



HAL
open science

Handling Non-Visible Referents in Situated Visualizations

Ambre Assor, Arnaud Prouzeau, Martin Hachet, Pierre Dragicevic

► **To cite this version:**

Ambre Assor, Arnaud Prouzeau, Martin Hachet, Pierre Dragicevic. Handling Non-Visible Referents in Situated Visualizations. Transactions on Visualization & Computer Graphics (TVCG), 2023. hal-04176381

HAL Id: hal-04176381

<https://inria.hal.science/hal-04176381>

Submitted on 3 Aug 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Handling Non-Visible Referents in Situated Visualizations

Ambre Assor , Arnaud Prouzeau , Martin Hachet , Pierre Dragicevic 

Abstract—Situated visualizations are a type of visualization where data is presented next to its physical referent (i.e., the physical object, space, or person it refers to), often using augmented-reality displays. While situated visualizations can be beneficial in various contexts and have received research attention, they are typically designed with the assumption that the physical referent is visible. However, in practice, a physical referent may be obscured by another object, such as a wall, or may be outside the user's visual field. In this paper, we propose a conceptual framework and a design space to help researchers and user interface designers handle non-visible referents in situated visualizations. We first provide an overview of techniques proposed in the past for dealing with non-visible objects in the areas of 3D user interfaces, 3D visualization, and mixed reality. From this overview, we derive a design space that applies to situated visualizations and employ it to examine various trade-offs, challenges, and opportunities for future research in this area.

Index Terms—Taxonomy, Models, Frameworks, Theory ; Mobile, AR/VR/Immersive, Specialized Input/Display Hardware

1 INTRODUCTION

Situated visualizations are data visualizations that are displayed close to the data's *physical referent*, i.e., the object, space or person the data is about [96]. Compared to ordinary computer-supported visualizations, situated visualizations let users explore data in-context, which can facilitate sense-making and decision making [88].

Consider for example a search-and-rescue scenario similar to the one envisioned by Willett et al. [96], where a fire team needs to rescue people from a burning house. Imagine the firefighters are equipped with augmented reality (AR) displays which show them data about the house (e.g., location of emergency exits), the fire (e.g., temperature and CO₂ level in different rooms), and people (e.g., location of victims and their vitals). All this data is directly overlaid on the physical objects it refers to (e.g., rooms, victims). The firefighters would be able to use this data in combination with perceptual cues informing them about their surroundings (e.g., did parts of the house collapse? Do victims appear distressed or wounded?). A situated information visualization system like this would likely help the fire team reach faster and more reliable decisions than if they had to examine the data on separate displays.

Although this scenario is hypothetical, AR-based situated visualizations have been demonstrated in various areas such as building maintenance [74], urban planning [93], and botany [94], and we will likely see more and more of them. However, designing effective situated visualizations needs a good understanding of their design space. This space is vast and complex and despite recent advances in exploring and understanding it [19, 74, 88, 96], large areas remain poorly understood. In particular, previous work has almost always considered situated visualizations where the physical referent is in sight. However, in the real world, objects are often hidden or out of sight, and thus we cannot always assume that physical referents are visible.

Back to the search-and-rescue scenario, victims may be invisible to the rescue team because they are in another room, occluded by fragments, flames or smoke. Yet they may need assistance, so knowing that they exist, where they are, and what information is available about them is important for the rescue team. The same applies to elements of the environment such as emergency exits. However, it is unclear how to visualize information about hidden physical referents like these and what interaction techniques to use. Since most previous situated visualization systems have been designed with visible referents in mind, they cannot provide us with much guidance for such questions.

We fill this gap by laying out a conceptual framework consisting of a terminology, a typology of techniques, and a design space for handling non-visible physical referents in situated visualizations. Although visual occlusion has been rarely discussed in the literature on situated visualizations, it is a known problem in the areas of 3D UIs and 3D visualization [29], as well as in AR and VR [37]. We survey the techniques proposed in these research areas, and classify them into eight families. We then use those families and previously published conceptual frameworks to build a design space. We illustrate the use of our conceptual framework in three application areas, discuss design trade-offs, limitations, and suggest several directions for future work.

2 BACKGROUND

2.1 Related Work

In this section we discuss related work, including work on situated visualizations and 3D occlusion management.

Situated visualizations. A situated visualization is a data representation located close to the data's *physical referent(s)*, meaning close to the space, object, or person the data is about [96]. Situated visualizations are thus data visualizations that are integrated (or appear to be integrated, in the case of augmented-reality displays) into the physical environment in a meaningful way. For example, a visualization of data about a house for sale is not situated if it shown on a desktop computer, but becomes situated if it is shown next to – or inside – the house.

Various techniques exist for displaying and situating data visualizations into the physical world. One of the simplest approaches is to place regular computer displays close to physical referents [96]. Alternatively, physical data representations (also called data physicalizations [53]) can be used instead of regular displays, and placed next to referents (e.g., placing a robotic plant next to garbage bins to convey data about recycling behavior [49]). The most common approach, however, is to use augmented reality (AR) to display visualizations in a way that makes them appear to be close to their physical referents or overlaid on top of them. AR-based situated visualizations have been explored in various domains such as health (e.g., to show vital information about a patient during a surgery [6]), or aviation (e.g., to display flight information directly in front of the pilot [38]). We will mostly focus on such types of situated visualizations in this paper.

Most situated visualizations are designed for the specific case where the user can simultaneously see the visualization and the physical referent. However, as explained by Ens et al. [31], one of the main challenges of situated visualization is that the data is often displayed in real-world environments that cannot be fully controlled, with cluttered spaces and constantly moving objects and people. As a consequence, in many scenarios there is a high chance that the physical referent is not visible while the user needs information about it.

Taxonomies of occlusion management. As in the physical world, occlusion is a major concern in virtual worlds and in 3D user interfaces and visualizations. As a result, various techniques have been proposed for handling occlusion, either by moving the user's point of view, or

-
- Ambre Assor E-mail: ambre.assor@inria.fr.
 - Arnaud Prouzeau E-mail: arnaud.prouzeau@inria.fr.
 - Martin Hachet E-mail: martin.hachet@inria.fr.
 - Pierre Dragicevic E-mail: pierre.dragicevic@inria.fr.

All authors are with Inria, CNRS, Université de Bordeaux.

by altering the rendering, position or shape of the occluding objects. Elmqvist and Tzigas [29] proposed a taxonomy of occlusion management techniques in 3D visualizations. They classify techniques in terms of their purpose, the arrangement of the visualization, and the user’s status (active/passive). Examples of techniques include x-ray vision (i.e., making occluding content transparent) and multiple views. The paper discusses a few examples of AR visualizations but they are not situated, and therefore the paper does not discuss how to handle the occlusion of physical referents. More recently, Macedo and Apolinario [37] surveyed AR techniques for handling objects that are physically hidden in real environments. In particular, they reviewed existing x-ray vision techniques that overlay hidden objects in front of occluding surfaces as if the user could see through them. Much of their paper focuses on how to fine-tune the rendering techniques to make the x-ray vision experience as believable as possible. The paper does not discuss data visualization. We borrow from this previous work to discuss how to handle non-visible physical referents in situated visualizations.

Situated visualization frameworks addressing non-visible referents. Closer to our own work, two conceptual frameworks have been recently published that address specific aspects of handling non-visible referents in situated visualizations. Lin et al. [61] reviewed labeling techniques to support visual search for situated visualizations where physical referents are out-of-view. We build on this survey and cover labeling approaches, but we also cover a range of other techniques for addressing non-visible referents, including techniques that convey information about referents in a way that does not require people to find them. Another closely related contribution is the work on ProxSituating Visualizations [77], which focuses on the use of proxies for physical referents – i.e., virtual representations that replicate physical referents. The paper extends Willett et al.’s [96] model to account for proxies, and observes that a visualization can be considered situated if it is close to a proxy, even if the physical referent is far away – a situation we will often encounter in this paper. Again, this paper focuses on specific techniques for handling non-visible referents. We do not build on this work as this is parallel work, but we will contrast their design space with ours in detail in [subsection 5.2](#). Despite these two prior conceptual frameworks, to our knowledge, no research has offered a broad and comprehensive survey of different ways non-visible physical referents can be handled in situated visualizations. This paper fills this gap by taking inspiration from the many occlusion management techniques proposed for VR, AR, and 3D user interfaces, and generalizing them to situated visualizations to build a comprehensive conceptual framework.

We now describe two short scenarios we will use throughout this paper, in order to illustrate our conceptual framework. The first scenario elaborates over the fire rescue example we mentioned in the introduction, while the second scenario covers a more casual situation.

Scenario 1. Alice is part of a firefighter team. Just before dinner time, the team receives a fire alert involving a nearby house. They know from the caller that the household consists of two parents, two children and an elder, and that one parent and the firstborn are out of the house and on their way to the incident. Soon, the fire truck speeds on the road and inside, the team prepares for the operation: they check the house layout on a 2D map, focusing especially on fire exits. They call the safe parent to gather more information and then assign each other tasks. Alice is in charge of rescuing the younger child, and she identified in the map that his bedroom is on the second floor. The parent confirmed that the child is sleeping in his bedroom because the room is equipped with a video baby phone connected to his smartphone. Some of Alice’s teammates have been assigned the rescue of the other family members, while others are in charge of extinguishing the fire. Two members are in charge of staying near the truck to coordinate the whole operation, give instructions and monitor information such as the amount of water left and the vitals of the team members captured in real time by wearable physiological sensors. Once the team arrives, the rescue operation goes relatively smoothly aside from a few unexpected events. For safety reasons, the rescue team first enters the basement and cuts the electricity. From there, they use stairs to reach the rest of the house. They manage to save everyone but not without difficulty due to the obscurity, the presence of heavy smoke and the important




Fig. 1: Diagram summarizing the key terminology introduced in Section 2.3. The pictures refer to Scenario 1, where a firefighter needs to evaluate the state of the victim hidden behind a wall.

number of rooms making it hard to see anything.

Scenario 2. Jessie gets home where her husband Max has been waiting for her. These days, they are trying to change their cooking and shopping habits in order to waste less. They took the habit to regularly log their groceries and the meals they prepare in a phone app. Tonight, Max is thinking about what he is going to prepare for dinner. He opens the fridge and sees that the bottom drawer is full with vegetables and fruits, and that the leeks must be used soon. Max opens the cooking book and chooses a leek pie. He leaves the book open on the counter and looks at the ingredients which are all listed with small illustrations along with the cooking instructions. He goes back-and-forth between the kitchen worktop and the counter to read the different steps. He realizes two items are missing: one they thought they had but forgot to buy, and another one which has expired. However, they are not essential and Max decides to continue. He puts the pie in the oven, starts to wash the dishes, and after a while he smells a slight burnt odor. He turns around and sees it is high time to take the leek pie out of the oven. Fortunately, the pie turned out pretty good at the end.

These two scenarios are very different but they both involve situations where people are engaged in tasks that require access to information about objects in the physical environment, and these objects may be hidden from sight. This information may be something that can be directly seen as long as the object is visible (e.g., whether a vegetable is ripe), or something that needs to be conveyed by a visualization (e.g., the firefighters’ vitals). In this article, we are interested in cases where data is visualized, and where seeing objects in the environment may provide complementary information [96].

2.2 Basic Terminology

In this section we introduce the basic terminology that underlies our conceptual framework, and which will be useful for describing the design space. [Figure 1](#) gives a visual overview of our terminology.

An **object of interest** is an entity about which the user wants to learn. For example, in Scenario 1, a person to rescue, an emergency exit, or another firefighter can all be objects of interest for a firefighter. In Scenario 2, an object of interest can be a grocery item necessary for the recipe. Whether an object is of interest depends on the user’s task and intent, which can change over time. However, we will be using the term “object of interest” liberally, to refer to objects that are currently of interest but also to objects that are expected to become of interest at some point in time. Note that an object of interest can be a physical referent, but it is a more general term as it does not indicate that any data visualization is attached to the object.

A **representation** of an object of interest is a visual depiction that stands for this object. It can reproduce the object’s visual appearance or not. For example, in Scenario 1, the image of the child that appears on

the parent’s smartphone is a relatively detailed and faithful representation of the actual child. In contrast, elements depicted on the house map like emergency exit icons are purely symbolic representations. The drawings of the ingredients in the cooking book from Scenario 2 are an intermediary example: while the drawings approximate the appearance of the actual ingredients, they are generic drawings which do not show for example whether the actual leaks in the fridge are overripe. This notion of representation faithfulness will be further discussed in the design space section.

Data is information about an object of interest that is not necessarily conveyed by its visual appearance. For example, in our Scenario 1, information like the composition of the household, the firefighters’ vitals, or the amount of water left is information that exists but that cannot be accessed through mere visual inspection. In Scenario 2, examples of data include the ingredients’ expiry date, the list and quantity of ingredients in a recipe, or the history of Jessie and Max’s meals. The location of an object of interest is also data. Data can be conveyed through data **visualizations**. Following previous terminology [96], whenever data is visualized about an object of interest, the data’s **physical referent** is this object of interest. In addition, when the visualization is near (or appears to be near) the object of interest, it is **situated**.

We refer to an object of interest as **non-visible** if it is visually imperceptible by the user at a given time. For example, in Scenario 1, while Alice is in the basement, the child located in the second floor is non-visible. In Scenario 2, most cooking ingredients are non-visible unless the fridge and kitchen cupboards are opened. We identify three main reasons why an object of interest may be non-visible:

1. Following previous terminology [62], an object is **out-of-view** if it is not within in the user’s visual field but can be seen if the observer turns their head. For example, in Scenario 2, the pie in the oven is out-of-view while Max is washing the dishes.

2. An object is **occluded** if its sight is blocked by another object. For example, in Scenario 1, the child is occluded by walls. We more specifically refer to an object as **hidden** if it is in the user’s vicinity but occluded from all reachable locations and viewpoints. For example, in Scenario 2, as long as the fridge’s door is closed, there is no viewpoint from which Max can see the food it contains.

3. An object is **illegible** if none of its useful details is visible to the user, for reasons other than being out-of-view or occluded. For example, the object can be too small, too far, moving too fast, or insufficiently illuminated. In Scenario 2, the cooking book is too far for Max to be able to read it, forcing him to move back-and-forth between the book and the kitchen worktop.

3 HANDLING NON-VISIBLE OBJECTS IN IMMERSIVE ENVIRONMENTS: A SURVEY

Although there has been little work on handling non-visible referents in situated visualizations, there has been a substantial body of research in VR and AR on techniques for handling non-visible objects of interest. In this section, we present a survey of such techniques and classify them into eight categories. This survey will inform the design space we lay out in the next section.

3.1 Methodology and Scope

In order to collect possibly relevant papers, we conducted multiple keyword searches on several scientific search engines: Google scholar, Scopus, ACM DL and IEEE Xplore. Our search terms consisted of *occlusion*, *off-view*, *distant*, *hidden*, and *awareness* combined with either *augmented reality*, *mixed reality*, or *immersive*. In addition, we added papers cited in the 2008 taxonomy of 3D occlusion management techniques from Elmqvist and Tsigas [29].

In order to keep the size and complexity of the survey and of the resulting design space within reasonable limits, we iterated over a set of criteria to restrict the scope of our survey. We converged towards three major criteria; For a paper to be included, the technique it presents had to meet the following requirements:

1. **Visual only.** The technique needs to use visual representations. Techniques that increase the user’s awareness about non-visible objects

of interest through non-visual channels (e.g., haptic or auditory) were not included in this survey.

2. **Object-aware.** The technique requires that an object of interest is defined. Techniques that are object-agnostic – e.g., techniques that broaden the field of view or magnify vision – were not included. This choice is consistent with our focus on situated visualizations, where the notion of object of interest is central (objects of interest are physical referents, i.e., objects about which data is visualized).

3. **No movement required.** The technique should be usable from any location in space. Techniques that do not provide information about the object of interest but instead help the user access it by locomotion or body movements were excluded.

We ended up with a total of 107 papers, which we divided into eight families of techniques. We give an overview of the eight families here. The full list of papers is available as supplementary material.

3.2 Worlds-In-Miniature




Fig. 2: A world-in-miniature showing a scaled-down version of the physical room where the user is located, overlaid in their visual field using augmented reality [11].

A World-In-Miniature (or WIM) is a virtual scaled-down replica of the environment the user is in, embedded within the users’ view of their immediate surroundings. The scaled-down environment is often virtual. Figure 2 shows an example of a WIM of a real environment (the laboratory room the user stands in), displayed in AR. The purpose of a WIM is to help users discover and learn about the environment around them, including parts that may be occluded, out-of-view or illegible.

WIMs have been used in various domains, including for medical imaging [25], building exploration for tourist guides or situated urban planning [87], air traffic visualization [28], US Forest Service’s forest inventory data visualization [71], and interacting with smart home and IoT appliances [2, 79]. In both AR and VR, WIMs have been mainly explored for supporting navigation [72], but they can also be used to get information about hidden objects of interest.

WIMs are more or less faithful depictions of the full-scale environment, depending on the implementation. In VR, WIMs can be exact copies. In AR, WIMs are typically simplified versions of the real scene (e.g., Figure 2), but some WIMs are photorealistic copies [16] (here, a minimap).

For a WIM to be usable, it is important that it can be freely manipulated. Although the first implementations only allowed users to move their head closer or further away from the WIM [10], further implementations allowed them to grab the WIM and rotate it [3], and even re-scale it [27]. Direct manipulation of elements inside the WIM has also been explored – for example, one application allows users to move furniture in a virtual room by moving their replica inside the WIM [15]. In another application showing building consumption in a university campus, users can get additional info about buildings by grabbing them within the WIM [32]. Some variations of WIMs even address occlusion problems with the WIM itself, for example by allowing users to make walls transparent [90], or to select non-visible regions of interest [3].




Fig. 3: A virtual twin of a smart factory shown in augmented reality, allowing the user to see the state the factory in real-time and remote-control it [59].

3.3 Virtual Twins

A virtual twin¹ is a dynamic 3D model that conveys the visual appearance as well as the key properties (e.g., state, features, physical measurements) of a physical system or object. Figure 3 shows an example of the virtual twin of an illegible factory.

Conceptually, a virtual twin is similar to a WIM, except it replicates an object instead of a spatial environment, and it focuses on replicating its key properties. Although a virtual twin can faithfully convey the visual appearance of the object of interest [86], often the focus is on conveying its important properties and associated data. Virtual twins have been suggested for a range of applications, including for surgery [60], radiology [82], assembly tutoring [75] and for sharing content during teleconferencing [57]. They can help support maintenance [59], learning [75], and collaboration [57] tasks.

3.4 Physical Replica




Fig. 4: Tangible interface for visualizing the energy consumption of a university [32].

A physical replica is a physical representation of an object of interest. As an example, Figure 4 shows a tangible model of a campus augmented with energy consumption data. The model is installed in a room within the campus itself, and thus it allows users to get an overview of the whole campus, which is out-of-view for the most part.

Physical replica are conceptually similar to WIMs and to virtual twins, except they are physical. A benefit of physical replicas is that they can be designed to be easy to grab and manipulate, as in the Uplift system [32] (Figure 4). However, accurate and dynamic physical replicas are difficult to build. As a consequence, the purpose of most physical replicas that have been developed so far is not to convey detailed and faithful information about the object of interest, but rather to serve as support surfaces on which to display digital information about the object, or as tangible controllers for navigating this information. Examples include the Tangible Globe [78] which users can use as an input device to browse information about the Earth, the head probe from Hinckley et al. [47] which users can use to control a virtual cutting

¹There are many definitions of virtual and digital twin in the literature. Here we consider that a virtual twin is a digital twin with a 3D visual representation.

plane on a separate screen, or the humanoid puppet Teegi which has a display integrated showing EEG data [39].

3.5 Telepresence Systems




Fig. 5: Example of a teleexistence/telepresence system. User A on the left sees a virtual representation of the persons B and C they want to interact with. The user can also remotely control a robot as if they were physically in the same location [85].

Telepresence or teleexistence broadly refers to techniques that give the user the impression of being in a different physical location, typically by providing an immersive view of the remote environment or object, and/or by allowing users to directly interact with it from a first-person perspective [80, 85]. Figure 5 shows a system allowing a user to see and interact with remote people as if the user was in the same room.

Telepresence systems sometimes let users physical interact with remote objects of interest using physical actuators or robots – for example, physically operating a remote cockpit through a robot [84]. In many systems focusing on social telepresence, the objects of interest are other people, who are rendered virtually as if they were near the end user [80]. Often these representations do not resemble the people themselves, but research has started to explore automatic 3D reconstruction [40, 65]. Alternatively the representation of remote people can be physical, like in the case of mutual teleexistence systems where remote users are represented by robots [58]. Overall, telepresence covers situations of teleconferencing [65], visualization sharing at a distance [92], tutoring [89]. School and Healthcare have been also identified as promising application areas [58].

3.6 X-Ray Vision

X-ray vision techniques emulate the fictional superhuman ability to see through occluding surfaces [70, 95].

Applications of x-ray vision range from visualizing the heat distribution of buildings [8] to assisting underground exploration [34]. A range of rendering techniques have been proposed to emulate x-ray vision in augmented reality (for reviews, see [63, 70]). In x-ray vision systems, objects of interest may be reconstructed based on 3D models [29, 56] or on video images, for example as captured by a drone [33]. Although most implementations of x-ray vision do not focus on interaction aspects, two systems have been proposed that use a torch metaphor to specify the region to reveal: one allows users to explore virtual annotations left in different rooms of a workspace [20], while the other one allows users to control home appliances such as lights and TVs even when they are in other rooms [41]. Explosion diagrams are closely related to x-ray vision, but they displace occluding surfaces instead of removing them or rendering them transparent [29, 56].

3.7 Out-Of-View Pointing

Out-of-view pointing techniques give information about the presence and location of an object of interest that is outside the user's field of view, typically by indicating the object's direction on the edge of the display. Figure 6 shows an example in AR.




Fig. 6: Out-of-view pointing in SidebARs [81]. On the top right, arrows point toward locations of fire companies.

Out-of-view pointing techniques were originally developed for regular computer displays [9, 17, 46], and later adapted to VR [23, 44] and AR [43, 45] displays. Some of the techniques augment the display hardware itself, e.g., with LEDs [64]. In the vast majority of the cases, no information is given about the object of interest apart from its existence and its approximate location. Most techniques only give direction clues [44, 45] while others also convey the distance of the object of interest [54]. A few applications show data visualizations about the object of interest for notifications [51], for fire companies [81] (figure 6) or in order to discover a city [36]. But in general, the amount of information shown with this technique is limited.

3.8 Labelling




Fig. 7: The stereo panorama of the Luxor Temple was provided by Tom DeFanti, Greg Wickham, and Adel Saad, with stereo rendering by Dick Ainsworth for AR labelling [76].

Labelling in AR applications consists of displaying additional textual information about objects of interest, which can be visible or not [62]. Figure 7 shows an AR application that displays information boxes linked to the physical objects they belong to.

Labels are often displayed near the objects they refer to, and sometimes visual links are drawn between the labels and the objects to highlight their association [1]. In order to position the labels at the right place, a specific layout tracking or processing system is often used [42] along with an adequate view management system [10] [42]. Recent work has introduced the use of labels in the specific case of out-of-view objects, and highlighted their benefits [62]. In terms of interaction, the user generally does not have explicit control over the position of the labels. However, in order to make sure the labels are visible and readable, their position is often automatically adjusted to the position and orientation of the users' head [62]. Although labels can sometimes provide rich information as in Figure 7, in most cases, labels only show very basic information about objects of interest. This makes them similar to out-of-view pointing techniques.

3.9 Virtual Mirrors

A virtual mirror is a virtual tool that behaves like a physical mirror, i.e., it reflects nearby objects. Figure 8 shows an example.

Users can generally manipulate a virtual mirror as they would manipulate any real mirror [13]. The mirror's surface shows the simulated reflected image that allows the user to see non-visible content, such as hidden faces of objects of interest or objects outside one's field of




Fig. 8: Example of a virtual mirror used to see the hidden side of a spine [12].

view. A variation of the virtual mirror are physical mirrors whose image is digitally-generated and therefore virtual. For example, in the Mind-Mirror project [66], users are equipped with an AR display and a real mirror, and can see their own electroencephalography data through the real mirror. One of the major benefits of mirror-based techniques is that the mirror metaphor is very familiar to most users.

4 DESIGN SPACE

To lay out possible techniques that can handle non-visible referents in situated visualizations, we describe 14 design dimensions broken down into three categories, depending on whether they are relative to the representation, to the visualization, or to the general situated visualization system (see Table 1 for an overview).

We will reuse the two scenarios introduced in Section 2.1 to illustrate the different values the dimensions can take. In contrast to the initial scenarios, we will assume that Alice and Max are now equipped with lightweight AR head-up displays. We will ignore technological limitations and implementation issues – for example, if Alice uses X-ray vision to see behind a wall, we will not specify how the hidden environment is visually reconstructed: it could be from an array of surveillance cameras, from drones equipped with depth sensors, etc.

4.1 Process

This design space is the product of several (about 15) meetings between the co-authors on the topic of handling hidden referents in situated visualizations, and an analysis of the literature. Our first meetings were dedicated to identifying different types of non-visibility, from which we derived the typology reported in Section 2.2. In the next meetings, we fleshed out a design space of occlusion management techniques, starting with the dimensions from Elmqvist et al.'s taxonomy of 3D occlusion management for visualization [29]. We kept iterating on our design space in the subsequent meetings, and informed our discussions by examining other related design spaces and taxonomies from the literature (see Section 5.2 for details). We identified aspects from these conceptual frameworks that were missing or under-discussed to be able to cover situated visualization systems – for example, the need to distinguish between the representation of the physical referent and the visualization, the notions of spatial and temporal indirection, and the use of mixed-reality displays. The survey of techniques from section 3 was developed in parallel to our design space, and the two informed each other. In particular, we used the families of techniques to identify missing dimensions from our design space, and we used our design space to refine the definitions of our families of techniques. Finally, once our set of dimensions was stabilized, we generated three use cases to test and illustrate our design space (reported in subsection 5.1).

4.2 Category: General

The dimensions in this category capture general characteristics of the situated visualization system that do not specifically apply to the representation of the physical referent or to its visualization.

Non-visibility support: Refers to the types of non-visible physical referents that are supported by the situated visualization system (see subsection 2.2 for our terminology on the subject). In contrast to the other dimensions, the values here are not mutually exclusive.

Table 1: Summary of all 14 design space dimensions. The dimensions are splitted into the three categories of Column 1. Their name appears Column 2, followed by a synthesized descriptions of each Column 3 and the different values they can take Column 4.

Category	Dimension	Description	Values		
General	Non-visibility support	The types of non-visible physical referents that are supported.	Out-of-view	Occluded	Illegible
	Interactivity	The extent to which the user can interact with the representation or the visualization.	Low	Medium	High
	Interaction realism	The extent to which the interaction techniques are reality-based.	Low	Medium	High
	Actionability	The extent to which the user can act on the non-visible physical referent.	Low	Medium	High
Representation	Representation fidelity	The level of detail with which the physical referent is represented.	Low	Medium	High
	Representation refresh rate	How often the representation of the referent is updated to reflect the real referent.	Low	Medium	High
	Surroundings inclusion	How much of the environment around the referent is included in the representation.	Low	Medium	High
	Representation medium	Whether the representation of the physical referent is digital or physical.	Digital	Mixed	Physical
	Representation integration	To what extent the representation appears to be integrated in the user's environment.	Low	Medium	High
Visualization	Location information	The amount of information shown about the physical referent's location.	Low	Medium	High
	Visualization richness	The amount of data visualized about the physical referent.	Low	Medium	High
	Visualization refresh rate	How often the visualizations are updated to reflect data about the physical referent.	Low	Medium	High
	Visualization medium	Whether the visualizations are digital or physical.	Digital	Mixed	Physical
	Visualization integration	To what extent the visualizations appear to be integrated in the user's environment.	Low	Medium	High

- *Out-of-view*: The technique allows the user to see information about physical referents that are outside their field of view. Example: In Scenario 1, Alice stands in front of the door of the bedroom where the child is stranded, thinking about what to do next. Meanwhile, she talks to her partner behind her. Even though she is not facing him, a virtual rearview mirror allows her to see him.
- *Occluded*: The technique allows the user to see information about occluded referents. Example: Alice can see the child and information about her through the bedroom door using an X-ray vision display.
- *Illegible*: The technique allows the user to see information about illegible referents. Example: Alice zooms into the child's face to see if she is still conscious.

Interactivity: The extent to which the situated visualization system allows the user to interact with – or reconfigure – the representation of the physical referent or the data visualization (see again Figure 1 for our terminology). A technique like World-In-Miniature is typically interactive, because users can rotate and rescale the representations. However, the same technique can have very different levels of interactivity.

- *Low*: The representation and the visualization are non-interactive. Example: In Scenario 1, Alice searches for the stranded child. She sees a video of the child with basic information about her vitals and an arrow pointing to her direction, but there is no control to change the camera settings or get further details about the child's state.
- *Medium*: The user can interact with the representation or the visualization to a limited degree. Example: Alice can turn the camera image and the visualization on or off.
- *High*: The representation and the visualization are highly interactive. Example: Alice can finely tune the camera image and the arrow representation, and explore many different types of data dimensions and visualizations for the vitals.

Interaction realism: The extent to which the interaction techniques are reality-based, i.e., “increase the realism of interface objects and allow users to interact even more directly with them—using actions that correspond to daily practices within the non-digital world” [52]. In the literature on immersive analytics, such interactions are also called embodied [26]. Examples of reality-based interactions are virtual mirrors, as they are used in a very similar way to real mirrors.

- *Low*: The interactions are not reality-based. Example: Alice is looking for the child's bedroom. She opens a virtual menu and selects an option that displays a 3D arrow pointing at the bedroom. Once in front of the bedroom door, she uses another menu to activate an X-ray display that reveals the interior of the room. She then activates visualizations of the room's temperature and air quality, and of the child's vitals.
- *Medium*: The interactions are weakly reality-based or combine elements that are reality-based and others that are not.
- *High*: The interactions are reality-based. Example: Alice looks around, which activates the 3D arrow. Once she sees the door, she grabs a physical controller and uses it like a torch to activate the X-ray display (as in [41]). The closer she moves to the door, the more detailed information she gets about the room and the child.

Actionability: The extent to which the user can act on the non-visible physical referent. For example, IoT and smart home applications can give users control over lights or home appliances (e.g., [41]), and thus they support actionability. In contrast, techniques such as out-of-view pointing and labelling rarely support actionability.

- *Low*: The situated visualization system does not let the user act on the physical referent. Example: in Scenario 2, while Max washes the dishes, his heads-up display shows a heatmap visualization of the pie's temperature. The visualization indicates that the pie is about to burn, but Max has to walk to the oven to turn it off.
- *Medium*: The technique supports basic actions on the physical referent such as flipping electronic states on and off, and other actions typically supported by remote controllers. Example: Max has an option to turn the oven off without having to walk to the oven.
- *High*: The technique allows rich manipulation of the physical referent, which include mechanical actions. Maximum actionability is achieved when the user can manipulate the object as if they were standing in front of it. Example: Max can open the oven's door and pull the oven rack out from a distance.

4.3 Category: Representation

The dimensions here capture characteristics of the representation of the physical referent.

Representation fidelity: Refers to the level of detail with which the physical referent is represented. It can go from a simple symbol to a highly-detailed model of the physical referent.

- *Low*: The representation of the physical referent is either non-existent or does not convey information on the visual appearance of the referent (i.e., it is a symbol [55]). Example: Alice sees a “plus” symbol indicating the location of the child behind the wall.
- *Medium*: The representation conveys some elements of the visual appearance of the physical referent (i.e., it is a simple icon [55]). Example: Alice sees the icon of a girl standing for the child.
- *High*: The representation is detailed and matches the visual appearance of the physical referent reasonably well. Examples: Alice sees a photo or a 3D model of the child.

Representation refresh rate: How often the representation of the physical referent is updated to reflect the real referent.

- *Low*: The representation of the physical referent is generated once and never updated. Example: in Scenario 2, Max asks the situated visualization system to show him the fruits and vegetables he bought a week ago at the market. The system shows him photos of the produce taken before they were put in the fridge, which tells about their number, type and size, but not whether they are about to go bad.
- *Medium*: The representation of the physical referent is occasionally updated. Example: Max asks the situated visualization system to show him the interior of the fridge. The system takes a photo every morning, and shows him the last one. Max can see the fruits and vegetables, and which of them should be eaten soon. However, he does not see that Alice ate all the remaining grapes just an hour ago.

- *High*: The representation of the referent is updated in real time. Example: Max sees a real-time video of the interior of the fridge.

Surroundings inclusion: How much of the physical environment surrounding the referent is included in the representation.

- *Low*: The surroundings of the physical referent are not represented. Example: Thanks to the x-ray feature, Alice sees the child through the bedroom door, but only the child is represented. She cannot see whether there are hazardous elements around the child.
- *Medium*: Several elements of the physical referent's surroundings are included in the representation. Example: The x-ray feature is able to capture and render the entire bedroom, but Alice can only see the portion that is visible through the virtual hole in the door.
- *High*: The representation includes a large part of the referent's surroundings. Example: Alice switches from the x-ray representation to a world-in-miniature replica of the bedroom, where she can see the child as well as all details of her physical environment.

Representation medium: Whether the representation of the physical referent is digital or physical. It can go from completely digital to completely physical, or can be anywhere in-between [53, 67].

- *Digital*: The representation is a computer-generated image that is immaterial. Example: While the truck is speeding to the house on fire, the firefighting team downloads a floor plan of the house from a central database, and video-projects it on the truck's floor.
- *Mixed*: The representation combines physical with digital elements. Example: The firefighting team uploads a floor plan of the house on a shape display, i.e., a device with motorized pins [53]. The shape displays morphs into a coarse physical replica of the house's floors. Other details (doors, fire exits, etc.) are video-projected.
- *Physical*: The representation is purely physical. Example: a high-resolution shape display whose pins can change color morphs into a detailed physical replica of the house's floors.

Representation integration: The extent to which the representation appears to be integrated in the user's physical environment. If the representation is physical (see above), its integration with the physical world is necessarily high. If it is digital, it depends on the visual metaphors and rendering techniques used to display it. One key factor is whether the representation is egocentric (it follows the user) or exocentric (it is attached to the physical referent or its environment).

- *Low*: The representation of the referent does not appear integrated in the physical environment at all, and it is very clear that this representation is a virtual image. Example: Alice sees a 3D model of the child floating in the air.
- *Medium*: There is some sense that the representation is integrated in the physical environment, but this integration is not fully believable. Example: Alice sees the child through the door but it is clear that the representation has been added, because of imperfections of the AR rendering, and because it is not possible to see through physical objects in reality.
- *High*: The representation appears to be an integral part of the physical environment. Example: The situated visualization system renders a perfect full-size replica of the child's bedroom in front of Alice, as if the door and the walls have been completely removed.

4.4 Category: Visualization

The dimensions in this category capture the characteristics of the visualizations of data about the physical referent.

Location information: The extent to which the situated visualization system visualizes information to help the user know where the referent is located. This dimension can be low or high depending on how precise and complete the information is.

- *Low*: The situated visualization system gives no information about the location of the physical referent. Example: Max needs to assemble the ingredients listed in the cook book. The system shows him the ingredients he currently has but not where to find them.
- *Medium*: The system gives hints about the location of the physical referent: for example, its distance or its direction. Example: The system adds an arrow next to each ingredient listed in the cookbook, that tells Max in which direction to go to find the ingredient.

- *High*: The systems gives enough information to be able to quickly find the physical referent. Example: a combination of arrows, labels and x-ray representations tell Max exactly where to find each ingredient.

Visualization richness: The amount of information shown about the physical referent. It can range from low (e.g., as in simple labelling systems) to high (e.g., as in many virtual twins). This dimension is analogous to representation fidelity, except representation fidelity is about the visual appearance of the physical referent, whereas visualization richness is about its associated data. In addition, there is an upper limit to representation fidelity, but no such limit for visualization richness.

- *Low*: The situated visualization system shows minimal information, such as the name or the location (see before) of the physical referent. Example: Alice sees a label next to the child behind the door, indicating her name and age.
- *Medium*: The system visualizes some data about the physical referent. Example: next to the child, the system displays her body temperature, breathing rate and heart rate.
- *High*: The system visualizes lots of data about the physical referent, allowing the user to perform analytical tasks. Example: Alice sees detailed health data about the child, allowing Alice to estimate how much time is left before the child starts suffocating.

Visualization refresh rate: How often the visualization is updated to reflect data about the physical referent. Refresh rate can be low either because the visualization is not updated frequently enough, or because the data is not updated or transmitted frequently enough. This dimension is analogous to the representation refresh rate, but applied to the referent's data.

- *Low*: The data is collected and visualized once, and never updated. Example: Alice sees a visualization of the planned trajectories of all her firefighting partners, but does not see their current location or any new information collected since they entered the house.
- *Medium*: The data and its visualizations are occasionally updated.
- *High*: The data and its visualizations are updated in real time. Example: Alice sees the real-time location of her partners, which she can compare with the trajectories to see who deviates from the plan.

Visualization medium: Whether the visualizations are digital or physical. This dimension is analogous to the representation medium, but applied to the visualizations.

- *Digital*: The visualizations are computer-generated images that are immaterial. Example: In their truck, the firefighting team video-projects a floor plan of the house on the truck's floor, together with visualizations of temperature and air quality data.
- *Mixed*: The visualizations combine physical with digital elements. Example: A shape display morphs into a physical replica of the house floors, and temperature and air quality data are video-projected on top of this replica.
- *Physical*: The visualizations are purely physical. Example: on the shape display, the walls and doors on the scale model change their texture and color to reflect temperature and air quality data.

Visualization integration: The extent to which the visualizations appear to be integrated in the user's physical environment. This dimension is analogous to the representation integration, but applied to the visualizations. It is generally more difficult for visualizations to be perceived as highly integrated, because visualizations tend to use abstract encodings that do not resemble objects in the real world [21].

- *Low*: The visualizations do not appear integrated in the physical environment at all. Example: While facing the bedroom door, Alice sees data about the air quality inside, as 2D charts floating in the air.
- *Medium*: There is some sense that the visualizations are integrated in the physical environment, but this integration is not fully believable. Example: With her x-ray vision, Alice sees situated volumetric visualizations of air quality data through the door, but it is clear that those visualizations are not part of the physical world.
- *High*: The visualizations appear to be part of the physical environment. Example: the situated visualization system replicates the interior of the bedroom perfectly and at full scale, giving the illusion that walls and doors have been removed. Smoke of different colors can be seen in different areas of the bedroom, to convey high

concentrations of invisible but harmful gases and particles.

5 DISCUSSION

In this section, we illustrate the use of our design space in different application areas, discuss trade-offs and limitations and propose several directions for future work.

5.1 Use Cases and Design Trade-Offs

To illustrate the different concepts from our framework (terms highlighted in bold), we discuss three examples of application areas for situated visualizations. For each of them, we cite an existing research prototype from the visualization literature, and discuss different ways in which this prototype could handle non-visible referents.

1. Building Management. Some buildings like factories and power plants contain many sensors producing large quantities of **data** [74]. Technicians need to understand this data to be able to deal with failures or optimize the building's behaviour (e.g., its energy consumption). It can be useful to have AR systems that can position data **visualizations** directly near their **physical referents**, but often the referents are **non-visible**. Consider an air conditioning system: **objects of interests** may include air handling units, air return fans, dampers, pipes and vents. Most of these are non-visible due to **occlusion**: they are either behind walls, in the ceiling or in the floor. Suppose we want to 1) help technicians diagnose air conditioning problems from anywhere in the building, 2) help them locate and fix the problem once they are in the room. For task 1), the AR system could use a virtual twin, as in the Corsican Twin [74], which uses virtual twins for building management. In such a case, it is important to provide as much useful information as possible about objects of interest, which could imply supporting a high **representation fidelity**, a high **visualization richness**, and **surroundings inclusion**. Now if we turn to task 2), surroundings inclusion is much less important, since the technicians can now see the physical environment. However, choosing a technique with high **integration** like x-ray vision is desirable, as technicians may need to quickly test and fix broken equipment. While equipment repair often requires physical presence, other building maintenance tasks can be done through the remote control of electrical and mechanical equipment. In these cases, integrating remote control features directly in the situated visualization system itself can help support **actionability**. Furthermore, the system could let users control remote objects of interest directly through their representation, as if they were actually manipulating the objects themselves, and thus support high **interaction realism**.

2. Studying Plants. Situated visualizations can help study plants: for example, Virtual Vouchers [94] is a mobile AR system that helps botanists recognize and study plant species using a large dataset consisting of photographs of dead leaves called vouchers. Now, commercial mobile apps (e.g., <https://plantnet.org>) can automatically recognize plants and display information about them. In all such cases, the **objects of interest** are plants, and AR can be used to display **data** about those plants through **visualizations**, next to the plant themselves. Examples of data include geo-temporal records of the species [94], or the place and time the same species was identified. While current apps assume users are always facing the plant they want to learn about, there are possible scenarios where the plants of interest cannot be seen: For example, a botanist, a forest conservationist, or a nature enthusiast may want to go to a location known for the presence of a rare plant species, or may go on a hike and wish to learn about interesting plants nearby. In such cases, a situated visualization may use a geographical database of plant species distributions [18] or individual specimens [14], and show data about plants that are in proximity but cannot be seen. The system could additionally help users orient themselves using techniques with high **representation integration** and low **surroundings inclusion** (e.g., off-view pointing techniques or labels) or techniques making the opposite trade-off (e.g., worlds-in-miniature). For plants, which do not change rapidly, the use of a high **representation refresh rate** or **visualization refresh rate** is not necessary. Having a high **representation fidelity** may be useful however, but is very difficult to do technically, unless the geographical database of plant species is regularly updated by volunteers taking photos or scans of plants.

3. Caregiving. Situated visualizations can be used to assist caregivers, nurses, and doctors by providing them with information about patients in a medical facility or a care home. For example, Situated Glyphs [91] provide caregivers with situated medical information about patients, such as records of their recent incidents (e.g., a fall) or their schedule (e.g., lunchtime, recreational activities, family visitation). In this context, the **objects of interest** are the patients, and the patients may be **non-visible** if they are in a different room than the caregiver. Yet caregivers may still need to access the data, especially for monitoring purposes or in a presence of a suspected emergency. Despite all their advantages, situated physical displays like Situated Glyphs have a poor **non-visibility support**. Medical facilities often have control rooms for monitoring patient vitals, but when caregivers are neither in the control room nor in the patient's room (e.g., they are moving between patients), they lose access to the information. Therefore, an alternative could be to equip caregivers with lightweight AR displays where they could see both a **representation** of the patients and **visualizations** conveying important information about them. In the context of caregiving, a high **visualization refresh rate** is important for patients under intensive care or who are at a risk of accident. If keeping an overview of all patients at all times is a desirable option, then the system would need to use a low **visualization richness** to prevent clutter, and offer the user the option to increase visualization richness for specific patients of interest. In this case, having a high **representation fidelity** as well as a high **representation refresh rate** could potentially convey important qualitative information that complements the quantitative information captured by sensors.

5.2 Link to Other Conceptual Frameworks

In this section, we discuss the differences and similarities between our terminology and design space, and other conceptual frameworks.

Non-visibility support. In the ProxSituating Visualization paper [77] we already mentioned in Section 2.1, the authors relate the visibility of the referent to *spatial indirection*, i.e., the extent to which the physical referent is far from the visualization and the user. Our terminology on non-visibility (Section 2.2) is more precise and considers whether the physical referent can actually be seen irrespective of spatial indirection, and if not, why. Our *non-visibility support* dimension also relates to the *scope* dimension of Willett et al.'s [95] visualization-as-superpower framework: the more the kinds of non-visibility supported, the larger the scope of the situated visualization system.

Interactivity and interaction realism. Those dimensions were initially inspired by the *interaction model* dimension in Elmqvist's [29] taxonomy of 3D occlusion management for visualization. Like their taxonomy, our dimensions aim to capture how the user interacts with the system: are they active or passive, what kind of gestures do they use, are they in control? We first reused Elmqvist's *active/passive* dimension. However, in XR environments, users are rarely completely passive – we capture this with our *interactivity* dimension, which lies on a continuum. This dimension also maps to the *degree of control* dimension in Willett et al.'s [95] visualization-as-superpower framework. In addition to the level of interactivity, users can be active in very different ways; In XR research, interactions that go beyond classical interaction styles are often qualified as embodied [26] or reality-based [52]. The two terms are similar but the former emphasizes the use of the body, while the latter emphasizes interactions that are close to real-world actions. We chose the latter term for our dimension *interaction realism*, because we found that it made it easier to classify techniques.

Actionability. In contrast with Elmqvist's model, we make the distinction between interacting with visualizations (*interactivity*) and interacting with physical referents (*actionability*). A similar distinction is made in Thomas et al.'s [88] characterization of situated visualizations, which distinguishes between three levels of interaction: altering the visualization pipeline, directly altering the visualization's physical presentation, and directly altering the visualization's physical referent. Actionability also relates to the notion of *pragmatic superpower* in Willett et al.'s [95] visualization-as-superpower framework.

Representation fidelity and representation / visualization medium. These three dimensions were initially inspired by a design

space of worlds-in-miniature introduced by Danyluk and al. [27], and in particular by its *abstraction* dimension. The ProxSituating framework [77], despite having been developed independently from ours, also includes the related dimension of *representation form* and in particular *tangibility*, which maps to our *representation / visualization medium* dimensions. The dimensions in our category *visualization* were also broadly inspired by Ens and al.'s [31] paper on Grand Challenges in Immersive Analytics, in which they discuss data representation in immersive situated visualizations.

Surroundings inclusion. This dimension maps to the concept of *environment* in the ProxSituating visualization framework [77].

Visualization richness. This dimension maps to the *information richness* dimension in Willett et al.'s [95] visualization-as-superpower framework, and the notion of *information bandwidth* discussed in Ens and al.'s [31] Grand Challenges paper. It is also close to the *semantic relevance* dimension in the taxonomy of annotation in outdoor augmented reality [97], and the concept of *value* in Lin et al.'s [62] taxonomy of labelling techniques.

Representation / visualization refresh rate. These dimensions relate to the notion of *temporal indirection* in the embedded representation [96] and the ProxSituating [77] frameworks, and to the notion of *temporal relevance* in the visualization-as-superpower framework [95].

Visualization integration. This dimension relates to the distinction between situated and embedded visualization in the embedded representation framework [96]: embedded visualizations are higher on the integration dimension than non-embedded situated visualizations.

Dimensions values. The choice to use three values (low, medium, high) in most of our dimensions was inspired by the design space of anthropographics by Morais and colleagues [69].

5.3 Descriptive and Generative Power

Our conceptual framework (terminology, typology of techniques, and design space) can help describe and compare situated visualization systems in terms of their support for non-visible referents – including the extent of this support, and how this support is achieved. We successfully applied our design space to characterize the few existing situated visualization systems we know of that have some support for non-visible referents (see table in supplementary material). The General and Representation categories of our design space (see Table 1) can also be used to characterize the mixed-reality systems we surveyed in section 3, which support non-visible objects of interest but generally include no data visualization. We however envision that our conceptual framework will be most useful in the hands of designers of new situated visualizations, to help them choose what types of non-visible referents they want to support and how. This ability to guide the design of new situated visualizations confers to our framework some generative power, which we illustrated through use cases in subsection 5.1. Our framework is probably less suited for designing novel techniques for handling hidden referents, although it provides an overview of existing techniques that can help identify gaps. Future work could involve a validation of the framework's descriptive power, for example by asking designers to prototype situated visualizations with and without our framework.

5.4 Techniques not Discussed

In Section 3, we defined a scope for our survey and our design space, which helped us keep our framework simple but inevitably led us to exclude potentially interesting techniques. We discuss them here, by revisiting our three main criteria for inclusion we listed in Section 3.

1. Visual only. We only considered the use of visual techniques to handle non-visible objects of interest. Although we are not aware of non-visual techniques in the context of situated visualisation, Prouzeau et al. [73] proposed to use vibrotactile feedback to manage occlusion and overplotting in immersive 3D scatterplots. Their technique encodes occluded information (e.g., the local density of points) with the strength of vibration of a VR controller. Their study suggested that the use of haptic feedback is beneficial compared to a classic visual technique (in that case, a cutting plane) in cases where the scene is already visually complex. Nonetheless, the paper also showed that the haptic channel

has a limited bandwidth compared to the visual channel. The use of non-visual channels, such as haptic and audio, deserve to be investigated to convey limited amounts of information about non-visible referents in situated visualizations, especially when the scenes are visually cluttered. Future work in this area would need to investigate how to link this non-visual information to the physical referents.

2. Object-aware. We assumed the existence of an object of interest that the system knows about, and we therefore excluded all object-agnostic approaches. Nevertheless, object-agnostic approaches could still be helpful in dealing with non-visible referents in situated visualizations. For example, an x-ray vision technique that is object-agnostic and renders anything that is occluded could still help users see potential physical referents or their physical surroundings, even though it may not be able to link visualizations to their physical referents. Similarly, a situated visualization system could display feeds of video cameras placed in hidden areas where physical referents are likely to be situated, such as areas under surveillance [7], places where maintenance tasks are being carried out [98], or the parts of a body that are undergoing surgery [50]. Among object-agnostic techniques, some techniques exist that widen the user's field of view, sometimes providing a 360° field of view [5]. Some of these techniques combine wide-angle views with regular views [35, 68]. All such techniques can help users see physical referents, although again, they cannot help them find referents or relate them with their corresponding visualizations.

3. No movement required. We only considered systems allowing users to see information about non-visible referents from where they stand, and therefore excluded systems requiring them to move. Such systems could nevertheless facilitate the use of situated visualizations where physical referents are not always in view. For example, tour planning techniques can show users how to get to a specific location [4, 22, 24, 30], and could potentially guide them to a non-visible physical referent; once the user is there, the referent stops being non-visible. The out-of-view pointing techniques we discussed in subsection 3.7 sometimes serve the same function, but we included them nonetheless because the arrow does give immediate information about the physical referent (its location), albeit very limited information. Attention guiding techniques are similar to tour planning, except they do not involve locomotion – only posture changes [83]. Another approach is to directly move the user to the non-visible object of interest. This is possible in virtual worlds using teleportation techniques (e.g., [48]), but such motions are virtual. In situated visualization systems, actually teleporting users is not possible, but autonomous vehicles could in principle be used to physically bring users closer to physical referents of interest.

6 CONCLUSION

We proposed a conceptual framework to help researchers and user interface designers think about how to handle non-visible physical referents in situated visualizations. This framework consists of 1) a terminology, 2) a typology of techniques, 3) a design space. Our framework reflects the diversity of mixed-reality techniques that can be used to address a variety of situations in which physical referents are not visible. It also highlights areas that remains unexplored and opportunities for future research. However, our framework is far from capturing all aspects relevant to the design of situated visualizations that can handle hidden referents, nor does it cover all possible designs. There are many considerations that go into the design of such interactive systems, and a single design space cannot possibly capture them all. However, we believe our conceptual framework is a useful starting point for discussion and reflection. In particular, we hope that our terms and concepts will make it easier for researchers and practitioners to think and discuss about the issue of non-visible physical referents in situated visualizations, and that they will be extended to cover broader cases. On a general level, we hope that this work will open interesting discussions and inspire more research about this neglected topic, and more broadly, about the question of how to take into account the complexity of physical-world constraints when designing situated visualizations.

7 ACKNOWLEDGEMENTS

This work was partly supported by the Agence Nationale de la Recherche (ANR), grant number ANR-19-CE33-0012.

REFERENCES

- [1] A. Ajanki, M. Billinghurst, H. Gamper, T. Järvenpää, M. Kandemir, S. Kaski, M. Koskela, M. Kurimo, J. Laaksonen, K. Puolamäki, et al. An augmented reality interface to contextual information. *Virtual reality*, 15:161–173, 2011. 5
- [2] G. Alce, M. Roszko, H. Edlund, S. Olsson, J. Svedberg, and M. Walbergård. AR as a User Interface for The Internet of Things: Comparing Three Interaction Models. In *ISMAR*, 2017. doi: 10.1109/ISMAR-Adjunct.2017.37 3
- [3] C. Andujar, F. Argelaguet, and R. Trueba. Hand-Based Disocclusion for the World-in-Miniature Metaphor. *Presence: Teleoperators and Virtual Environments*, 19(6), 2010. doi: 10.1162/pres_a_00018 3
- [4] C. Andujar, P. Vázquez, and M. Fairén. Way-Finder: guided tours through complex walkthrough models. *Computer Graphics Forum*, 23(3), 2004. doi: 10.1111/j.1467-8659.2004.00781.x 9
- [5] J. Ardouin, A. Lécuyer, M. Marchal, C. Riant, and E. Marchand. FlyVIZ: a novel display device to provide humans with 360 vision by coupling catadioptric camera with hmd. In *Proc. VRST*, 2012. 9
- [6] P. Arpaia, E. De Benedetto, C. A. Dodaro, L. Duraccio, and G. Servillo. Metrology-Based Design of a Wearable Augmented Reality System for Monitoring Patient’s Vitals in Real Time. *IEEE Sensors Journal*, 21(9), 2021. doi: 10.1109/JSEN.2021.3059636 1
- [7] B. Avery, W. Piekarski, and B. H. Thomas. Visualizing Occluded Physical Objects in Unfamiliar Outdoor Augmented Reality Environments. In *Proc. of ISMAR*, 2007. doi: 10.1109/ISMAR.2007.4538869 9
- [8] R. Bane and T. Hollerer. Interactive tools for virtual x-ray vision in mobile augmented reality. In *ISMAR*, 2004. doi: 10.1109/ISMAR.2004.36 4
- [9] P. Baudisch and R. Rosenholtz. Halo: a Technique for Visualizing Off-Screen Locations. *NEW HORIZONS*, (5):8, 2003. doi: 10.1145/642611.642695 5
- [10] B. Bell, S. Feiner, and T. Höllerer. View management for virtual and augmented reality. In *Proc. UIST*, 2001. doi: 10.1145/502348.502363 3, 5
- [11] B. Bell, T. Höllerer, and S. Feiner. An annotated situation-awareness aid for augmented reality. In *UIST*, 2002. doi: 10.1145/571985.572017 3
- [12] C. Bichlmeier, S. M. Heining, M. Feuerstein, and N. Navab. The Virtual Mirror: A New Interaction Paradigm for Augmented Reality Environments. *IEEE TMI*, 28(9), 2009. doi: 10.1109/TMI.2009.2018622 5
- [13] C. Bichlmeier, S. M. Heining, M. Rustae, and N. Navab. Virtually Extended Surgical Drilling Device: Virtual Mirror for Navigated Spine Surgery. In *Proc. MICCAI*, 2007. doi: 10.1007/978-3-540-75757-3_53 5
- [14] R. Blanco, Z. Salazar, and T. Isenberg. Exploring carnivorous plant habitats based on images from social media. In *IEEE VIS 2019*, 2019. 8
- [15] A. Bluff and A. Johnston. Don’t Panic: Recursive Interactions in a Miniature Metaworld. In *Proc. of VRCAI*, 2019. 3
- [16] F. Bork, U. Eck, and N. Navab. Birds vs. Fish: Visualizing Out-of-View Objects in Augmented Reality using 3D Minimaps. In *Proc. ISMAR*, 2019. doi: 10.1109/ISMAR-Adjunct.2019.00-28 3
- [17] F. Bork, C. Schnelzer, U. Eck, and N. Navab. Towards Efficient Visual Guidance in Limited Field-of-View Head-Mounted Displays. *IEEE TVCG*, 24(11), 2018. doi: 10.1109/TVCG.2018.2868584 5
- [18] J. A. Boudell and B. A. Middleton. Plot locator: An app for locating plots in the field. *Applications in plant sciences*, 7(12):e11311, 2019. doi: 10.1002/aps3.11311 8
- [19] N. Bressa, H. Korsgaard, A. Tabard, S. Houben, and J. Vermeulen. What’s the situation with situated visualization? a survey and perspectives on situatedness. *IEEE TVCG*, 28(1):107–117, 2021. 1
- [20] G. Caggianese, V. Colonnese, and L. Gallo. Situated Visualization in Augmented Reality: Exploring Information Seeking Strategies. In *Proc. of SITIS*, 2019. doi: 10.1109/SITIS.2019.00069 4
- [21] M. Card. *Readings in information visualization: using vision to think*. Morgan Kaufmann, 1999. doi: 10.5555/300679 7
- [22] I. Chaturvedi, F. H. Bijarbooneh, T. Braud, and P. Hui. Peripheral vision: a new killer app for smart glasses. In *Proc. of IUI*, 2019. doi: 10.1145/3301275.3302263 9
- [23] L. Chittaro and S. Burigat. 3D location-pointing as a navigation aid in Virtual Environments. In *Proc. of AVI*, 2004. doi: 10.1145/989863.989910 5
- [24] L. Chittaro, R. Ranon, and L. Ieronutti. Guiding visitors of Web3D worlds through automatically generated tours. In *Proc. WEB 3D*, 2003. doi: 10.1145/636593.636598 9
- [25] D. Coffey, N. Malbraaten, T. Le, I. Borazjani, F. Sotiropoulos, and D. F. Keefe. Slice WIM: a multi-surface, multi-touch interface for overview+detail exploration of volume datasets in virtual reality. In *Proc. of I3D*, 2011. doi: 10.1145/1944745.1944777 3
- [26] M. Cordeil, A. Cunningham, T. Dwyer, B. H. Thomas, and K. Marriott. Imaxes: Immersive axes as embodied affordances for interactive multi-variate data visualisation. In *Proc. of UIST*, 2017. doi: 10.1145/3126594.3126613 6, 8
- [27] K. Danyluk, B. Ens, B. Jenny, and W. Willett. A design space exploration of worlds in miniature. In *Proc. CHI*, 2021. doi: 10.1145/3411764.3445098 3, 9
- [28] T. Drey. Data Visualisation using Augmented Reality with a Focus set on Head Mounted Displays and Collaborative Tasks. 2018. 3
- [29] N. Elmqvist and P. Tsigas. A Taxonomy of 3D Occlusion Management for Visualization. *IEEE TVCG*, 14(5), 2008. doi: 10.1109/TVCG.2008.59 1, 2, 3, 4, 5, 8
- [30] T. Elvins, D. Nadeau, R. Schul, and D. Kirsh. Worldlets: 3D Thumbnails for Wayfinding in Virtual Environments. *Journal of the ACM*, pp. 21–30, Aug. 1997. doi: 10.1145/263407.263504 9
- [31] B. Ens, B. Bach, M. Cordeil, U. Engelke, M. Serrano, W. Willett, A. Prouzeau, C. Anthes, W. Büschel, C. Dunne, et al. Grand challenges in immersive analytics. In *Proc. of CHI*, pp. 1–17, 2021. 1, 9
- [32] B. Ens, S. Goodwin, A. Prouzeau, F. Anderson, F. Wang, S. Gratzl, Z. Lucarelli, B. Moyle, J. Smiley, and T. Dwyer. Uplift: A Tangible and Immersive Tabletop System for Casual Collaborative Visual Analytics. *IEEE TVCG*, 2020. doi: 10.1109/TVCG.2020.3030334 3, 4
- [33] O. Erat, W. A. Isop, D. Kalkofen, and D. Schmalstieg. Drone-Augmented Human Vision: Exocentric Control for Drones Exploring Hidden Areas. *IEEE TVCG*, 24(4), 2018. doi: 10.1109/TVCG.2018.2794058 4
- [34] M. T. Eren and S. Balcisoy. Evaluation of X-ray visualization techniques for vertical depth judgments in underground exploration. *The Visual Computer*, 34(3), 2018. doi: 10.1007/s00371-016-1346-5 4
- [35] K. Fan, J. Huber, S. Nanayakkara, and M. Inami. Spidervision: Extending the human field of view for augmented awareness. In *Proc. of AH*, 2014. doi: 10.1145/2582051.2582100 9
- [36] A. Fedosov and S. Misslinger. Location based experience design for mobile augmented reality. In *Proc. EICS*, 2014. 5
- [37] L. F.Macedo. Occlusion Handling in Augmented Reality: Past, Present and Future. *IEEE TVCG*, 2021. 1, 2
- [38] D. C. Foyle, A. D. Andre, and B. L. Hooley. Situation Awareness in an Augmented Reality Cockpit: Design, Viewpoints and Cognitive Glue. In *Proc of HCI Int.*, 2005. 1
- [39] J. Frey, R. Gervais, S. Fleck, F. Lotte, and M. Hachet. Teegi: tangible EEG interface. In *Proc. of UIST*, 2014. doi: 10.1145/2642918.2647368 4
- [40] H. Fuchs, A. State, and J.-C. Bazin. Immersive 3D Telepresence. *Computer*, 47(7), 2014. doi: 10.1109/MC.2014.185 4
- [41] S. Ghosh, G. Bailly, R. Despouys, E. Lecolinet, and R. Sharrock. Super-Vision: Spatial Control of Connected Objects in a Smart Home. In *Proc of CHI*, 2015. 4, 6
- [42] R. Grasset, T. Langlotz, D. Kalkofen, M. Tatzgern, and D. Schmalstieg. Image-driven view management for augmented reality browsers. In *2012 IEEE ISMAR*, pp. 177–186. IEEE, 2012. 5
- [43] U. Gruenefeld, A. E. Ali, S. Boll, and W. Heuten. Beyond Halo and Wedge: visualizing out-of-view objects on head-mounted virtual and augmented reality devices. In *Proc. MobileHCI*. ACM, Sept. 2018. 5
- [44] U. Gruenefeld, T. C. Stratmann, A. E. Ali, S. Boll, and W. Heuten. RadialLight: exploring radial peripheral LEDs for directional cues in head-mounted displays. In *MobileHCI*, 2018. 5
- [45] U. Gruenefeld, T. C. Stratmann, L. Prädell, and W. Heuten. MonocularAR: a radial light display to point towards out-of-view objects on augmented reality devices. In *MobileHCI*, 2018. doi: 10.1145/3236112.3236115 5
- [46] S. Gustafson, P. Baudisch, C. Gutwin, and P. Irani. Wedge: clutter-free visualization of off-screen locations. In *Proc. CHI*, 2008. 5
- [47] K. Hinckley, R. Pausch, J. C. Goble, and N. F. Kassell. Passive real-world interface props for neurosurgical visualization. In *Proc. of CHI*, 1994. doi: 10.1145/191666.191821 4
- [48] K. Hirose, T. Ogawa, K. Kiyokawa, and H. Takemura. Interactive Reconfiguration Techniques of Reference Frame Hierarchy in the Multi-viewport Interface. In *Proc. of 3DUI*, 2006. doi: 10.1109/VR.2006.89 9
- [49] D. Holstius, J. Kembel, A. Hurst, P.-H. Wan, and J. Forlizzi. Infotropism:

- living and robotic plants as interactive displays. In *Proc. DIS*, 2004. doi: 10.1145/1013115.1013145 1
- [50] A. Ilie, K.-L. Low, G. Welch, A. Lastra, H. Fuchs, and B. Cairns. Combining Head-Mounted and Projector-Based Displays for Surgical Training. *Presence: Teleoperators and Virtual Environments*, 13(2), 2004. 9
- [51] Y. Ishiguro and J. Rekimoto. Peripheral vision annotation: noninterference information presentation method for mobile augmented reality. In *Proc. of AH*. Tokyo, Japan, 2011. doi: 10.1145/1959826.1959834 5
- [52] R. Jacob, A. Girouard, L. Hirshfield, M. Horn, O. Shaer, E. Solovey, and J. Zigelbaum. Reality-based interaction: A framework for post-WIMP interfaces. In *Proc. of CHI*, 2008. doi: 10.1145/1357054.1357089 6, 8
- [53] Y. Jansen, P. Dragicevic, P. Isenberg, J. Alexander, A. Karnik, J. Kildal, S. Subramanian, and K. Hornbæk. Opportunities and challenges for data physicalization. In *Proc. of CHI*, 2015. 1, 7
- [54] H. Jo, S. Hwang, H. Park, and J.-h. Ryu. Aroundplot: Focus+context interface for off-screen objects in 3D environments. *Computers & Graphics*, 35(4), 2011. doi: 10.1016/j.cag.2011.04.005 5
- [55] J. D. Johansen and S. E. Larsen. *Signs in use: an introduction to semiotics*. Psychology Press, 2002. 6
- [56] D. Kalkofen, M. Tatzgern, and D. Schmalstieg. Explosion Diagrams in Augmented Reality. In *IEEEVR*, 2009. doi: 10.1109/VR.2009.4811001 4
- [57] T. Kantonen, C. Woodward, and N. Katz. Mixed reality in virtual world teleconferencing. In *Proc of IEEE VR*, 2010. 4
- [58] A. Kristoffersson, S. Coradeschi, and A. Loutfi. A Review of Mobile Robotic Telepresence. *Advances in Human-Computer Interaction*, 2013. doi: 10.1155/2013/902316 4
- [59] M. Kritzler, M. Funk, F. Michahelles, and W. Rohde. The virtual twin: controlling smart factories using a spatially-correct augmented reality representation. In *Proc. of IoT*, 2017. doi: 10.1145/3131542.3140274 4
- [60] H. Laaki, Y. Miche, and K. Tammi. Prototyping a Digital Twin for Real Time Remote Control Over Mobile Networks: Application of Remote Surgery. *IEEE Access*, 7, 2019. doi: 10.1109/ACCESS.2019.2897018 4
- [61] T. Lin, Y. Yang, J. Beyer, and H. Pfister. Labeling out-of-view objects in immersive analytics to support situated visual searching. *CoRR*, 2021. doi: 10.1109/TVCG.2021.3133511 2
- [62] T. Lin, Y. Yang, J. Beyer, and H. Pfister. Labeling out-of-view objects in immersive analytics to support situated visual searching. *IEEE TVCG*, 2021. doi: 10.1109/TVCG.2021.3133511 3, 5, 9
- [63] M. A. Livingston, A. Dey, C. Sandor, and B. H. Thomas. *Pursuit of "X-Ray Vision" for Augmented Reality*, chap. Human Factors in Augmented Reality Environments. 2013. doi: 10.1007/978-1-4614-4205-9_4 4
- [64] P. Lubos, G. Bruder, O. Ariza, and F. Steinicke. Ambiculus: LED-based low-resolution peripheral display extension for immersive head-mounted displays. In *Proc. VRIC*, 2016. doi: 10.1145/2927929.2927939 5
- [65] A. Maimone, X. Yang, N. Dierk, A. State, M. Dou, and H. Fuchs. General-purpose telepresence with head-worn optical see-through displays and projector-based lighting. In *Proc. of IEEE VR*, 2013. doi: 10.1109/VR.2013.6549352 4
- [66] J. Mercier-Ganady, F. Lotte, E. Loup-Escande, M. Marchal, and A. Lécuyer. The Mind-Mirror: See your brain in action in your head using EEG and augmented reality. In *IEEE VR*, 2014. doi: 10.1109/VR.2014.6802047 5
- [67] P. Milgram and F. Kishino. A Taxonomy of Mixed Reality Visual Displays. *IEICE Trans. Information Systems*, 1994. 7
- [68] F. Mizuno, T. Hayasaka, and T. Yamaguchi. Virtual chameleon- a system to provide different views to both eyes. In *Proc. of IFMB*, 2009. 9
- [69] L. Morais, Y. Jansen, N. Andrade, and P. Dragicevic. Showing data about people: A design space of anthropographics. *IEEE TVCG*, 28(3):1661–1679, 2020. doi: 10.1109/TVCG.2020.3023013 9
- [70] S. Mori, S. Ikeda, and H. Saito. A survey of diminished reality: Techniques for visually concealing, eliminating, and seeing through real objects. *IPSJ TCVA*, 9(1), 2017. doi: 10.1186/s41074-017-0028-1 4
- [71] J. W. Nam, K. McCullough, J. Tveite, M. M. Espinosa, C. H. Perry, B. T. Wilson, and D. F. Keefe. Worlds-in-Wedges: Combining Worlds-in-Miniature and Portals to Support Comparative Immersive Visualization of Forestry Data. In *IEEE (VR)*, 2019. doi: 10.1109/VR.2019.8797871 3
- [72] R. Pausch, T. Burnette, D. Brockway, and M. E. Weiblen. Navigation and locomotion in virtual worlds via flight into hand-held miniatures. In *Proc. of SIGGRAPH*, 1995. doi: 10.1145/218380.218495 3
- [73] A. Prouzeau, M. Cordeil, C. Robin, B. Ens, B. H. Thomas, and T. Dwyer. Sceptics and highlight-planes: Immersive interaction techniques for finding occluded features in 3d scatterplots. In *Proc. of CHI*, 2019. 9
- [74] A. Prouzeau, Y. Wang, B. Ens, W. Willett, and T. Dwyer. Corsican Twin: Authoring In Situ Augmented Reality Visualisations in Virtual Reality. In *Proc. of AVI*, 2020. doi: 10.1145/3399715.3399743 1, 8
- [75] C. Qiu, S. Zhou, Z. Liu, Q. Gao, and J. Tan. Digital assembly technology based on augmented reality and digital twins: a review. *Virtual Reality & Intelligent Hardware*, 1(6), 2019. doi: 10.1016/j.vrih.2019.10.002 4
- [76] D. Ren, T. Goldschwendt, Y. Chang, and T. Höllerer. Evaluating wide-field-of-view augmented reality with mixed reality simulation. In *Proc. IEEE VR*, 2016. doi: 10.1109/VR.2016.7504692 5
- [77] K. Satriadi, A. Cunningham, R. Smith, T. Dwyer, A. Drogemuller, and B. Thomas. Proxysituated visualization: An extended model of situated visualization using proxies for physical referents. In *CHI*, 2023. 2, 8, 9
- [78] K. A. Satriadi, J. Smiley, B. Ens, M. Cordeil, T. Czauderna, B. Lee, Y. Yang, T. Dwyer, and B. Jenny. Tangible globes for data visualisation in augmented reality. In *Proc. of CHI*, 2022. 4
- [79] D. W. Seo, H. Kim, J. S. Kim, and J. Y. Lee. Hybrid reality-based user experience and evaluation of a context-aware smart home. *Computers in Industry*, 76, 2016. doi: 10.1016/j.compind.2015.11.003 3
- [80] W. R. Sherman and A. B. Craig. Understanding Virtual Reality—Interface, Application, and Design. *Presence: Teleoperators & Virtual Environments*, 12(4), 2003. 4
- [81] T. Siu and V. Herskovic. SideARs: improving awareness of off-screen elements in mobile augmented reality. In *Proc. of CHI*. Chile, 2013. 5
- [82] M. Sousa, D. Mendes, S. Paulo, N. Matela, J. Jorge, and D. S. Lopes. VRRRoom: Virtual Reality for Radiologists in the Reading Room. In *Proc. of CHI*, 2017. doi: 10.1145/3025453.3025566 4
- [83] M. Sukan, C. Elvezio, O. Oda, S. Feiner, and B. Tversky. ParaFrustum: visualization techniques for guiding a user to a constrained set of viewing positions and orientations. In *Proc. UIST*, 2014. 9
- [84] S. Tachi, K. Komoriya, K. Sawada, T. Nishiyama, T. Itoko, M. Kobayashi, and K. Inoue. Telexistence cockpit for humanoid robot control. *Advanced Robotics*, 17(3), 2003. doi: 10.1163/156855303764018468 4
- [85] S. Tachi, K. Watanabe, K. Takeshita, K. Minamizawa, T. Yoshida, and K. Sato. Mutual telexistence surrogate system: ESAR4 - telexistence in real environments using autostereoscopic immersive display -. In *IEEE/RSJ*, pp. 157–162, 2011. doi: 10.1109/IROS.2011.6094543 4
- [86] F. Tao, M. Zhang, and A. Y. C. Nee. Chapter 11 - Digital Twin and Virtual Reality and Augmented Reality/Mixed Reality. In *Digital Twin Driven Smart Manufacturing*. 2019. 4
- [87] M. Tatzgern, R. Grasset, E. Veas, D. Kalkofen, H. Seichter, and D. Schmalstieg. Exploring real world points of interest: Design and evaluation of object-centric exploration techniques for augmented reality. *PMC*, 18, 2015. doi: 10.1016/j.pmcj.2014.08.010 3
- [88] B. H. Thomas, G. F. Welch, P. Dragicevic, N. Elmqvist, P. Irani, Y. Jansen, D. Schmalstieg, A. Tabard, N. A. ElSayed, R. T. Smith, et al. Situated analytics. In *Immersive analytics*, pp. 185–220. 2018. 1, 8
- [89] B. Thoravi Kumaravel, F. Anderson, G. Fitzmaurice, B. Hartmann, and T. Grossman. Loki:Facilitating Remote Instruction of Physical Tasks Using Bi-Directional Mixed-Reality Telepresence. In *Proc. UIST*, 2019. doi: 10.1145/3332165.3347872 4
- [90] S. Truman and S. von Mammen. An Integrated Design of World-in-Miniature Navigation in Virtual Reality. In *Proc. of FDG*, number 69. 2020. doi: 10.1145/3402942.3402994 3
- [91] J. Vermeulen, F. Kawsar, A. L. Simeone, G. Kortuem, K. Luyten, and K. Coninx. Informing the design of situated glyphs for a care facility. In *Proc. of VL/HCC*, 2012. doi: 10.1007/978-3-642-21726-5_5 8
- [92] X. Wang, P. E. D. Love, M. J. Kim, and W. Wang. Mutual awareness in collaborative design: An Augmented Reality integrated telepresence system. *Comp. in Ind.*, 2014. doi: 10.1016/j.compind.2013.11.012 4
- [93] S. White and S. Feiner. SiteLens: situated visualization techniques for urban site visits. In *Proc. of CHI*, 2009. doi: 10.1145/1518701.1518871 1
- [94] S. White, S. Feiner, and J. Kopylec. Virtual Vouchers: Prototyping a Mobile Augmented Reality User Interface for Botanical Species Identification. In *Proc. of 3DUI*, 2006. doi: 10.1109/VR.2006.145 1, 8
- [95] W. Willett, B. A. Aseniero, S. Carpendale, P. Dragicevic, Y. Jansen, L. Oehlberg, and P. Isenberg. Superpowers as inspiration for visualization. *IEEE TVCG*, 2021, 2021. doi: 10.1109/TVCG.2021.3114844 4, 8, 9
- [96] W. Willett, Y. Jansen, and P. Dragicevic. Embedded Data Representations. *IEEE TVCG*, 23(1), 2017. doi: 10.1109/TVCG.2016.2598608 1, 2, 3, 9
- [97] J. Wither, S. DiVerdi, and T. Höllerer. Annotation in outdoor augmented reality. *C&G*, 33(6), 2009. 9
- [98] M.-L. Wu and V. Popescu. Efficient VR and AR Navigation Through Multiperspective Occlusion Management. *IEEE TVCG*, 24(12), 2018. doi: 10.1109/TVCG.2017.2778249 9