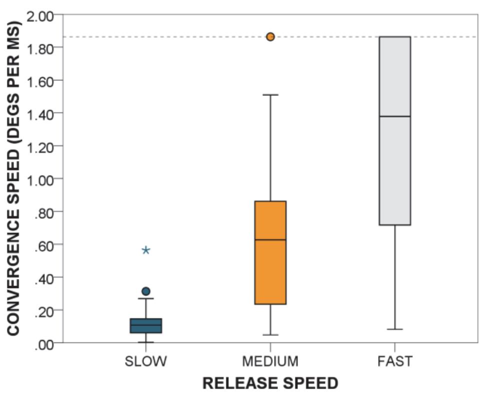

Figure 4: Tuned highest convergence speeds for "easy release" (box-and-whiskers plots, circles and asterisks denote outliers). Dotted line indicates maximum allowed value (1.863° per ms).

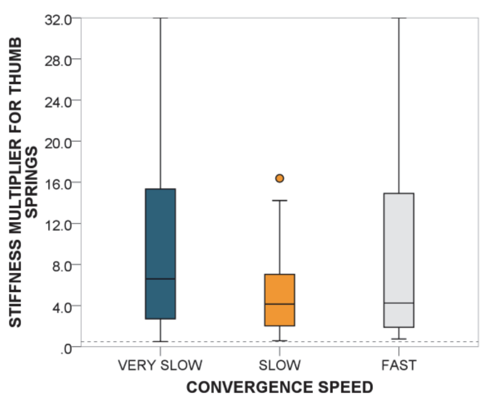

Figure 5: Tuned stiffness values for "preferred hand motion after release". Dotted line indicates minimum allowed value (0.5).

5 Conclusion and Implications

We described benefits of separating a virtual hand from the real hand configuration and why subsequent convergence is important. We conducted the first grasping-focused studies to understand convergence motion effects and to optimize finger convergence speed. The new study provides understanding of performance effects, a relationship to user motion speeds, and additional insight into motion discrepancy, while the initial study focused on subjective effects. In contrast to Burns's principles [3], the studies suggest user detection of motion discrepancy appears secondary, and convergence speed should first avoid sticking during release and effects of excessive convergence time. New results, in combination with initial results, show that the default speed of 0.045° per ms avoids these artifacts and also provides good objective performance overall. This speed's convergence time, around 371 ms for normal release [2], is short enough to avoid grasp problems at normal speeds. It is close to the subjectively-best convergence speed from the initial study. It is also known to give better objective and subjective performance than spring-based grasping without the release mechanism [2]. Overall, we expect the default speed to be near optimal for general interaction.

We suggest adaptive convergence to further optimize behaviors by basing convergence speed on detected release speed. For example, the default speed may be low for fast regrasp sequences (e.g., multiple grasps in quick succession). We found that fast release increases convergence time. Per release speed, convergence speed could be set considering the highest value below notable performance drops. High convergence speed seems to reduce detectable motion artifacts. Future work can focus on an adaptive approach guided by our data.

Overall, the studies suggest users are usually insensitive to the motion discrepancy. The new study suggests remaining motion discrepancy can be improved by increasing thumb spring stiffness relative to other digits to address sensitivity to thumb motion and to provide digit motions matching user preference. Adjustment can also consider stiffness effects on grasped object position, release accuracy, and overall subjective quality, as studied in [11].

References

[1] M. Prachyabrued and C. W. Borst. Visual Interpenetration Tradeoffs in Whole-Hand Virtual Grasping. In *IEEE 3DUI*, pp. 39-42, 2012.

[2] M. Prachyabrued and C. W. Borst. Virtual Grasp Release Method and Evaluation. *International Journal of Human-Computer Studies*, 70(11):828-848, 2012.

[3] E. Burns. *MACBETH: Management of Avatar Conflict by Employment of a Technique Hybrid*. PhD Dissertation, University of North Carolina at Chapel Hill, 2006.

[4] A. Pusch, O. Martin, and S. Coquillart. HEMP – Hand-Displacement-Based Pseudo-Haptics: A Study of a Force Field Application. In *IEEE 3DUI*, pp. 59-66, 2008.

[5] G. Zachmann and A. Rettig. Natural and Robust Interaction in Virtual Assembly Simulation. In *ISPE International Conference on Concurrent Engineering: Research and Applications*, 2001.

[6] R. W. Lindeman, J. L. Sibert, and J. N. Templeman. The Effect of 3D Widget Representation and Simulated Surface Constraints on Interaction in Virtual Environments. In *IEEE VR*, pp. 141-148, 2001.

[7] J. Jacobs, M. Stengel, and B. Froehlich. A Generalized God-Object Method for Plausible Finger-Based Interactions in Virtual Environments. In *IEEE 3DUI*, pp. 43-51, 2012.

[8] O. Hilliges, D. Kim, S. Izadi, M. Weiss, and A. D. Wilson. Holodesk: Direct 3D Interactions with a Situated See-Through Display. In *ACM CHI*, pp. 2421-2430, 2012.

[9] M. Moehring and B. Froehlich. Enabling Functional Validation of Virtual Cars through Natural Interaction Metaphors. In *IEEE VR*, pp. 27-34, 2010.

[10] C. W. Borst and A. P. Indugula. Realistic Virtual Grasping. In *IEEE VR*, pp. 91-98, 2005.

[11] C. W. Borst and M. Prachyabrued. Nonuniform and Adaptive Coupling Stiffness for Virtual Grasping. In *IEEE VR*, 2013.