Physical Guides: An Analysis of 3D Sketching Performance on Physical Objects in Augmented Reality

Philipp Wacker
RWTH Aachen University
52056 Aachen, Germany
wacker@cs.rwth-aachen.de

Adrian Wagner
RWTH Aachen University
52056 Aachen, Germany
adrian.wagner@rwth-aachen.de

Simon Voelker
RWTH Aachen University
52056 Aachen, Germany
voelker@cs.rwth-aachen.de

Jan Borchers
RWTH Aachen University
52056 Aachen, Germany
borchers@cs.rwth-aachen.de

Abstract

Augmented Reality (AR) lets users sketch 3D designs directly attached to existing physical objects. These objects provide natural haptic feedback whenever the pen touches them, and, unlike in VR, there is no need to digitize the physical object first. Especially in Personal Fabrication, this lets non-professional designers quickly create simple 3D models that fit existing physical objects. We studied how accurately visual lines and concave/convex surfaces let users draw 3D shapes attached to physical vs. virtual objects in AR. Results show that tracing physical objects is 48% more accurate, but takes longer than tracing virtual objects. Concave physical edges further improve accuracy due to their haptic guidance. Our findings provide initial metrics when designing AR sketching systems.

Author Keywords

Augmented reality; 3D sketching; tracing; motor ability; physical objects; guide classification.

ACM Classification Keywords

H.5.1. [Multimedia Information Systems]: Artificial, augmented, and virtual realities

Introduction

The recent advancement of Virtual Reality (VR) and Augmented Reality (AR) technology has rekindled an interest

---

CHI'18 Extended Abstracts, April 21–26, 2018, Montreal, QC, Canada
ACM 978-1-4503-5621-3/18/04.
https://doi.org/10.1145/3170427.3188493

Constraint	Virtual	Physical
No Guide	Virtual table surface	Real table surface
Visual	Pen stroke on a virtual sketch	Waterline in a bottle
Concave	Inside of a virtual bucket	Intersection of shelf & wall
Convex	Edge of a virtual desk	Opening of a wine glass

Figure 1: Examples for all combinations of initial constraint & subsequent limitation to a line.

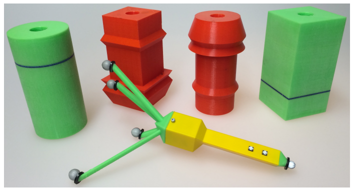

Figure 2: Front: Custom pen used in the study. Back: Physical objects used in the study. For each *shape* condition, we had two physical objects to reflect the *no guide* and *visual* condition as well as the *convex* and *concave* condition.

in using these techniques in 3D design tasks. Numerous research projects [2, 5, 7] and products like Tiltbrush[1] or Gravity Sketch VR[2] focus on sketching in mid-air to create and view 3D models in both VR and AR.

However, only AR enables sketching directly on existing physical objects in the real world, whether to extend them or to design a new, matching object. This is particularly helpful when designing an object that has to fit an existing object, such as a snap-on handle for a drinks container. Being able to achieve this quickly and easily without extensive knowledge of professional modeling tools opens up 3D object design to novices. In VR, the user would first need to digitize the existing model. In AR, she can use the physical object directly instead. This also adds the benefit of haptic feedback from the physical object surfaces as guidance during the modeling task.

Sketching planar shapes in VR is more accurate when a flat physical surface is provided [1]. However, real objects, like a water bottle, have more complex surface shapes and offer a variety of guidance elements both visual (such as a printed line) and haptic (such as curves and edges). To explore this space of designing objects that can fit existing objects using AR, we provide a first classification of the types of guidance that physical and virtual objects offer, and study their impact on 3D sketching precision in AR.

Related Work

Immersive Modeling looks at interaction techniques to directly create 3D models in mid-air. This requires tracking input devices in space and aligning the visualization to create the impression of drawing strokes in mid-air [5]. Jackson and Keefe [2] enable users to create models in VR by letting them "lift" individual strokes from 2D sketches and pull out surfaces between these strokes. WireDraw [7] supports creating a physical wireframe of a known 3D model with a 3D extruder pen, by displaying stroke guides in AR.

While performance of freehand sketches improves with practice [6], the missing haptic feedback still affects their precision. For VR, Arora et al. [1] recently showed that displaying visual guidelines already improves the drawing accuracy of freehand strokes, while providing a physical surface as a guide improves it dramatically.

Using **Physical Objects as Guides in Personal Fabrication** requires aligning virtual models and physical objects, for example when designing an object to 3D-print that should fit around or inside an existing object.

Zhu et al. [8] use physical objects such as pens during the 3D printing process to create exact cutouts on printed objects. MixFab [4] lets users place small physical objects behind a see-through display and create virtual models aligned to them, e.g., to cut holes in the virtual object.

Thus, while researchers have explored sketching in AR from several directions, there has been no quantitative study of how different types of physical and virtual guidance affect the precision when users sketch in AR.

Classification of Guidance Types

To structure our study, we first introduce a classification of the types of guidance that objects may offer. These guides can assist a user in creating strokes that are aligned to the object surface by tracing them with a pen. While many factors affect drawing on physical objects, from material properties such as hard, soft, rough or smooth, to object size, location, and orientation, we focus on local object features that may guide a stroke.

---

[1] www.tiltbrush.com

[2] www.gravitysketch.com

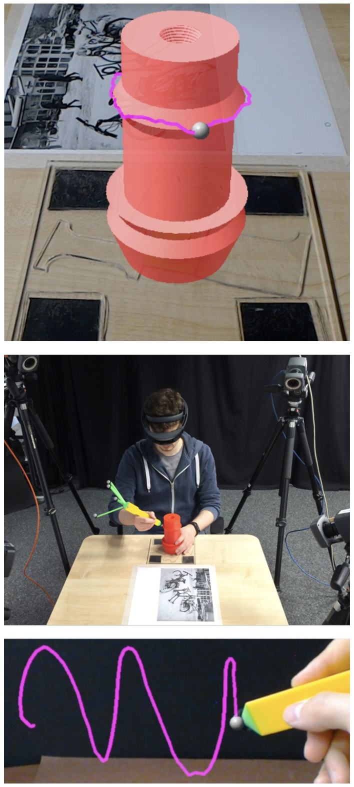

Figure 3: After aligning the coordinate systems, points reported by the VICON can be rendered into the viewing area of the HoloLens (bottom). In our study, participants were asked to draw a stroke around virtual objects (top) and physical objects (middle).

The guidance types on an object can be seen as limiting the degrees of freedom for sketching. The first limitation guides the free-hand movement to a surface by constraining one degree of freedom (*surface guidance*). On *physical* objects, this constraint is hard for one half degree of freedom (can't push inside the object) and soft for the other half (can lift pen off surface). On *virtual* objects, it is entirely a soft constraint. Movement can be limited further to tracing a 1D line or curve by *visual* guides and object shape (*line guidance*). Curvature increases guidance, with *concave* and *convex* edges as extremes. *Concave* shapes provide more guidance, reducing the degree of freedom more than *convex* shapes. This results in 8 combinations (Fig. 1).

Tracing Study

Our study was designed to quantify the effect of different guidance types on tracing time and accuracy. This knowledge can be useful to determine the necessary strengths of smoothing algorithms in different guidance situations. For this, we chose the basic task of drawing on the surfaces of fixed, upright objects. We chose a cylinder and a cuboid as our basic object types, around which the user had to draw a *circle* and *square*, respectively. We measured accuracy by comparing the trace to the optimal shape.

Experimental Design

For each object (cylinder and cuboid), we included both *physical* and *virtual* models as conditions to reflect the *surface guidance* described before. Our guidance types led to four conditions: *no guide* as baseline, *visual*, *concave*, and *convex*. For this study, we used edges as the most common and pronounced *convex* and *concave* surface features.

The result is a 2×2×4 design: *shape* (*circle* / *square*) × *surface guidance* (*physical* / *virtual*) × *line guidance* (*no guide* / *visual* / *concave* / *convex*). The sequence of all 16 conditions was counterbalanced with a Latin square. We recruited 16 participants (4 female, 19–29 years, M=24.9 years, SD=2.4 years, all able-bodied). No one had experience with AR drawing tools. Each performed 80 strokes, five strokes in each condition. They could practice each condition before the five strokes. A session took 45–60 min.

Apparatus

Our 3D-printed physical objects for tracing were 16 cm high, 8 cm wide, and could be mounted normally or upside down. A different *line guidance* was placed at 4 cm from each end. This let us combine the *no guide* and *visual* conditions on each green object, and the *convex* and *concave* conditions on each red object (Fig. 2, back). For each condition, the object was mounted on a plate in front of the participant. The plate was attached to the table to avoid movements.

Pen Design, Tracking & Rendering

We created a custom pen (Fig. 2, front) tracked by a VICON motion tracking system that used 6 high-speed infrared cameras to track reflective markers from different angles at 100 fps with sub-millimeter accuracy. Similar to other studies [1, 2], our pen featured markers at the end and tip. Using a spherical marker as the tip both improved tip tracking stability and prevented user confusion when the alignment of the virtual line rendering drifted slightly around the center of the physical tip. The pen included two buttons for inking and calibration, and a Bluetooth LE module to send their states to a server. While the inking button was pressed, the software recorded the pen tip position as a path, and forwarded it to a Microsoft HoloLens headset, which rendered the path into the user's view at 30 fps.

To roughly align the VICON and HoloLens coordinate systems, we used the VICON calibration wand and a visual marker that the HoloLens could track. Afterwards, we fine-aligned the coordinate systems manually by adjusting rota-

Condition	Significance	Mean	SD
Physical	A	4.93	1.98
Virtual	B	9.48	4.70
Convex	A	6.70	4.07
Visual	A	6.91	3.97
Concave	A	7.16	4.69
No Guide	B	8.09	4.18
Physical,	A	4.11	1.52
Physical, Visual	B	4.74	1.44
Physical, Convex	B	4.97	1.79
Physical, No	C	5.90	2.54
Virtual, Convex	D	8.48	4.90
Virtual, Visual	D	9.00	4.48
Virtual, Concave	E	10.13	4.83
Virtual, No Guide	E	10.25	4.35

Figure 4: Means and standard deviations of *mean deviation in 3D* for main effects and interactions (in mm, rounded to two decimal places). Rows represent conditions, horizontal lines separate main effects and interactions. Rows not connected by the same letter are significantly different.

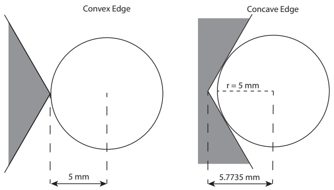

Figure 5: Optimizations made to the target stroke for the *physical* conditions to correct the inability to draw the perfect stroke.

tion and location of a digital model to fit a real-world counterpart placed at a known point.

Based on pilot tests, we occluded those parts of the path behind the model in both the physical and the virtual condition to preserve realism (Fig. 3, top).

Study Procedure

Participants sat at a table inside the VICON's tracking volume (Fig. 3, center). They were allowed to move their head and torso, but asked to remain seated. Each trial, we showed the object to trace around. In *physical* conditions, we screwed the object onto the mounting plate. We asked participants to rest their non-dominant hand on the mounting plate.

We asked our participants to draw around the object at normal speed, while keeping precision in mind. In the *virtual* conditions, participants were allowed to move their hand and the pen through the virtual object.

During the implementation of the system, we observed that the HoloLens occasionally adjusts its coordinate system due to updated environment tracking information. This caused the VICON and HoloLens coordinate systems to become misaligned, and the real and rendered pen tip would deviate from each other. We asked participants to mention any offset to us during the study (happened twice) and also inquired about the correct alignment throughout the session. In case of a misalignment, we re-synchronized the coordinate systems before repeating the last trial.

Measurements

Similar to Arora et al. [1], we processed the data for each stroke with a low-pass filter averaging over a 10 frame window, created a path from the resulting points, and then resampled the path to 100 equidistant points for evaluation. This removed a potential bias due to the higher point count in areas such as corners where participants slowed down. We compared these resampled points to the optimal stroke to calculate *Mean Deviation in 3D* by averaging the shortest distance from each point to the target shape. This represents the mean deviation from the target shape [1].

Since the line drawn originated from the center of the spherical marker at the pen tip, the marker displaced the user's input on physical surfaces by its radius of 5 mm, making it impossible to perfectly trace physical surfaces. To account for this, we virtually enlarged all physical target shapes by that radius for our calculations. The *physical*, *concave* condition required another enlargement to a total of 5.77 mm since the radius of the pen tip marker was larger than the concave opening, keeping it further away from the concave edge (Fig. 5). These adjustments were not necessary for the *virtual* conditions since the user could penetrate object surfaces to align her strokes. We also recorded the *Stroke Duration* for each inking operation.

Results

Due to measurement issues, we discarded 20 strokes of the 1280 recorded. We performed mixed-effect ANOVAs with the user as a random variable on the log-transformed measurements to analyze the effects of our experimental conditions. All post-hoc pairwise comparisons were performed using Tukey HSD tests.

The evaluation of *Mean Deviation in 3D* indicates that both guidance conditions (*surface guidance*: $F_{1,1229} = 910.802$, $p < .001$; *line guidance*: $F_{3,1229} = 20.511, p < .001$) as well as their interaction $(F_{3,1229} = 20.584, p < .001)$ had a significant effect. The means and results of the post-hoc tests can be seen in Table 4. *Physical* objects were traced 48% more accurate than *virtual* objects with *physical*, *concave* showing the least deviation and *virtual*, *no guide* as well as *virtual*, *concave* showing the most.

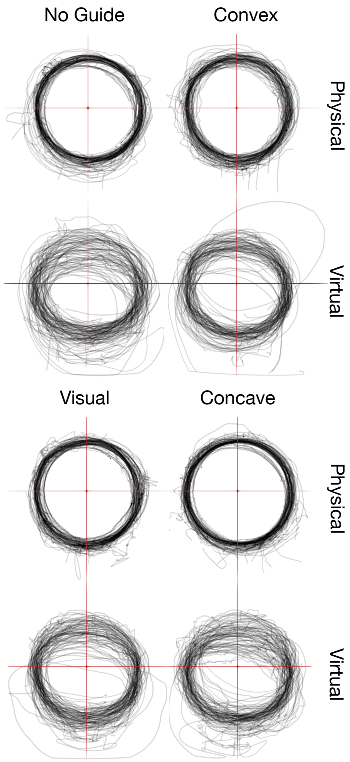

Figure 6: Top rendering of the interaction *surface guidance*×*line guidance*. The bottom half of each condition represents the front view. *Physical* strokes show less deviation than *virtual* strokes.

For the *Stroke Duration* we report two interesting findings. Participants took longer to draw around a *physical* ($M = 9.05s$, $SD = 4.97$) compared to a *virtual* object ($M = 8.8s$, $SD = 6.13$) ($F_{1,1229} = 11.253, p < .001$). The *physical*, *concave* object ($M = 7.96s$, $SD = 3.73$) had the lowest average duration and was significantly faster compared to the *physical*, *convex* object which was the slowest ($M = 10.08s$, $SD = 5.78$).

Front-to-Back Comparison

Studying aggregate renderings of the strokes performed, as in Fig. 6, we noticed a pattern that motivated us to compare accuracy on the front- and back-facing halves of each object, to evaluate how well users could continue strokes they could no longer see. We split recorded points into a front-facing and a back-facing half and computed the mean 3D deviation for both halves. We evaluated all interactions that involved the new variable *side* and the main effect.

We found an effect of *side* ($F_{1,2422} = 13.711, p < .001$). The means shows that *front* ($M = 6.84mm$, $SD = 4.32mm$) deviated less than *back* ($M = 7.27mm$, $SD = 4.14mm$). However, there were no interaction effects.

Another intriguing observation was that the *virtual square* shapes showed a slight counter-clockwise rotation (Fig. 8: diagonal lines and shifted corners especially on the back side of the virtual condition). Their sides were also traced more accurately than their front and back. Both observations appeared in all *virtual*, but no *physical* conditions.

Discussion

Our results indicate that both *surface guidance* and *line guidance* affect the performance of drawing on objects. In particular, *physical* objects improve drawing accuracy showing that the hard constraint of a surface supports the user more than its soft “lift-off” constraint. However, it took participants slightly longer to draw around *physical* than *virtual* objects, possibly because of the friction on the physical surface. Another possibility is that drawing on a *physical* object required hand movements that could be saved by moving through the *virtual* models.

The combination of *surface guidance* and *line guidance* is particularly interesting. *Physical*, *concave* guides lead to the most accurate traces, as one would expect, since they provide the strongest constraints. Stroke duration with these guides was also the fastest on average. However, the *virtual*, *concave* condition performed generally poorly and has the most deviations together with *virtual*, *no guide*. This might be due to the added difficulty of determining the correct depth ‘inside’ an object. Since the virtual pen tip got occluded behind an object, its absence while aligning the pen with the *concave* edge might also affect accuracy.

While our results show differences in the mm range, these already matter for certain modeling tasks, and previous studies indicate that such differences increase with the size of the target shape [1].

After our main evaluation, we also compared the performance on the *front* and *back* of the object. We found that participants generally deviated more on the *back* of the object. This is most likely due to the missing visual indication of where to trace, so that the user has to rely on physical guides and motor memory to continue the trace. Looking at the top-view renderings suggests that participants deviated more into the object on the back of the object, while drawing outside the object in front. This suggests further studies of these effects and how different path and guide visualizations affect tracing performance.

The qualitative observations about a higher accuracy on the sides of a *square* in the *virtual* condition might be ex-

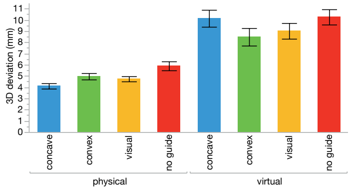

Figure 7: Interaction effect of *surface guidance*×*line guidance* on *mean deviation in 3D*, with 95% Confidence Intervals (CI). The *physical*, *concave* condition has the least deviation (significant). *Physical*, *visual* and *convex* are also significantly different from the other conditions.

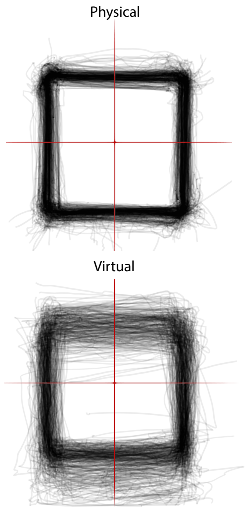

Figure 8: Top rendering of the accumulated strokes for the *square* separated by *surface guidance*. The bottom half of each condition represents the front view. *Virtual* strokes appear slightly rotated counter-clockwise and more spread out at the front and back than the sides.

plained by the issues humans have with detecting depth in virtual environments [3]. The sides present a clearer edge, so that participants could judge the correct position more easily. The front and back do not present such a clear border. In combination with overshoots when approaching a corner, this can also explain the ‘rotation’ of squares in the *virtual* condition for counter-clockwise strokes (75.6% of strokes performed). An intriguing question for future studies is whether the same phenomenon occurs when drawing in clockwise direction. Participants’ handedness is also a likely factor in this.

Conclusion & Future Work

Immersive modeling is a promising approach to make generating simple 3D models more accessible to nonprofessionals. This is especially beneficial in the area of Personal Fabrication, when designing models that should attach to or align with existing objects.

In our lab study, we quantified the effects different guidance types have on both accuracy and time to complete a stroke when drawing around an object. We found that users’ accuracy was highest for the *physical* conditions, in particular the *physical*, *concave* condition and lowest for *virtual*, *concave* and *virtual*, *no guide* objects. Furthermore, participants drew more slowly on *physical* than *virtual* objects, with *physical convex* objects being the slowest.

Our study covers only a small area in the design space of drawing interactions with physical and digital objects in Augmented Reality. We focused on a single-handed task of drawing around a static object. In further studies, we want to look at bi-manual tracing and the effect of different object orientations, sizes, surface structures, materials, and pen styles on drawing performance.

REFERENCES

1. Rahul Arora, Rubaiat Habib Kazi, Fraser Anderson, Tovi Grossman, Karan Singh, and George Fitzmaurice. 2017. Experimental Evaluation of Sketching on Surfaces in VR. In *Proc. CHI '17*. ACM, 5643–5654.
2. Bret Jackson and Daniel F. Keefe. 2016. Lift-Off: Using Reference Imagery and Freehand Sketching to Create 3D Models in VR. *IEEE Transactions on Visualization and Computer Graphics* 22, 4 (2016), 1442–1451.
3. J. Adam Jones, J. Edward Swan, II, Gurjot Singh, Eric Kolstad, and Stephen R. Ellis. 2008. The Effects of Virtual Reality, Augmented Reality, and Motion Parallax on Egocentric Depth Perception. In *Proc. APGV ’08*. ACM, 9–14.
4. Christian Weichel, Manfred Lau, David Kim, Nicolas Villar, and Hans W. Gellersen. 2014. MixFab: A Mixed-Reality Environment for Personal Fabrication. In *Proc. CHI'14*. ACM, 3855–3864.
5. Gerold Wesche and Hans-Peter Seidel. FreeDrawer - A Free-Form Sketching System on the Responsive Workbench. In *Proc. VRST'01*. ACM, 167–174.
6. Eva Wiese, Johann Habakuk Israel, Achim Meyer, and Sara Bongartz. 2010. Investigating the Learnability of Immersive Free-hand Sketching. In *Proc. SBIM ’10*. Eurographics Association, 135–142.
7. Ya-Ting Yue, Xiaolong Zhang, Yongliang Yang, Gang Ren, Yi-King Choi, and Wenping Wang. 2017. WireDraw: 3D Wire Sculpturing Guided with Mixed Reality. In *Proc. CHI'17*. ACM, 3693–3704.
8. Kening Zhu, Alexandru Dancu, and Shengdong (Shen) Zhao. 2016. FusePrint: A DIY 2.5D Printing Technique Embracing Everyday Artifacts. In *Proc. DIS'16*. ACM, 146–157.