Corsican Twin: Authoring In Situ Augmented Reality Visualisations in Virtual Reality

Arnaud Prouzeau, Yuchen Wang, Barrett Ens, Wesley Willett, Tim Dwyer

► To cite this version:

Arnaud Prouzeau, Yuchen Wang, Barrett Ens, Wesley Willett, Tim Dwyer. Corsican Twin: Authoring In Situ Augmented Reality Visualisations in Virtual Reality. International Conference on Advanced Visual Interfaces, Sep 2020, Island of Ischia, Italy. 10.1145/3399715.3399743 . hal-02614521

HAL Id: hal-02614521

<https://hal.science/hal-02614521>

Submitted on 10 Jun 2020

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.

Corsican Twin: Authoring In Situ Augmented Reality Visualisations in Virtual Reality

Arnaud Prouzeau¹, Yuchen Wang¹, Barrett Ens¹, Wesley Willett², Tim Dwyer¹

¹ Monash University, Australia, ² University of Calgary, Canada

{arnaud.prouzeau,barrett.ens,tim.dwyer}@monash.edu,ywan715@student.monash.edu,wesley.willett@ucalgary.ca

Figure 1: The Corsican Twin is an immersive authoring tool for authoring AR in situ visualisations. Middle: Embedded and situated visualisations are authored in a digital twin in VR. Right: The visualisations are displayed in their locations in AR.

ABSTRACT

We introduce Corsican Twin, a tool to author augmented reality data visualisations in virtual reality using digital twins. The system provides users with the necessary contextual information needed to design embedded and situated data visualisations in a safe and convenient remote setting. The system was co-designed with and for people with little or no programming experience. Using the system, we illustrate three potential use cases for situated visualizations in the context of building maintenance, including: (1) on-site equipment debugging and diagnosis; (2) remote incident playback; and (3) operations simulations for future buildings. From feedback gathered during formative evaluations of our prototype tool with domain experts, we discuss implications, opportunities, and challenges for future in situ visualisation design tools.

CCS CONCEPTS

• **Human-centered computing** → **Visualization techniques**; *Interaction techniques*.

KEYWORDS

Immersive Analytics, Authoring Tool, In situ Visualisation

ACM Reference Format:

Arnaud Prouzeau¹, Yuchen Wang¹, Barrett Ens¹, Wesley Willett², Tim Dwyer¹. 2020. Corsican Twin: Authoring In Situ Augmented Reality Visualisations in Virtual Reality. In *International Conference on Advanced Visual Interfaces (AVI '20)*, September 28–October 2, 2020, Salerno, Italy. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3399715.3399743>

1 INTRODUCTION

The rapid development of Augmented Reality (AR) technologies is providing new opportunities to integrate data visualisations directly into physical environments, where they can help viewers solve complex situated problems and perform data-related tasks. Displaying applications and data *in situ* [14, 15] provides the necessary context for users to understand problems and take action based on the data. Moreover, integrating data representations tightly with their physical source (or referent), known as *situated analytics* [14, 33], can reinforce this spatial awareness. Situated analytics shows promise in domains like: building maintenance where there are clear benefits of overlaying sensor data on facilities and equipment [20]; construction where there is a need for simulations of data alongside future infrastructure [18]; or in educational settings where data can demonstrate invisible phenomena directly in real spaces [12].

Situated visualisations can be created manually by programming using AR libraries (Vuforia, AR Core, etc.). However, this requires programming skill and makes rapid production and iteration of designs difficult. Recently, several toolkits to author such visualisations have been introduced. Recent toolkits such as IATK provide prepackaged visualisation primitives and a grammar of graphics approach to authoring in the Unity desktop environment [11]. While this is easy to use and efficient, performing this kind of design work at the desktop can make it difficult to accurately place the visualisations. Moreover, the transition from small-scale design on computer screen to real-life scale visualisation in AR can lead to unexpected issues (visualisation too small, real life object hiding the

Figure 2: Timeline of the design phase of the system with the co-design activities: Int: Interview, Obs: Observation, FS: Feedback Session, US: Usability Study

visualisation, etc.). On the other hand, recent tools like DXR [31] or MARViT [7] allow users to author visualisations directly in AR. While this solves the issues of position and scale of the visualisation, the designer has to be on site, which can be problematic in some contexts (e.g. plant rooms accessible only to accredited technicians, very large environments). The limited interaction vocabulary of current AR headsets can also make complex interaction difficult.

In this paper, we present the Corsican Twin, a Virtual Reality (VR) tool to design AR visualisation for large and complex environments (Figure 1). This name refers to Alexandre Dumas’s novel, the Corsican Brothers [13], in which the pain felt by each of two twin brothers is experienced identically by his twin. Similarly, our system uses digital twins of physical environments to support authoring of visualisations in VR, which can then be experienced in the original environments using AR. By using a 3D model of the environment (3D CAD Model, or Photogrammetry scans), we allow the designers to safely and easily create real-scale situated visualisations for specific target environments without visiting the site. We propose a virtual reality authoring tool that allows designers to author situated and embedded AR visualisations for real-world locations and augment them with active and proxemic interactions.

We used a co-design process (Fig. 2) involving experts from the building domain. Expert feedback led to a set of Design Goals (§3.2). The system stemming from these guidelines (§4) allows users to create both (1) situated visualisations, which are located close to their physical referent, and (2) embedded visualisations, which have a one to one correspondence with a physical referent [38]. To our knowledge, our prototype is the first system to provide authors with the substantial spatial awareness needed to integrate data visualisations tightly with their physical referents, and to link these with data from real systems data in an industrial environment. Feedback sessions with domain experts using the system (§5) suggest that Corsican Twin allows end users to easily design efficient in situ visualisations without being on site. Finally (§6), we discuss the potential evolution of the system and the future research directions emerging from these sessions.

2 RELATED WORK

Our research bridges prior work on *in situ* prototyping of data visualisations and immersive authoring tools.

2.1 Prototyping Visualisations in Situ

AR is not the only method to display visualisations in situ. They can also be displayed using a smartphone, a large screen, or be printed on paper. A comprehensive survey of in situ visualisations

is provided by Willet et al. [38], but very few of these expose how they were designed. When the spatial context is relevant to the data visualisation, it is important to design in that context [3]. This concept has generally been called *Bodystorming* [5] and may involve elaborate acting [4, 29]. According to Oulasvirta et al., actually being in the place encourages designers to consider contextual elements, is more inspiring and makes ideas more memorable [30]. Prototyping visualisations in context has been applied to street infographics [9, 10] and hospitals [35].

However, it is not always possible to have site access, especially in the design phase before construction. A site may also be located in a remote location or access-restricted for safety reasons. In lieu of physical access, Eriksson et al. highlight the importance of using physical maps [16]. Hansen and Dalsgaard had end users play out use case scenarios with puppets and blueprints [19]. Korsgaard et al. organised site tours of future buildings using 3D models on a 2D screen [23]. It is our contention that VR and AR offer a better solution to bridge the gap between remote access and the need for in situ visualisation prototyping and authoring.

2.2 AR Authoring

Desktop tools for authoring AR have been developed for some time [8, 25, 37]. However, this approach creates a potential gap between the design and final result. With iaTAR [24], Lee et al. propose that authoring be done directly using the AR device, allowing users to directly experience their application. They defined this concept as WYXIWYG: What You eXperience Is What You Get. A user study showed this method is faster and more preferred by participants compared to a 2D GUI-based authoring tool. Similarly, in Farrago [39], Wozniowski and Warne use a smartphone application to place 3D objects, add textures to them and associate them with a tracking marker. In the Reality Editor, Huen et al. use a smartphone interface to author data connections between smart objects [21]. Vera et al. associate text or images to GPS positions to facilitate the creation of outdoor AR applications [34].

In contrast, our work explores authoring AR interfaces in the comfort, safety, and convenience of a digital twin in VR. Only very recently has there been some initial exploration of VR authoring for AR. For example, CAVE-AR [6] is a CAVE 2 VR system which allows AR design by placing virtual objects in a 3D scene. Similarly, Microsoft Layout [28] allows developers to provide the layout of a room in VR and subsequently augment the same space in AR. However, these systems involve simple elements without the precise contextual information required for in situ data visualisation. There has been some recent research on authoring or prototyping of in situ visualisation in AR. For example DXR [31] lets users author their visualisation from the desktop and then modify the visualisation position and encoding in AR. MARViT allows authoring of tablet-based Augmented Reality visualisations [7], but mostly for the purposes of static data storytelling. Our goal is high-fidelity digital-twin VR authoring of AR data visualisation by and for industry domain users.

3 AUTHORIZING IN SITU AR VIS IN VR

Digital twins are virtual replicas of physical assets, facilities, or environments that leverage the real-time data capabilities of the

Figure 3: Corsican Twin workflow. a) An author creates data visualisations and interactions in a digital replica using VR. b) The author places a virtual marker in the scene. c) A physical tracking marker is put at the same location in the real environment. d) The real environment is populated with in situ data visualisations.

Internet of Things. They are becoming increasingly adopted by businesses to save costs by testing new procedures before deploying them or preventing equipment issues that would otherwise require expensive maintenance.

With the Corsican Twin, we propose using digital twins, viewed through VR, as proxy environments for authoring in situ data visualisations, as illustrated in Figure 3. After recognising a need for in situ data visualisations, an *author* creates them using a simple VR interface in a digital twin. The author then places a virtual and a physical tracking marker (this could be a standard QR code, or a custom marker designed by the author) at the same location in both scenes. The visualisations can then be viewed – just as they appeared in the digital twin – by an *end user* wearing an AR display.

3.1 Building Management Domain Use Cases

While the ideas may be more broadly applicable, we focus our co-design efforts on the building management domain. Our design method consist of 2 iterations (See Figure 2), each informed by interviews and observations with experts from this domain (building managers and maintenance technicians). The first iteration was evaluated in a feedback session with Building Management System (BMS) engineers and a usability study with non experts users (students from our department). The second iteration was evaluated in three feedback sessions with maintenance technicians, building managers, and BMS engineers, they are presented in §5. This process helped us derived the following scenarios and design goals.

On site maintenance – Alice is a maintenance engineer who looks after a university’s campus facilities. She receives a call about a room occupants found too hot. She takes her AR headset and heads

directly to the room. At the door, she scans the QR code of the room and a simplified version of the BIM model for the HVAC system is shown in AR. This visualisation helps her to understand quickly where to find the Air Handling Unit (AHU - which is behind one wall), the fans, the vents, etc. The vent colours are depicted as a function of the air temperature propelled into the room, and the colour of the AHU encodes the speed of the fan. When she enters the room, a large situated line graph shows the evolution of the room temperature over the last 24 hours. She can see that the temperature has continued to increase even after the target temperature was reached. By looking at the AHU, she can see that it is running as the fan speed seems to be over 50%. Finally, she sees that the air supply coming from the vents is hot. As she walks closer to the AHU, two new line graphs appear showing the opening and closing of the heating and cooling coils. She sees that the cooling coil is closed and the heating one opened which is probably due to a Building Management System software issue.

Incident playback – Bob is a chiller technician working in a company servicing the university’s facilities. He is called on Monday morning because a tube failure alarm has been raised during the night. When arriving in the plant, Bob puts on his AR headset and scans the QR code. Embedded and situated visualisations appear showing the state of the chiller and the water pressure and temperature in the pipe. the chiller seems to be working correctly now, so Bob makes a calendar chooser appear by saying “calendar” and selects the time of the alarm he received that morning. The visualisations update and show the state of the chiller at that moment. He notices that the pressure in the pipe bringing the water to the chiller was higher than usual, and had been building up for the previous 24 hours. Bob knows that this is a common cause of tube failure and was probably the cause in this case. He can now continue his investigation to understand why the pressure was at this level.

Remote maintenance – Nathalie is a Building Management System engineer at the university. She receives a call about a meeting room which has been warmer than usual for the past 3 months, since it was refurbished. The room is on a different campus, which is 2 hours away by car. Before sending someone there, Nathalie puts on her VR headset and enters a digital twin of the room, where she can see the BIM model of the HVAC system and a photogrammetry scan of this room, taken after refurbishment. When she enters the room, a bar chart showing the temperature in the last 48 hours appears on a nearby wall. She notices that the temperature has been warmer than it should have been during the day, but does not seem to be affected outside of working hours. By viewing the photogrammetry scan, she sees a large wall display composed of 24 screens. The HVAC system for the room seems to have been designed for a regular meeting room and can not cope with the heat produced by the wall display. Realising that there is no quick fix for this situation, she saves herself the trip to the campus, and makes a note to discuss solutions at the next regular meeting.

3.2 Design Goals

Before discussing our prototype development, we discuss the goals we aim to achieve. These design goals for in situ authoring are drawn from existing literature on authoring tools for AR and

Figure 4: Photogrammetry scan of the facility from Fig. 1.

situated visualisation, and from information learned during our interviews with domain experts.

DG1: Allow Design in Context of the Environment and remotely. From our interviews with BMS engineers we found that while they had to go on site to gather information about the building and discuss with the building managers, the design of building management system is mostly done offsite. The spatial contexts of building management and maintenance can vary widely and it is very important to provide it to design appropriate in situ visualisations (§2). In addition, Spaces like meeting rooms and classrooms are often very accessible, whereas plant facilities containing transformers, chillers, and boilers may limit access to just a small number of personnel who must wear safety equipment. Finally, if visualisations are required for a building in construction, the context simply does not exist yet. While the perfect situation would be to design the visualisation using an AR device in the intended context, this is not always possible.

DG2: Enable experts with no programming experience. It is critical that visualisations are designed in collaboration with building managers and maintenance technicians. One of our interviewees mentioned that, while most of their applications are personalised for the university, some are not adapted to their tasks. Thus, the authoring tool should allow building manager and maintenance workers to test, but also prototype visualisations. In the past, in situ visualisations have required co-design with programmers (§2). An authoring tool that allows domain practitioners to author in situ visualisations themselves is better. The usability study of the first prototype showed the importance of allowing for direct interactions to create and manipulate visualisations and to provide menus to change the settings of the visualisations.

DG3: Provide a Variety Visualisations for Different Needs. When on site, our users use exclusively *situated visualisations* (i.e. visualisations close to their physical referent [38]). Interviews and Observations of maintenance technicians showed they bring a laptop connected to the Building Management System with them and install it close to the plant or assets they are working on. They are, most of the time, interested in two types of information: First, (1) what is the current state of the system, for instance what is the temperature of the air leaving a specific Air Handling Unit and is the fan running or not. They are also interested in (2) the temporal evolution of specific variables, for instance, how was the pressure in the damper in the last two hours before the failure. These help a building manager or maintenance worker to diagnose a failure when on site, to periodically check operation of the assets, and to assess the effect of their actions on the system. AR allows for

Figure 5: Left: CAD + BIM model. Right: CAD + BIM + scaffolds generated automatically by the system.

the use of *situated visualisations* as well as *embedded visualisations*, which are visualisations that are overlaid directly on their physical referents [38]. With *embedded visualisations*, the added spatial context should lead to a better understanding of the visualisation and reduce the risk of errors, but also constrains the design of the visualisation to the location and shape of the physical object.

DG4: Support Interactions to Show/Hide Visualisations. In the current building management system of the university, an initial visualisation shows an overview of the system with only information regarding the current state. To access more information (e.g. temporal evolution of a variable, for instance the temperature in the room), building managers need to click on a specific link. As the detailed information is not useful all the time, they can choose when they need to access it. In AR, this is even more important, as the space around the user would become visually cluttered if everything is visualised concurrently at all times. It is important to provide a mechanism to add interactions to the visualisation to allow the users to get detailed information only when needed.

DG5: Support flexible methods to link physical objects with their data. With the development of Building Information Modelling (BIM), assets are often tagged with their location and shape. It is then possible to link each asset with the data coming from the sensors associated with them, and then automatically create a 3D overlay for this asset, which is essential for embedded visualisations. However, interviews with building managers and BMS engineers pointed out that BIM models were far to be generalised for all current buildings (most of the old buildings does not have one), thus, our system cannot rely only on this technology. It is important to provide methods to manually create the overlays for the different assets, tag it and link it to the appropriate data source. Building managers also pointed out that such models were rarely updated when a modification was done in a room. This means that even when a BIM model is provided, it is important to be able to modify the information regarding the overlays.

The 5 design goals defined in this section have been used in design and implementation of the authoring tool. They guided our design choices and allowed for quick evaluation of the tool.

4 CORSICAN TWIN PROTOTYPE

Guided by the above design goals, we implement a working prototype of the Corsican Twin using Unity 3D with support for a HTC Vive Pro VR headset and Microsoft HoloLens AR display.

4.1 Overview

To provide context about a physical site (**DG1**), the Corsican Twin immerses authors in a life-size, 3D model of the building interior. We use either 3D photogrammetry scans of environments created with a Matterport [26] camera (Figure 4) or 3D CAD and BIM

Figure 6: Authoring in situ visualisations. Left: The author creates a viewport to display a situated visualisation (a) and associates it with a temporal data source (b). Right: The author creates a scaffold to display an embedded visualisation for this electric cabinet (c), associates it with a data source and sets its encoding (d).

models (Figure 5). These models provide detailed contextual information about the physical spaces, allowing authors to tailor visualisations to those specific sites. Moreover, authors can quickly teleport throughout the model instead of walking, reducing time and effort during design sessions.

Another advantage of VR is that it allows for direct interaction in the 3D space. As most VR headsets support 6 Degree-Of-Freedom controllers, we take advantages of these to support direct manipulation [2] for creating and manipulating visualisations. With this approach we aim to mimic natural interactions with physical objects as much as possible, providing an interface that is simple to use and easy to understand (DG2). We also propose a laser pointer to select object at a distance. System commands and other more abstract operations are completed via menus.

Our prototype supports several different classes of in situ data displays [38], including both *situated* and *embedded* visualisations (DG3).

Situated visualisations are familiar 2D visualisations that are ‘spatially situated’ [17] in 3D space near their physical referent. Our system uses standard visualisation types such as line graphs or bar charts, which are useful for analysing the kinds of time series data that are common in building management.

Embedded visualisations integrate data-driven visual marks directly into the environment, displaying representations of individual data points close to their physical referents. In our system, these visualisations encode data onto the 3D geometry of the relevant equipment by varying visual attributes including colour and size. For time series data these attributes can vary over time to reflect changing data values.

To reduce visual clutter, the Corsican Twin also lets authors add both explicit and implicit interactions to show or hide visualisations (DG4). When enabled, end-users using the Microsoft HoloLens can use the air-tap gesture to show or hide individual visualisations. Authors can also define spatial volumes that show or hide a visualisation when entered, allowing them to toggle visualisations automatically when users enter or leave a certain part of the space.

4.2 VR Authoring Interface

4.2.1 Menus –. Global commands (create visualisation, create interaction, main settings, etc.) are accessed through a global menu on the author’s non-dominant hand [22]. Commands related to

individual visualisations and their encodings are accessed through panels that appear next to them (Figure 6-b, -d) when selected.

4.2.2 Visualisation creation and scaling –. To create a visualisation, authors first choose either a situated or embedded element from the global menu. Figure 6-Left shows the process of creating a situated visualisation. First, the author draws the chart canvas using their controller. A control panel then appears, which the author can use to associate the new visualisation with an physical asset (such as an air handling unit) and a data set (like fan speed or temperature). The author can also use the panel to set the chart type and adjust settings including axis and encoding parameters.

Figure 6-Right shows the process of creating an embedded visualisation. The author first chooses a “scaffold” shape (cube, sphere or cylinder) that best matches the geometry of the relevant physical asset (usually a piece of equipment). Then they use the controller to place and size the scaffold so that it surrounds or intersects the asset. Once placed, the author can bind visual attributes of the scaffold like colour and vibration to show data from the asset. If a BIM model is available, authors can skip the drawing step and instead select geometry directly from the model (DG5). The system then automatically rescales the geometry to create the scaffold for a new visualisation (as in Figure 5-Right).

Once created, authors can continue to select and move visualisations using both direct manipulation or using distant pointing [2]. When a visualisation is selected, handles appear at the corners of the shape to allow reshaping and scaling.

4.2.3 World-In-Miniature (WiM) –. Large environments may cause challenges in the Corsican Twin since it may be difficult for authors to reach and scale scaffolds for bigger assets. To provide an overview of the scene, authors can view a WiM [32], a miniature copy of the life-sized virtual scene. The author can manipulate the WiM model to view it from different perspectives, as well as to create, select and move visualisations. Changes are immediately reflected in the life-sized scene (Figure 7).

4.2.4 Showing/Hiding Visualisations –. Authors can also create volumes that act as triggers for show/hide events. These can be strategically placed in the environment to enable proxemic interactions [1], for example revealing a visualisation only when a viewer approaches a related piece of equipment. To create these triggers,

Figure 7: The author views a World-in-Miniature model of the virtual environment. The miniature BIM model can be seen in the model along with the world-scale BIM model in the background.

authors first draw and scale a trigger region then connect it to visualisations by drawing arrows (Figure 8-Left). Authors can choose one of three different trigger types: *Click*, *Hover* and *Position*. *Click* triggers allow end users to hide and show visualisations manually using the HoloLens air tap gesture. *Hover* triggers show any linked visualisations when a viewer’s head is aimed at the trigger region and hides them when they look away. *Position* triggers cause linked visualisations to appear when the viewer enters the region defined by the scaffold (Figure 8), and hides them on exit.

4.3 End User AR Application

End users can view the visualisations in the real-world environment using a Microsoft HoloLens [27]. The spatial location of the AR components are calibrated using a single QR code placed in the physical environment. In the VR application, the author places a virtual QR Code in a suitable location. A matching code is then placed at the same location in the real room. Using Vuforia, the HoloLens detects this marker and then aligns the authored visualisations in their relative locations.

5 EXPERT FEEDBACK SESSIONS

To evaluate our system, we performed three feedback sessions with groups of expert users using an HTC Vive VR headset and a Microsoft HoloLens AR headset. We conducted the first two sessions in an on-campus research space for which we have a variety of sensor data, a photogrammetry scan, an architectural model and a mechanical BIM model. The last session was conducted in an office space with the same models. We streamed both views on a large display to allow all participants to see the current user’s view.

5.1 Procedure

To demonstrate the full functionality of the prototype, we showed the experts a complex set of situated visualizations tailored uniquely to the space (Figure 5). We first added a simplified version of the mechanical BIM model which incorporated simple geometric models of several pieces of equipment situated inside and outside the space—including an air handling unit (AHU), air return fan, dampers, pipes, and vents. On top of these we visualised 15 different datasets drawn from the building management system, including data from

Figure 8: Left: Authoring a proxemic interaction. The two situated visualisations will appear when a user enters the red rectangle and disappear when they leave it. Right: The situated visualisation appears to the AR user when they enter the corresponding location.

pieces of equipment in the room and environmental data about the space itself. We used colour and size encodings to show the return and supply air temperature of the air handling unit and speed of the return fan. We also added a vibration encoding to both units that would trigger in response to a “fail to start” alarm on either device. Similarly, we used a colour encoding on the room’s vents to represent the temperature of the incoming air. Finally, we used situated line graphs to show longer-term changes in room temperature as well as changes in the air handling unit’s heating and cooling coil. To reduce the amount of visual clutter in the room, we included interactions for toggling on and off these situated charts. Viewers could toggle the temperature chart by clicking on any of the vents, or toggle the heating and cooling coil chart by standing close to the air handling unit.

During each session, we first showed the AR application to give participants a better idea of what AR visualisations could reveal in the space. Next, we discussed the capabilities of the AR tool and detailed each of its features. After this, we showed participants the VR authoring tool and demonstrated its functionality. Here we first introduced the VR environment and navigation controls, then the general features such as the photogrammetry scan, the world in miniature, and changing the opacity of the BIM model. Finally, we showed participants how to author both situated and embedded visualisations and add interactions to them. At each step, we first demonstrated the functionality then encouraged participants to repeat the task, reserving time afterwards to discuss the feature.

To ensure that the data contained interesting events that might trigger discussion, we used data from an instance (01/01/2018 at 09h39) when a “Fail To Start” alarm occurred in the space. During the entire session, we encouraged participants to think aloud while interacting with the prototype and to comment as they watched others using the tool. We audio and video recorded the sessions then transcribed them for analysis. Overall, each session lasted roughly 1 hour and 30 minutes.

5.2 Participants

For the first session, we recruited a group of 3 building management system engineers from our university (Table 1-Left), whose primary job is to monitor the mechanical assets in campus buildings and to perform maintenance in case of failure. The participants in our second session included 4 members of the university’s project

Group 1 - Engineers		Group 2 - Project Planners		Group 3 - UX Team	
P1	BMS Engineer	P4	Sustainability Analyst	P8	UX Designer
P2	BMS Engineer	P5	Project Engineer	P9	Marketing Director
P3	BMS Technician	P6	Asset Planning Engineer	P10	UX Designer
		P7	Strategic Asset Planner	P11	Software Developer
				P12	UX Designer
				P13	UX Designer
				P14	Software Developer
				P15	UX Designer

Table 1: Participants of the feedback sessions and their roles.

planning team (Table 1-Middle), who focus on the construction of new campus buildings, retrofitting existing ones, and managing large changes to campus infrastructure. Finally in the last session, we recruited a group of 8 participants from the UX design team of a large international company that specialises in Building Management Systems (Table 1-Right). None of the participants in the first group had previously used a VR headset, compared to 3/4 of in the second group, and 8/8 in the third. However, none had substantial experience with either AR or VR technologies.

5.3 Reflections

In the 3 sessions, experts gave detailed feedback on the prototype and possible building maintenance and planning applications. Though some participants had never used VR before, all were able to successfully use the system after a short round of training. The main usability challenge for participants was using the grasping metaphor in menus, and most relied on the laser pointer instead.

5.3.1 Using AR During Maintenance and Planning. The experts in all three sessions found the AR prototype very easy to use and required very little training to grasp the functionality of the implemented tools. Participants in the first group emphasised the potential utility of AR visualisations for maintenance, especially in plant rooms where lots of data is available, and engineers need to account for many different parameters to understand a situation. They also liked the fact data visualisations could be situated near their physical referents as this feature could avoid confusion between similar assets. P2 stressed this point, noting that he “*could walk up and look at a valve and press a button and it showed me the trend log on that valve.*”. In the third session P9 suggested that AR could provide new information, saying “*[it will] make them super smart. They will see the flow [overlaid on the pipes]. This information is not even in the books.*”. Similarly, the project planners in the second session suggested that there is considerable potential to show live data in AR for higher-level building management. However, one participant mentioned that it would be useful to be able to use a phone as an AR device, as AR headset are not yet common. The experts also highlighted the fact that interactions to show visualisations can be used to allow users to “dig in” to required information only when needed.

5.3.2 Model Quality. Despite visible gaps and missing elements in the 3D photogrammetry scan of the room, participants responded positively, noting that it was good enough to give a sense of the space, especially for someone not familiar with the building. The engineers in the first session emphasised that models of this quality could give a good overall sense of the types of equipment found in each room, and might be considerably more useful than floorplans which rarely provide this information. One expert also suggested

that a such models could be combined with thermal simulations showing air flow to diagnose problems with heating and cooling systems, making it possible to “*visually see the draft [of air]*” (P1).

5.3.3 Visualisations. Participants in all groups agreed that allowing different end-users to **author and customise** their own visualisations was important, since each technician is sensitive to the particular types of information they need when diagnosing a failure. Engineers in the first group contrasted this against traditional building management systems, which do not allow for interface customisation and often require multiple clicks to access important pieces of information. Our prototype, by contrast, allowed them to view personally-relevant visualisations of task-specific data in the appropriate spaces. Participants in the second group further emphasised the benefit of allowing users to author their own sets of visualisations for a space, noting that personnel in different roles (mechanical, electrical, etc.) can have distinctly different information needs. They also stressed the value of interaction to hide and show visualisations to help reduce information overload, with P6 emphasising “*When you are in front of this equipment, you know this is what you want to focus on.*”

Meanwhile, experts in all sessions emphasised the potential value of situated and glanceable encodings for **monitoring** tasks. P2 emphasised how persistent visualisations could help with common heating diagnostic tasks where “*there is the common [visualisation] that you want. You are looking at a valve and you show the speed of the fan, the boiler temperature and the chiller temperature. So you know that stuff is running.*”. In the second session, P4 stressed the value of visually striking encodings like vibration for monitoring noting that, “*You can walk into rooms and go—I know something is not working.*”. Designers in the third group also suggested focusing technicians’ attention by revealing visualisations only when they show abnormal values. They also proposed allowing visualisations to pivot dynamically, so that they always face the viewer—although doing so could result in occlusions and may be problematic during collaboration, where multiple users’ views may not align.

Experts in the second session also discussed the potential for situated visualisations to serve as entry points for other kinds of **documentation**. In particular, they suggested linking reference documentation and information regarding the history of the equipment directly to the situated visualisations, noting that “*You could show all the history of an asset without having to search for it*” (P5).

In addition to information directly linked to the equipment, the engineers also highlighted the importance of **showing data about more distant systems**, since these can also provide valuable context. P1 illustrated the value of visualising information about the location of related systems elsewhere the building with an anecdote: “*For example, recently, I’ve got an alarm in building 89 and I know a room is too hot, but where that your chilled water come from? It doesn’t tell on the BMS. You’ve got to look at how you embed all that sort of stuff into these models so that people can find the fault easily.*”.

5.3.4 Visualisation Placement and Layout. Overall, participants tended to create and place visualisations directly in the virtual world. However, several also chose to use the WiM to create and place larger visualisations. Several participants in the third group also chose to place their situated visualisations on walls or other flat surfaces rather than in space. P14 suggested that snapping

visualisations to walls “like in Powerpoint” would help make them more manageable. Another participant (P10) extended this idea, suggesting that the system pre-populate spaces by creating and snapping a default templated visualisation to each asset in the model, then allowing designers to customise it. One participant (P11) also suggested placing embedded visualisations on the sensors rather than the equipment, noting that in the case of a failure, technicians are likely to focus their attention there first.

5.3.5 Other uses of VR. The three groups of experts were divided over the value of using the VR interface for **remote monitoring**. The engineers in the first session found remote visualisation in VR unappealing, since they assumed that in the event of a fault they would need to go to the site anyway. However, the planners in the second session were more interested in this functionality, remarking that “If I don’t have to go into the plant room, it is good to be able to log in in a virtual world” (P6). The UX team in the third group were the most able to foresee practical benefits of VR.

Groups agreed that the virtual environment could be useful when **planning** new buildings or retrofitting existing ones, where the virtual system could allow them to see the impact of different equipment configurations. Participants in the first group described the potential for using a VR tool to assess whether a cooling system could cope with the additional heat generated by adding new freezers to a lab space. In the second group, P4 echoed this statement, noting that “if we’re going to put any generator in...you can load the model with the specs of what you’re looking at and put it into the room and see if that works.”

6 DISCUSSION AND CONCLUSION

Our Findings suggest that the Corsican Twin allows users to easily create useful in situ visualisations. They also highlight how a user-centered design process could improve efficiency of the visualisations, avoid errors in the design, and support customisation. Authoring visualisations in VR using 3D models of the target environments can provide users with an awareness of the context and a strong sense of space. However, as is well known, VR can be overwhelming for inexperienced users or those predisposed to simulator sickness. To ease the transition between desktop and VR could be valuable for opening up the authoring process to them.

6.0.1 Visualisation authoring approaches. In our discussions, UX designers in the third group explained that when designing dashboards for building managers, they don’t start from scratch. Instead, they typically start with a generic template containing a wide variety of visualisations, then customise it to their particular use case. However, understanding how to place a large number of initial visualisations in a unique space while avoiding occlusion remains an interesting and challenging direction for future research. We noticed that precise tasks like using menus to set up and design the 2D visualisations were challenging in VR. A hybrid approach where designers create and place visualisations in VR but configure them via a desktop interface may be more efficient. This process could even be collaborative, with one user in VR and another on the desktop doing complementary tasks. However, the impact of transitioning between or synchronising these two separate interfaces would require further investigation.

6.0.2 Visualisation placement. The position of embedded visualisations is by definition constrained by the positions of the physical objects to which they are linked. On the other hand, it is possible to place situated visualisations anywhere in the space. Participants in our third session tended to place situated visualisations on walls or other flat surfaces. However, this may be due to the fact that all of our situated visualisations were 2D displays. This suggests that snapping and alignment tools for placing these kinds of visualisations may be a useful addition to future systems. Tools for authoring and placing three dimensional situated visualisations that more fully utilise the empty space in technicians’ work environments also represent a promising opportunity.

6.0.3 Binding data to physical referents. Associating data to their physical referents is an essential part of our solution. Our prototype allows users to extract these associations from BIM models, as well manually define referents for cases where a BIM model is not available. Another possible solution could be to use machine learning to detect and classify the equipment from photogrammetry scans. Initial research in this area is promising [36], but does not yet deal with the kinds of complex geometry found in machine rooms and other spaces with large amounts of equipment. Semi-automated approaches, in which technicians label each piece of equipment with a QR before the scan and use image recognition to identify those codes in the mesh, could help address this complexity.

6.0.4 Working with changing spaces. Spaces undergo constant change. This is a well known challenge in building management and often leads to technical issues — for example adding new equipment to a room can overstretch the capabilities of the cooling system. In our case, changing environments mean that the contextual information the Corsican Twin provides to the user can become outdated and potentially misleading. One solution would be to regularly check the BIM model or periodically re-scan every room to detect changes. Providing technicians with simple mechanisms for highlighting and correcting mismatches between the model and the real environment could also make it easier to deal with small changes or anticipate larger redesigns.

6.0.5 Handling more complex data. The data our Corsican Twin supports are primarily time series. However, more complex data can be useful in building management, as well as in other domains. During our feedback sessions, participants highlighted the value of integrating these kinds of historical displays with simulations and other predictive tools. Data about related equipment in other spaces (for example on another floor) may also be useful to display, especially if they are part of the same larger systems. World-in-miniature displays or virtual portals to other parts of a building could help spatially connect these other pieces of equipment to the current space and also provide visual referents around which to display these kinds of related data.

6.0.6 Application to other domains. While we focus on building maintenance, in situ visualisations have potential in other domains. For example, classrooms, museums, and other public spaces all present opportunities for AR content which connects to the objects and equipment in the environment [12]. Here, systems like the Corsican Twin could help educators or curators design new interactive AR content without requiring access to the spaces themselves.

REFERENCES

- [1] Till Ballendat, Nicolai Marquardt, and Saul Greenberg. 2010. Proxemic interaction: designing for a proximity and orientation-aware environment. In *ACM International Conference on Interactive Tabletops and Surfaces*. ACM, 121–130.
- [2] Doug A. Bowman, Ernst Kruijff, Joseph J. LaViola, and Ivan Poupyrev. 2004. *3D User Interfaces: Theory and Practice*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA.
- [3] Nathalie Bressa, Kendra Wannamaker, Henrik Korsgaard, Wesley Willett, and Jo Vermeulen. 2019. Sketching and Ideation Activities for Situated Visualization Design. In *Proceedings of the 2019 on Designing Interactive Systems Conference (DIS '19)*. ACM, New York, NY, USA, 173–185. <https://doi.org/10.1145/3322276.3322326>
- [4] Marion Buchenau and Jane Fulton Suri. 2000. Experience Prototyping. In *Proceedings of the 3rd Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques (DIS '00)*. ACM, New York, NY, USA, 424–433. <https://doi.org/10.1145/347642.347802>
- [5] Colin Burns, Eric Dishman, William Verplank, and Bud Lassiter. 1994. Actors, Hairdos & Videotape—Informance Design. In *Conference Companion on Human Factors in Computing Systems (CHI '94)*. ACM, New York, NY, USA, 119–120. <https://doi.org/10.1145/259963.260102>
- [6] M. Cavallo and A. G. Forbes. 2019. CAVE-AR: A VR Authoring System to Interactively Design, Simulate, and Debug Multi-user AR Experiences. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 872–873. <https://doi.org/10.1109/VR.2019.8798148>
- [7] Z. Chen, Y. Su, Y. Wang, Q. Wang, H. Qu, and Y. Wu. 2019. MARViT: Authoring Glyph-based Visualization in Mobile Augmented Reality. *IEEE Transactions on Visualization and Computer Graphics* (2019), 1–1. <https://doi.org/10.1109/TVCG.2019.2892415>
- [8] J. Choi, Y. Kim, M. Lee, G. J. Kim, Y. Nam, and Y. Kwon. 2010. k-MART: Authoring tool for mixed reality contents. In *2010 IEEE International Symposium on Mixed and Augmented Reality*. 219–220. <https://doi.org/10.1109/ISMAR.2010.5643576>
- [9] Sandy Claes, Jorgos Coenen, and Andrew Vande Moere. 2018. Conveying a Civic Issue Through Data via Spatially Distributed Public Visualization and Polling Displays. In *Proceedings of the 10th Nordic Conference on Human-Computer Interaction (NordiCHI '18)*. ACM, New York, NY, USA, 597–608. <https://doi.org/10.1145/3240167.3240206>
- [10] Sandy Claes and Andrew Vande Moere. 2013. Street Infographics: Raising Awareness of Local Issues Through a Situated Urban Visualization. In *Proceedings of the 2Nd ACM International Symposium on Pervasive Displays (PerDis '13)*. ACM, New York, NY, USA, 133–138. <https://doi.org/10.1145/2491568.2491597>
- [11] M. Cordeil, A. Cunningham, B. Bach, C. Hurter, B. H. Thomas, K. Marriott, and T. Dwyer. 2019. IATK: An Immersive Analytics Toolkit. In *2019 IEEE Conference on Virtual Reality and 3D User Interfaces (VR)*. 200–209. <https://doi.org/10.1109/VR.2019.8797978>
- [12] F. M. Dinis, A. S. Guimarães, B. R. Carvalho, and J. P. P. Martins. 2017. Virtual and augmented reality game-based applications to civil engineering education. In *2017 IEEE Global Engineering Education Conference (EDUCON)*. 1683–1688. <https://doi.org/10.1109/EDUCON.2017.7943075>
- [13] Alexandre Dumas. 1844. *The Corsican Brothers*.
- [14] N. ElSayed, B. Thomas, K. Marriott, J. Piantadosi, and R. Smith. 2015. Situated Analytics. In *2015 Big Data Visual Analytics (BDVA)*. 1–8. <https://doi.org/10.1109/BDVA.2015.7314302>
- [15] Barrett Ens and Pourang Irani. 2016. Spatial analytic interfaces: Spatial user interfaces for in situ visual analytics. *IEEE computer graphics and applications* 37, 2 (2016), 66–79.
- [16] Eva Eriksson, Martin Ludvigsen, A Lykke-Olesen, and R Nielsen. 2006. Bthere or be Square: A Method for Extreme Contextualization of Design. (01 2006).
- [17] George W Fitzmaurice. 1993. Situated information spaces and spatially aware palmtop computers. *Commun. ACM* 36, 7 (1993), 38–50.
- [18] Mani Golparvar-Fard, Feniosky Peña-Mora, and Silvio Savarese. 2009. D4AR—a 4-dimensional augmented reality model for automating construction progress monitoring data collection, processing and communication. *Journal of information technology in construction* 14, 13 (2009), 129–153.
- [19] Nicolai Brodersen Hansen and Peter Dalsgaard. 2012. The Productive Role of Material Design Artefacts in Participatory Design Events. In *Proceedings of the 7th Nordic Conference on Human-Computer Interaction: Making Sense Through Design (NordiCHI '12)*. ACM, New York, NY, USA, 665–674. <https://doi.org/10.1145/2399016.2399117>
- [20] S. Henderson and S. Feiner. 2011. Exploring the Benefits of Augmented Reality Documentation for Maintenance and Repair. *IEEE Transactions on Visualization and Computer Graphics* 17, 10 (Oct 2011), 1355–1368. <https://doi.org/10.1109/TVCG.2010.245>
- [21] Valentin Heun, James Hobin, and Pattie Maes. 2013. Reality editor: programming smarter objects. In *Proceedings of the 2013 ACM conference on Pervasive and ubiquitous computing adjunct publication*. ACM, 307–310.
- [22] Ken Hinckley, Randy Pausch, Dennis Proffitt, James Patten, and Neal Kassell. 1997. Cooperative bimanual action. In *CHI*, Vol. 97. 27–34.
- [23] Henrik Korsgaard, Nicolai Brodersen Hansen, Ditte Basballe, Peter Dalsgaard, and Kim Halskov. 2012. Odenplan: A Media FaÇade Design Process. In *Proceedings of the 4th Media Architecture Biennale Conference: Participation (MAB '12)*. ACM, New York, NY, USA, 23–32. <https://doi.org/10.1145/2421076.2421081>
- [24] Gun A. Lee, Gerard J. Kim, and Mark Billinghurst. 2005. Immersive Authoring: What You eXperience Is What You Get (WYXIWYG). *Commun. ACM* 48, 7 (July 2005), 76–81. <https://doi.org/10.1145/1070838.1070840>
- [25] B. MacIntyre, M. Gandy, J. Bolter, S. Dow, and B. Hannigan. 2003. DART: the Designer's Augmented Reality Toolkit. In *The Second IEEE and ACM International Symposium on Mixed and Augmented Reality, 2003. Proceedings*. 329–330. <https://doi.org/10.1109/ISMAR.2003.1240744>
- [26] Matterport. 2019. *Matterport*. <https://matterport.com/>
- [27] Microsoft. 2019. *Microsoft HoloLens*. <https://www.microsoft.com/en-us/hololens>
- [28] Microsoft. 2019. *Microsoft Layout*. <https://dynamics.microsoft.com/en-au/mixed-reality/layout/>
- [29] William Odom, John Zimmerman, Scott Davidoff, Jodi Forlizzi, Anind K. Dey, and Min Kyung Lee. 2012. A Fieldwork of the Future with User Enactments. In *Proceedings of the Designing Interactive Systems Conference (DIS '12)*. ACM, New York, NY, USA, 338–347. <https://doi.org/10.1145/2317956.2318008>
- [30] Antti Oulasvirta, Esko Kurvinen, and Tomi Kankainen. 2003. Understanding contexts by being there: case studies in bodystorming. *Personal and Ubiquitous Computing* 7, 2 (01 Jul 2003), 125–134. <https://doi.org/10.1007/s00779-003-0238-7>
- [31] R. Sicat, J. Li, J. Choi, M. Cordeil, W. Jeong, B. Bach, and H. Pfister. 2019. DXR: A Toolkit for Building Immersive Data Visualizations. *IEEE Transactions on Visualization and Computer Graphics* 25, 1 (Jan 2019), 715–725. <https://doi.org/10.1109/TVCG.2018.2865152>
- [32] Richard Stoakley, Matthew Conway, and Y Pausch. 1995. Virtual Reality on a WIM: Interactive Worlds in Miniature. (02 1995). <https://doi.org/10.1145/223904.223938>
- [33] Bruce H Thomas, Gregory F Welch, Pierre Dragicevic, Niklas Elmqvist, Pourang Irani, Yvonne Jansen, Dieter Schmalstieg, Aurélien Tabard, Neven AM ElSayed, Ross T Smith, et al. 2018. Situated Analytics.
- [34] Fernando Vera, J. Alfredo Sánchez, and Ofelia Cervantes. 2017. A Platform for Creating Augmented Reality Content by End Users. In *Applications for Future Internet*, Enrique Sucar, Oscar Mayora, and Enrique Munoz de Cote (Eds.). Springer International Publishing, Cham, 167–171.
- [35] J. Vermeulen, F. Kawsar, A. L. Simeone, G. Kortuem, K. Luyten, and K. Coninx. 2012. Informing the design of situated glyphs for a care facility. In *2012 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. 89–96. <https://doi.org/10.1109/VLHCC.2012.6344490>
- [36] Rebekka Volk, Julian Stengel, and Frank Schultmann. 2014. Building Information Modeling (BIM) for existing buildings – Literature review and future needs. *Automation in Construction* 38 (2014), 109 – 127. <https://doi.org/10.1016/j.autcon.2013.10.023>
- [37] Yuan Xing Wang, Tobias Langlotz, Mark Billinghurst, and Tobin Bell. 2009. An Authoring Tool for Mobile Phone AR Environments.
- [38] W. Willett, Y. Jansen, and P. Dragicevic. 2017. Embedded Data Representations. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (Jan 2017), 461–470. <https://doi.org/10.1109/TVCG.2016.2598608>
- [39] Mike Wozniowski and Paul Warne. 2011. Towards in situ authoring of augmented reality content. In *ISMAR*.