Don't Pull the Balrog — Lessons Learned from Designing Wizualization: a Magic-inspired Data Analytics System in XR

Peter W. S. Butcher* Bangor University
Andrea Batch† U.S. Bureau of Economic Analysis
Panagiotis D. Ritsos‡ Bangor University
Niklas Elmqvist§ Aarhus University

Figure 1: **Designing Wizualization.** Building the Wizualization system was a quest spanning many arduous months in the wilderness. Here we hope to share a tale of our travails for the benefit of those who want to follow in our footsteps. (Image by MidJourney v3.)

Abstract

This paper presents lessons learned in the design and development of Wizualization, a ubiquitous analytics system for authoring visualizations in WebXR using a magic metaphor. The system is based on a fundamentally hybrid and multimodal approach utilizing AR/XR, gestures, sound, and speech to support the mobile setting. Our lessons include how to overcome mostly technical challenges, such as view management and combining multiple sessions in the same analytical 3D space, but also user-based, design-oriented, and even social ones. Our intention in sharing these teachings is to help fellow travellers navigate the same troubled waters we have traversed.

Index Terms: H.1.2 [Models and Principles]: User/Machine Systems—Human factors; H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Artificial, augmented, and virtual realities; H.5.2 [Information Interfaces and Presentation]: User Interfaces—Graphical user interfaces.

1 Introduction

Everday wizardry is full of practical tips and tricks that the observant magic user will pick up during their life's journey or from fellow practitioners of the Art. Never pick toadstool after midnight lest the potency of the mushroom be diminished. Water spells are best cast with a strong wind at your back and not your front. Long hair and—for the male wizard—beard may be photogenic and suitably arcane, but can be hazardous during the evocation of fire without appropriate hairnets. And, indeed, don't pull the Balrog. He has a mean temper on account of the orcs' incessant drumming as well as the endless debates about whether his wings are figurative or literal.

*e-mail: p.butcher@bangor.ac.uk
†e-mail: andrea.c.batch@gmail.com
‡e-mail: p.ritsos@bangor.ac.uk
§e-mail: elm@cs.au.dk

In this paper, we attempt to share the same kind of practical tips and tricks from our experiences developing WIZUALIZATION [1], a WebXR system for visualization authoring built using a magic-inspired metaphor. Designed to support ubiquitous [9, 10], immersive [6], and situated analytics [11, 19], the Wizualization system supports both current Augmented Reality (AR) head-mounted displays (HMDs) such as the Microsoft HoloLens 2 and the upcoming Apple Vision Pro, as well as handheld Augmented Reality [1]. Because of the mobile and multimodal usage context of the system, it is based on a hybrid user interface approach [12] where multiple devices such as a handheld mobile phone or fixed static displays can augment the HMD.

The project was a significant engineering effort involving tens of thousands of lines of code written over the span of two years, and it is still ongoing. Akin to the experiences of fictional magic users above, we encountered many challenges during the design and development of the system: some of a technical nature, but also those of a more scientific, design-based, or even social one. Moreover, our project was carried out in parallel with conducting an extensive survey on the situated analytics landscape [19], and right after we developed and evaluated a series of view-management techniques, implemented in handheld AR using web technologies [1]. Both these activities contributed to the design of Wizualization, and consequently, bestow teachings to our lessons learned that we present here. The goal of this paper is to share these hard-won lessons in the hope that fellow ubiquitous and immersive analytics (UA/IA) practitioners can benefit from our experience.

2 Related Work

Ubiquitous analytics (UA) [10] is the overarching research area of data visualization conducted using post-WIMP devices and displays, such as mobile devices, large-display environments, and extended reality (XR) technologies [9, 17]. Immersive analytics (IA) [6] is a

subset of this area that deals specifically with immersive technologies, such as Mixed, Augmented, and Virtual Reality (MR/AR/VR). Situated analytics (SA) [11, 19] is a form of analysis conducted in physical space contextual to the data, typically using MR/AR. Finally, mobile data visualization [16] is a cross-cutting topic on how to use mobile devices for effective data representation.

Several UA/IA/SA system have been developed over the years; here follows a representative sampling. ImAxes [8] uses direct manipulation of multidimensional axes to allow users to create data visualizations in VR. ART [5] supports collaborative analysis using a 3D parallel coordinate application. ImEcon [2] extends the ImAxes system for an economics use case. MARVIS [15] is a framework combining mobile devices and AR HMDs for analysis. ReLive [14] combines ex-situ (2D desktop) and in-situ (3D Virtual Reality) views in the same tool; VRxD [18] adds collaborative mechanisms to this cross-platform setting. Finally, DXR [20], VRIA [4], and IATK [7] are toolkits for building immersive analytics applications that inspired the design and implementation of Wizualization.

3 Wizualization: Data Analytics in WebXR

Our starting point in designing Wizualization was to envision a user in a mobile setting who wants to analyze *in situ* data relevant to a location. In such a setting, the user would need to keep at least one or both hands free. They would also want to mostly keep their eyes on the world surrounding them, and not on a handheld display. This means that most **input** would presumably consist of gestures, voice commands, and simple touch interaction on handheld mobile devices. For **output** we would want to use an MR/AR head-mounted display (HMD), such as a Microsoft HoloLens 2, to enable integrating visual representations of data into the user's own surroundings. The device would be networked to enable the user to **collaborate** with other analysts sharing the same physical space or distributed across time and space. In other words, our vision for this ubiquitous analytics environment is fundamentally a *hybrid user interface* [12] consisting of multiple device modalities complementing each other.

The actual design of Wizualization was based on the observation that a user authoring visualizations and analyzing data in a mobile setting using an AR HMD would inevitably be behaving very much like a wizard conjuring visual constructs out of thin air. Hence, the tool is based on a "hard magic" metaphor; one where the rules of creating visualizations and connecting them to data is made explicit using a grammar (called Optomancy) that is similar to existing visualization grammars such as Vega/Vega-Lite, or those implemented in VRIA [4], and DXR [20]. Specifications using this grammar—aptly called *spells*—are built using a succession of voice commands, gestures, and touch input. Finally, the tool is realized solely using Web technologies, using the WebXR Device API [21] as well as React. Our paper on the Wizualization system [1] provides full technical details of the system. Figure 2 provides an overview of the main components of Wizualization.

4 Lessons Learned

The design and development of the Wizualization system spanned well over two years. During this time, the first two co-authors built the software under frequent consultations with the full author team. Several hurdles were encountered and overcome during this time; most of a technical nature, but some more social or even political. Here we list the lessons learned during this design process. Minor teachings are collected first; the larger lessons get a subsection each.

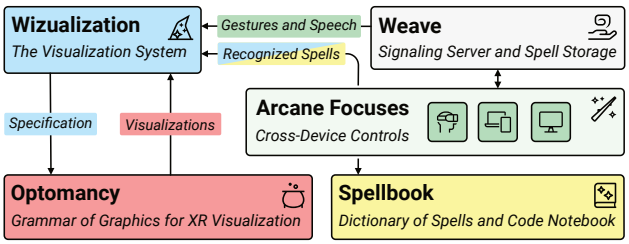

Figure 2: **Wizualization System Overview.** Main components: (i) Arcane Focuses provide support for different devices and user system control; (ii) Weave synchronizes the state of all connected clients (Arcane Focuses) and stores the state of cast spells in each room; (iii) Optomancy, the grammar of graphics we use to define visualizations in XR; and (iv) Spellbook provides a dictionary of spells, spell primitives, and mappings to Optomancy specification code blocks.

4.1 Minor Teachings

We encountered several minor challenges during the design process. While these teachings are by no means novel or original, some are not common to academic research projects, or possibly restricted to work that is predominately an 'engineering' contribution.

- **Professional software engineering management practices:** The meticulous use of version control, frequent backups, agile design processes, and strict typing yield robustness, even for a research project. This becomes even more important when implementation teams rely on less experienced developers (not in our case), or collaborations are done across institutions, in different locations and different timezones (as in our case).
- **Applications provide grounding:** Designing toolkits and libraries is notoriously difficult, but building real tools that use the toolkits at the same time can help ground the design.
- **Beyond the proof of concept:** Working on a proof of concept is often sufficient for research. Nevertheless, the applied nature of interface design, especially for visualization, does require careful consideration of how the systems we develop are used in practice, and what are the parameters that affect the resulting user experience. The experience of client developers that are to utilize a toolkit is particularly important.
- **Appeasing reviewers:** The applied nature of engineering work does not always go down well with scientific reviewers. The complexity of getting systems papers published in HCI and related disciplines is well known and documented.

4.2 The Problem with (the) Edge

Having access to the latest and greatest devices and APIs is a privilege for researchers, but oftentimes pushing these bleeding-edge technologies to their limits puts you firmly in uncharted territory. This can be both an exciting and an unpleasant place to be.

- **API Volatility:** Have you ever woken up to find everything in your room has moved, or that some things are missing entirely? For most people, this would merely be the stuff of nightmares, but for software developers living at the edge, this can be a regular occurrence as entire swathes of API endpoints can change, be removed, or break at the whim of their maintainers.
- **Device Support:** You've managed to get hold of the latest piece of tech, it's in your hands, you've written some code, and that concept you wanted to test works! But then, the powers

that be decide to lay off the workforce who were tasked to maintain the software on your new hardware, and uncontrollably, the usefulness of your hardware degrades from within, rendering your hard work unusable.

While this may sound extreme, it happened. Towards the end of development of the publication version of Wizualization, Microsoft released an update for its Edge browser on HoloLens 2 which degraded hand tracking capabilities to the point where finger and joint pose data was no longer accessible in the browser. A seemingly innocuous update to Edge, with no mention of changes to the WebXR Device API in the changelog, had knocked out a core feature of Wizualization. Spell casting and recording were no longer possible on the HoloLens.

Our team posted on the Microsoft Tech Community forum searching for answers, directly messaged Microsoft Employees on the WebXR Discord server, and even took to Twitter to try and raise awareness of the problem.

At the time of writing, hand pose data is still unavailable in the stable release of Edge on the Hololens 2. Forcing the device to use an older version of the browser is currently the only known way to restore hand tracking features.

This all raises worrying questions about the ability (or willingness) of industry, including corporations such as Microsoft, to sustainably maintain cutting edge APIs on their devices into the future.

4.3 More Danger Than Fiction

Metaphors have always been important in HCI; just witness the persistence of the desktop metaphor for graphical user interfaces, where documents, files, trashcans, and the desktop itself serve as useful and familiar shortcuts to computer concepts. It turns out that superpowers may be a useful metaphor and inspiration for traditional data visualization about amplifying cognition and making the invisible visible [22]. Similarly, in the area of ubiquitous and immersive analytics, we noted that a mobile approach using HMDs would be very much akin to a wizard weaving mysterious gestures in mid-air and mumbling voice commands, thus inspiring us to adapt a magic approach for our data analytics system.

The magic metaphor has been useful in our work because it provides a familiar mental image for both users and developers. We have let it guide not just the terminology and the interactions in the system, but also used it to recast our system architecture. Visualization specifications are *spells*, and spells can either be one off or saved in a *spellbook*—essentially a computational notebook for reusable specs. The devices that a user may add to the hybrid interface, such as a smartphone or a handheld controller, are *arcane focuses* that channel and concentrate the user's magic in an artifact.

At the same time, there are drawbacks to adopting a somewhat whimsical metaphor such as magic in a scientific research paper. We found out early on that not all reviewers appreciate levity in a research paper, and at least one of our two rejections of the original paper was likely caused by a reviewer protesting against some of our more lighthearted passages in the submission.

While we recognize that humor certainly is a subjective and even divisive issue, we still found this somewhat disappointing. In fact, recent research has shown that humorous paper titles are more likely to be cited than those without [13]. In addition, magic is now mainstream. The fantastical has shed its disreputable and lurid character of the 1940s and 50s. Brandon Sanderson, who coined the term "hard magic" that we used in the Wizualization paper [3], is a bestselling mainstream author whose books have sold 21 million copies; for the *The Lord of the Rings* the number is 150 million, making it the third most bestselling book of fiction of all time. Regardless of whether you stan superheroes or not, the fact that the Marvel Cinematic Universe and its ilk currently dominates cinema is not up for debate. We now feel vindicated in our approach.

4.4 Opportunities for Remote Evaluation

The development of Wizualization was carried out after our evaluation of view management techniques in handheld AR, also implemented using web technologies. Inevitably, much of our experience in engineering aspects for that project influenced aspects of the implementation of Wizualization. Nevertheless, one of the important elements of our view management evaluation was the fact it was carried out remotely (due to the pandemic restrictions) in a variety of supported hardware, in true open-standards web-based fashion.

Indeed, the main purpose of using web technologies is to access even advanced AR software in the browser, without requiring local installation. Two of the authors encountered this with the ReLive system [14]. Our experiences with our view-management evaluation show that web-based immersive evaluations can facilitate remote evaluation. This is beneficial compared to the more traditional game engine implementations for deployment and data collection.

On the other hand, web-based technologies are often challenging to harness due to compatibility issues, such as device and browser support for WebXR, limited rendering capacity and quality when compared to game engine,[1] as well as differences in the local networks of the participants. Nevertheless, we believe that carefully designed, standards-compliant web-based implementations can enable remote evaluation of VR and AR/MR implementations, providing improved versatility—and possibly improved ecological validity—when compared to game-engine systems.

5 Research Vision

As part of our background work for our view-management evaluation and the development of Wizualization, we reviewed the literature in situated visualization and situated analytics. Our goal was to define the design space for situated analytics, derive and classify design dimensions and identify usage archetypes. An important finding of this process is that the vast majority of the existing SA systems offer limited support for the sensemaking loop beyond reading and exploring information, with very few systems supporting schematizing and none supporting reporting mechanisms. Consequently, there is an opportunity to apply frameworks such as Wizualization to deep analysis scenarios.

In the future, we anticipate conducting laboratory and field studies to validate the utility of the Wizualization system. There is also opportunity for new visual representations, new mid-air interaction techniques, and new analytics methods specialized for mobile analytics. We look forward to building on the Wizualization system as we continue to answer such questions in our future research.

6 Conclusion

In this paper, we discuss our experiences from designing and implementing Wizualization [1], a WebXR system for visualization authoring built using a magic-inspired metaphor. Wizualization is designed to support ubiquitous, immersive, and situated analytics, through AR HMDs such as the Microsoft HoloLens 2 and the upcoming Apple Vision Pro, as well as handheld devices. Our experience demonstrates that although using Web-technologies for such an endeavor is challenging, the potential for synergies with other systems, domains, and disciplines is very high. Moreover, drawing inspiration from unlikely or unexpected domains such as popular fiction literature provides inspiration, prompt us to question 'conceptional' practices of the past, and makes use of the often untapped potential

[1] Although performance gaps are diminishing!

of human inspiration that finds place in books, movies, comics, and gaming (some serious, and some less so).

ACKNOWLEDGMENTS

This work was supported by the grant IIS-1901485 from the U.S. National Science Foundation as well as the DSP Centre funded by the ERDF through the Welsh Government. Any opinions, findings, and conclusions expressed in this material are those of the authors and do not necessarily reflect the views of the funding agency.

REFERENCES

[1] A. Batch, P. W. S. Butcher, P. D. Ritsos, and N. Elmqvist. Wizualization: A "hard magic" visualization system for immersive and ubiquitous analytics. *IEEE Transactions on Visualization and Computer Graphics*, 30(1), 2024. to appear.

[2] A. Batch, A. Cunningham, M. Cordeil, N. Elmqvist, T. Dwyer, B. H. Thomas, and K. Marriott. There is no spoon: Evaluating performance, space use, and presence with expert domain users in immersive analytics. *IEEE Transactions on Visualization and Computer Graphics*, 26(1):536–546, 2020. doi: 10.1109/TVCG.2019.2934803

[3] A. Batch, S. Shin, J. Liu, P. W. S. Butcher, P. D. Ritsos, and N. Elmqvist. Evaluating view management for situated visualization in web-based handheld AR. *Computer Graphics Forum*, 42(3):349–360, 2023. doi: 10.1111/cgf.14835

[4] P. W. S. Butcher, N. W. John, and P. D. Ritsos. VRIA: A web-based framework for creating immersive analytics experiences. *IEEE Transactions on Visualization and Computer Graphics*, 27(7):3213–3225, 2021. doi: 10.1109/TVCG.2020.2965109

[5] S. Butscher, S. Hubenschmid, J. Müller, J. Fuchs, and H. Reiterer. Clusters, trends, and outliers: How immersive technologies can facilitate the collaborative analysis of multidimensional data. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, p. 1–12. ACM, New York, NY, USA, 2018.

[6] T. Chandler, M. Cordeil, T. Czauderna, T. Dwyer, J. Glowacki, C. Goncu, M. Klapperstueck, K. Klein, F. Schreiber, and E. Wilson. Immersive analytics. In *Proceedings of the International Symposium on Big Data Visual Analytics*, pp. 1–8. IEEE, Piscataway, NJ, USA, 2015. doi: 10.1109/BDVA.2015.7314296

[7] M. Cordeil, A. Cunningham, B. Bach, C. Hurter, B. H. Thomas, K. Marriott, and T. Dwyer. IATK: an immersive analytics toolkit. In *Proceedings of the IEEE Conference on Virtual Reality*, pp. 200–209. IEEE Computer Society, Los Alamitos, CA, USA, 2019. doi: 10.1109/VR.2019.8797978

[8] M. Cordeil, A. Cunningham, T. Dwyer, B. H. Thomas, and K. Marriott. ImAxes: Immersive axes as embodied affordances for interactive multivariate data visualisation. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pp. 71–83. ACM, New York, NY, USA, 2017. doi: 10.1145/3126594.3126613

[9] N. Elmqvist. Anywhere & everywhere: A mobile, immersive, and ubiquitous vision for data analytics. *Communications of the ACM*, 2023. To appear.

[10] N. Elmqvist and P. Irani. Ubiquitous Analytics: Interacting with Big Data Anywhere, Anytime. *Computer*, 46(4):86–89, 2013.

[11] N. A. M. ElSayed, B. H. Thomas, K. Marriott, J. Piantadosi, and R. T. Smith. Situated analytics: Demonstrating immersive analytical tools with augmented reality. *Journal of Visual Languages & Computing*, 36:13–23, 2016. doi: 10.1016/j.jvlc.2016.07.006

[12] S. Feiner and A. Shamash. Hybrid user interfaces: breeding virtually bigger interfaces for physically smaller computers. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pp. 9–17. ACM, New York, NY, USA, 1991. doi: 10.1145/120782.120783

[13] S. B. Heard, C. A. Cull, and E. R. White. If this title is funny, will you cite me? citation impacts of humour and other features of article titles in ecology and evolution. *FACETS*, 8:1–15, 2023. doi: 10.1139/facets-2022-0079

[14] S. Hubenschmid, J. Wieland, D. I. Fink, A. Batch, J. Zagermann, N. Elmqvist, and H. Reiterer. ReLive: Bridging in-situ and ex-situ visual analytics for analyzing mixed reality user studies. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 24:1–24:20. ACM, New York, NY, USA, 2022. doi: 10.1145/3491102.3517550

[15] R. Langner, M. Satkowski, W. Büschel, and R. Dachselt. MARVIS: Combining mobile devices and augmented reality for visual data analysis. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. ACM, New York, NY, USA, 5 2021. doi: 10.1145/3411764.3445593

[16] B. Lee, R. Dachselt, P. Isenberg, and E. K. Choe. *Mobile Data Visualization*. AK Peters, Boca Raton, FL, USA, 2022.

[17] J. C. Roberts, P. D. Ritsos, S. K. Badam, D. Brodbeck, J. Kennedy, and N. Elmqvist. Visualization beyond the desktop – the next big thing. *IEEE Computer Graphics and Applications*, 34(6):26–34, Nov. 2014. doi: 10.1109/MCG.2014.82

[18] D. Saffo, A. Batch, C. Dunne, and N. Elmqvist. Through their eyes and in their shoes: Providing group awareness during collaboration across virtual reality and desktop platforms. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 383:1–383:15. ACM, New York, NY, USA, 2023. doi: 10.1145/3544548.3581093

[19] S. Shin, A. Batch, P. W. S. Butcher, P. D. Ritsos, and N. Elmqvist. The reality of the situation: A survey of situated analytics. *IEEE Transactions on Visualization and Computer Graphics*, 30(1), 2024. to appear.

[20] R. Sicat, J. Li, J. Choi, M. Cordeil, W.-K. Jeong, B. Bach, and H. Pfister. DXR: a toolkit for building immersive data visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):715–725, 2019. doi: 10.1109/TVCG.2018.2865152

[21] W3C. *WebXR*. [Online]. Available `https://www.w3.org/TR/webxr/`, 2017. Accessed: 07/11/22.

[22] W. Willett, B. A. Aseniero, S. Carpendale, P. Dragicevic, Y. Jansen, L. Oehlberg, and P. Isenberg. Perception! Immersion! Empowerment! Superpowers as inspiration for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):22–32, 2022. doi: 10.1109/TVCG.2021.3114844