Multi-Touch Querying on Data Physicalizations in Immersive AR

BRIDGER HERMAN, MAXWELL OMDAL, STEPHANIE ZELLER, CLARA A. RICHTER, FRANCESCA SAMSEL, GREG ABRAM, and DANIEL F. KEEFE

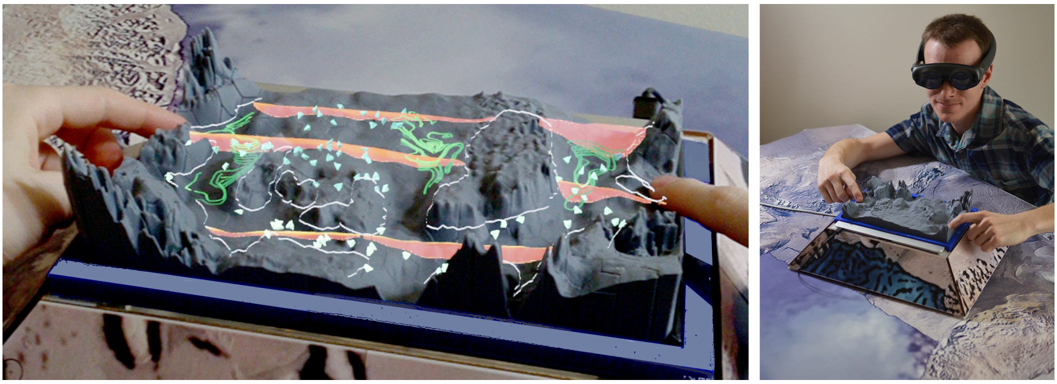

Fig. 1. Adding direct multi-touch input sensing to a 3D physical printout and combining this with a perspective-tracked, stereoscopic AR display creates a hybrid visualization system that benefits from both the high-resolution, passive haptic feedback of the physical data printout and the interactive data exploration made possible by digital displays. Here, multi-touch queries create streamlines and cutting planes to explore a supercomputer simulation of ocean currents and ice melting under Antarctica's Filchner-Ronne ice shelf.

Data physicalizations (3D printed terrain models, anatomical scans, or even abstract data) can naturally engage both the visual and haptic senses in ways that are difficult or impossible to do with traditional planar touch screens and even immersive digital displays. Yet, the rigid 3D physicalizations produced with today's most common 3D printers are fundamentally limited for data exploration and querying tasks that require dynamic input (e.g., touch sensing) and output (e.g., animation), functions that are easily handled with digital displays. We introduce a novel style of hybrid virtual + physical visualization designed specifically to support interactive data exploration tasks. Working toward a "best of both worlds" solution, our approach fuses immersive AR, physical 3D data printouts, and touch sensing *through* the physicalization. We demonstrate that this solution can support three of the most common spatial data querying interactions used in scientific visualization (streamline seeding, dynamic cutting places, and world-in-miniature visualization). Finally, we present quantitative performance data and describe a first application to exploratory visualization of an actively studied supercomputer climate simulation data with feedback from domain scientists.

Authors' addresses: B. Herman , M. Omdal, D. F. Keefe, {herma582, omdal009, dfk}@umn.edu, Department of Computer Science and Engineering, University of Minnesota, Twin Cities, 200 Union St SE, Minneapolis, MN 55455, USA. S. Zeller, G. Abram, F. Samsel, {szeller, gda, fsamsel}@tacc.utexas.edu, Texas Advanced Computing Center, Advanced Computing Building, J.J. Pickle Research Campus, Building 205, 10100 Burnet Rd, Austin, TX 78758, USA. C. A. Richter, richt22c@mtholyoke.edu, Mt. Holyoke College, 50 College St South Hadley, MA 01075, USA.

2573-0142/2021/11-ART497 $15.00
https://doi.org/10.1145/3488542

CCS Concepts: • **Human-centered computing → Interaction devices**; **Visualization systems and tools**; **Scientific visualization**.

Additional Key Words and Phrases: Data physicalization;tangible user interfaces;mixed reality

ACM Reference Format:
Bridger Herman, Maxwell Omdal, Stephanie Zeller, Clara A. Richter, Francesca Samsel, Greg Abram, and Daniel F. Keefe. 2021. Multi-Touch Querying on Data Physicalizations in Immersive AR. *Proc. ACM Hum.-Comput. Interact.* 5, ISS, Article 497 (November 2021), 19 pages. https://doi.org/10.1145/3488542

1 INTRODUCTION

For thousands of years, humans have articulated, transcribed, and catalogued information through tangible and/or haptic channels. Though the term "data physicalization" is relatively new, material properties such as size, form, and surface texture have been used throughout history for communication and education, including in the current digital age [36]. Most modern data physicalizations leverage our natural predisposition for "passive haptic" data exploration: touching, holding, and rotating 3D data representations manually. This is something we cannot yet replicate with digital technologies. Even today's best active-haptic immersive displays (e.g., head-tracked, stereoscopic VR coupled with a glove robotic input/output device to provide touch feedback at the fingertips [22]) are incapable of effectively reproducing the haptic benefits of running multiple fingers over the surface of a data physicalization such as using a finger to mark a point while comparing other areas of interest, or to follow a path on the physical surface [35]. If we wish to leverage our innate proclivity for learning and communicating through physical touch in data visualization (while avoiding an additional layer of input/output devices that distance users from their data [36]), investigating the myriad forms of digital fabrication and hybrid display [29] is essential to visualization research.

Despite numerous successful examples of data physicalization, however, we notice what may be a fundamental limitation with respect to exploratory visualization. While recent physicalization work clearly demonstrates the ability to convey data from a variety of sources including melting glaciers [61], the predicted impact of future climate scenarios on native bird populations [64], global terrorism statistics [46], and biological interactions of marine ecosystems [47], none of these examples involve interactive data exploration. This raises the concern: Does the typically static, pre-fabricated property of physicalizations preclude their use for interactive tasks such as data querying, filtering, and calling up details-on-demand?

Researchers have explored two main approaches to overcome this seemingly inherent limitation. The first is mechanical: including a system of sensors to detect touch and linear actuators to physically move the 3D "pixels" of a display surface up and down in real-time [45]. These impressive actuated displays may indeed represent the future of physicalization work but presently remain costly and challenging to construct and transport. The second main approach is to create a hybrid display, combining physical elements with virtual elements to provide both interactivity and dynamic content (e.g., projection [73]). For this paper, we define the scope of *hybrid visualization* [35] (also known as *augmented physicalization* [21]) to include any visualization that contains physical components but adds interactivity using digital input and output mechanisms.

Building on this work, in this paper we introduce a novel style of hybrid virtual + physical visualization designed specifically to support interactive data exploration tasks. We target the use case of multivariate, time-varying spatial data because this is a long-standing challenge for exploratory visualization research. This is one area where physicalizations have yet to be widely adopted, but have the strong potential to be beneficial because they make it possible to perceive data via multiple sensory channels simultaneously [21].

Our high-level goal is to combine the complementary strengths of digital visualization and physicalization while overcoming the individual limitations of each. Since physicalizations are meant to be touched with the fingers, we aim to support a level of interactivity that is similar to a digital multi-touch screen. Our first technical requirement is, therefore, to sense and respond to simultaneous touch input with multiple fingers placed directly on the surface of the physicalization. Further, the accuracy and responsiveness of this sensing should be similar to that of now-ubiquitous, multi-touch displays. However, unlike those displays, this input device should provide an added benefit of passive haptic feedback during interaction. Our second technical requirement is to robustly modify the visual display in response to the user input. Since we aim to support users analyzing complex spatial relationships in their data, meeting this requirement necessitates high-accuracy and low-latency alignment between the physicalization and the virtual data display so that users can make spatial judgements between the two (e.g., relating a flow feature displayed virtually to ridges in the bathymetry data displayed physically).

This research is significant because whether on planar surfaces, in immersive environments, or even hybrid virtual + physical systems, interactive techniques are essential to supporting exploratory visualization tasks, and so far these have been missing from data physicalizations. For 3D and 4D data, such traditional interactive techniques include user interfaces for 3D cutting planes, particle emitters, streamline seeding widgets, and multi-scale views [49]. While implementations of these techniques exist for planar-display desktop setups (e.g., [4, 15]), projection-based hybrid virtual + physical visualizations (e.g., [5, 53, 80]) and immersive environments, including via multi-touch input (e.g. [12, 17, 44]), our solution is the first to enable these interactive data exploration techniques in an immersive hybrid virtual + physical visualization.

Thus, this paper contributes a novel approach to virtual + physical, hybrid, interactive visualization based on combining perspective-tracked stereoscopic Augmented Reality (i.e., immersive AR) with a creative approach to sensing touch *through* 3D printed data physicalizations. This work represents the first example of such a hybrid visualization that supports exploratory visualization tasks using classic interactive widgets for spatial visualization, such as particle emitters, streamline seeding widgets, and multi-scale views. The specific design developed (e.g., the custom tracking base) is justified via a combination of discussions of related work and results from formative experiments in the lab. The results are applied to create an exploratory visualization of Antarctic ice melt simulations actively studied by climate science collaborators (see Figure 1) who provide early feedback on the utility of the tool. This evaluation is complemented by a lab user study to characterize the lower-level accuracy of the physicalization touch input pipeline.

2 RELATED WORK

2.1 Representing Data in the Physical World

Several advantages of data physicalization as compared to digital visualization have been demonstrated in the literature, including more efficient comprehension of data values when compared to 2D and 3D stereoscopic virtual displays of information [35] and preference for working with physical visualizations over immersive VR visualizations [19]. Many physicalizations are also works of art; though their primary goal is often to engage and inspire, these artworks can also improve data understanding relative to more traditional visualizations [21]. Examples of physicalizations that are both artistic and improve data understanding range from 3D printing of volumetric datasets [7] to massive tripod data sculptures depicting climate change effects on native bird populations [64]. Researchers and artists believe that these may provide audiences with increased data engagement, with the added potential of emotional connection that may be an even more valuable outcome than traditional metrics of speed and accuracy in data comprehension [77].

Other data physicalizations exist solely for the pragmatic purposes of data analysis and interpretation. Such physicalizations have been applied to biological and medical data. Examples of this approach include the physicalization of molecular models [14] and the 3D printing of multi-color surfaces for analyzing the multivariate properties of pain phenotypes [76]. Geospatial data are also commonly encoded with physicalization, with works ranging from tactile maps [56] to strategies for 3D printing an approximation of the Earth by dividing it into interchangeable, hierarchical cells [20]. While there are some indications that physicalizations are helpful for these pragmatic, exploratory visualizations (e.g., [19, 35]), there is also evidence to the contrary (e.g., [8]). In fact, Weiß et al. argue that there may be no significant difference in user engagement between web-based, VR, and hybrid physical visualizations for certain tasks [78]. Although additional research is needed to continue to clarify the possible benefits, risks, and tradeoffs with data physicalization, it is quite clear that physicalization is not the right fit for every data analysis task; this further motivates our desire to explore the potential of hybrid visualization systems.

2.2 Exploring Data in the Virtual World

There is a rich history of research on interactive techniques for exploring spatial data using spatial and surface-based displays. Immersive data querying interfaces include flow visualization widgets, such as direct-manipulation 3D placement of particle emitters and streamline rakes [12, 44], interactive world-in-miniature widgets for improved navigation and manipulation [3, 16], bimanual and multi-touch interfaces for slicing plane manipulation [17, 60], freehand and haptic-aided 3D sketching for data selection [31, 40, 65], multi-touch convex hull selection on slicing planes [17], mid-air volumetric selection of sub-volumes of interest [38], and many more.

Many of these techniques translate directly from VR environments to AR, and both styles of immersive visualization have a history of successful application to geographic or topographic data (e.g., [10, 27]), as in our application, which combines flow and terrain data. Our approach therefore adopts many of the lessons learned from prior work, such as the importance of situating data queries relative to features in the spatial terrain and supporting the most common immersive flow visualization widgets. Of course, even the best of these purely virtual immersive environments lacks the real-world tangibility provided by data physicalizations [36].

2.3 Early Tangible Interactions with Visualizations

Many visualization systems have built upon early research in Tangible User Interfaces (TUIs) (e.g., [30]) to make the data interactions described in the previous section even more natural or intuitive [63]. Our work shares a clear motivation with TUIs to bridge the gap between digital and physical elements. However, the "tangibles" in TUIs are typically different than the data physicalizations studied in this paper. Most TUIs use comparatively simple tangible objects that act as proxies (e.g., pucks, tiles, bricks, pointers, or a doll's head). These have regularly been used to represent or select data; for example, using physical pointers to select virtual points on an map with an immersive AR overlay [10], exploring neural fiber data using a rolled-up piece of paper and stereoscopic panel display [33], and performing query tasks in AR on 3D scatter plots with direct manipulation of physical props [6]. However, with a few noted exceptions (e.g. [23, 71]), users are not intended to perceive more than the position or orientation of the data by touching or viewing the tangibles. With data physicalizations, the role of the physical component shifts significantly to prioritize reading data, often using both the haptic and visual senses.

TUIs often integrate well with AR and VR displays [9], and the technologies used to sense how the tangibles are touched and manipulated are also relevant to our work. For detecting touch with TUIs, there are two primary options: optical and direct. The optical approach, using cameras to detect touch contacts on the surface of a tangible, suffers from a well-known limitation that the precise moment of contact is difficult to detect [79]. The precision of data exploration that is possible via this strategy is therefore limited, which is problematic when working with complex physical surfaces like geographic and bathymetric data. More direct touch sensing technologies (e.g., capacitive [58] and pneumatic [75]) are therefore preferred from a standpoint of accuracy, but these more accurate sensing technologies are often difficult to integrate with digital fabrication techniques. For example, although flexible multi-touch sensors can be rolled to fit onto the surface of cylindrical tangible prop [66], they cannot currently be stretched to fit over a more complex surface, even a sphere. We seek a solution that is direct-touch and high-accuracy but still integrates well with digital fabrication techniques.

2.4 Hybrid Virtual + Physical Systems

Moving beyond simpler tangibles to physical models that users touch and observe to read data values, the work that most closely relates to our own is in *augmenting* physicalizations with electromechanical or virtual components, an active topic of research also known as hybrid virtual + physical visualization. There are four categories of augmented physicalizations: surface augmentation, active physicalization, six-degree-of-freedom (6DoF) augmentation with a 2D display and 6DoF augmentation with an immersive 3D display [21].

In the first category, prior work using surface augmentation includes combining physical landscapes with GIS software [73], a table-scale physicalization paired with a digital energy data visualization [41], and area-preserving displays for geovisualization [18]. These surface-augmented physicalizations are often accomplished via projection (e.g. [72, 73]), though other technologies have been utilized (e.g., embedded LEDs [71]). Real-time registration between virtual and physical spaces is not usually required; often a single calibration is sufficient each time the system is assembled [73]. However, projection-based augmented physicalizations sometimes require projection mapping to ensure that no distortion occurs on the digital image – it has been reported as necessary in physicalizations over 6cm in height [21]. Surface augmentation is not ideal for our applications since we require understanding volumetric data, such as ocean currents at multiple depths.

Systems in the second category, active physicalizations, tend to require specialized hardware, and are usually limited to 2.5D applications where the mechanical elements are low resolution and move along a single axis. For example, Tangible CityScape [72] allows users to visualize multiple variables in the context of an urban environment physicalized with actuated columns, while EMERGE [71] enables visualization tasks like annotation, filtering, and navigation on 3D bar charts with actuated columns and color-changing LEDs. This approach is powerful and will perhaps one day evolve to create an ultimate display for surface-based 3D data. However, the current technology is limited in its ability to represent fine detail and in the prohibitive complexity and cost required to implement such displays. These factors along with a desire to complement physicalized data with spatially registered mid-air data steer us toward a hybrid solution using some form of AR.

In the third category, mobile phones, tablets, and desktop systems have been used to create a graphics viewed on the device's screen that overlay virtual content with a view of the physicalization captured by the device's camera. Examples include the tablet-based PLANWELL [51] and desktop-based AR systems for analyzing structural molecular biology [24, 25]. Unlike projection-augmented physicalizations, but similar to our work, these planar-display hybrid visualizations require real-time registration between virtual and physical spaces to allow a user to manipulate both the 3D viewpoint and the physical model itself. In these situations, inside-out image-based tracking is often used (e.g. [25, 51]), though other approaches exist (e.g. optical outside-in tracking [8, 55]). Techniques in this category can take advantage of the robust multi-touch input provided by the mobile device to support a variety of data interactions; however, they are performed on the mobile

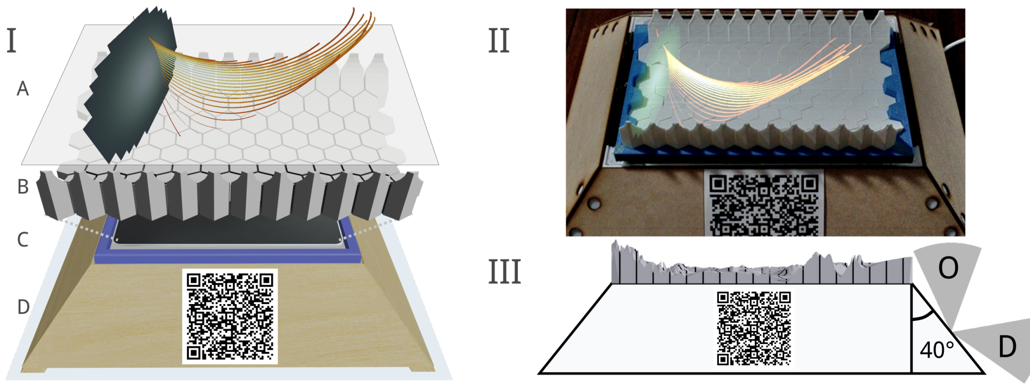

Fig. 2. (I) An augmented reality visualization (A) is superimposed on top of a hexagonally segmented data physicalization (B). Force from a user's touches is transferred through the hexagonal columns to the force-sensing surface (C). The virtual visualization is spatially co-located with the physicalization using an image-based tracking system that determines 3D positions and rotations of image targets (D). (II) An example hybrid visualization containing a hexagonally segmented physicalization of a pipe and immersive augmented reality content showing user-chosen streamlines seeded based on a generated velocity field and a cutting plane selected with multi-touch querying. (III) The side angles of the truncated pyramidal base are angled at 40° to optimize tracking accuracy when viewing the physicalization from the side to for details (D) and the top-down overview (O). Adapted from [1].

device's screen *not* on the physicalization itself, so the potential passive haptic benefits of data physicalization are lost.

The fourth category of augmented physicalizations pairs 6DoF-tracked physical content with immersive, head-tracked stereoscopic AR virtual content. Prior work includes Augmented Maps [10] and a hybrid visualization of flood dynamics using 3D-printed topography [81]. Systems in this style require robust, high-accuracy and low-latency 6DoF tracking for both the physicalization and the immersive AR device so that the virtual visualization remains spatially registered to the physical reference frame. We take this as a key requirement for our work and build upon the prior uses of image-based tracking methods (e.g., [6]). Although some prior work exists in this category, none meets our goal and requirement of supporting interaction through direct touch on the physicalization at a level similar to what can be provided by a typical multi-touch display.

3 HARDWARE AND STRATEGIES FOR ADDRESSING TECHNICAL CHALLENGES

Figure 2 presents an overview of our technical approach. By placing the physicalization on a custom base, we minimize tracking latency and maximize accuracy. By dividing the physical print into a regular grid of pieces that can move independently in the vertical direction, a pressure-sensing pad placed under the physicalization is able to sense touch events *through* the physicalization.

3.1 Continuous Real-Time Registration of Virtual and Physical Displays

Our strategy for consistently registering the position and rotation of the data physicalization relative to the user's eye uses a passive inside-out optical image tracking system. Although this is the best match with the Magic Leap AR hardware we used, other 6 DoF tracking solutions were considered, including electromagnetic, passive outside-in, and inside-out 3D model (as opposed to

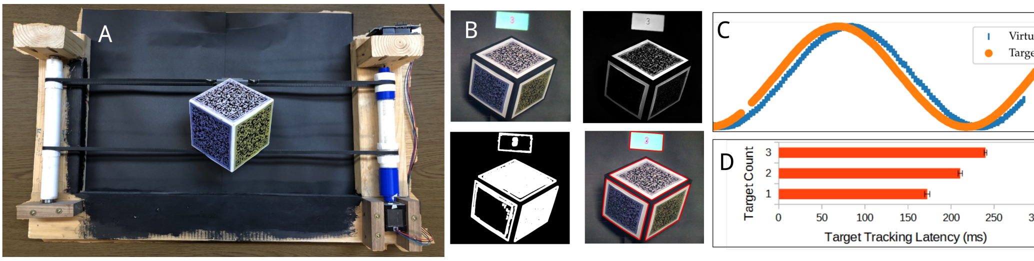

Fig. 3. The experimental setup to optimize latency for image-based tracking includes (A), an oscillating image target, (B) boundary extraction for the virtual and physical object positions, and (C) positional samples of virtual and physical target used to determine latency. Image (D) shows the latency results based on number of targets.

image) tracking. Passive inside-out tracking based on planar image targets is the most portable, only requiring a single camera for capturing the scene and images affixed to each object to be tracked. Its sole requirement is a target image placed near the physicalization such that it is in the line of sight of the camera simultaneously with the physicalization, as in the image target in Figure 2(I-D). Coverage and portability of inside-out image-based optical tracking also surpasses that of its outside-in counterparts because it is not bounded by a "tracking volume" – the space where all external cameras have a line of sight – and instead, the images and cameras are connected directly to the physicalization and the viewer, and thus will always be available without setup. This also means that it is not necessary to track both the user's head position and the physicalization; the relative offset from the image target to the tracking camera is sufficient to place the physicalization in 3D space. It is possible to perform an initial registration with image targets, then switch to the AR HMD's onboard 6DoF tracking to maintain the alignment. The onboard tracking is lower-latency but it does not allow a user to manipulate the physical model interactively, for example, to spin it around to view the other side. In practice, we found that the Magic Leap's internal room-scale tracking system was not accurate enough to maintain sufficient alignment between the virtual and physical content of the visualization without continuously enabled image tracking.

The key to optimizing an image-based tracking system is to understand the trade-offs in design choices for the number of images used to track objects and their placement in physical space. Thus, our design decisions focus on these two issues.

To optimize the number of images used to track the physicalization, we conducted an experiment to measure latency of the tracking system based on the number of image targets in view of the tracking camera. Figure 3A shows the approximate view of the AR device facing the sinusoidally oscillating image tracking assembly with a variable number of image targets attached. Video was captured through the left eye of the AR device of a virtual object's position tracked with the physical object's location, and computer vision methods were used to extract the true positions of the physical and virtual objects in the scene (Figure 3B). Lastly, following Steed's Sine Fitting method, both the virtual and physical object trajectories were fit to a sine curve, and the latency of the system was calculated as the phase distance between the two curves [67]. We found that the latency of the image tracking system does indeed increase with the number of simultaneously visible targets; we found a latency of 172ms, 211ms, and 240ms for 1, 2, and 3 targets, respectively. This suggests the physicalization tracking will have the lowest latency when tracked by just a single image. Using positional accuracy measurement methods based on Groves et al. [26], we verified that tracking just a single image is sufficiently accurate, providing sub-millimeter positional accuracy for a comfortable viewing distance of 30cm. However, we found it impossible to find a position where a single image could be placed next to the physicalization and force-sensing tablet while also providing reliable tracking from all viewing angles.

This led us to consider a stage or stand for the physicalization that provides a good tracking target from any viewpoint but limits the view to one (ideal) or at most two image targets at a time. Our starting point was a box with one image on each vertical side. However, noting that prior work has shown that image targets track most accurately when viewed from non-perpendicular angles [1], we further refined the design to the truncated pyramid shape pictured in Figure 2. The four-sided, truncated-pyramid base minimizes the number of targets simultaneously visible from any viewpoint, while tilting the sides of the base to an angle of 40° from vertical, which, based on the previously reported data [1], should optimize tracking for the top-down and side-on perspectives, as illustrated in sub-figure (III).

Our implementation uses commercially available software and hardware. A Magic Leap is used to display the augmented reality content with a resolution of 1280 x 960 pixels per eye, a horizontal field of view of 50°, and a 120Hz refresh rate. Optical tracking uses Vuforia [54]. Graphics are rendered using the Unity game engine, and a custom shader is used to occlude the augmented reality graphics with the physical model, as shown in Figure 1 *right* (though a user's fingers do not occlude the graphics). It should be noted that the images in this paper containing AR visualizations were generated using the Magic Leap's built-in screen capture system, which had a screen space offset between real-world and virtual content of approximately -0.037 units in the x direction and -0.006 units in the y direction. This was compensated for by visually calibrating the scene to a physical target before making the screen captures.

3.2 Sensing Multi-Touch Input on 3D Physicalizations

Our design for sensing touch through a 3D printed terrain-like surface builds upon related work and workshops on the theme of "Touching the 3rd Dimension" [68]. Extending the fluidity of multi-touch interaction off the surface of the screen and into/onto spatial displays is a longstanding goal of this research, which includes work on perceptual illusion [11], tracking 3D gestures anchored to a surface [34], redirected touching in purely virtual environments [43], and even dynamically inflatable multi-touch surfaces [69].

Figure 2 shows our solution for translating physical touch into virtual environments. To enable multi-touch interactions directly on the physicalization, it is split into a series of hexagonal columns with a radius of 0.75cm. This size of hexagon is chosen to be small enough to approximate the radius of a human fingertip yet maintain a manageable number of columns for the entire physicalization. Tiled hexagons are used to avoid the rigid visual structure that a regular square or rectangular grid imposes on the physicalization.

It should be noted that while the resolution of each column is fairly coarse, sub-hexagon touch precision is achieved by placing the partitioned physicalization atop a high-resolution force-sensitive surface, a Sensel Morph [62]. The hexagonal columns have flat bottoms which enables our software to interpolate a user's touch position within a hexagon. The Sensel Morph has a size of 24cm x 13.85cm and contains 185 x 105 force sensing elements which detect forces 0.049N-49N.

Using polygonal columns on such a sensitive device presents its own series of challenges. Firstly, as the user presses down on a particular column, multiple touch points may be registered by the device from the column's edges contacting the surface. In order to remove these "ghost" touches and ensure there is only one touch point registered for each finger, configurable parameters are used. Since the physicalization is divided into 0.75cm columns, an assumption is made that two fingers will never touch the same column – if a cluster of touch points within a diameter of 1.5cm is registered, the point inside this cluster with the greatest force is kept and the remainder are culled. To remove any other ghost touches remaining from the weight of the physicalization or

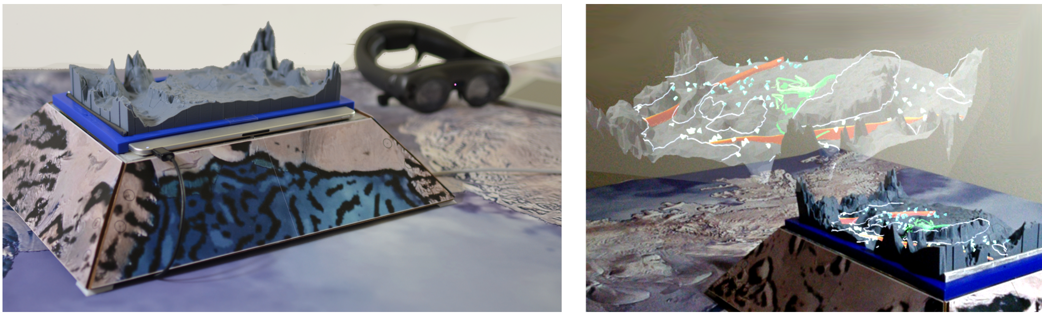

Fig. 4. (Left image) The physicalization and force-sensing surface are placed on a truncated pyramidal base. The base is situated onto a large fabric poster with contextual NASA satellite imagery around the Filchner-Ronne ice shelf in the Antarctic. *The imagery is not to scale with the physicalization.* (Right image) Worlds-in-miniature allow the user to view their data at multiple scales. All interactions still take place on the physicalization; when users are focused on the larger world, the physicalization acts as a haptic proxy.

other physical factors, the touch points are thresholded such that those with at least 5x less force than the others and those with an absolute force of less than 0.98N are ignored. Once the touch points are filtered, they are made available on a network using VRPN [74] in order to be accessed by other devices such as the augmented reality display.

In addition to the well-developed application to Antarctic ocean current simulations described in detail later, we have also successfully applied these strategies for multi-touch input and tracking to other datasets, such as the flow field in a pipe, pictured in Figure 2.

4 SYSTEM AND INTERACTION TECHNIQUES FOR MULTI-TOUCH QUERYING ON DATA PHYSICALIZATIONS

This section describes how three classic interactive data querying techniques can be re-implemented in the new context of hybrid virtual+physical visualization. Before describing these in detail, we begin with a brief overview of the system architecture used in our implementation.

4.1 System Overview

Our system architecture includes four components: the immersive augmented reality visualization, the multi-touch physicalization attached to a computer, the data host, and the visualization host.

Let us motivate this multi-component system by following the steps a user might take when making queries into a visualization. The user begins by donning an AR device in a space near the data physicalization. Figure 1 shows an approximate view of what the user might see here; the immersive AR content for that visualization is spatially registered to the physicalization using optical image-based tracking of the target images shown in Figure 4 (left). When the user touches the physicalization, this input is sensed as described in Section 3.2 then relayed to the AR device by a VRPN server on the touch host. The AR headset in turn calculates the proper virtual position of each touch point. After converting the touch points from room-space to data-space using a series of affine transformation matrices [32], a VTK data query is synthesized and sent to the data host as an HTTP request. Query results are calculated on the data host using VTK [59], and a geometric representation of the data is sent back to the AR headset and displayed in the visualization.

While not the focus of our human-computer interaction research, a companion touch screen interface is also supported in this architecture for input related to system control (e.g., specifying a timestep to view or switching between query modes). Additionally, we provide a touch-enabled

Fig. 5. Streamlines are interactively seeded by moving a finger around on the surface of the physicalization and releasing (left), while cutting planes are created by using two fingers (right).

design user interface that allows all data-to-visual mappings on the AR visualization to be modified interactively [28]. The AR visualization is rendered using the Artifact-Based Rendering (ABR) technique [39], delivering visually rich graphics to complement the data physicalization.

4.2 Streamline Seeding Interaction

Figure 5 (left) depicts a demonstration of an interactive stream rake widget that seeds streamlines along a line segment defined by a point on the surface of the physicalization and a second point directly above it. Streamlines are one of the essential tools for exploring fluid flow data from computational fluid dynamics and other numerical simulations [49]. In particular, stream rakes are used to create and view a collection of flow lines that pass through multiple seed points that are evenly spaced along a user-defined line segment. Variation between the resulting flow lines can then show researchers how flow patterns change along the axis defined by the line (ocean depth in our case).

In our implementation, the base of the rake is defined by a user's touch point on the physicalization and the up vector. As the user moves their finger along the physicalization, a preview of the seeding line is displayed on the AR device and updates in real time. The streamline resolution, number of seed points, and length are all adjustable parameters. When the streamline rake is selected, its coordinates are converted into data-space. Coordinates are sent to the data server, where VTK components are used to advect the streamline geometry. The streamlines are then sent back to the immersive AR visualization application, as shown in Figure 5 (left).

4.3 Cutting Plane Interactions

Figure 5 (right) demonstrates a cutting plane query defined by a user's two touch points on the physicalization. Interactively creating cutting planes is another useful tool for exploring 3D volumetric data [42]; they allow a user interactively define a 2D slice into the data.

In our system, the user's two touch points and the up vector define a basis for the cutting plane. Again, an interactive preview of the plane is updated in real-time as the user moves their fingers. When these touches are released, the coordinates are converted into data-space and sent to the data server. VTK components are used to create the cutting plane geometry which is then returned to the immersive AR app.

4.4 World-in-Miniature Widget

The ability to explore data at multiple scales is another indispensable technique in scientific visualization [49]. Worlds-in-miniature [70] is a technique that makes it possible to view a virtual environment at different scales while simultaneously interacting with a proxy that is at a comfortable

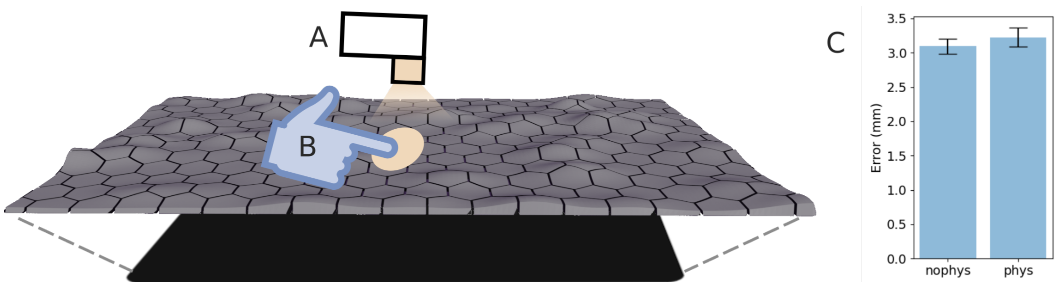

Fig. 6. Touch accuracy was measured with and without the tiled physicalization present. Random Cartesian coordinates are projected on the bare force-sensing surface or physicalization (A), and a participant touches these points (B). Image (C) shows positional error across conditions. Error bars show CI95.

scale. This technique has been used extensively in visualization research, for instance to compare forest plots over time [48] and to explore flow data via multimodal displays [16].

Figure 4 (right) shows our implementation of worlds-in-miniature using the hybrid virtual + physical visualization model. The hybrid worlds-in-miniature technique uses the data physicalization as both a primary view for the data and as a haptic proxy [50] for a larger AR visualization.

5 EVALUATION AND APPLICATION

Returning to the goals and requirements stated earlier, supporting exploratory visualization tasks requires a robust, high-accuracy pipeline for capturing multi-touch input and a responsive display that can usefully update the digital aspects of the visualization while maintaining accurate registration of the digital and virtual components. This section reports on two complementary evaluations: 1. a quantitative lab study on touch accuracy, and 2. an application to Antarctic climate science data and feedback from a domain scientist.

5.1 Characterizing Touch Accuracy

We conducted a user performance study to better understand the low-level capability of sensing multi-touch input *through* a data physicalization. During the touch study, the AR device was not included because the touch coordinates reported to the Touch Host and AR device are identical. Touch data were collected using the experimental setup pictured in Figure 6, which was designed to detect touches at specific (x,y) coordinates with and without the hexagonally tiled physicalization present. The study was designed as a within-subjects, repeated measures A vs. B comparison of two conditions: touch input sensed through a hexagonally-tiled physicalization (*phys*) and touch input sensed directly on the Sensel Morph's touch pad (*nophys*). From our experience, we hypothesized that the *nophys* condition would be close, but slightly more accurate than the *phys* condition.

5.1.1 Methodology. Participants ($n = 8$) were recruited from the local computer science student population. All participants were right-handed and all reported having normal or corrected-to-normal vision. After informed consent, the participants completed a series of pointing tasks in which they were instructed to touch a target as accurately as possible with their dominant-hand index finger. The tasks were presented in three blocks of 30 trials each. The first block was used for training and consisted of 30 trials in conditions alternating between *phys* and *nophys*. During the second and third blocks, participants touched points on the bare force-sensing tablet (condition *nophys*) and the physicalization atop the tablet (condition *phys*). The order of these two blocks was counter-balanced across subjects to ensure minimal learning effects.

Target touch points were generated using a random uniform distribution across the 24cm x 13.85cm physical surface. Each random point was then displayed sequentially to the participant via a projector aligned with the force-sensing surface (see Figure 6 A); each point was revealed after the participant had completed the touch action for the prior point, and points were visualized with a diameter of 5mm. The participant's touch points were recorded after filtering to remove ghost touches (using the strategy described in Section 3), and additionally adjusting for the Sensel Morph's hardware offset (x: 3.2mm, y: -2.6mm) and scale factor of 0.97.

The physicalization used in the study was of a synthetic terrain dataset and was the same for all participants. The terrain data were generated by adding 50 randomly distributed Gaussian peaks to create a bumpy surface with a maximum height of 8mm (Figure 6 B). These synthetic data were used, rather than the pipe or Antarctic data presented elsewhere in the paper, because we determined in pilot testing that a flatter surface was needed to avoid distortion effects when projecting the target points on the physicalization.

5.1.2 Results. Touch data were recorded for *8 users* $\times$ *30 points*, totalling 240 data points for each condition. 18 duplicate touches (touches that registered within 0.5s of each other) were registered for *nophys* and 38 were registered for *phys*; as is common in touch-tracking software, these duplicates were averaged and counted as one touch. Additionally, 10 expected touches for *nophys* and 4 expected touches for *phys* were not detected as they were either within the dead zone surrounding the force-sensing tablet or below the force threshold in our touch reporting software.

Touch accuracy was measured by calculating the `touch_error` as the Euclidean distance between the ground truth Cartesian coordinates and the coordinates of a participant's touches. Points with an x or y coordinate more than 3σ away from the mean were considered outliers and discarded.

As shown in Figure 6C, participants selected points more accurately under the *nophys* condition (μ = 3.09mm, σ = 1.64mm) as compared to the *phys* condition (μ = 3.22mm, σ = 2.1mm). A Shapiro-Wilk test showed that the data were not normally distributed ($p < 0.0001$). Thus, a Wilcoxon signed-rank was used to analyze significance. The two-sided Wilcoxon signed-rank test indicated that the physicalization did not have a significant effect on `touch_error` ($p = 0.753$); thus, we fail to reject the null hypothesis that there is no difference in mean `touch_error` between conditions.

5.1.3 Interpretation and Discussion. Comparing the results of `touch_error` between conditions indicates that the average Euclidean distance between a participant's touch points and the ground truth points was slightly greater in the *phys* (3.22mm) condition than in *nophys* (3.09mm). We did anticipate some difference in this direction since the rigidity of the physicalization hexagons certainly contributes some error. However, after observing participants during the study, we also suspect that some error was due to shadows in the projection used to display touch point as well as the physical pixel size (about 0.5mm x 0.5mm).

These combine to leave us surprised that the study was unable to detect a significant difference in touch accuracy across approximately 240 trials. Other factors could also impact this result. For example, the number of participants was small and their backgrounds rather homogeneous (limitations necessitated by the current constraints of the pandemic).

Despite these factors, we take the study as a promising initial characterization of the feasibility of the core concept behind our tiled physicalizations–that the force of a participant's touch can be transferred directly through the column to the tablet, opening the door for sensing touch through the type of rigid 3D printed physicalizations that are most easily created today.

5.2 Antarctic Climate Science: Application and Feedback

To assess the potential of multi-touch querying on data physicalizations in a real-world scientific context, we constructed the augmented physicalization pictured in the examples throughout this

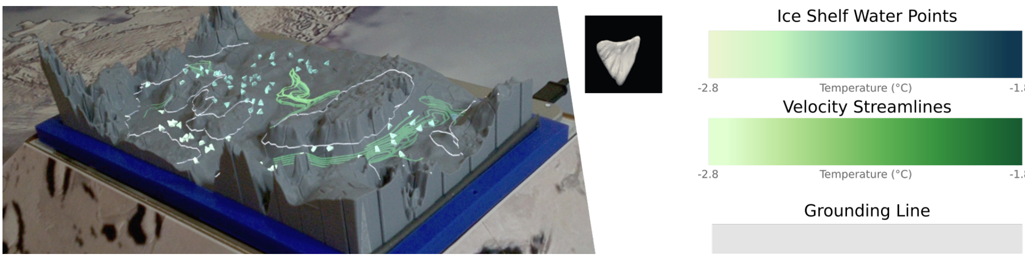

Fig. 7. Streamlines are seeded based on the velocity of the ocean currents underneath the Filchner-Ronne ice shelf, and are colored by the ocean temperature. Glyphs are distributed based on a density-based point sampling of the Ice Shelf Water mass. The grounding line where the ice intersects with the water and land is shown as a ribbon on the visualization.

paper and conducted a user interview with a climate scientist collaborator actively studying the phenomena. The Filchner-Ronne ice shelf in Antarctica is of significant interest to climate scientists because ice shelves extending from the Antarctic continent prevent ice sheets on the continent itself from melting rapidly and flowing into the sea (see Figure 7). The water content of the Antarctic ice sheets and shelves combined has the potential to inundate many of the planet's major coastal population centers if completely melted. The melt rates beneath the ice shelves exceed the surface melt rates and are difficult to monitor, as water masses and currents of varying temperature, salinity, flow rate, and more must be tracked and are invisible from the surface. A more detailed understanding of the currents, their intensity, and their temperature will enable scientists to more accurately predict the melt rate of the ice shelves, which play a critical role in sea-level rise as well as climate patterns across the globe [2, 37, 52].

Technical results of the application to climate science in Antarctica's Filchner-Ronne ice shelf are pictured throughout the paper as well as in the accompanying video. All three of the classic data querying interaction techniques described in Section 4 are applicable to these data and are controllable via multi-touch queries made on top of the Filchner-Ronne ice shelf 3D print.

5.2.1 Data Conversion. Due to the physical constraints of the direct-touch sensing mechanism, a subsection of the Filchner-Ronne ice shelf was selected based on conversations with collaborating climate scientists about areas of interest within this region. The subsection measures 945 km by 545 km, and underwent transformations from its original coordinate space in the climate simulations, including a stereographic projection, a vertical exaggeration of 60x, clipping and scaling to the size of the force-sensing tablet, and finally segmenting into hexagonal columns as discussed in Section 3.2. Special care was taken to preserve the original directionality of the ocean current flow vectors when applying the stereographic projection and performing affine transformations on the data. Hexagonal columns of the surface data were 3D printed with a layer height of 0.15mm, taking a total print time of about two days for the 247 hexagons in the segmented surface data.

Figure 4 (left) shows the 3D printed region of interest along with NASA satellite imagery from the Filchner-Ronne ice shelf. Tracking images that align the virtual and physical visualizations are adapted from NASA satellite imagery, which provides additional context for the position and orientation of the data and the physical model relative to a bird's-eye view of Antarctica. The images are affixed to the truncated pyramidal tracking base, and the context they provide is enhanced by placing the entire assembly on top of a large 2D poster, extending the Antarctic imagery.

5.2.2 Climate Scientist Interview. The interview was conducted with one investigator physically present and two remote investigators joining via video conferencing. We began the interview process by asking the climate scientist to explore the system, using a “think-aloud” approach as we directed him through a series of introductory tasks designed to demonstrate the capabilities of our system. These tasks included seeding streamlines, creating cutting planes through the volumetric data, and using the multi-scale view to view details of the ocean current mixing. We then prompted the scientist throughout his exploration process with questions designed to encourage both constructive criticism about the system as a whole and feedback on the scientist’s usage of the system. Questions included, for example, “How does it feel to explore your data in this format?” and “How do you think this configuration could play a role in your collaboration with peers?”

When setting up the physicalization for the interview, we noted that the orientation of the physical model relative to the user is a critical consideration [57]. The bathymetric physicalization was oriented with North facing away from the user, a common perspective that the scientists use when working with these data.

The interview revealed several preliminary insights. Following the introduction tasks, the climate scientist conducted a rapid “lay-of-the-land” exploration by touching the printed bathymetry in several locations, seeding clusters of streamlines to understand the velocity, temperature, and direction of interesting flow patterns. He noted our system’s utility in selecting streamlines precisely on the surface and added that automatically seeding an initial set of streamlines at the start would help narrow down areas of interest in the data.

The climate scientist reacted to the physicality of the data visualization, saying “*The fact that the data overlays something physically is really wonderful...I’ve never experienced this before. There’s a really intuitive feel to it.*” and “*[The printed bathymetry] really grounds you and provides useful context for the data itself.*” He noted the system’s potential for improving his group’s research and inter-scientist communication more generally, particularly if a future iteration of the system acted as a centerpiece that team members could walk around, examine, and reference interactively on a tabletop. He further indicated that a hybrid system might increase public interest and improve communication of critical but complex concepts in this field.

In addition to these overarching comments, the climate scientist provided useful feedback on several specific aspects of the design. He observed of the hexagonally-segmented physicalization itself, “*I like the hexagons... it’s much more stable [than squares], and I think having squares instead of hexagons would be much more distracting.*” The scientist expressed particular interest in the grounding line – the line where water, depressed by surface ice, meets the land. This feedback prompted the grounding line’s inclusion as shown in Figure 7. He also stated that a larger physicalization of the bathymetry would have been extremely useful, even if it wasn’t all touch-sensitive, adding that he found himself primarily using the printed bathymetry as an orientation device as opposed to searching for small details in the physical model. This observation was striking. Counter to our expectations, the physical model became a spatial marker in three dimensions that acted as a peripheral canvas for the data of greater interest as opposed to a primary data object itself.

Although the subjective user feedback presented here should be treated only as an early data point in understanding the potential of this approach, the interviewed scientist and other collaborators noted several use cases where there is strong potential for hybrid virtual + physical visualization. Speaking on the subject of collaboration, one scientist remarked “*Having it available to work with and show other scientists – I think there’s a lot of potential there.*” Furthermore, “*There’s a lot of potential for people to see and get excited about the science because it’s something new and interesting.*”

6 DISCUSSION

6.1 Limitations

Our approach to hybrid visualization, while a novel step forward, has two principal limitations. First, for some workflows, the digital fabrication time may be prohibitive. Printing time can be reduced by increasing the layer height, but this comes at the cost of a lower-resolution model. However, our scientist feedback interview indicates potential for some flexibility in resolution, depending on how a user identifies and works off of the physical model.

Another issue that emerged with the hexagonal column design is the challenge of keeping all the tiles linked together. Researchers needed to be cautious when moving the physicalization from the 3D printer to the touch surface, as the hexagons could become a 3D puzzle very quickly if dropped. Currently, we mitigate this issue with the 3D printed frame shown in blue on Figure 2 I-C, which nicely holds the tiles in place on the surface but will not prevent them from scattering if turned upside down. Although several prototypes for linkage between hexagons together while preserving complete touch sensitivity have been made, we have not yet found a foolproof resolution.

6.2 Potential for Other Applications

Our solution to hybrid virtual + physical visualization has been applied to several datasets so far, including those shown in this paper, though our solution is not limited to such data. We have considered, for example, applying this technique to brain microstructure data as well as treatment-planning data from radiation therapy procedures. In these anatomical cases, we would physicalize a slice or surface out of a CT scan and segment it, enabling the domain scientists to create cutting planes and interactively seed directional glyphs at points of interest to further explore their data.

The addition of a force-sensitive surface underneath the physicalization also opens up new possibilities for interactive techniques that not only leverage a user's *touch* but also the touch's *pressure*. This is not a new area of research [13], but its application to 3D exploratory visualization could inspire new data querying widgets, such as sampling glyphs with density determined by pressure or using pressure to control the number and length of streamlines Further, such systems present new and highly complex visual design challenges and open new avenues for experimentation, including the materiality, appearance, and scale of the physical models and their interactions with the virtual visualization.

6.3 Further Touch Sensing Evaluations

One of the most interesting discussion points raised by the touch accuracy study is the issue of what accurate touch means when sensing touch *through* a physicalization. The results from the study in Section 5 suggest that there are many contributing factors, both hardware and software. While the parameters of the touch-filtering algorithm remained constant throughout the study and application, these are all individually adjustable to achieve a more accurate touch and further avoid duplicate and missed touches. Such tuneable parameters may include the clustering diameter, minimum absolute force threshold, and force threshold relative to other touch points on the surface. Beyond these software parameters, we must also consider the hardware setup – what is the effect of polygon size on touch accuracy? Polygon shape? Height of the columns? The durability and linkage between columns? This last item proved to be of particular interest; during the touch study, several participants voiced concerns about "breaking" the physicalization during the *phys* condition, and one managed to unseat a hexagonal tile while touching the physicalization. These are all factors that have the potential to impact accuracy of a multi-touch physicalization, and thus need to be evaluated in future studies. Additional studies must also be run (e.g. in the style of [35]) to test the exploratory capabilities and interactive techniques used in the hybrid visualization system.

7 CONCLUSION

The motivation for the work was to fuse the "best of" physicalizations and virtual displays; specifically our goal was to create a form of data physicalization that can be used for interactive, exploratory visualization. We demonstrated that this is indeed possible; in fact, we showed for the first time what three classic spatial data exploration widgets might look like when designed for use with a data physicalization in conjunction with immersive AR. The preliminary quantitative lab study of touch accuracy lends support to the feasibility of the core concept of sensing touch input through a tiled physicalization surface, and the application to exploratory visualization of Antarctic climate simulations demonstrates how the sensing and interaction techniques can be used with real data. We believe the sensing solutions and interactive techniques presented here are a major step forward for the future of hybrid visualization platforms. This exploratory research provides evidence for significant future potential of interactive, surface-based hybrid visualization.

ACKNOWLEDGMENTS

The authors wish to thank Josh Rector, Cullen Jackson, and the rest of the Sculpting Vis Collective, as well as the reviewers whose feedback helped improve this paper. This work was supported in part by the National Science Foundation (IIS-1704604 & IIS-1704904). MPAS-Ocean simulations were conducted by Mark Petersen, Phillip Wolfram, Mathew Maltrud and Xylar Asay-Davis as part of the Energy Exascale Earth System Model project, funded by the U.S. Department of Energy, Office of Science, Office of Biological and Environmental Research.

REFERENCES

[1] Daniel F Abawi, Joachim Bienwald, and Ralf Dorner. 2004. Accuracy in optical tracking with fiducial markers: an accuracy function for ARToolKit. In *Third IEEE and ACM International symposium on mixed and augmented reality*. IEEE, 260–261.

[2] Greg Abram, Francesca Samsel, Mark R. Petersen, Xylar Asay-Davis, Darin Comeau, Stephen F. Price, and Mike Potel. 2021. Antarctic Water Masses and Ice Shelves: Visualizing the Physics. *IEEE Computer Graphics and Applications* 41, 1 (2021), 35–41. https://doi.org/10.1109/MCG.2020.3044228

[3] Daniel Acevedo, Eileen Vote, David H Laidlaw, and Martha S Joukowsky. 2001. Archaeological data visualization in VR: Analysis of lamp finds at the Great Temple of Petra, a case study. In *Proceedings Visualization, 2001. VIS'01.* IEEE, 493–597.

[4] James Ahrens, Berk Geveci, and Charles Law. 2005. Paraview: An end-user tool for large data visualization. *The visualization handbook* 717, 8 (2005).

[5] Luis Alonso, Yan Ryan Zhang, Arnaud Grignard, Ariel Noyman, Yasushi Sakai, Markus ElKatsha, Ronan Doorley, and Kent Larson. 2018. Cityscope: a data-driven interactive simulation tool for urban design. Use case volpe. In *International conference on complex systems*. Springer, 253–261.

[6] Benjamin Bach, Ronell Sicat, Johanna Beyer, Maxime Cordeil, and Hanspeter Pfister. 2017. The hologram in my hand: How effective is interactive exploration of 3d visualizations in immersive tangible augmented reality? *IEEE transactions on visualization and computer graphics* 24, 1 (2017), 457–467.

[7] Christoph Bader, Dominik Kolb, James C Weaver, Sunanda Sharma, Ahmed Hosny, João Costa, and Neri Oxman. 2018. Making data matter: Voxel printing for the digital fabrication of data across scales and domains. *Science advances* 4, 5 (2018), eaas8652.

[8] Francois Berard and Thibault Louis. 2017. The object inside: Assessing 3d examination with a spherical handheld perspective-corrected display. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. 4396–4404.

[9] Mark Billinghurst, Hirokazu Kato, Ivan Poupyrev, et al. 2008. Tangible augmented reality. *Acm siggraph asia* 7, 2 (2008), 1–10.

[10] Joachim Bobrich and Steffen Otto. 2002. Augmented maps. *International Archives of Photogrammetry Remote Sensing and Spatial Information Sciences* 34, 4 (2002), 502–505.

[11] Gerd Bruder, Frank Steinicke, Phil Wieland, and Markus Lappe. 2011. Tuning self-motion perception in virtual reality with visual illusions. *IEEE Transactions on Visualization and Computer Graphics* 18, 7 (2011), 1068–1078.

[12] Steve Bryson. 1996. Virtual reality in scientific visualization. *Commun. ACM* 39, 5 (1996), 62–71.

[13] William Buxton, Ralph Hill, and Peter Rowley. 1985. Issues and techniques in touch-sensitive tablet input. In *Proceedings of the 12th annual conference on Computer graphics and interactive techniques*. 215–224.

[14] Felix A Carroll and David N Blauch. 2017. 3D printing of molecular models with calculated geometries and p orbital isosurfaces. *Journal of Chemical Education* 94, 7 (2017), 886–891.

[15] Hank Childs. 2012. VisIt: An end-user tool for visualizing and analyzing very large data. (2012).

[16] Dane Coffey, Nicholas Malbraaten, Trung Le, Iman Borazjani, Fotis Sotiropoulos, and Daniel F Keefe. 2011. Slice WIM: a multi-surface, multi-touch interface for overview+ detail exploration of volume datasets in virtual reality. In *Symposium on Interactive 3D Graphics and Games*. 191–198.

[17] Dane Coffey, Nicholas Malbraaten, Trung Bao Le, Iman Borazjani, Fotis Sotiropoulos, Arthur G Erdman, and Daniel F Keefe. 2011. Interactive slice WIM: Navigating and interrogating volume data sets using a multisurface, multitouch VR interface. *IEEE Transactions on Visualization and Computer Graphics* 18, 10 (2011), 1614–1626.

[18] Shima Dadkhahfard, Katayoon Etemad, John Brosz, and Faramarz Samavati. 2018. Area Preserving Dynamic Geospatial Visualization on Physical Globe.. In *VISIGRAPP (3: IVAPP)*. 309–318.

[19] Kurtis Danyluk, Teoman Tomo Ulusoy, Wei Wei, and Wesley Willett. 2020. Touch and Beyond: Comparing Physical and Virtual Reality Visualizations. *IEEE Transactions on Visualization and Computer Graphics* (2020).

[20] Hessam Djavaherpour, Ali Mahdavi-Amiri, and Faramarz F Samavati. 2017. Physical visualization of geospatial datasets. *IEEE computer graphics and applications* 37, 3 (2017), 61–69.

[21] Hessam Djavaherpour, Faramarz Samavati, Ali Mahdavi-Amiri, Fatemeh Yazdanbakhsh, Samuel Huron, Richard Levy, Yvonne Jansen, and Lora Oehlberg. 2021. Data to Physicalization: A Survey of the Physical Rendering Process. *arXiv preprint arXiv:2102.11175* (2021).

[22] Bin Fang, Di Guo, Fuchun Sun, Huaping Liu, and Yupei Wu. 2015. A robotic hand-arm teleoperation system using human arm/hand with a novel data glove. In *2015 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2483–2488.

[23] Sean Follmer, Daniel Leithinger, Alex Olwal, Akimitsu Hogge, and Hiroshi Ishii. 2013. inFORM: dynamic physical affordances and constraints through shape and object actuation.. In *Uist*, Vol. 13.

[24] Alexandre Gillet, Michel Sanner, Daniel Stoffler, David Goodsell, and Arthur Olson. 2004. Augmented reality with tangible auto-fabricated models for molecular biology applications. In *IEEE Visualization 2004*. IEEE, 235–241.

[25] Alexandre Gillet, Michel Sanner, Daniel Stoffler, and Arthur Olson. 2005. Tangible interfaces for structural molecular biology. *Structure* 13, 3 (2005), 483–491.

[26] Leah A Groves, Patrick Carnahan, Daniel R Allen, Rankin Adam, Terry M Peters, and Elvis CS Chen. 2019. Accuracy assessment for the co-registration between optical and VIVE head-mounted display tracking. *International journal of computer assisted radiology and surgery* 14, 7 (2019), 1207–1215.

[27] Nicholas R Hedley, Mark Billinghurst, Lori Postner, Richard May, and Hirokazu Kato. 2002. Explorations in the use of augmented reality for geographic visualization. *Presence: Teleoperators & Virtual Environments* 11, 2 (2002), 119–133.

[28] Bridger Herman, Francesca Samsel, Annie Bares, Seth Johnson, Greg Abram, and Daniel F Keefe. 2020. Printmaking, Puzzles, and Studio Closets: Using Artistic Metaphors to Reimagine the User Interface for Designing Immersive Visualizations. In *2020 IEEE VIS Arts Program (VISAP)*. IEEE, 19–28.

[29] Trevor Hogan and Eva Hornecker. 2016. Towards a design space for multisensory data representation. *Interacting with Computers* 29, 2 (2016), 147–167.

[30] Hiroshi Ishii and Brygg Ullmer. 1997. Tangible bits: towards seamless interfaces between people, bits and atoms. In *Proceedings of the ACM SIGCHI Conference on Human factors in computing systems*. 234–241.

[31] Bret Jackson, Dane Coffey, and Daniel F Keefe. 2012. Force Brushes: Progressive data-driven haptic selection and filtering for multi-variate flow visualizations. *Proceedings of EuroVis 2012* (2012), 7–11.

[32] Bret Jackson and Daniel F Keefe. 2019. From Painting to Widgets, 6-DOF and Bimanual Input Beyond Pointing.

[33] Bret Jackson, Tung Yuen Lau, David Schroeder, Kimani C Toussaint, and Daniel F Keefe. 2013. A lightweight tangible 3D interface for interactive visualization of thin fiber structures. *IEEE transactions on visualization and computer graphics* 19, 12 (2013), 2802–2809.

[34] Bret Jackson, David Schroeder, and Daniel F Keefe. 2012. Nailing down multi-touch: anchored above the surface interaction for 3D modeling and navigation. In *Proceedings of Graphics Interface 2012*. Citeseer, 181–184.

[35] Yvonne Jansen, Pierre Dragicevic, and Jean-Daniel Fekete. 2013. Evaluating the efficiency of physical visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 2593–2602.

[36] Yvonne Jansen, Pierre Dragicevic, Petra Isenberg, Jason Alexander, Abhijit Karnik, Johan Kildal, Sriram Subramanian, and Kasper Hornbæk. 2015. Opportunities and challenges for data physicalization. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 3227–3236.

[37] Hyein Jeong, Xylar S. Asay-Davis, Adrian K. Turner, Darin S. Comeau, et al. 2020. Impacts of Ice-Shelf Melting on Water-Mass Transformation in the Southern Ocean from E3SM Simulations. *Journal of Climate* 33,

13 (06 2020), 5787–5807. https://doi.org/10.1175/JCLI-D-19-0683.1 arXiv:https://journals.ametsoc.org/jcli/article-pdf/33/13/5787/4955899/jclid190683.pdf

[38] Seth Johnson, Daniel Orban, Hakizumwami Birali Runesha, Lingyu Meng, Bethany Juhnke, Arthur Erdman, Francesca Samsel, and Daniel F Keefe. 2019. Bento box: An interactive and zoomable small multiples technique for visualizing 4d simulation ensembles in virtual reality. *Frontiers in Robotics and AI* 6 (2019), 61.

[39] Seth Johnson, Francesca Samsel, Gregory Abram, Daniel Olson, Andrew J Solis, Bridger Herman, Phillip J Wolfram, Christophe Lenglet, and Daniel F Keefe. 2019. Artifact-based rendering: Harnessing natural and traditional visual media for more expressive and engaging 3D visualizations. *IEEE transactions on visualization and computer graphics* 26, 1 (2019), 492–502.

[40] Daniel F Keefe, Robert C Zeleznik, and David H Laidlaw. 2008. Tech-note: Dynamic dragging for input of 3D trajectories. In *2008 IEEE Symposium on 3D User Interfaces*. IEEE, 51–54.

[41] Nurit Kirshenbaum, James Hutchison, Ryan Theriot, Dylan Kobayashi, and Jason Leigh. 2020. Data in Context: Engaging Audiences with 3D Physical Geo-Visualization. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–9.

[42] Tijmen Klein, Florimond Guéniat, Luc Pastur, Frédéric Vernier, and Tobias Isenberg. 2012. A design study of direct-touch interaction for exploratory 3D scientific visualization. In *Computer Graphics Forum*, Vol. 31. Wiley Online Library, 1225–1234.

[43] Luv Kohli. 2010. Redirected touching: Warping space to remap passive haptics. In *2010 IEEE Symposium on 3D User Interfaces (3DUI)*. IEEE, 129–130.

[44] Joseph LaViola and Robert Zeleznik. 1999. Flex and pinch: A case study of whole hand input design for virtual environment interaction. In *Proceedings of the second IASTED international conference on computer graphics and imaging*. 221–225.

[45] Daniel Leithinger, David Lakatos, Anthony DeVincenzi, Matthew Blackshaw, and Hiroshi Ishii. 2011. Direct and gestural interaction with relief: a 2.5 D shape display. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*. 541–548.

[46] Loren Madsen. 2018. https://www.facebook.com/loren.madsen.9/posts/10156981297933841. Accessed July 2020.

[47] Nathalie Miebach. [n.d.]. Changing Waters. http://nathaliemiebach.com/waters.html. Accessed March 2021.

[48] Jung Who Nam, Krista McCullough, Joshua Tveite, Maria Molina Espinosa, Charles H Perry, Barry T Wilson, and Daniel F Keefe. 2019. Worlds-in-Wedges: Combining Worlds-in-Miniature and Portals to Support Comparative Immersive Visualization of Forestry Data. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. IEEE, 747–755.

[49] Gregory Nielson, Hans Hagen, and Heinrich Muller. 1997. Scientific visualization. Institute of Electrical & Electronics Engineers.

[50] Niels Christian Nilsson, Andre Zenner, and Adalberto L Simeone. 2020. Haptic Proxies for Virtual Reality: Success Criteria and Taxonomy. In *Workshop on Everyday Proxy Objects for Virtual Reality at CHI'20*.

[51] Aditya Shekhar Nittala, Nico Li, Stephen Cartwright, Kazuki Takashima, Ehud Sharlin, and Mario Costa Sousa. 2015. PLANWELL: Spatial User Interface for Collaborative Petroleum Well-Planning. In *SIGGRAPH Asia 2015 Mobile Graphics and Interactive Applications* (Kobe, Japan) *(SA '15)*. Association for Computing Machinery, New York, NY, USA, Article 19, 8 pages. https://doi.org/10.1145/2818427.2818443

[52] Mark R. Petersen, Xylar S. Asay-Davis, et al. 2019. An Evaluation of the Ocean and Sea Ice Climate of E3SM Using MPAS and Interannual CORE-II Forcing. *Journal of Advances in Modeling Earth Systems* 11, 5 (2019), 1438–1458. https://doi.org/10.1029/2018MS001373

[53] Ben Piper, Carlo Ratti, and Hiroshi Ishii. 2002. Illuminating clay: a 3-D tangible interface for landscape analysis. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. 355–362.

[54] PTC. [n.d.]. Vuforia. https://www.ptc.com/en/products/vuforia/vuforia-engine

[55] Daniel Pustka, Jan-Patrick Hülß, Jochen Willneff, Frieder Pankratz, Manuel Huber, and Gudrun Klinker. 2012. Optical outside-in tracking using unmodified mobile phones. In *2012 IEEE International Symposium on Mixed and Augmented Reality (ISMAR)*. IEEE, 81–89.

[56] Wolf-Dieter Rase. 2011. Creating physical 3D maps using rapid prototyping techniques. In *True-3D in Cartography*. Springer, 119–134.

[57] Kim Sauvé, Dominic Potts, Jason Alexander, and Steven Houben. 2020. A Change of Perspective: How User Orientation Influences the Perception of Physicalizations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. 1–12.

[58] Valkyrie Savage, Xiaohan Zhang, and Björn Hartmann. 2012. Midas: fabricating custom capacitive touch sensors to prototype interactive objects. In *Proceedings of the 25th annual ACM symposium on User interface software and technology*. 579–588.

[59] Will J Schroeder, Bill Lorensen, and Ken Martin. 2004. *The visualization toolkit: an object-oriented approach to 3D graphics.* Kitware.

[60] Jürgen Schulze-Döbold, Uwe Wössner, Steffen P Walz, and Ulrich Lang. 2001. Volume rendering in a virtual environment. In *Immersive Projection Technology and Virtual Environments 2001.* Springer, 187–198.

[61] Adrien Segal. 2015. Grewingk Glacier. https://www.adriensegal.com/grewingk-glacier. Accessed July 2020.

[62] Sensel, Inc. [n.d.]. The Sensel Morph. https://morph.sensel.com/. Accessed March 2021.

[63] Orit Shaer and Eva Hornecker. 2010. *Tangible user interfaces: past, present, and future directions.* Now Publishers Inc.

[64] Gary Singh. 2018. Wearing Multiple Hats. *IEEE Computer Graphics and Applications* 38, 4 (2018), 6–8.

[65] Jason S Sobel, Andrew S Forsberg, David H Laidlaw, Robert C Zeleznik, Daniel F Keefe, Igor Pivkin, George E Karniadakis, Peter Richardson, and Sharon Swartz. 2004. Particle flurries. *IEEE Computer Graphics and Applications* 24, 2 (2004), 76–85.

[66] Hyunyoung Song, Hrvoje Benko, Francois Guimbretiere, Shahram Izadi, Xiang Cao, and Ken Hinckley. 2011. Grips and gestures on a multi-touch pen. In *Proceedings of the SIGCHI conference on Human Factors in computing systems.* 1323–1332.

[67] Anthony Steed. 2008. A simple method for estimating the latency of interactive, real-time graphics simulations. In *Proceedings of the 2008 ACM symposium on Virtual reality software and technology.* 123–129.

[68] Frank Steinicke, Hrvoje Benko, Antonio Krüger, Daniel Keefe, Jean-Baptiste de la Riviére, Ken Anderson, Jonna Häkkilä, Leena Arhippainen, and Minna Pakanen. 2012. The 3rd dimension of CHI (3DCHI) touching and designing 3D user interfaces. In *CHI'12 Extended Abstracts on Human Factors in Computing Systems.* 2695–2698.

[69] Andrew Stevenson, Christopher Perez, and Roel Vertegaal. 2010. An inflatable hemispherical multi-touch display. In *Proceedings of the fifth international conference on Tangible, embedded, and embodied interaction.* 289–292.

[70] Richard Stoakley, Matthew J Conway, and Randy Pausch. 1995. Virtual reality on a WIM: interactive worlds in miniature. In *Proceedings of the SIGCHI conference on Human factors in computing systems.* 265–272.

[71] Faisal Taher, John Hardy, Abhijit Karnik, Christian Weichel, Yvonne Jansen, Kasper Hornbæk, and Jason Alexander. 2015. Exploring interactions with physically dynamic bar charts. In *Proceedings of the 33rd annual acm conference on human factors in computing systems.* 3237–3246.

[72] Sheng Kai Tang, Yusuke Sekikawa, Daniel Leithinger, Sean Follmer, and Hirishi Ishii. [n.d.]. Tangible CityScape. http://tangible.media.mit.edu/project/tangible-cityscape. Accessed March 2021.

[73] Laura Tateosian, Helena Mitasova, Brendan Harmon, Brent Fogleman, Katherine Weaver, and Russel Harmon. 2010. TanGeoMS: Tangible geospatial modeling system. *IEEE transactions on visualization and computer graphics* 16, 6 (2010), 1605–1612.

[74] Russell M Taylor, Thomas C Hudson, Adam Seeger, Hans Weber, Jeffrey Juliano, and Aron T Helser. 2001. VRPN: a device-independent, network-transparent VR peripheral system. In *Proceedings of the ACM symposium on Virtual reality software and technology.* 55–61.

[75] Carlos E Tejada, Raf Ramakers, Sebastian Boring, and Daniel Ashbrook. 2020. AirTouch: 3D-printed Touch-Sensitive Objects Using Pneumatic Sensing. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems.* 1–10.

[76] Michael C Thrun and Florian Lerch. 2016. Visualization and 3D printing of multivariate data of biomarkers. *WSCG* (2016).

[77] Yun Wang, Adrien Segal, Roberta Klatzky, Daniel F Keefe, Petra Isenberg, Jörn Hurtienne, Eva Hornecker, Tim Dwyer, and Stephen Barrass. 2019. An emotional response to the value of visualization. *IEEE computer graphics and applications* 39, 5 (2019), 8–17.

[78] Maximilian Weiß, Katrin Angerbauer, Alexandra Voit, Magdalena Schwarzl, Michael Sedlmair, and Sven Mayer. 2020. Revisited: Comparison of empirical methods to evaluate visualizations supporting crafting and assembly purposes. *IEEE Transactions on Visualization and Computer Graphics* 27, 2 (2020), 1204–1213.

[79] Andrew D Wilson. 2010. Using a depth camera as a touch sensor. In *ACM international conference on interactive tabletops and surfaces.* 69–72.

[80] Terri L Woods, Sarah Reed, Sherry Hsi, John A Woods, and Michael R Woods. 2016. Pilot study using the augmented reality sandbox to teach topographic maps and surficial processes in introductory geology labs. *Journal of Geoscience Education* 64, 3 (2016), 199–214.

[81] Guoyong Zhang, Jianhua Gong, Yi Li, Jun Sun, Bingli Xu, Dong Zhang, Jieping Zhou, Ling Guo, Shen Shen, and Bingxiao Yin. 2020. An efficient flood dynamic visualization approach based on 3D printing and augmented reality. *International Journal of Digital Earth* 13, 11 (2020), 1302–1320.

Received July 2021; revised September 2021; accepted September 2021