Drawing into the AR-CANVAS: Designing Embedded Visualizations for Augmented Reality

Benjamin Bach* University of Edinburgh, UK
Ronell Sicat† Harvard University, MA
Hanspeter Pfister‡ Harvard University, MA
Aaron Quigley§ University of St. Andrews, UK

ABSTRACT

With AR-CANVAS we introduce the notion of the augmented reality canvas for information visualization. This is beyond the traditional empty (white), rectangular, and flat-dimensional canvas seen in traditional information visualization. Instead, the AR-CANVAS describes the part of a viewer's field-of-view where information visualization is rendered *in-situ* with respect to visible and potentially invisible real-world objects. The visual and spatial complexity of the canvas requires rethinking how to design visualizations for augmented reality. Based on an example of a library exploration scenario, we describe the essential aspects of the AR-CANVAS as well as dimensions for visualization design. We conclude with a brief discussion of challenges in designing visualizations into such a canvas.

1 INTRODUCTION

Embedding data visualization into the physical world surrounding us has recently come into focus as one class of visualization with potentially broad applications [2, 20]. Examples include, monitoring production in a chemical plant, accessing on-line manuals during assembly in an airline factory [18], studying data flows in a laboratory, or exploring a book collection in a library or bookstore. Broadly speaking, any application requiring real-time data and respective "context data" [13] to support users *in-situ* activity, might eventually benefit from embedded data visualizations.

Embedding visualizations has two key aims: showing data in relation to its *physical referents* [20], and preserving a user's flow state while concentrating on both information and handling tools [14] or navigation at the same time. Through the emergence of hands-free augmented reality technology (e.g., HoloLens, Meta), designing data visualizations for such scenarios becomes both feasible and pressing. However, while placing abstract information visualizations in-situ is conceptually straightforward, the reality is quite different. Designing visualizations to be situated into the real world poses a new set of challenges compared to designing information visualizations for on-screen use alone. A few challenges and considerations include: *Where to place visualizations?, How to avoid visual clutter?, What visualization types to choose from?, How to maintain faithful perception?, How to provide for interaction?, When to update a visualization?,* etc.

Traditional information visualization design has built theory and practice around designing for an (i) *empty* and monochrome, (ii) *two-dimensional* and in most cases (iii) *limited rectangular canvas* such as paper or onscreen. Such a canvas is reserved entirely for a visualization design, i.e, designers can define shapes for visual elements and assign visual attributes solely based on design perceptual laws and visual guidelines. Similarly, designers can decide how to layout elements, e.g., based on similarity (e.g., MDS), connectivity (e.g., force-directed layout), grouping (e.g., treemap) or data dimensions (e.g., time).

If visualizations are to be designed and hence situated into the real world, then the viewer's field-of-view becomes the new canvas, which we call the AR-CANVAS. This canvas is (i) *populated* (not blank) as the physical world contains physical objects with position and size as well as a background which, in most cases, is neither monochrome white or black; (ii) an AR-canvas can be *dynamic* as the position of objects or the observer can change; (iii) the canvas effectively is *three-dimensional* and objects and visualizations can occlude each other resulting in invisibility and clutter; (iv) in case of multiple observers, each one obtains a *different perspective* on the same scene. Some data physicalizations deal with similar issues, yet as we aim to show here, the challenges and the visualization design space on an AR-canvas are larger. Previous work on embedded data visualizations has started surveying the concept of visualizations in the physical world [20] and at the time of finalizing this article, Bauer makes a very important point on "*Silent Augmented Reality*" [3]. However, no systematic description of the design space and the resulting implications exist, and adressing this is what we aim to contribute here.

In this paper, we present a framework to explore and ignite the design of embedded data visualizations in augmented reality. We focus on augmented reality as provided by head-attached displays such as HoloLens or Meta: users can operate hands-free, digital content can be shown in stereoscopy, and users can freely move their head and body inside the scene. We exclude hand-held (e.g. ARKit) and spatial see-through augmented reality (e.g. public display AR) techniques, for now, in our framework [6]. We assume an *optimal* display technology, neglecting, for example, the currently small field-of-view or low graphical performance of current devices and poor fiducial marker tracking [1].

We define the key components in an AR-CANVAS as (**C**ontext-data, **A**rtifact, **N**avigator, **V**isualization, **A**ctivity, **S**cene) as well as characterize their potential relationships. Based on a motivational example of a library exploration scenario, we describe a preliminary design space for visualization design, taking into account our five components of the canvas. This design space aims at supporting visualization designers with an initial framework of how to approach the design of embedded visualizations and to highlight respective design challenges. We expect this design space will evolve as designers and developers create designs and clever solutions for the initial challenges, as well as informing the structured evaluation of designs and techniques.

2 MOTIVATIONAL SCENARIO

Libraries and bookstores are great places to make discoveries—intended or serendipitously [16]. They provide an inviting atmosphere, a curated and purposely organized collection of books, an informative librarian, and most importantly, an instant visual and physical overview of the available material. Unlike virtual books in on-line libraries—most of the time hidden from the user and accessible only through initial queries or selected examples—physical books can provide tangible information through their size, appearance (e.g., can show age, and how much they are used), the design of their cover, hardcover or softcover, etc. One touches books and

*e-mail: benj.bach@gmail.com
†e-mail: sicat@g.harvard.edu
‡e-mail:pfister@g.harvard.edu
§e-mail:aquigley@st-andrews.ac.uk

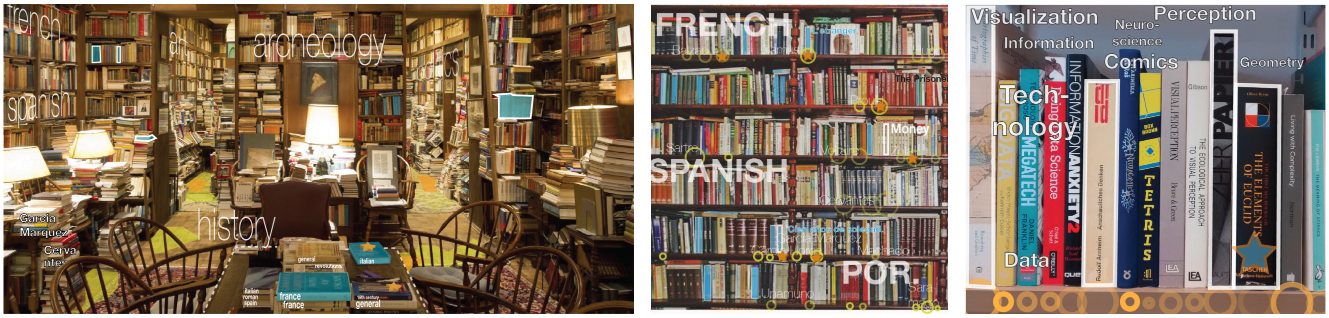

Figure 1: Visualizations on AR-canvases showing information depending on the viewer's distance: a) library room, b) shelf, c) books.

starts reading, or browsing through all the pages at will, observing illustrations and footnotes. All this can provide a rich and immersive experience of browsing and discovery. However, there is a lot of information available in databases, largely inaccessible in the physical library setup: ratings, critiques, citations, keywords, books by the same author, level of required pre-knowledge, edition, as well as characters, times, and places in the story [16]. Embedding visualizations into libraries and bookstores could support visitors by combining the best of both worlds: browsing a physical environment with physical information about books coupled with rich digital data.

Figure 1 shows a set of mock-up designs for different levels of detail, depending on where a user stands and what is their field of view. Figure 1-left is what a user may see when entering the library. Words on the shelves indicate high-level categories and can describe distant rooms. Books closer to the visitor may reveal topics and authors (center front in Figure 1-left). As an observer approaches a shelf, detailed information indicate the subcategories on each shelf. Figure 1(center) shows categories for books and highlights specific subcategories (e.g., important authors). For some selected books, this view shows green circles ○: their SIZE can indicate some sort of popularity (e.g., user rating, sold units), transparency can indicate a second dimension such as time of publication. Books in their native language bear a blue overlay on their spines, while books recommended by a particular agent are highlighted with a white frame. In Figure 1-right the circles for all books are shown on the shelf in order to not obstruct the books. Some keywords can be placed close to the books they describe.

In creating the mock-up designs, we drew from experiments and our experience with visualization on the traditional and the AR-CANVAS (e.g., [1]). We very carefully chose which information to show, how to visualize it, where to place it, and which information to show depending on where a user is situated. Note that the 2D pictures in this paper can only provide a limited representation of the actual impression and visual scene in AR; stereoscopy and visual fidelity allows humans to better distinguish between AR and actual content of the scene, improving the readability of information in Figure 1.

3 Related Research

Augmented reality, as a concept, dates back to the earliest days of computer graphics, with several applications in libraries [7]. Besides labeling and support for navigating the user to their destination, few projects in research and design demonstrates how to design visualizations on an AR-CANVAS. Projects are limited in that they often do not consider or discuss the relations between and across users, real-world physical objects, and virtual content, as well as the implications of their design choices. Other research related to visualization in AR has been concerned with legible and dynamic placements of labels [15], highlighting patterns [9], dynamic adjustment of content with respect to user distance [4], and interactive focus+context for scientific visualization [11]. Most recently, Bauer advocates for the careful blending of artificial content into the real world [3].

Few examples exist that visualize information for exploration on an AR-CANVAS, these include, air pollution in 3D environments using floating bubbles, clouds, and other glyphs [19]; corrosion by applying color maps directly to large structures such as building walls [17]; abstract and temporal data on buildings, overlaying bar-charts and line graphs on their walls [21]; visualizations overlaid on shelved product packaging, which the authors called *situated analytics* [8, 10]. Most related is a categorization by Bell et al. [4], describing object properties in AR through *visibility*, *position*, *size*, *transparency*, and *priority*. We draw from previous approaches, but aim to deliver a more visualization-focused perspective.

4 Elements of the AR-Canvas

The term AR-CANVAS describes an acronym of key terms, including: **C**ontext-data, **A**rtifact, **N**avigator, **V**isualization, **A**ctivity, **S**cene.

CONTEXT DATA—Data can refer to objects, tasks, or anything else of interest to the user in performing their task in a particular context [13]. We differentiate between data associated to present objects (*referred data*) and data not associated to objects. We also focus on visualization for data related to objects, however similar concerns and designs may be valid for data not related to objects. Data can be of different **types** (e.g., attributes, states, categorical, temporal, relational). Data can be associated to an individual object or **aggregated** for a group of objects (a set of books in a shelf). Data can further be **static or dynamic**, i.e., updating over time intended for real-time monitoring.

ARTIFACTS—The artifact is a physical body in the real world (object) that can be associated with data in the respective scenario, meaning it can be a *physical referent* [20]. An object has physical characteristics such as **size, color, texture, shape, rigidness, transparency**, etc. It can bear textual **labels** and detailed information written on it such as on product packaging. Objects can be **static**, **moveable** (e.g. by the user), or **dynamic** in that they move without the user's interaction (e.g. cars in the street). Objects have a distance to the user and can have different **spatial relationships** to other objects such as distance, nestedness, or stacking. Specific spatial properties can imply a **grouping** of objects, such as proximity or nestedness. Eventually, certain objects can serve as containers (cups, petri dishes, rooms) for **content** that is hard to translate into an augmented reality environment (e.g. liquids), or which content is invisible to the human but data exists (e.g. bacteria). Artifacts are the primary objects in each AR-CANVAS, i.e., their physical properties, their dynamicity, their spatial relationships as well as their distance to the user will greatly determine any visualization design.

NAVIGATOR—The navigator is the human being perceiving and interacting with the visualization. The navigator is physically present in the scene, with a specific location and free to walk and leave the scene. A navigator's location defines a **distance** to all artifacts as well as how much is **visible** of each artifact. A navigator has a **viewing direction** and a **field-of-view**, both determining the size and location of the personal AR-canvas. Both can change as a navigator moves, eventually changing artifact visibility. Potentially, there can be multiple navigators collaborating within the same setup. Each has their own AR-CANVAS which can show different artifacts and data. Eventually, any physically present navigator becomes a dynamic object, visually interfering with visualizations displayed in others' personal AR-CANVAS.

VISUALIZATION—A visualization refers to every data-representing visual element projected onto one's AR-canvas. Visualizations have a **type** (e.g. node-link diagrams, parallel coordinates, scatterplot, etc.), they contain **data marks** (e.g. points, crosses, lines, and other shapes) and **visualization marks** such as axes, labels, and grids, and employ **visual variables** such as color, texture, size, and position. A visualization can be as simple as indicating the state of an object by applying a single color to its surface or can be more complex by incorporating multiple attributes about or summarizing data about a group of objects.

ACTIVITY—Activities with visualizations and artifacts can involve look-ups and search, browsing and exploration, or complex analytical tasks. Some of these tasks rely more on perception, while others require more interaction with objects and visualizations [1]. In some cases, users may need access to the data plus the artifact (to understand the state of an object); they may need only the object (e.g. move the object, group objects); or they need only the data (e.g. find maximum value in the data). Depending on the use case, designs for an AR-CANVAS must support such a range of tasks.

SCENE—Artifacts and navigators are situated in a scene. A scene can be of different **size**, e.g., a single table, a room, a street or an entire city with only a fraction visible to the navigator at any time. The scene is essentially the background or backplane for the AR-CANVAS, i.e., it will determine how and where visualizations can be displayed and how readable they can be. A scene can provide **empty spaces** (walls, surfaces, sky, floor, etc.). It has an intrinsic **lighting** which can change as the sun starts shining or gets covered by clouds; light intensity and color also changes over the course of the day. Any light source can introduce shadows, shadings, reflections, and refractions on artifacts. A scene can **allow the user to move** around, or **move artifacts** themselves. Eventually, some artifacts or other parts of the scene may move by themselves (e.g, a passing car, another navigator).

Using our five elements we can describe an AR-CANVAS for many specific visualization scenarios. For instance, in our library example, we can describe *context-data* (spatial data, ratings, publication data, keywords, authors, social data, schedules, language, prizes, prices, etc.), *artifacts* (books, magazines, media items), the *navigator* (a visitor interested in certain topics, having some pre-knowledge, knowing some authors, etc.), *visualizations* (text labels, circles, rectangular spine overlays, links), *activity* (serendipitous book browsing, focused search), and the *scene* (different rooms, books organized into shelves, sometimes stacked to towers, lying on tables, etc.).

5 Design Space

This section describes the parameters and design space to consider when designing *visualizations* for an AR-CANVAS. Each design decision depends on some or all of the elements of the AR-CANVAS described in the previous section.

VISUAL MARKS—*What shapes are used to encode information?* On the traditional canvas, data elements are represented by (2-dimensional) visual marks: **circles**, **squares**, or **other shapes**. In our example, a book can have a circle mark whose size is mapped to an attribute, e.g., user rating. An AR-CANVAS can incorporate both 2D and 3D visual marks [19]. Moreover, in an AR-CANVAS, artifacts can serve as visual marks whose visual attributes are altered. In our library example, in some cases, a book's spine can be *altered* with rectangular highlights and text labels (Figure 1-center). Other book categorizations, e.g., grouping by language, are indicated through text labels, frames, or superimposed glyphs (e.g., stars).

LOCATION—*Where is the visualization / are the visual marks located with respect to artifacts and the scene?* Visual marks can be **overlaid** onto an object (e.g., blue highlights in Figure 1(center), also [8]), it can be **beside** or **around** an object (green circles in Figure 1(center)). In case of groups of objects, a visualization can additionally be **in-between** this set of objects (white keyword labels in Figure 1(right)). A visualization can be **detached** from the objects it is referring to, e.g. by floating in the air [19], being anchored onto some wall or other surface [21], or being placed at some deliberate location close to the user. In this case, the design requires linking indications that link a visualization (or parts of) to the objects it is referring to. Eventually, a visualization can be **attached to the screen**, i.e., always visible and follows the user when moving and turning his/her head. Screen-attached visualizations are practically detached from artifacts and require linking indications as well. Bell et al. [4] mention the possibility of a maximal distance of a virtual object from its physical referent. They also describe automatic approaches to place virtual content into an AR scene, similar to force-directed label placement in networks and other approaches common in AR [12, 15].

DIMENSIONALITY—*Which dimensionality does the visualization has?* A visualization can be entirely 2D and always facing the user, be relief-like (i.e., the main layout and layout of the visualization happens in a 2D-space, but some elements may be rendered 3D but without encoding specific information), or it can be entirely 3D [2, 19]. There may also be intermediate cases where a visualization is intended to be mapped onto a non-flat surface such as the surface of a cup, a sculpture, or walls in a library (Figure 1-left, also [21]).

ORIENTATION—*How is the visualization oriented with respect to the artifact(s) it is referring to and the human (navigator) observing it?* A (2D) visualization, or parts of it, can be designed to always face the user, e.g., to increase readability of textual labels [4], and reduce estimation error in visualizations through perspective distortion [5]. On the other side, labels and shapes oriented away from the navigator will not clutter their field-of-view but subtly come into focus as the navigator turns towards them. This can mean that the information displayed is important just at a specific position, e.g., when directly facing a shelf (Figure 1-left→center).

VISIBILITY—*Under which conditions is a visualization visible?* A designer may want a specific visualization to be **always visible**, independent on whether the respective artifacts are visible to the user or not [11]. Or, the designer decides to remove a visualization from the navigator's field-of-view whenever the respective artifact is not visible (e.g., **hidden** by another artifact or scene objects). **Orientation** is one means to control the visibility of a visualization depending on the navigator's location. Alternatively, visibility can be defined through the **distance** between the navigator and the artifact or visualization or their respective head-pointing direction. Our library example will show more information about more objects as the navigator comes closer (Figure 1-center→right). Deciding on visibility requires consideration of a navigator's activities, the complexity of the scene (visual distraction and noise), and factors such as number of concurrent visualizations and their respective complexity.

STYLING—*What is the visual appearance of the visualization?* The visual style of a visualization includes all its visual parameters such as **coloring**, **texture**, **size**, or **transparency**. Style defines how a visualization blends with its environment. In our example, we have decided to make categorical labels (languages, authors) semi-transparent to decrease visual disruption. A visualization can be perceived as part-of-the-environment, e.g., subtle strokes on a glass or light waves on a surface. On the other side, a visualization can be high in contrast with its environment and hence be clearly distinguishable and which comes into focus (recommended white framed books in Fig.1-center).

VISUALIZATION-DESIGN—*What type of visual representation is the visualization?* The visual representation (or type) of a visualization implies how it is being read and decoded. Traditional visualization has described types such as barcharts, node-link diagrams, parallel-coordinates plots, or flow lines. In AR, a visualization can be a simple encoding of an attribute or a state for an individual artifact (as in our example). Visualizations can encode multiple attributes on one object (size and transparency of our circles), or data and relationships about multiple objects (tags in Figure 1(right)). More complex visualizations have been projected on walls [12], but would treemaps, networks, and other complex visualization work the same way? The critical point is how to relate the given layout or placement of artifacts in the scene to a layout of data marks in the visualization. Traditional visualizations explicitly create data marks to represent objects in the data (e.g., points in a scatterplot, lines in a PCP, links in a node-link diagram, etc.) and maps attributes of that data object to visual attributes of the data mark. The positions of these marks is then defined by the specific visualization layout determined by spatial dimensions (e.g., scatterplot), geographic positions (e.g., maps) or optimization techniques (e.g., force-directed graphs). However, positions of marks in an AR-CANVAS cannot be arbitrarily determined by such since their corresponding artifacts already have scene-dependent positions (this problem is related to showing networks on geographical maps). The trade-offs between visualizations that map data mark positions to their artifacts' positions and those with traditional layouts but link data marks with artifacts, e.g., via arrows, have to be considered.

6 Designing Embedded Visualization

Our visualization design space is not meant to be complete and final but we believe it can serve as an initial framework to describe some of the decisions a designer has when designing visualizations for an AR-CANVAS. Now, we discuss challenges and concepts related to employing the design space for an AR-CANVAS and which are based on general knowledge and summarize our experience and discussions on previous and ongoing AR-visualization work.

2D vs. 3D visualization: 3D visualizations are often discussed with a focus on 2D monoscopic screens. In an AR environment, a visualization can still use a 2D layout while being projected on a real-world surface, eventually being integrated into a 3D environment. This can lead to perspective distortion and the misinterpretation of visual variables [5]. Similarly, if 2D visualizations are projected onto non-flat surfaces, perspective issues may arise. Pure 3D visualizations, however, may emerge as a means to blend naturally in the environment and may invite for interaction and manipulation. 3D visualizations can show intrinsic 3D data or fill the space between and around objects (see 3D stars in Figure 1(left)). However 3D visualizations may require careful solutions to relate to the (also 3D) artifacts they are referring to.

Visual Clutter: The main source of visual clutter will most likely come from i) the scene and its respective background, ii) potentially overlapping visualizations and visual marks, and iii) visualizations potentially hiding real-world information and artifacts of use to the navigator. A crowded scene will make it difficult to find space for displaying visualizations without hiding other information. Visualizations could be designed with the background in mind so that they blend seamlessly in color and style; visualizations could be projected onto object surfaces and blend into the respective nature of the object. In cases where objects must not be cluttered, carefully chosen empty spaces around the objects can be used. If an object is not visible to the user, visualizations associated to it could be removed. Surfaces most likely to host more complex visualizations will be walls, tables, and floors. Alternatively, respective empty screens such as paper, whiteboards, or projection screens can be integrated into the scene and explicitly used for visualization display.

Adaptivity: Besides static designs to reduce visual clutter, visualizations can adapt to a wide range of conditions: a user's distance to an object, visibility of objects, a user's task at hand or a current selection as well as any other measure of relative importance. Visualizations can change style and level-of-detail, e.g., they could be subtle and almost hidden indicating that there is (more) information to display on demand; when a user approaches they can expand and unfold in the space, involve color, and reveal full detail. Visualizations could be marked important and stick to a viewer's field-of-view or remain visible even if its artifacts are not.

Labels and Legends: Labels are important to name elements—artifacts in the scene or visual marks in the visualization. Readability of text on an AR-CANVAS will be constrained by the respective background, a label's relative size as well as whether a label is facing the user or not. Layout techniques have been presented for labels in 3D and AR environments [12,15] yet, labels can be adaptive or designed to be visible only under certain circumstances such as in our example as too many labels may clutter the scene.

Relations between artifacts or visualizations: In some cases, e.g. the keywords in Figure 1-right, it might be useful to relate objects in space. Relations can be indicated through graphical lines or visual marks of same color (e.g., all blue books are in their native language). Perhaps a good solution is to show relations *on demand*, as a navigator is close to an object or otherwise implicitly or explicitly signaled interest. The open problem at this point is how to effectively interact with visualizations and artifacts in AR.

Data not related to artifacts: Eventually, a scene, or an activity may need access to data and information not related to any artifacts (e.g., email, surveilling a distant set of artifacts). Such information could possibly be shown in empty spaces in the scene.

7 Discussion

The ultimate goal of embedded visualizations is to computationally support individual and collective human thinking where it happens—in our minds. By focusing our attention and concentrating on a particular problem we can exhibit a deep mental involvement with the task and data. In this, computation should augment, not replace our thinking. When we examine the elements or structure of a problem, situation, or object in the real world, we want to be able to draw in new information which we don't currently know. The new information or computational process should be available in such a fluid manner that we don't need to expend additional mental effort to access it. Embedded visualizations are able to provide this fluidity by seamlessly blending the needed digital information with the real world into a single task-adaptive and cohesive view—a truly data augmented environment in an AR-CANVAS.

We believe designing embedded visualizations on an AR-CANVAS requires rethinking certain design dimensions that we have learned to master on a traditional canvas. We are also just starting to build an understanding of design rules and what factors influence the effectiveness of visualization designs in such an augmented environment. This paper is meant to gather ideas, questions, and possibilities to design for the AR-canvas, and most importantly, to open a discussion and collect examples of the few existing designs. The many open questions are about visualization types (*what novel types of visualizations emerge in augmented reality?*), usability (*what makes a good and usable AR visualization?*), efficiency of visualizations (*how efficient is a specific visual encoding for a particular task?*), about tasks (*what visual analysis can be performed better in AR than on standard environments?*), interaction (*what hybrid interactions with artifacts and visualizations can be discovered?*), and scenarios (*what scenarios can be supported through embedded visualizations?*). With this paper, we hope to encourage more visualization designs in augmented reality, eventually showing the strengths and benefits of this novel field.

References

[1] B. Bach, R. Sicat, J. Beyer, M. Cordeil, and H. Pfister. The hologram in my hand: How effective is interactive exploration of 3d visualizations in augmented reality? In *IEEE TVCG*, 2018.

[2] D. Bauer. Can augmented reality solve mobile visualization? https://hackernoon.com/can-augmented-reality-solve-mobile-visualization-f06c008f8f84, 2017.

[3] D. Bauer. Silent augmented reality. https://hackernoon.com/silent-augmented-reality-f0f7614cab32, 2017.

[4] B. Bell, S. Feiner, and T. Höllerer. View management for virtual and augmented reality. In *Proc. of UIST*, pp. 101–110. ACM, 2001.

[5] A. Bezerianos and P. Isenberg. Perception of visual variables on tiled wall-sized displays for information visualization applications. *IEEE TVCG*, 18(12):2516–2525, 2012.

[6] O. Bimber and R. Raskar. *Spatial augmented reality: merging real and virtual worlds*. CRC press, 2005.

[7] D. Chen, S. Tsai, C.-H. Hsu, S. J. Pal, and G. Bernd. Mobile augmented reality for books on a shelf. In *2011 IEEE International Conference on Multimedia and Expo*, pp. 1–6, 2011.

[8] N. ElSayed, B. Thomas, K. Marriott, J. Piantadosi, and R. Smith. Situated analytics. In *Proc. of BDVA*, pp. 1–8. IEEE, 2015.

[9] N. A. ElSayed, R. T. Smith, and B. H. Thomas. Horus eye: See the invisible bird and snake vision for augmented reality information visualization. In *Proc. of IEEE ISMAR*, pp. 203–208, 2016.

[10] N. A. ElSayed, R. T. Smith, and B. H. Thomas. Situated analytics: Where to put the abstract data? 2016.

[11] D. Kalkofen, C. Sandor, S. White, and D. Schmalstieg. Visualization techniques for augmented reality. In *Handbook of Augmented Reality*, pp. 65–98. Springer, 2011.

[12] T. Langlotz, C. Degendorfer, A. Mulloni, G. Schall, G. Reitmayr, and D. Schmalstieg. Robust detection and tracking of annotations for outdoor augmented reality browsing. *Computers & graphics*, 35(4):831–840, 2011.

[13] A. Quigley. From GUI to UUI: Interfaces for ubiquitous computing. In J. Krumm, ed., *Ubiquitous Computing Fundamentals*, chap. 6. Chapman & Hall/CRC, 1st ed., 2009.

[14] A. Tabard, J.-D. Hincapié-Ramos, M. Esbensen, and J. E. Bardram. The elabbench: an interactive tabletop system for the biology laboratory. In *Proc. of ITS*, pp. 202–211. ACM, 2011.

[15] M. Tatzgern, D. Kalkofen, R. Grasset, and D. Schmalstieg. Hedgehog labeling: View management techniques for external labels in 3d space. In *Virtual Reality (VR), 2014 iEEE*, pp. 27–32. IEEE, 2014.

[16] A. Thudt, U. Hinrichs, and S. Carpendale. The bohemian bookshelf: supporting serendipitous book discoveries through information visualization. In *Proc. of CHI*, pp. 1461–1470. ACM, 2012.

[17] J. A. Walsh and B. H. Thomas. Visualising environmental corrosion in outdoor augmented reality. In *Proc. of Australasian User Interface Conference*, pp. 39–46, 2011.

[18] X. Wang, S. Ong, and A. Nee. A comprehensive survey of augmented reality assembly research. *Advances in Manufacturing*, 4(1):1–22, 2016.

[19] S. White and S. Feiner. Sitelens: situated visualization techniques for urban site visits. In *Proc. of CHI*, pp. 1117–1120. ACM, 2009.

[20] W. Willett, Y. Jansen, and P. Dragicevic. Embedded data representations. *IEEE TVCG*, 23(1):461–470, 2017.

[21] S. Zollmann, D. Kalkofen, C. Hoppe, S. Kluckner, H. Bischof, and G. Reitmayr. Interactive 4d overview and detail visualization in augmented reality. In *Proc. of IEEE ISMAR*, pp. 167–176, 2012.