

Tangible Volumetric Brushing in Augmented Reality

Stéphane Gosset*
Université Paris-Saclay,
CNRS, Inria, LRI, France

Mickael Sereno†
Université Paris-Saclay,
CNRS, Inria, LRI, France

Lonni Besançon‡
Linköping University,
Sweden

Tobias Isenberg§
Université Paris-Saclay,
CNRS, Inria, LRI, France

ABSTRACT

We discuss a 3D volume selection technique supported by Augmented Reality (AR) using both a tangible multi-touch device and an AR Head-Mounted Display (AR-HMD). While the HMD stereoscopically renders the dataset, the tablet allows the user to draw a lasso that can be extruded using the tablet’s relative or absolute position and orientation. This interaction is inspired by the original Tangible Brush technique. We aim, with this setup, to understand the implications of a tangible device where its absolute position now has meaning in the user’s output space compared to traditional implementations where only the relative movements are captured.

Index Terms: Human-centered computing—Human computer interaction (HCI)—Interaction paradigms—Mixed / augmented reality; Human-centered computing—Visualization—Visualization application domains—Scientific visualization

1 INTRODUCTION AND BACKGROUND

Augmented Reality (AR) allows users to see, in their real environment, 3D-spatially aligned scientific datasets such as volume data or particle physics simulations. However, AR-based scientific visualization received little attention from researchers [2], even though AR technologies have been available since the late 1990s. With AR head-mounted displays (AR-HMD) becoming increasingly powerful, designers can now render complex visualizations in realtime. While this new computing power is available, it remains unclear how to best interact with such visualizations. We focus in this work on volumetric selection, a fundamental interaction task for visualization.

Compared to Virtual Reality (VR), AR allows users to couple their AR-HMD view with their physical tools such as a workstation or a tablet. We argue that, for users to be efficient with AR-HMDs in their immersive analytics applications, they need another device for interactions. Indeed, as mid-air gestures is tiring and speech recognition is not enough on its own, this complementary device would allow users to interact as they are used today with their datasets. To not reduce the user’s mobility, we use a multi-touch tablet to fit that need. Several researchers already explored the use of multi-touch tablets in VR (e. g., [4]), but this tablet use in VR leads to a different perceived resolution, compared to a normal display of a tablet.

We thus use these devices (Fig. 1) together with an AR-HMD and explore which interaction design works best for volumetric selections. Besançon et al.’s [1] study showed benefits of using a tangible device to select manually a part of a 3D data. Their setup relied on a tangible tablet displaying the visualization in an orthographic projection. Another static screen allowed the user to perceive the virtual scene using two perspective projections. The first one displayed what the tablet would have shown in a perspective view, and the second one displayed a static birds-eye-view of the scene. Besançon et al.’s setup, however, decoupled the users’ interactive space and visualization space. Compared to their solution, we want to understand the benefits and limitations of using a tangible device where its absolute position has a meaning in the users’ environment. This raises multiple concerns that we discuss in this extended abstract.




Figure 1: The system from an external and internal point of view.

First, we want to understand what visuals the screen of the tablet should display. Second, similar to López et al.’s [3] setup, we assume that users should input their lasso on a static view to reduce input noise. Then, even if the tablet’s 3D position is now meaningful to the user, we want to understand what positional mapping we should use for the tablet during the selection—an absolute or a relative one.

2 CONCEPTS

We examine these questions as we develop our interaction technique. As a primary hypothesis, performing a tangible selection in AR should increase the users’ performance because (1) more depth cues are available to the users and (2) their movements now physically match what they see. To perform the selection, a user draws a lasso on the tablet and moves it around. Displaying the dataset in a static orthographic view on the tablet as in Tangible Brush is thus essential to ensure precise input by allowing the users to relate the lasso to the dataset in which they want to select a region.

However, the absolute position and size of the tablet for selections have limitations. First, when a user is manipulating the tablet inside a visualization, the user will no more be able to see its screen because the HMD renders data around it. Second, the dataset, to be visualized efficiently, has a certain size that might be beyond the physical size of the tablet. This leads us to examine a virtual tablet whose size might differ from its real counterpart. Third, to still be visualized comfortably, the visualization is usually vertically close to the user’s eyes, which would force the user to move the tablet with both arms raised, increasing the user’s fatigue during the interaction.

These concerns stated, we propose two modes for the selection. One relies on the absolute position of the tablet. The other uses relative motions, thus decoupling what users physically manipulate from their virtual actions, similar to mouse input and to Tangible Brush’s relative movements. The relative movements may allow users to move the tablet in comfortable positions, to not be disturbed by the visualization surrounding the tablet during the interaction, and might make more sense in the user’s mental model when the need to increase the size of the tablet arises. For the relative mode, questions arise about what gain factors should be applied during the interaction and what should the coordinate system of the tablet be. We considered either to align the tablet’s coordinate system with the user’s environment coordinate system (“Relative-Aligned” mode), or to define a new coordinate system based on the tablet’s normal (“Relative-Full” mode). While we only discuss design concerns here, a study is needed to assess the user’s performance depending on the gain factor being used, the coordinate system axis orientation (whether absolute or defined by the normal axis of the tablet) of the tablet, and whether the absolute position of the tablet performs better than using its relative movements. To study these cases, we implemented a prototype that we describe next.

*e-mail: stephane.gosset@gmail.com

†e-mail: serenomickael@gmail.com

‡e-mail: lonni.besancon@gmail.com

§e-mail: tobias.isenberg@inria.fr




Figure 2: The tablet interface during the interaction.

3 IMPLEMENTATION AND INTERACTION PROCESS

Our system relies on four components: a multi-touch tablet for the tangible interaction, a HoloLens 2 as an AR-HMD, a VICON as an external tracking system, and a server workstation for the communications. While we use an external tracking system to track the tablet, it is worth noting that some existing systems (e.g., the Oculus Quest) are standalone and track their input devices (e.g., remote controllers) without external sensors. The server runs on a Linux workstation and handles all the communications between the components. We designed it to handle multiple HMDs and multiple tablets. We focus, in this project, to only one HMD and one tablet. Once connected to the VICON using the VRPN protocol, the server continuously streams to the multi-touch tablet its virtual position and orientation in the HoloLens' coordinate system via TCP/IP.

The tablet (Fig. 2) acts as a remote controller and initializes all possible actions. It relies on Android's SDK and NDK. It renders, in an OpenGL window, the orthographic projection of the scene with respect to the position of the virtual tablet in the 3D space. A slider allows users to resize the orthographic view and the virtual tablet. To start the interaction, the user presses the button "Start Selection," which magnifies the OpenGL view (i.e., hides the widgets not needed for the interaction). Two buttons allow them to move the tablet. A first one ("Position Button") allows them to place the tablet in space before performing a selection, useful for relative movements where the user needs to define an origin. A second one ("Tangible Button") allows the user to build the selection mesh by extruding the lasso drawn. For this second button, the virtual position and orientation of the tablet are functions of the mode the user chose. In the "Relative-Full" mode, a new coordinate system is defined each time the user presses the "Tangible Button," where we store the orientation of the tablet o_{beg} . There, we consider only orientation deltas between the current orientation of the tablet and the stored orientation o_{beg} , and apply that offset to the orientation defined by the user when using the "Position Button." Updating the view only when users presses the buttons allows them to enter a precise lasso once they found a suitable orientation, position, and scaling. This design also allows them to create as many selection meshes as needed. Indeed, each time the "Tangible Button" is pressed, a new selection mesh is created to not connect end-points that should not be. An "End Selection" button allows users to stop the interaction process, where they can (in)validate their selections.

The HoloLens 2 (Fig. 3) uses Unity 2019 and MRTK v2.2. When the user starts the selection process, the HMD displays a wireframe representing the virtual size, position, and orientation of the virtual tablet, and draws the lasso on this frame. Then, each time the user presses the "Tangible Button," the HMD initializes a mesh to be updated. For each update of the position and orientation of the tablet, the HMD computes the new position of every point of the lasso and connects the previous computed lasso with the new one in the current mesh. Using 3D meshes allows the user to quickly grasp what the system has captured, before (in)validating the selection.




Figure 3: The interface of the HoloLens 2 during and after the selection of the red galaxy. While the 3D mesh is not lighted, users have the motion parallax and stereoscopic perspective as useful depth cues.

4 DISCUSSION AND CONCLUSION

Based on our past experience and our testing of the new design, we discuss the pros and cons we hypothesize for the described system. First, as discussed in Sect. 2, the absolute mode (1) forces users to raise their arms high, (2) hides the tablet by the HMD displaying data between tablet and the user's eyes, and (3) makes the user loose the overall context (i.e., the scene does not fit the user's field of view), which would be essential during the selection.

Second, we argue that users need to interact with relative coordinate systems and not absolute ones. We think that the relative coordinate system should be defined by the screen on which the user focuses. If this screen is the tablet screen, then the "Relative-Full" mode might be preferred, as the user will no longer look at the virtual scene through the HMD. We think, however, that the user's focus will most often be the AR view generated by the HMD because the tablet shows the dataset orthographically projected, while the user needs depth cues during extrusion to comfortably catch the desired regions with the drawn lasso. The HMD provides these cues with its stereoscopic perspective projection that matches the human visual system and also adds motion parallax.

Based on these hypotheses, we think that the "Relative-Aligned" might be the mode that will yield the best performance and comfort in an experiment. Users would see the scene entirely, might mentally understand the difference between size of the real and the virtual tablet, and all the interactions can be done relatively to where the user looks (i.e., toward the visualization). We are in the process to prepare formal user study to test these hypotheses. Moreover, this experiment will also give us hints about how the users perceive and use the screen of the tablet, allowing us to improve the design of the visual representation on the tablet.

In conclusion, we adapted an existing interaction technique [1] to AR environments, allowing users to select subsets of a 3D dataset using a tablet to extrude a volume in space. Compared to Tangible Brush [1], we use an AR-based setup instead of another 2D screen to grasp depth, which we argued improves the user's performance beyond the already established benefits of such manual interaction. Our setting raises multiple research questions based on human perception and comfort that we propose to investigate in future work.

REFERENCES

- [1] L. Besançon, M. Sereno, M. Ammi, L. Yu, and T. Isenberg. Hybrid touch/tangible spatial 3D data selection. *Computer Graphics Forum*, 38(3):553–567, June 2019. doi: 10.1111/cgf.13710
- [2] A. Fonnnet and Y. Prié. Survey of immersive analytics. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–1, 2019. Early Access. doi: 10.1109/TVCG.2019.2929033
- [3] D. López, L. Oehlberg, C. Doger, and T. Isenberg. Towards an understanding of mobile touch navigation in a stereoscopic viewing environment for 3D data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 22(5):1616–1629, May 2016. doi: 10.1109/TVCG.2015.2440233
- [4] H. B. Surale, A. Gupta, M. Hancock, and D. Vogel. TabletInVR: Exploring the design space for using a multi-touch tablet in virtual reality. In *Proc. CHI*, pp. 13:1–13:13. ACM, New York, 2019. doi: 10.1145/3290605.3300243