

Attention-Aware Visualization: Tracking and Responding to User Perception Over Time

Srinivasan, Arvind; Ellemose, Johannes ; Butcher, Peter; Ritsos, Panagiotis D.; Elmqvist, Niklas

IEEE Transactions on visualization and computer graphics

DOI:
[10.1109/TVCG.2024.3456300](https://doi.org/10.1109/TVCG.2024.3456300)

E-pub ahead of print: 09/09/2024

Peer reviewed version

[Cyswllt i'r cyhoeddiad / Link to publication](#)

Dyfyniad o'r fersiwn a gyhoeddwyd / Citation for published version (APA):
Srinivasan, A., Ellemose, J., Butcher, P., Ritsos, P. D., & Elmqvist, N. (2024). Attention-Aware Visualization: Tracking and Responding to User Perception Over Time. *IEEE Transactions on visualization and computer graphics*, 1-11. Advance online publication. <https://doi.org/10.1109/TVCG.2024.3456300>

Hawliau Cyffredinol / General rights






Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Attention-Aware Visualization: Tracking and Responding to User Perception Over Time

Arvind Srinivasan* , Johannes Ellemose* , Peter W. S. Butcher , Panagiotis D. Ritsos , Niklas Elmqvist 




Fig. 1: **Attention-aware visualization.** Our work proposes tracking, recording, and revisualizing a user's attention on a visualization to aid them in understanding what they have seen and not seen, as well as cause the visualization to respond to the user's gaze. Our two implementations—3D and 2D, respectively—make use of the head direction of a Meta Quest 3 or similar XR HMD, the mouse pointer, or an embodied eye-tracker to measure the user's gaze on the visualization. We also present several methods for revisualization, including heatmaps and contour plots overlaid on the chart, or compact visualizations on the border of the viewport.

Abstract—We propose the notion of *attention-aware visualizations* (AAVs) that track the user's perception of a visual representation over time and feed this information back to the visualization. Such context awareness is particularly useful for ubiquitous and immersive analytics where knowing which embedded visualizations the user is looking at can be used to make visualizations react appropriately to the user's attention: for example, by highlighting data the user has not yet seen. We can separate the approach into three components: (1) measuring the user's gaze on a visualization and its parts; (2) tracking the user's attention over time; and (3) reactively modifying the visual representation based on the current attention metric. In this paper, we present two separate implementations of AAV: a 2D data-agnostic method for web-based visualizations that can use an embodied eyetracker to capture the user's gaze, and a 3D data-aware one that uses the stencil buffer to track the visibility of each individual mark in a visualization. Both methods provide similar mechanisms for accumulating attention over time and changing the appearance of marks in response. We also present results from a qualitative evaluation studying visual feedback and triggering mechanisms for capturing and revisualizing attention.

Index Terms—Attention tracking, eyetracking, immersive analytics, ubiquitous analytics, post-WIMP interaction.

1 INTRODUCTION

The art and science of visualization design is primarily guided by a plethora of principles and guidelines—some empirically validated, others more anecdotal—in the pursuit of crafting visual representations that are not only *expressive and effective* [30, 32] but also *visually salient* [25, 31]. This extensive body of knowledge spans both the practical realm of visualization practitioners and the burgeoning academic discipline. However, these visualization design principles primarily focus on creating static visual artifacts that only respond to direct user interaction, such as explicit clicks and taps. What if we could transcend this passive paradigm to design visualizations that actively engage with the viewer's attention? Imagine visualizations imbued with the capability to track and respond to the viewer's gaze, dynamically adapting their visual attributes in real time to optimally convey the data.

We introduce the concept of *attention-aware visualizations* (AAVs):

contextually aware visual representations that track user attention on their surface over time and modify their visual attributes accordingly. Conceptually, this approach consists of three steps (Fig. 1):

1. **Measuring user attention** on a visualization and its marks. This could be as simple as estimating attention based on the scroll location or mouse cursor [53], the user's head orientation, or their exact gaze location using an eyetracker.
2. **Recording attention over time** as the user's perception (i.e. gaze) moves over the visualization over time. This also includes modeling the user's short-term memory as attention smoothly decaying.
3. **Modifying the visualization** in response to user attention. This could be used as a diagnostic—allowing the user to see which part of visualization they have surveyed and which they have not—or as a subtle visual cue for read wear [22].

While there are numerous potential approaches to designing AAVs, this paper focuses on the fundamental idea of showing users their own attention on a visualization. Our goal is to encourage users to explore areas of a visualization they may not have previously investigated as well as to reflect on the areas they have, thus enhancing their overall understanding and interaction with the data.

The contributions of this paper are as follows: (1) the concept of attention-aware visualization as a context-aware [41] method for creating visualizations responsive to implicit gaze and not just explicit interaction; (2) a library implementing attention-awareness for 2D visualizations based on numeric integration; (3) an implementation of attention-awareness and occlusion [16] for 3D visualizations using the

* *Equal contribution.*

• Arvind Srinivasan, Johannes Ellemose, and Niklas Elmqvist are with Aarhus University in Aarhus, Denmark. E-mail: {arvind, ellemose, elm}@cs.au.dk.
• Peter W. S. Butcher and Panagiotis Ritsos are with Bangor University, Bangor, United Kingdom. E-mail: {p.butcher, p.ritsos}@bangor.ac.uk.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org.
Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

stencil buffer; and (4) results from a qualitative evaluation investigating the approach for both 2D (eye-tracker) and 3D settings (XR HMD). Refer to [Sec. 9](#) for the supplementary materials.

2 BACKGROUND

In this section we discuss our inspirations from prior research for this work, which include fundamental concepts such as context awareness, low-level aspects of attention in visualization, eyetracking as a measure of attention in data visualization, and the use of gaze for interaction in general HCI systems. Note that in this paper we do not consider attention and awareness beyond the individual [20], such as for collaboration, territoriality, and conflict management in group work.

2.1 Context Awareness

One of the primary inspirations for this work is the notion of context-aware (CA) systems [41]. CA systems monitor and react to a user's changing context to promote and mediate said user's interaction with the system itself, as well as other systems and users. The term *context* is defined in the literature in several ways [4], with the most popular definition from Dey [14, p. 5], as “any information that can be used to characterize the situation of an entity,” where “an entity can be a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and applications themselves.” More importantly, for the context (pun intended) of visualization explored in this work, Dey [14, p. 5] defines a system as CA if “it uses context to provide relevant information and/or services to the user, where relevancy depends on the user's task.”

Examples of such context range from tracking location, environmental factors, usage patterns, and user behavior. Such context-aware systems are particularly central within the areas of mobile computing [5], the internet of things [35], augmented reality [19], and wearable ubiquitous computing [1, 52]. This includes ubiquitous [15] and immersive/situated analytics [42], which is why we explore context awareness as a concept in both 2D and 3D in this paper.

2.2 Attention in Visualization

The human perceptual and cognitive systems generate a wealth of information in real time, all of which cannot be processed simultaneously. *Attention* is the selective concentration of mental awareness on a subset of this information, excluding the remaining information [46]. Early work on cognitive processes in attention noted that for highly complex attentive processes—such as visual search—to be possible, they must be organized in both low-level parallel—so called *pre-attentive*—and complex serial—*attentive*—processes [23, 33]. While subsequent work has criticized the existence of a clear division between pre-attentiveness and attentiveness [54], it remains a useful construct for understanding operational aspects of human perception and attention, e.g., for visual search optimization [29].

For data visualization, the work by Healey et al. [21] has been instrumental in understanding the effective use of color, shape, and motion in preattentive processing of complex visual representations. Colin Ware [50] emphasize the cognitive limits of visual attention and memory, and provide guidelines for designing more effective visualizations that cater to human perception. Several specialized topics related to attention has been studied within the visualization field, including change blindness [34], visual salience [25, 31], and the use of eye tracking for visualization. We discuss the latter in more detail next.

2.3 Eye Tracking for Visualization

Eye tracking is a useful tool for visualization researchers [12]; it can help identify visual exploration behaviors [13], assess perceptual and cognitive aspects of engaging with visualizations [26], and even be used in combination with other usability metrics to understand user behavior [7]. Examples include the work of Burch et al. [13] on tree-drawing perception, Alam et al. [2] on eye tracking analysis in the data space instead of the commonly used image space, and Polatsek et al. [39] on evaluation of salience models for visual analysis. On the other hand, Kim et al. [26] assess the limitations of eye trackers in information visualization studies and find that in certain situations

inferences about the underlying cognitive processes may be misleading. Kurzahls et al. [27] present a survey of eye tracking technologies for the evaluation of visualization techniques, identifying future directions at the time of publication. In more recent work Shin et al. [43] review prior work on eye tracking and propose a novel approach of creating a virtual eye tracker that produces gaze heatmaps using deep learning.

Another take on visualization and eye tracking is the use of visualization for eye tracking data. Blascheck et al. provide a somewhat dated survey on this practice from 2014 [8]. Notable examples include Tsang et al.'s approach for visualizing fixation sequences [47], Alam and Jianu [3] use visualization to connect gaze patterns with visual elements on a screen, and the VETA system [18] combines eye tracking visualization with video of the underlying scene. Our work in this paper both uses eye tracking (in addition to other user input, such as touch events and mouse actions) to collect the context from the user as well as visualize attention collected from eye tracking data on a visualization. Our attention-aware visualization approach both builds on this past work and extends it significantly by using the eye tracking data as a first-class input channel and not merely a diagnostic tool; more below.

2.4 Gaze as Input

The use of gaze as a method for computer input is marked by pioneering efforts to establish its viability and utility. Bolt [9] demonstrated one of the earliest implementations, showing how dynamic window systems could be orchestrated through gaze. This foundational work led to followup efforts, including Ware and Mikaelian [51], who evaluated eye tracking for computer input, and Jacob's proposal of eye movement-based techniques that link gaze direction with interface control [24].

As gaze tracking technology matured, researchers began exploring its integration with other input modalities to create richer, more nuanced interaction paradigms. Zhai et al.'s MAGIC pointing [55], Salvucci and Anderson's gaze-added interfaces [40], and Fono and Versteeg's Eye-Windows [17] each contribute to this domain by cascading gaze with manual inputs to streamline selection tasks and manage focus within interfaces. Pfeuffer et al. introduced gaze-touch in 2014 [36], combining gaze with multi-touch on the same surface and thereby illustrating the synergistic potential of gaze with other interaction techniques.

Over the past decade, there has been significant progress in gaze-supported multimodal interactions, particularly for exploring large datasets and improving target acquisition. Stellmach et al. [45] and Stellmach and Dachsel [44] explored gaze-supported interactions for large image collection exploration and target acquisition, respectively. Turner et al. [48], Pfeuffer et al. [36, 37], and Velloso et al. [49] further investigated the integration of gaze with touch and mid-air gestures, demonstrating the potential for seamless and efficient interactions across devices and interfaces. Pfeuffer et al.'s work on combining gaze with pinch interactions in virtual reality [38] represents the latest evolution of this research direction, highlighting the growing sophistication and breadth of gaze-supported interactions.¹ Although this research is closely related to AAVs, our approach focuses less on utilizing direct gaze as an input to control an interface, as explored in [28], and more on the retrospective analysis of gaze patterns over time.

3 DESIGN: ATTENTION-AWARE VISUALIZATION

The intellectual contribution proposed in this work is the concept of *attention-aware visualizations* (AAV): data visualizations that track the viewer's attention on their surface over time and react accordingly. We see several benefits to this approach, ranging from creating a more responsive and engaging visual environment, to showing data coverage, illuminating missed outliers, and even predicting attention patterns. Although numerous possibilities exist for how a visualization could respond to a viewer's attention, this paper focuses solely on displaying the viewer's attention on the visualization. We can summarize our approach using the following research questions:

RQ1 How should a viewer's attention on a visualization be **measured in real time**?

¹Incidentally, gaze and pinch has also been adopted by the Apple Vision Pro, making it the first commercial device using dedicated gaze interaction.

RQ2 How should a viewer’s attention on a visualization be **recorded over time**?

RQ3 How should a visualization be **modified to display** the viewer’s attention over time?

Here we propose our design framework for attention-aware visualizations that addresses each of these research questions in turn.

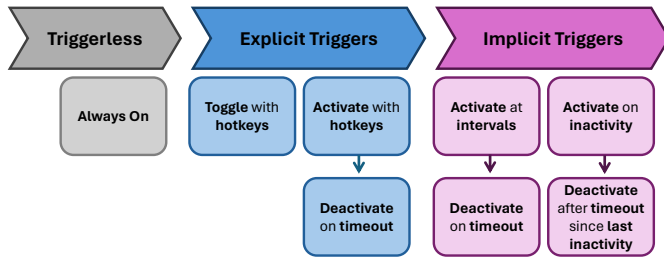


Fig. 2: **Triggering strategies.** Overview of the different triggering strategies for our attention-aware visualization interventions.

3.1 Measuring Attention

First of all, what is *attention*? As discussed in the Background (Sec. 2), there are several interpretations. From a cognitive perspective, attention refers to the selective focusing of mental resources on specific information or stimuli. Even though this view encompasses all senses, we focus on vision here since most interactive visualizations primarily leverage the visual system. Attention, then, determines which aspects of the visual scene are prioritized for mental processing, guiding where and how we allocate our visual resources.

Given all that, how should we measure it (RQ1)? Since it is a purely cognitive process, attention cannot be measured in its own right—at least not with widely available equipment. However, depending on the context when using a computing device, there are several situations when interaction with the device can be used to measure attention:

- 👉 *Taps on a touch screen* signifies an interest in that area (if the touch was intentional);
- 🖱️ *Pointer movement* on the area surrounding the pointer or its destination (if the pointer is being used deliberately); and
- 👁️ *Eye movement* on the screen (if the user is viewing the device actively and not preoccupied by some other task).

We have experimented with all forms of attention measurement. They all have their particular uses, especially since some of them are less available than others. For example, eye movement is indisputably the most useful metric, but capturing user gaze requires a dedicated eyetracker to achieve high accuracy, and these are typically expensive and not widely available. Mouse movement, on the other hand, is trivially available on most computing platforms, but may not always be deliberate—and it can be hard to determine when this is the case.

3.2 Recording Attention

We assume a reliable attention measurement mechanism that reports the user’s attention on a visualization in real time. We further assume that the attention focus is a specific point in the user’s point of vision, presumably with a specific area surrounding it and with some degree of uncertainty depending on the measurement device (e.g., an embodied eyetracker being much more accurate than a webcam-based eyetracker).

Now, how do we record this attention over time (RQ2)? At this point, we can identify two separate approaches to recording attention depending on whether we have knowledge of the visualization or not:

- **Data-agnostic:** We do not have specific knowledge of the underlying data in the visualization being tracked—in other words, we do not know what the pixels of the visualization represent, only that the user’s attention is moving over it; and

- **Data-aware:** We have specific knowledge of exactly what parts—what marks representing specific data items—of the visualization that the user is focusing their attention on.

Data-agnostic. Since we have no knowledge of the underlying visualization or its data, we cannot directly associate attention with a specific mark or data item. Rather, this approach requires subdividing the visual space of the visualization into discrete bins—typically square cells on a regular grid—that we can use to accumulate attention over time. Unfortunately, this limits our chances of detecting when the user’s attention is drifting because we cannot detect situations when they are fixating on empty space in a visualization.

Data-aware. Because we can associate points on the visualization with visual marks, we can directly accumulate user attention on these marks (or even parts of marks). We can also discard attention that is spent on empty parts of a visualization—although we naturally still cannot reliably understand when the user is absent-mindedly viewing the visualization without actually seeing it.

Representing attention. Each *attention target*—a cell in the agnostic case or a visual mark in the aware case—can be assigned an attention value that starts at zero and then accumulates for every time unit the user’s area of attention overlaps with the cell or mark. If more than one attention metric is active at the same time—say, pointer and gaze—then each target has two attention values. Finally, attention should be normalized or capped across the entire visualization so that viewing one target for a long time will not overshadow all other targets.

Attention decay. People have limited short-memories and thus attention decays over time. In other words, just because a person has spent a long time looking at a specific set of bars in a barchart does not mean that they will be indelibly fixed in the person’s mind. After a while looking at other parts of the chart, high attention values for specific attention targets should decay towards zero over time.

3.3 Revisualizing Attention

How do we actually use the accumulated attention metric to change the visualization (RQ3)? One option is to not change the visualization at all, but to merely use the data for evaluation purposes where the viewers themselves do not ever see it. However, in this paper, we are interested in investigating how this data can be utilized by the visualization itself, e.g. to encourage exploration, or facilitate introspection. We call this *revisualizing* the attention on the original visualization.

Presentation. There are several potential methods of presenting the accumulated attention metric. They differ in how invasive they are to the target visualization:

- **Marks:** Modify the original marks to convey attention;
- 📄 **Overlay:** Visualize attention on an overlay on top of the view;
- 🏠 **Border:** Use view borders to visualize attention; and
- 📍 **Minimap:** Visualize attention on a separate window.

There are many options for the visual representation for the attention metric. Their availability depends vastly on whether the approach is data-agnostic or data-aware. For the former, the marks themselves cannot be changed since we have no knowledge of the underlying visualization. Instead, we have to use image processing operations on cells to convey the data, such as color transforms (hue rotate, desaturation, and luminance), blur, transparency, heatmaps, etc. For the latter, we can modify the visual attributes of the marks using object-level operations.

Triggering. Finally, because visualizing attention while tracking attention easily becomes a circular and self-reinforcing system, we must be careful to consider when and how to trigger these techniques. Figure 2 gives an overview of our proposed triggering strategies. A fully naïve solution would be triggerless and “always-on,” both for the attention collection as well as the revisualization. However, this would invariably mean that as the revisualization updates in response to the user’s attention, the change will draw the user’s attention, and so on. This will lead to an unstable and potentially confusing display.

We further delineate two additional strategies: *explicit vs. implicit* triggers. The former would require an explicit action to activate on behalf of the user; for example, checking a box or pressing a hotkey. This has the disadvantage that the trigger is outside the main workflow of the visualization and may thus be forgotten and never triggered. The latter approach would activate not by user action, but by some external implicit event, such as user inactivity, regular intervals, or perhaps the presence of consistently overlooked parts of the data. It makes sense that for both these triggers, attention data collection is disabled while the visualization is active to prevent reinforcing loops.

4 IMPLEMENTING ATTENTION-AWARE VISUALIZATION

We have implemented two versions of the attention-aware visualization method (Sec. 3), one for 2D settings using the data-agnostic approach, the other for 3D and the data-aware approach. Providing two separate implementations enables us to explore two quadrants in the design space in Tab. 1. The other two quadrants in the design space have not been implemented, but the extension should be trivial.

At the core of this endeavor lies a common model for measuring, recording, and revisualizing attention for both implementations. In this section, we describe this common model. In the following two sections, we describe the 2D and 3D implementations in detail, respectively.

Table 1: **AAV implementations.** We provide two AAV implementations; a data-agnostic one for 2D, and a data-aware one for 3D.

	Data-agnostic	Data-aware
2D	AAV2D (Sec. 5)	–
3D	–	AAV3D (Sec. 6)

4.1 Attention Model

Drawing on Sec. 3, the common attention model underpinning both of our implementations conceptualizes attention as a transient, ephemeral attribute that mirrors the characteristics of human short-term memory. This model is based on measuring the user’s gaze at specific elements within the visualization—whether these are discrete *bins* in a data-agnostic approach, or distinct visual *marks* (representing data points) in a data-aware framework. The attention allocated to these elements will accordingly escalate at a specific rate. The increment reflects the elements’ entry into the viewer’s short-term memory, akin to how prolonged observation may enhance memorability. However, once the viewer’s gaze shifts away, the model dictates that attention diminishes at a predetermined decay rate, analogous to the fading of details from short-term memory.

In practice, our model manages two *attention maps* for the visualization: (1) the global cumulative attention, and (2) the current short-term attention. For a data-aware implementation, an attention map is an associative array recording attention value per data point. For a data-agnostic one, an attention map is a spatial data structure (2D planar or 3D volumetric grid) recording attention per spatial cell in the visualization substrate. The global cumulative attention map does not decay over time, but keeps track of the attention measure over the entire lifecycle of the visualization. It is used primarily as a diagnostic and is not directly revisualized to the user (except potentially at the end of a session). We use the cumulative attention map to update the short-term attention map; these current attention values will ebb and flow as the user’s gaze moves across different parts of the visualization. This short-term attention map is used to determine both **when** (triggering) and **how** (revisualization) to render the user’s attention on the visualization.

4.2 Triggering Mechanisms

As noted in Sec. 3.3, we identify three strategies for triggering the revisualization of the visualization: ALWAYS-ON, EXPLICIT triggering, and IMPLICIT triggering (Fig. 2). ALWAYS-ON is constantly active. Depending on the revisualization approach, this behavior may lead to the revisualization itself affecting the user’s attention.

For EXPLICIT triggering, we activate revisualization when the user presses a hotkey or controller button—potentially spring-loaded, so

that when the button or key is released the revisualization disappears. Furthermore, as long as the revisualization is visible, we temporarily disable capturing attention to avoid self-reinforcing behavior.

Finally, for IMPLICIT triggering, we distinguish between two forms of revisualization: *emphasis* and *de-emphasis*. When the short-term attention for an entity (bin or mark) has decayed below a lower threshold, this means that there is a risk that the user is not seeing the entity, so the emphasis revisualization is activated. Analogously, when the short-term attention has increased above an upper threshold, the de-emphasis revisualization will be activated to discourage further study. Once the value rises above the lower or dips below the upper thresholds, the revisualization is de-activated. Specific implementations may choose to disable the emphasis or de-emphasis revisualizations as needed.

4.3 Revisualizing Attention

Finally, we use a common revisualization approach for both our 2D and 3D implementations. Common between them is the use of a straightforward heatmap with a quantitative color scale to revisualize attention. Furthermore, as discussed above, the notion of emphasis and de-emphasis is used for implicit triggering, and is similar across both implementations. We use *color saturation* as a common denominator for both: reducing saturation to de-emphasize and increasing it to emphasize. In addition, as discussed below, in the 2D case we also support blur to de-emphasize a segment of a visualization.

In the 3D implementation, we apply all of our revisualizations on the 3D scene itself; see Sec. 6 for detail. However, in the case of the 2D implementation, the visualizations are not immersive but bounded in 2D space, which allows us to revisualize attention also outside the visualization itself. This has the added benefit of allowing us to disable capturing attention when the user’s gaze is directed outside of the visualization itself. See Sec. 5 for more details.

The use of color to encode the revisualization of course means that color cannot be used to simultaneously encode dimensions of the data, which is not desired. However for immersive environments, using e.g. scale to denote attention would skew the perception of the data. For this paper we consider it a necessary compromise, and using an EXPLICIT triggering approach still allows for encoding data using color.

5 DATA-AGNOSTIC AAV IN 2D

AAV2D is our 2D implementation of attention-aware visualization. It is a data-agnostic JavaScript package made available on npm. Because it is data-agnostic, it can decorate any given HTML element in a DOM and decorate it with attention capture and visualization. Below we discuss attention capture, revisualization methods, and implementation.

5.1 Attention Capture

The AAV2D package maintains both a global cumulative and a short-term current attention map, as described in Sec. 4.1, in the form of a 2D grid fitted to the attached element. The size of individual grid cells can be controlled through programming-level and user-level options; see below. The AAV2D package installs input event handlers on the element, including both mouse and eye-tracker events (if available). This makes it possible to track the user’s gaze or mouse pointer position whenever it moves within the bounds of the element. For both approaches, the user’s attention point is treated like the center of an “attention circle,” and the package accumulates attention to all cells intersected by this circle. Attention data outside the region is ignored. This 2D grid emits events at the individual cell level, enabling the users of this library to create custom renderers if they so desire, similar to ones used in the study with the corresponding triggering mechanisms.

5.2 Revisualization

Our AAV2D implementation uses a “picture framing” metaphor for its revisualization mechanism, supporting several of the presentation methods in Sec. 4.3. More specifically, attaching AAV2D to an HTML element is seen as *mounting* a picture. The Overlay presentation is a *glaze* applied on top of the picture. The library also provides a Border presentation that serves as the *decorative frame* of the picture. Finally, the *Minimap* is provided as another decoration on the frame. The frame




Fig. 3: **AAV2D picture framing metaphor.** AAV2D provides several methods for revisualizing attention as an (a) **OVERLAY**, either on the **MOUNT** that holds the visualization or on a (b) **MINIMAP** around the (c) decorated **BORDER** that revisualizes attention.

is also used for showing an options panel that can be toggled to be visible or not. Figure 3 shows a visual overview of this method.

Furthermore, AAV2D supports several revisualization approaches depending on the presentation method. For both the glaze overlay and border minimap, the approach can use a heatmap, contour plot, or mesh (Fig. 4); in the case of the overlay, this revisualization is shown on top of the element, whereas for the minimap, it is shown on the miniature. For the frame border, we support bar charts, area charts, and heatmaps that show the short-term or cumulative attention on the plot for a specific x or y coordinate; see Fig. 5 for examples.

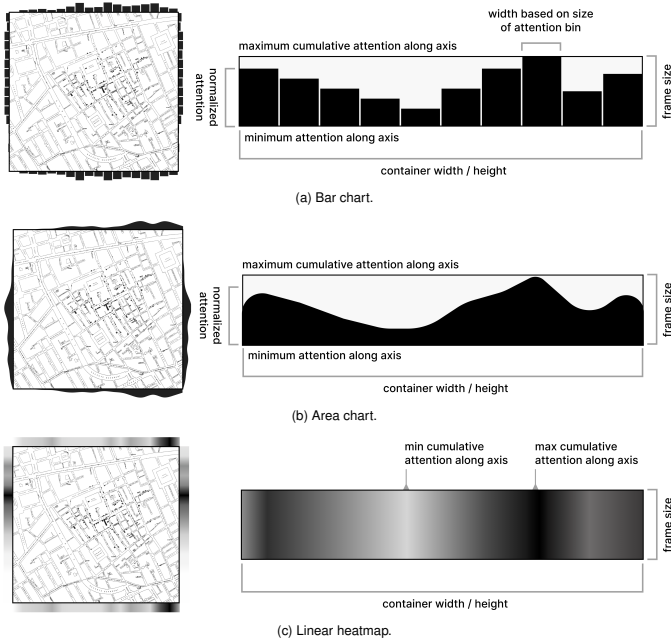


Fig. 5: **AAV2D border revisualizations.** AAV2D provides several methods for revisualizing attention on the decorative border of the element, each of which can be tailored to fit the mounted visualization.

5.3 Implementation Notes

We implemented AAV2D using vanilla JavaScript without the use of specialized libraries or bindings. D3² [10] was used for composing Scalable Vector Graphics (SVG) of contour and heatmaps. The mesh implementation uses CSS image filters to perform some of its visual transformations on the underlying HTML element, such as blur, saturation, and hue-rotation. All input events are intercepted by AAV2D event handlers, but are then propagated to the underlying element, meaning that it can be used even for interactive visualizations and interfaces.

²<https://d3js.org>




Fig. 4: **AAV2D overlay (glaze) revisualizations.** The three primary revisualization mechanisms include (a) **heatmap**, (b) **contour plot** and (c) **mesh**. The colors shown in the figure are tailored to fit the corresponding mounted visualization, which in this case is *John Snow's Cholera Map*.




Fig. 6: **AAV3D revisualizations.** AAV3D provides several methods for revisualizing attention on the visual marks themselves.

6 DATA-AWARE AAV IN 3D

Here we present AAV3D, our 3D data-aware implementation of attention-aware visualizations suitable for immersive analytics settings. Below we discuss the attention capture and triggering, revisualization, and notes about our implementation.

6.1 Attention Capture

The implementation has two strategies for capturing attention: using screen coordinates, and using the center of the screen. The screen coordinates strategy allows attention to be captured for elements currently rendered at any coordinate of the screen. This is useful for a non-immersive setting, where the interactive 3D scene is rendered on a normal display. In such a setting, the screen coordinates can be found with an eye-tracker, or by using the cursor as a proxy for gaze.

The second strategy only apply attention to elements in the center of the screen. This strategy is specifically designed for immersive environments using XR headsets that do not include eyetracking, such as the Meta Quest 3 (which was the target implementation platform for AAV3D). In such immersive environments, the user's head orientation thus acts as a proxy for their attention.

Common for both of these strategies is that they apply attention to all elements that are within a specified distance from the user's gaze position or center of the screen. This allows attention to be captured for the general area that the user's attention is directed. By carefully selecting an appropriate distance, we can approximate the user's attention from their head orientation in XR.

Beyond recording attention to individual visual marks, we can capture attention at the level of individual 3D triangles based on their visibility to the user. This can be useful to deal with occlusion arising from other objects in the scene, or from hidden surfaces facing away from the user. We go into details on how this is done in Sec. 6.3.

6.2 Revisualization

Our implementation currently supports two types of visualizations:

- **Heatmap** based on the normalized, cumulative attention measured to all parts of the visualization. This revisualization recolors each face of the 3D geometries in the visualization based on the cumulative amount of attention the face has received, normalized between zero and the maximum amount of attention.
- **Emphasis/de-emphasis** based on increasing and decreasing the color saturation of marks in the visualization. The emphasis effect increases saliency by recoloring the mark into a bright yellow, while the de-emphasis effect lowers visual saliency by desaturating the mark through a filter. Our implementation currently uses the short-term attention for this approach.

The 3D implementation follows the triggering strategies laid out in Sec. 4.2, and each of the revisualizations could be applied to each of the triggering strategies. However, we have currently only implemented emphasis/de-emphasis for IMPLICIT triggering.

6.3 Implementation Notes

Our 3D visualization implements bespoke visualizations in a WebXR application using the Three.js library.³ Attention is tracked through GPU color-picking on a per face basis; i.e., each triangle of every object’s geometry is assigned a unique RGB color in the picking scene. The red RGB colorspace is reserved for object IDs, while the green and blue RGB colorspace are used for faces. A benefit of using GPU-picking is that it is independent of the data geometry; this means we can even use custom 3D models. With this approach we also do not need to be concerned with the layout of the 3D scene, and we can more easily get objects from an area compared to using a raycaster.

The current implementation has a few limitations. Firstly, GPU picking is in itself an expensive operation. To improve performance, we only perform color-picking every 100 milliseconds, rather than every frame. Secondly, the use of colors limits the number, and complexities, of objects in the scene, and thus the size of the dataset to be visualized. This has not been a problem for the datasets we have used.

7 EVALUATION STUDY

To validate our proposed solution to the research questions RQ1–RQ3 (Sec. 3), we carried out a qualitative evaluation involving human participants, presented in this section. Our goal is to understand the design choices and deriving optimal strategies for visualizing attention as well as the associated triggering mechanisms.

7.1 Study Design Rationale

We made the following decisions for our design study:

- **No baseline.** We choose not to conduct a comparative or graphical perception experiment, partly because no suitable baseline comparison exists, and partly because we are interested in the richer guidance that a qualitative interview study provides.
- **Experienced participants.** We opted to recruit participants knowledgeable about data visualization to ensure that their attention would be representative of experienced visualization usage.
- **Convenience population.** Since our goal is to collect initial qualitative usage data of attention-aware visualizations, we feel that our specialized population (e.g., students in computer science and engineering) is of minor concern.
- **Between-participants on platform.** To keep sessions short, we split participants so that half did the study using the 2D data-agnostic AAV module, and half using the 3D data-aware one.

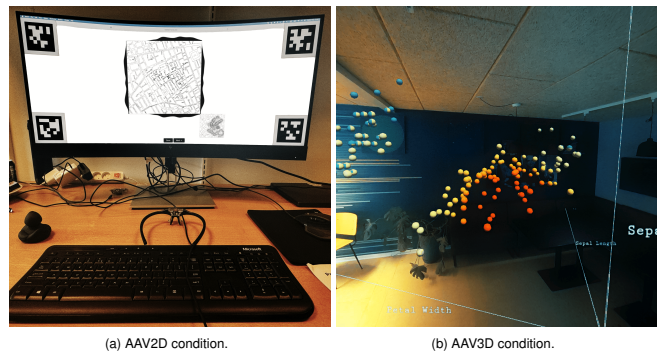


Fig. 7: **Experimental interface.** 2D and 3D user study setups.

7.2 Participants

We recruited 12 participants (9 male, 3 female) from a convenience population (students at two different universities). The participants were screened to be experienced computer users and knowledgeable about reading and interpreting data visualizations through a self assessment questionnaire. Their age ranged from 24 to 52 (mean 30), all were enrolled in a computer science program, and all had normal or corrected-to-normal vision with no color vision deficiencies.

7.3 Apparatus

For the 2D group, we used the data-agnostic implementation of attention-aware visualization in 2D (Sec. 5). We ran the experiment on an Apple MacBook Pro connected to a 34 inch ultra-wide monitor. The software was displayed in a maximized web browser in a 3024 × 1964 pixel native resolution. We used a Pupil Labs Neon embodied eye-tracker connected to an Android smartphone, to process gaze data that was then fed into the AAV2D module on the computer. Figure 7a shows an image of the experimental setup.


For the 3D group, we used the data-aware implementation of attention-aware visualization in 3D (Sec. 6). This study was conducted using a Meta Quest Pro head-mounted display in video passthrough mode. Figure 7b shows the experimental setup.

7.4 Procedure


Participants first provided informed consent at the beginning of their evaluation session. They filled out a demographic form and were given a brief introduction. They were then asked to complete three task blocks in random order. Each task block consisted of a visualization and a list of four open-ended questions. For the 2D tasks, participants engaged with 3 visualizations: Charles Minard’s Napoleon March (Fig. 8a), John Snow’s Cholera Map (Fig. 8b), and Florence Nightingale’s Rose Chart (Fig. 8c). They were asked questions warranting exploration, such as identifying river crossings by Napoleon’s army, counting the number of pumps contaminated by cholera, and comparing mortality rates from diseases across different years. For the 3D tasks, participants worked with a scatterplot of R.A. Fisher’s Iris data set (Fig. 8d), a barchart of Quinlan’s Auto MPG dataset from UC Irvine (Fig. 8e), and an elevation map of Mount St. Bruno (Fig. 8f). They were asked to analyze car cylinder counts over time, identify the largest flower sepals, and compare elevations on the terrain map. Detailed questions for these tasks can be found in the supplemental materials (see Sec. 9). Participants were assigned a random combination of attention presentation and triggering mechanism. Table 2 gives the experimental conditions, three for 2D and 3D each. Participants were given unlimited time to find answers to questions. Their performance was not timed and correctness was not measured. Participants were encouraged to verbalize their thinking, but did not follow a strict think-aloud protocol.

During each task block, the AAV software would continuously collect the user’s gaze using the Neon eyetracker or using GPU picking based on the Meta Quest head direction. The attention revisualization was displayed depending on the trigger mechanism randomly assigned

³<https://threejs.org>




(a) Charles Minard's *Napoleon's March* (English Translation) (AAV2D).




(b) Florence Nightingale's *Coxcomb Plots* (AAV2D).




(c) John Snow's *Cholera Map* (AAV2D).



(d) Fisher's *Iris flower dataset* as a 3D scatterplot (AAV3D).



(e) Quinlan's *Auto MPG dataset* as a 3D barchart (AAV3D).



(f) *Mt. Bruno elevation 3D surface* (AAV3D).

Fig. 8: Study visualizations. Each of these six visualizations (three for the 2D and three for the 3D platform) were associated with four tasks each (see supplemental material). All three of the 2D visualizations (top row) are in the public domain, and are slightly enhanced to improve readability. The 3D visualizations were our own designs.

to the participant: either manually or automatically. Participants were shown how to trigger and use the attention display.

After each block, participants were administered a brief System Usability Scale (SUS) [11] form and questionnaire on their experience and use of the attention-aware visualization. After the study, they were administered a summary questionnaire and were interviewed on their experience. Beyond these responses, we also captured attention metrics as well as usage statistics for each participant.

Table 2: Experimental conditions. Six combinations of presentation and triggers used for the experiment; three per platform (2D and 3D).

Condition	Platform	Trigger	Presentation
2D-Always	AAV2D	Always-on/Explicit	Contour
2D-Implicit	AAV2D	Implicit/Explicit	Mesh
2D-Explicit	AAV2D	Explicit/Always-on	Heatmap
3D-Always	AAV3D	Always-on	Heatmap
3D-Implicit	AAV3D	Implicit	Saturate
3D-Explicit	AAV3D	Explicit	Heatmap

7.5 Results

Overall, all participants understood the visualizations and were able to complete the tasks. Since we are not interested in completion time or task accuracy, we do not report such results here. Instead, here we summarize the results of the qualitative aspects of our evaluation, including observations, the subjective ratings, and themes from the spoken and written participant feedback.

7.5.1 Observations

All participants were able to complete all the tasks. We made no specific direct observations regarding task performance during either the 2D or 3D study. We observed that the Pupil Labs Neon eye-tracker used for the 2D study was not completely accurate. For some participants it measured slightly below their actual gaze, and for others slightly above, which likely skewed the attention tracking. Although this slightly offset measured attention on the visualization potentially led to inaccurate revisualizations, we did not observe this being an issue for our study,

given that the tasks were not dependent on revisualization accuracy, and the participants did not voice any accuracy concerns.

7.5.2 Subjective Ratings

The results for the SUS questionnaire were overall just above average [6], with the average score across participants of 70 (SD = 15). Per condition scores were all around 65 (below average SUS score), with one notable exception, the 3D visualization with EXPLICIT triggering, which scored 86 (excellent SUS score). When asked to rate the usefulness of the revisualizations, the responses varied significantly, with slightly more positive than negative responses. Fig. 9 illustrates the reported usefulness of all the revisualizations. These dispersed responses aligns with the SUS questionnaire results.

When asked to rank the triggering and revisualization techniques, participants using the 3D implementation had a clear preference for the EXPLICIT triggering condition and disliked the ALWAYS-ON heatmap. They preferred having control over when the revisualization was triggered, noting that the constant recoloring in the ALWAYS-ON condition was highly distracting, particularly because it overrode the colors of the visualization itself. This preference for EXPLICIT triggering is supported by the SUS score. The IMPLICIT triggering conditions were considered hard to understand, with one participant referring to it as “mysterious.” At the same time, other participants considered IMPLICIT triggering useful for certain tasks by helping to direct their attention.

The participants using the 2D implementation did not have a clear preference for revisualization and triggering. While some participants liked the contour map for its precise marking of attention and disliked the fuzzy nature of the heatmap, other participants disliked how the contour map cluttered the minimap with lines and liked how the heatmap was a more subtle color overlay. Most participants did not like the Mesh condition. Reasons ranged from being harder to understand what was seen, to being annoyed that it made the minimap hard to read, since it darkened seen areas too much in their opinion, and was imprecise due to the bin size used.

7.5.3 Participant Comments and Feedback

Our examination of participants’ experiences with revisualization tools in 2D and 3D contexts not only revealed the significant impact of revisualizations on user engagement, attention, and analytical depth but

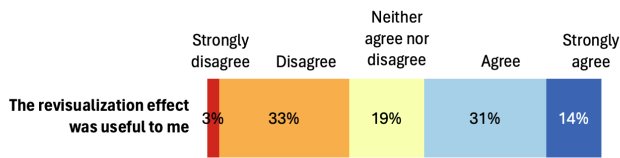


Fig. 9: **Subjective usefulness.** The reported usefulness of the revisualizations was widely dispersed, with slightly more in favor of the revisualizations while a large portion of the participants were neutral.

also shed light on design preferences and the nuanced effects of these tools on user interaction with data visualizations. We used thematic analysis to distill these insights into several key areas, incorporating participants’ voices to ground our findings with additional context.

Guiding Attention and Informing Curiosity. Revisualizations were predominantly used by participants to refocus and redirect their attention, often illuminating areas within the visual space that had either been overlooked or required further exploration. This utility of revisualizations bridged both 2D and 3D contexts, enhancing users’ exploration thoroughness by guiding them towards both explored and unexplored areas (P2–P5 in 2D; P1, P3, P5 in 3D). A participant (2D, P2) highlighted this dual functionality, noting, “*It was more that it guided me to where I had been looking and which, of course, did tell me where I haven’t been looking.*” Similarly, the stimulation of curiosity, especially when participants were seeking specific details or understanding the breadth of visualization techniques, was a recurring theme (P2, P5 in 2D; P2, P3 in 3D). Some participants in 3D shared that they “*treated it like a game*” (3D, P4) and “*found (themselves) wanting to go inside the visualization and see it from inside,*” (3D, P2), underscoring the depth of engagement revisualizations can foster.

Subtlety, Clarity, and Control. Participants leaned towards a preference for revisualizations characterized by subtlety, clarity, and the provision of EXPLICIT control. For instance, more than a couple of participants in 2D appreciated the subtlety of “*shading (frame) around the screen*” (P1), with some stating that it “*didn’t disturb in situations where (they) choose to not use them*” (P3, P5). Highly dynamic, vivid, or cluttered visualizations were often met with disapproval due to their propensity to distract or even obscure crucial data points (P5 in 2D; P2, P6 in 3D). For instance, more than a couple of participants in 2D remarked on the distracting nature of the Mesh implementation, stating how “*the (Mesh) was very contrasting and changing... added too much clutter...*” (P5) and “*occluded parts of the visualization*” (P2), even leading to some explicitly disliking it (P4). Similar responses were seen from the 3D participants where they stated that “*vibrant colors*” (P2) and “*overriding or changing colors frequently*” seemed “*harmful to experience*” (P6). The preference for EXPLICIT control mechanisms over ALWAYS-ON features was also notably prevalent, as they were seen to reduce distractions and enhance user autonomy in data exploration (P1, P3 in 2D; P1, P3, P6 in 3D). “*Explicit is not distracting at all; I had control,*” a participant (3D, P3) reflected, highlighting the value of user-driven engagement with revisualization tools.

Subjective Preferences. Participant feedback also pointed to the highly subjective nature of revisualization preferences and the critical role of task suitability in the effectiveness of these tools. This was evident in the 2D context, as some participants gravitated towards specific revisualization techniques such as heatmaps or contour maps based on their orientation towards precision and their detail-oriented utility, others valued minimaps for providing a holistic perspective of the data (P2–P6 in 2D; P3, P5, P6 in 3D) “*I did use the heatmap and the contour map more because I found them easier to understand how they can help me,*” noted a participant (2D, P3), exemplifying the varied preferences among users. This diversity underscores the need for adaptable revisualization techniques that can cater to the preferences of individual users and specific analytical tasks.

Learning and Adaptation Over Time. Participants increasingly engaged with the revisualization techniques as the study progressed. Initially drawn to these tools by their novelty, participants experienced a rise in their comfort and efficiency in using them over time (P4 in 2D; P1–P3, P5 in 3D). One participant in the 2D setting remarked, “*At first, I was really paying attention because it was new... but after getting familiar, it feels routine and I hardly give it a second thought*” (P4 in 2D). Similarly, in the 3D setting, another participant noted, “*Indeed, as I became more accustomed to the effects, I found my attention more effectively directed,*” highlighting the dynamic process of their interaction with the tools (3D, P1). This suggests that over time, users are likely to not just get used to the revisualization features but also incorporate them more fluidly into their analytical processes.

8 DISCUSSION

We have validated our approach to attention-aware visualizations both through two separate implementations—one in 2D and one in 3D—as well as through a qualitative evaluation study. Here we discuss what this work means for data visualization and context-aware designs within this field. We also enumerate some of the limitations of our work.

8.1 Explaining the Results

Overall, our findings mostly confirmed our expectations: participants perceived attention awareness to be both novel and useful, and confidently incorporated these mechanisms into their analytical workflow. As expected, however, the triggering and revisualization of attention is an important design challenge that requires careful attention (no pun). On the other hand, there were some nuanced differences observed between the 2D and 3D implementations that are worth closer scrutiny.

The immersive quality of 3D environments accentuates the need for revisualization mechanisms, likely due to the increased potential of the immersion to overwhelm or distract. “*Quite like being able to see 3D data from a new perspective inside the headset, the difference of being able to walk around to see things from different angles rather than rotating,*” remarked a participant (3D, P3), highlighting the new perceptual possibilities made accessible through the addition of depth. During its use, a participant commented, “*I feel like my attention is being guided around the visualizations; I have to actually move around to paint the scene, encouraging exploration*” (3D, P6), indicating that revisualization mechanisms can become essential for navigating and understanding data in 3D spaces by serving as a counterbalance to the potential for sensory overload and confusion. In contrast, in 2D environments, where depth is absent and the entire visualization is fully visible at all times, the need for revisualization appears less critical.

These insights suggest that the design of AAVs might benefit from a nuanced understanding of how users interact with and perceive visualizations differently in 2D vs. 3D. While the immersive complexity of 3D environments could amplify the need for user control and engaging interaction mechanisms, 2D spaces might call for simpler and more direct approaches. Anticipating and adapting to these differences could be key to unlocking the full potential of AAVs, ensuring that they enhance rather than complicate the user experience.

8.2 Design Implications

The introduction of attention-aware visualizations (AAVs) marks a shift towards creating more interactive and personalized data exploration experiences. By leveraging user attention data to dynamically adjust visualizations, we think that AAVs can make visual data exploration more tailored to individual users’ needs. However, our study highlights a critical balance that must be struck to harness the full potential of AAVs effectively. Continuous active monitoring and adjustment based on user gaze can be disruptive, creating a feedback loop that might skew the user’s natural data exploration process. Conversely, relying solely on user-initiated triggers for attention-aware adjustments risks underutilization due to forgetfulness or the extra cognitive load of remembering to activate the feature. This underscores the need for AAV systems to intelligently discern when to provide guidance without

overwhelming the user.⁴ Such a system would ideally analyze patterns of user engagement to anticipate the need for attention redirection, offering a seamless blend of automated and user-controlled interactions.

Even though we have only implemented and evaluated two of the settings that AAV could be useful in, we feel that AAV has potential in other settings as well, such as attention-aware 2D visualizations in immersive environments, or 3D visualizations viewed in a non-immersive environment, i.e., on a screen. Non-immersive 3D visualizations are arguably more common than their immersive counterparts, which is why our approach in AAV3D is designed to also work for such settings. We chose to focus on an immersive environment in this paper because we believe that AAVs can provide especially great value for such settings. When data is distributed all around the user, the cognitive load of keeping track of the data is high. In such scenarios, AAVs could be used to reduce the cognitive load, by pointing to unseen areas, or by reminding the user if they have already analyzed certain areas of the data. While a data-aware 2D implementation will not cause any difficulty, there are issues associated with a data-agnostic 3D implementation using the current approach. Our 2D implementation captures attention as a grid overlay. If the user is allowed to rotate view of the visualization and thereby the position of the marks on the screen, the stored attention will no longer correspond to the correct marks, since they can be at another cell in the overlay. This is a general issue if the user is able to change the representation. How to address this issue is left to future work.

Looking to the bigger picture, the integration of AAVs into visualization design may help improve user engagement with data. One potential direction involves the development of context-aware AAVs that adapt not just to where users look but also to the context of their interaction—what they are looking for and their task at hand. By understanding the user's intent and the context of their exploration, AAVs could provide more meaningful adjustments, such as highlighting relevant data points or suggesting alternative visualizations that might offer deeper insights into the data. This level of adaptation requires advancements in machine learning and artificial intelligence to interpret complex patterns of attention and interaction, paving the way for visualizations that are not only reactive but also proactive in facilitating data discovery.

Finally, the step to collaborative and educational uses for attention-aware visualization is not far. In situations where multiple users interact with the same visualization, AAVs could dynamically adjust to accommodate the focus and needs of different users, enhancing collaborative analysis and decision-making. In educational contexts, AAVs could tailor the presentation of information to the learner's pace and points of interest, providing a more engaging and effective learning experience.

8.3 Limitations

As discussed in our design framework (Sec. 3), one limitation of our approach is the potential for reinforcing behavior, where revisualizing mechanism based on user gaze may inadvertently influence the user's subsequent attention patterns. Such feedback could lead to a self-reinforcing cycle, where the visualization continuously emphasizes areas previously focused on by the user, potentially overshadowing other valuable but less attended data points. Understanding and mitigating the impact of such reinforcing behavior on the user's exploration and discovery process remains a challenge for future iterations of AAVs.

Another limitation concerns the choice of triggering to avoid disrupting the user experience. Finding the right balance between providing useful feedback and avoiding overwhelming the user with frequent or intrusive adjustments to the visualization is critical. This requires careful consideration of the timing, frequency, and extent of revisualizations to ensure they enhance rather than detract from the data exploration process. We leave such fine-tuning and evaluation to future work.

While eye tracking technology is common in the current generation of AR/VR headsets, it is still not commonplace in typical computing setups, thus limiting the adoption of AAVs. This restricts the immediate applicability of our proposed methods and limits the user base that can benefit from them. We speculate that the prevalence of gaze

interaction in, e.g., the Apple Vision Pro will lead to such technologies also becoming popular for regular computers and mobile devices.

The informal nature of our evaluation represents a methodological limitation. The evaluation's qualitative focus, lack of quantitative metrics, and absence of a baseline comparison restrict our ability to draw definitive conclusions about the efficacy of the framework. Further research involving more rigorous, quantitative evaluations and comparisons with existing approaches is necessary to fully understand the benefits and limitations of attention-aware visualization.

While we envision numerous futures for AAV, realizing them means developing algorithms capable of managing the delicate balance between being informative and non-disruptive. These algorithms are needed to ensure that AAVs enrich the user experience by making data visualizations not just tools for viewing but also guiding exploration.

9 CONCLUSION AND FUTURE WORK

We have presented attention-aware visualizations (AAVs), a context-aware data visualization paradigm that leverages user attention to dynamically adapt visual content. By tracking and analyzing where and how users allocate their attention over time, AAV not only offers a more personalized data exploration experience, but also detects situations where the user is not seeing everything in a given visualization. Our approach has been operationalized through two distinct implementations: a 2D data-agnostic method suitable for web-based visualizations and a 3D data-aware technique utilizing the stencil buffer for immersive analytics environments. Both methods underscore the potential of integrating attention metrics to enhance the interpretability and user engagement of visualizations. Accordingly, our evaluation of these implementations has provided insights into the effectiveness of visual feedback and attention-triggering mechanisms, highlighting the utility of AAVs in facilitating a deeper understanding of complex datasets.

We feel that this work points to future context-aware visualization techniques, where the focus is on developing visualization systems that respond to the needs of their users. Exploring the integration of AAVs with other sensory input methods, such as auditory feedback or haptic interfaces, could offer a richer, more multisensory data interaction experience. Additionally, the application of machine learning techniques to predict user attention shifts could refine the responsiveness of AAVs, making visualizations even more intuitive and user-centric. Another promising direction involves investigating the scalability of AAVs in handling larger, more complex datasets and their applicability across diverse fields such as medical imaging, financial analysis, and social network visualization. Finally, conducting extensive user studies across different domains could shed light on the universal applicability of AAVs and tailor them more closely to user needs and preferences.

SUPPLEMENTAL MATERIALS

All supplemental materials are available on OSF at <https://osf.io/8mfhp/>, released under a CC BY 4.0 license. In particular, they include the following: (1) Analysis files containing the final results of the quantitative System Usability Rating (SUS) analysis for both the 2D and 3D studies, along with a consolidated and summarized qualitative analysis with corresponding quotes from participants; (2) Data files containing the collected data from the participants of our evaluation. These files include experimental data in the form of data snapshots (only for 2D) and participant responses consolidated from block-level and study-level surveys, supported by notes and quotes from post-study interviews; (3) Materials used during the evaluation, including demographic forms and questionnaires; and (4) Zipped versions of the implementations of AAVs.

FIGURE CREDITS

Unless specified below, all figures are credited to the authors.

Figure 8a image credit: Modern redraw in English by Inigo Lopez Vazquez. May 3, 2015, CC BY-SA 4.0. Original by Charles Minard. November 20, 1869, public domain.

Figure 8b image credit: Florence Nightingale. 1858, public domain.

Figure 8c image credit: John Snow, 1854, public domain.

⁴After all, no one wants a visualization version of Clippy ☹️.

ACKNOWLEDGMENTS

This work was supported partly by Villum Investigator grant VL-54492 by Villum Fonden. 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] D. Abowd, A. K. Dey, R. Orr, and J. Brotherton. Context-awareness in wearable and ubiquitous computing. *Virtual Reality*, 3(3):200–211, sep 1998. doi: [10.1007/BF01408562](https://doi.org/10.1007/BF01408562) 2
- [2] S. S. Alam and R. Jianu. Analyzing eye-tracking information in visualization and data space: From where on the screen to what on the screen. *IEEE Transactions on Visualization and Computer Graphics*, 23(5):1492–1505, 2017. doi: [10.1109/TVCG.2016.2535340](https://doi.org/10.1109/TVCG.2016.2535340) 2
- [3] S. S. Alam and R. Jianu. Analyzing eye-tracking information in visualization and data space: From where on the screen to what on the screen. *IEEE Transactions on Visualization and Computer Graphics*, 23(5):1492–1505, 2017. doi: [10.1109/TVCG.2016.2535340](https://doi.org/10.1109/TVCG.2016.2535340) 2
- [4] U. Alegre, J. C. Augusto, and T. Clark. Engineering context-aware systems and applications: A survey. *Journal of Systems and Software*, 117:55–83, 2016. doi: [10.1016/j.jss.2016.02.010](https://doi.org/10.1016/j.jss.2016.02.010) 2
- [5] C. B. Anagnostopoulos, A. Tsounis, and S. Hadjiefthymiades. Context awareness in mobile computing environments. *Wireless Personal Communications*, 42(3):445–464, aug 2007. doi: [10.1007/s11277-006-9187-6](https://doi.org/10.1007/s11277-006-9187-6) 2
- [6] A. Bangor, P. Kortum, and J. Miller. Determining what individual sus scores mean: adding an adjective rating scale. *Journal of Usability Studies*, 4(3):114–123, may 2009. 7
- [7] T. Blascheck, M. John, S. Koch, L. Bruder, and T. Ertl. Triangulating user behavior using eye movement, interaction, and think aloud data. In *Proceedings of the ACM Symposium on Eye Tracking Research & Applications*, pp. 175–182. ACM, New York, NY, USA, 2016. doi: [10.1145/2857491.2857523](https://doi.org/10.1145/2857491.2857523) 2
- [8] T. Blascheck, K. Kurzhals, M. Raschke, M. Burch, D. Weiskopf, and T. Ertl. State-of-the-art of visualization for eye tracking data. In *State of the Art Reports of the Eurographics Conference on Visualization*. Eurographics Association, 2014. doi: [10.2312/EUROVISSTAR.20141173](https://doi.org/10.2312/EUROVISSTAR.20141173) 2
- [9] R. A. Bolt. Gaze-orchestrated dynamic windows. In *Proceedings of the ACM Conference on Computer Graphics and Interactive Techniques*, pp. 109–119. ACM, New York, NY, USA, 1981. doi: [10.1145/800224.806796](https://doi.org/10.1145/800224.806796) 2
- [10] M. Bostock, V. Ogievetsky, and J. Heer. D³ Data-Driven Documents. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2301–2309, Dec. 2011. doi: [10.1109/TVCG.2011.185](https://doi.org/10.1109/TVCG.2011.185) 5
- [11] J. Brooke et al. Sus-a quick and dirty usability scale. *Usability evaluation in industry*, 189(194):4–7, 1996. 7
- [12] M. Burch. *Eye Tracking and Visual Analytics*. River Publishers, Gistrup, Denmark, 2021. 2
- [13] M. Burch, N. Konevtsova, J. Heinrich, M. Hoferlin, and D. Weiskopf. Evaluation of traditional, orthogonal, and radial tree diagrams by an eye tracking study. *IEEE Transactions on Visualization and Computer Graphics*, 17(12):2440–2448, 2011. doi: [10.1109/TVCG.2011.193](https://doi.org/10.1109/TVCG.2011.193) 2
- [14] A. K. Dey. Understanding and using context. *Personal Ubiquitous Comput.*, 5(1):4–7, jan 2001. doi: [10.1007/s007790170019](https://doi.org/10.1007/s007790170019) 2
- [15] N. Elmqvist and P. Irani. Ubiquitous Analytics: Interacting with Big Data Anywhere, Anytime. *Computer*, 46(4):86–89, 2013. doi: [10.1109/MC.2013.147](https://doi.org/10.1109/MC.2013.147) 2
- [16] N. Elmqvist and P. Tsigas. A taxonomy of 3D occlusion management for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 14(5):1095–1109, 2008. doi: [10.1109/TVCG.2008.59](https://doi.org/10.1109/TVCG.2008.59) 1
- [17] D. Fono and R. Vertegaal. EyeWindows: evaluation of eye-controlled zooming windows for focus selection. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 151–160. ACM, New York, NY, USA, 2005. doi: [10.1145/1054972.1054994](https://doi.org/10.1145/1054972.1054994) 2
- [18] S. Goodwin, A. Prouzeau, R. Whitelock-Jones, C. Hurter, L. Lee, U. Afzal, and T. Dwyer. VETA: visual eye-tracking analytics for the exploration of gaze patterns and behaviours. *Visual Informatics*, 6(2):1–13, 2022. doi: [10.1016/j.visin.2022.02.004](https://doi.org/10.1016/j.visin.2022.02.004) 2
- [19] J. Grubert, T. Langlotz, S. Zollmann, and H. Regenbrecht. Towards pervasive augmented reality: Context-awareness in augmented reality. *IEEE Transactions on Visualization and Computer Graphics*, 23(6):1706–1724, 2017. doi: [10.1109/TVCG.2016.2543720](https://doi.org/10.1109/TVCG.2016.2543720) 2
- [20] C. Gutwin and S. Greenberg. Design for individuals, design for groups: Tradeoffs between power and workspace awareness. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work and Social Computing*, pp. 207–216. ACM, New York, NY, USA, 1998. doi: [10.1145/289444.289495](https://doi.org/10.1145/289444.289495) 2
- [21] C. G. Healey, K. S. Booth, and J. T. Enns. High-speed visual estimation using preattentive processing. *ACM Transactions on Computer-Human Interaction*, 6(2):107–135, 1999. doi: [10.1145/230562.230563](https://doi.org/10.1145/230562.230563) 2
- [22] W. C. Hill, J. D. Hollan, D. Wroblewski, and T. McCandless. Edit wear and read wear. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 3–9. ACM, New York, NY, USA, 1992. doi: [10.1145/142750.142751](https://doi.org/10.1145/142750.142751) 1
- [23] J. E. Hoffmann. A two-stage model of visual search. *Perception & Psychophysics*, 25:319–327, 1979. doi: [10.3758/BF03198811](https://doi.org/10.3758/BF03198811) 2
- [24] R. J. K. Jacob. What you look at is what you get: eye movement-based interaction techniques. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*. pp. 11–18. ACM, New York, NY, USA, 1990. doi: [10.1145/97243.97246](https://doi.org/10.1145/97243.97246) 2
- [25] H. Jänicke and M. Chen. A saliency-based quality metric for visualization. *Computer Graphics Forum*, 29(3):1183–1192, 2010. doi: [10.1111/J.1467-8659.2009.01667.X](https://doi.org/10.1111/J.1467-8659.2009.01667.X) 1, 2
- [26] S.-H. Kim, Z. Dong, H. Xian, B. Upatising, and J. S. Yi. Does an eye tracker tell the truth about visualizations?: Findings while investigating visualizations for decision making. *IEEE Transactions on Visualization and Computer Graphics*, 18(12):2421–2430, 2012. doi: [10.1109/TVCG.2012.215](https://doi.org/10.1109/TVCG.2012.215) 2
- [27] K. Kurzhals, B. Fisher, M. Burch, and D. Weiskopf. Eye tracking evaluation of visual analytics. *Information Visualization*, 15(4):340–358, 2016. doi: [10.1177/1473871615609787](https://doi.org/10.1177/1473871615609787) 2
- [28] F. Lu, S. Davari, L. Lisle, Y. Li, and D. A. Bowman. Glanceable ar: Evaluating information access methods for head-worn augmented reality. In *2020 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*, pp. 930–939. IEEE Computer Society, 2020. doi: [10.1109/VR46266.2020.00113](https://doi.org/10.1109/VR46266.2020.00113) 2
- [29] W. Lu, H. B.-L. Duh, S. Feiner, and Q. Zhao. Attributes of subtle cues for facilitating visual search in augmented reality. *IEEE Transactions on Visualization and Computer Graphics*, 20(3):404–412, 2014. doi: [10.1109/TVCG.2013.241](https://doi.org/10.1109/TVCG.2013.241) 2
- [30] J. D. Mackinlay. Automating the design of graphical presentations of relational information. *ACM Transactions on Graphics*, 5(2):110–141, 1986. doi: [10.1145/22949.22950](https://doi.org/10.1145/22949.22950) 1
- [31] L. E. Matzen, M. J. Haass, K. M. Divis, Z. Wang, and A. T. Wilson. Data visualization saliency model: A tool for evaluating abstract data visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):563–573, 2018. doi: [10.1109/TVCG.2017.2743939](https://doi.org/10.1109/TVCG.2017.2743939) 1, 2
- [32] T. Munzner. *Visualization Analysis and Design*. A K Peters, Boca Raton, FL, USA, 2014. 1
- [33] U. Neisser. *Cognitive Psychology*. Appleton-Century-Crofts, New York, NY, USA, 1967. 2
- [34] L. T. Nowell, E. G. Hetzler, and T. Tanasse. Change blindness in information visualization: A case study. In *Proceedings of the IEEE Symposium on Information Visualization*, pp. 15–22. IEEE Computer Society, Los Alamitos, CA, USA, 2001. doi: [10.1109/INFVIS.2001.963274](https://doi.org/10.1109/INFVIS.2001.963274) 2
- [35] C. Perera, A. Zaslavsky, P. Christen, and D. Georgakopoulos. Context aware computing for the internet of things: A survey. *IEEE Communications Surveys & Tutorials*, 16(1):414–454, 2014. doi: [10.1109/SURV.2013.042313.00197](https://doi.org/10.1109/SURV.2013.042313.00197) 2
- [36] K. Pfeuffer, J. Alexander, M. K. Chong, and H. Gellersen. Gaze-touch: combining gaze with multi-touch for interaction on the same surface. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pp. 509–518. ACM, New York, NY, USA, 2014. doi: [10.1145/2642918.2647397](https://doi.org/10.1145/2642918.2647397) 2
- [37] K. Pfeuffer and H. Gellersen. Gaze and touch interaction on tablets. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pp. 301–311. ACM, New York, NY, USA, 2016. doi: [10.1145/2984511.2984514](https://doi.org/10.1145/2984511.2984514) 2
- [38] K. Pfeuffer, B. Mayer, D. Mardanbegi, and H. Gellersen. Gaze + pinch interaction in virtual reality. In *Proceedings of the ACM Symposium on Spatial User Interaction*, pp. 99–108. ACM, New York, NY, USA, 2017. doi: [10.1145/3131277.3132180](https://doi.org/10.1145/3131277.3132180) 2
- [39] P. Polatsek, M. Waldner, I. Viola, P. Kapec, and W. Benesova. Exploring visual attention and saliency modeling for task-based visual analysis. *Computers & Graphics*, 72:26–38, 2018. doi: [10.1016/j.cag.2018.01.010](https://doi.org/10.1016/j.cag.2018.01.010)

- [40] D. D. Salvucci and J. R. Anderson. Intelligent gaze-added interfaces. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 273–280. ACM, New York, NY, USA, 2000. doi: 10.1145/332040.332444 2
- [41] B. Schilit, N. Adams, and R. Want. Context-aware computing applications. In *1994 First Workshop on Mobile Computing Systems and Applications*, pp. 85–90, 1994. doi: 10.1109/WMCSA.1994.16 1, 2
- [42] 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*, 2023. To appear. doi: 10.1109/TVCG.2023.3285546 2
- [43] S. Shin, S. Chung, S. Hong, and N. Elmqvist. A scanner deeply: Predicting gaze heatmaps on visualizations using crowdsourced eye movement data. *IEEE Transactions on Visualization and Computer Graphics*, 29(1):396–406, 2023. doi: 10.1109/TVCG.2022.3209472 2
- [44] S. Stellmach and R. Dachsel. Look & touch: gaze-supported target acquisition. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 2981–2990. ACM, New York, NY, USA, 2012. doi: 10.1145/2207676.2208709 2
- [45] S. Stellmach, S. Stober, A. Nürnberger, and R. Dachsel. Designing gaze-supported multimodal interactions for the exploration of large image collections. In *Proceedings of the Conference on Novel Gaze-Controlled Applications*, pp. 1:1–1:8. ACM, New York, NY, USA, 2011. doi: 10.1145/1983302.1983303 2
- [46] A. M. Treisman and G. Gelade. A feature-integration theory of attention. *Cognitive Psychology*, 12(1):97–136, 1980. doi: 10.1016/0010-0285(80)90005-5 2
- [47] H. Y. Tsang, M. Tory, and C. Swindells. eSeeTrack - visualizing sequential fixation patterns. *IEEE Transactions on Visualization and Computer Graphics*, 16(6):953–962, 2010. doi: 10.1109/TVCG.2010.149 2
- [48] J. Turner, J. Alexander, A. Bulling, D. Schmidt, and H. Gellersen. Eye pull, eye push: Moving objects between large screens and personal devices with gaze and touch. In *Proceedings of the IFIP Conference on Human-Computer Interaction*, vol. 8118 of *Lecture Notes in Computer Science*, pp. 170–186. Springer, New York, NY, USA, 2013. doi: 10.1007/978-3-642-40480-1_11 2
- [49] E. Velloso, J. Turner, J. Alexander, A. Bulling, and H. Gellersen. An empirical investigation of gaze selection in mid-air gestural 3D manipulation. In *Proceedings of the IFIP Conference on Human-Computer Interaction*, vol. 9297 of *Lecture Notes in Computer Science*, pp. 315–330. Springer, New York, NY, USA, 2015. doi: 10.1007/978-3-319-22668-2_25 2
- [50] C. Ware. *Information Visualization: Perception for Design*. Morgan Kaufmann, San Francisco, CA, USA, 3 ed., 2012. 2
- [51] C. Ware and H. H. Mikaelian. An evaluation of an eye tracker as a device for computer input. In *Proceedings of the SIGCHI/GI Conference on Human Factors in Computing Systems and Graphics Interface*, pp. 183–188. ACM, New York, NY, USA, 1987. doi: 10.1145/29933.275627 2
- [52] M. Weiser. The computer for the 21st Century. *Scientific American*, 265(3):94–104, 1991. doi: 10.1145/329124.329126 2
- [53] J. H. R. White and G. Buscher. User see, user point: Gaze and cursor alignment in web search. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 1341–1350. ACM, New York, NY, USA, 2012. doi: 10.1145/2207676.2208591 1
- [54] J. M. Wolfe. Guided search 2.0 a revised model of visual search. *Psychonomic Bulletin & Review*, 1(2):202–238, 1994. doi: 10.3758/BF03200774 2
- [55] S. Zhai, C. Morimoto, and S. Ihde. Manual and gaze input cascaded (MAGIC) pointing. In *Proceedings of the ACM Conference on Human Factors in Computing Systems*, pp. 246–253. ACM, New York, NY, USA, 1999. doi: 10.1145/302979.303053 2