LuxAR: A Direct Manipulation Projected Display to Extend and Augment Desktop Computing

Yuan Chen*
y2238che@uwaterloo.ca
School of Computer Science, University of Waterloo
Waterloo, Canada

Géry Casiez†‡
gery.casiez@univ-lille.fr
Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9189 CRIStAL
Lille, France

Sylvain Malacria§
sylvain.malacria@inria.fr
Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9189 CRIStAL
Lille, France

Daniel Vogel
dvogel@uwaterloo.ca
School of Computer Science, University of Waterloo
Waterloo, Canada

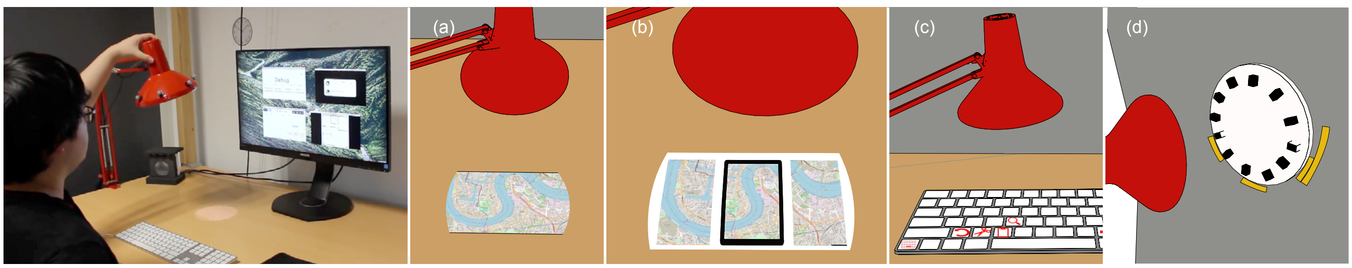

Figure 1: LuxAR is an architect lamp instrumented with a projector to extend direct manipulation interfaces into the physical surroundings, such as: (a) moving desktop windows onto desk surfaces; (b) extending interaction space of mobile devices; (c) augmenting keyboard with shortcut keys; (d) visualizing information related to nearby objects like a wall clock.

ABSTRACT

We prototype and evaluate a desktop input and output device in the form of an architect desk lamp. The bulb is replaced with a pico laser projector, a button replaces the switch, and the mechanical design allows it to remain in position between user manipulations. By also tracking lamp position and orientation, we explore novel interaction techniques to extend and augment conventional devices in a physical desktop environment. Content can be transferred from devices to the surrounding environment, and the representation of content can be adapted to surfaces, objects, and other devices using the lamp projection target and lamp proximity. A user study with the prototype and semi-structured interviews examine proposed interactions and consider potential scenarios and applications. Based on the results, we propose further design considerations for direct manipulation systems to extend and augment desktop computing.

KEYWORDS

Augmented Reality, Desktop Computing, Interaction Techniques.

ACM Reference Format:
Yuan Chen, Géry Casiez, Sylvain Malacria, and Daniel Vogel. . LuxAR: A Direct Manipulation Projected Display to Extend and Augment Desktop Computing. In *Proceedings of* . ACM, New York, NY, USA, 12 pages.

*Also with Univ. Lille, Inria, CNRS, Centrale Lille, UMR 9189 CRIStAL.
†Also with University of Waterloo.
‡Also with Institut Universitaire de France.
§Also with University of Waterloo.

, ,
.

1 INTRODUCTION

People often rely on multiple displays and devices to create an expanded desktop space for multitasking and information coordination [3, 9, 13, 47]. However, physical space constraints, optimal device placements, and cost often pose challenges [13, 47]. A mixed reality (MR) head-mounted display (HMD) can be used to expand desktop computing space [20, 25]. However, using an HMD creates an isolated digital workspace [25] and can induce fatigue [4].

In contrast, spatial augmented reality (SAR) uses one or more projector-camera units (pro-cams) for projection mapping on physical surfaces to transform them into interactive displays [14, 30]. SAR has been used to extend desktop computing to surfaces [11, 31, 40], but in a limited way with a single fixed pro-cam. Some SAR systems use multiple fixed-position pro-cams to cover more surfaces [14, 30, 31] for presenting information. However, projection space cannot be adjusted once the system is running. One way to overcome this limitation is to use steerable pro-cams [15, 44], but users cannot explicitly control the projection space. Handheld pro-cam systems supporting explicit and direct input have been proposed [2, 7, 12, 42], but these must be held in the air for aiming and input movements, fully occupying at least one hand and leading to fatigue. To support explicit and direct manipulation with a pro-cam, a method to hold the pro-cam in space would be ideal, especially to support long-duration tasks such as placing additional desktop information on surfaces.

An architect lamp has articulated, counterbalanced arms, that can support a pro-cam in space. This means the lamp can be aimed at a surface and it will maintain its pose after manipulation. The

Lantern demo [19] and the LuminAR prototype [22] both integrate a pro-cam in an architect lamp and leverage these adjustment and support properties. However, the Lantern requires a phone for all input relegating the role of the lamp to be only an adjustable display. LuminAR actuates the lamp using motors, with display position controlled indirectly through gestures performed in front of the lamp with touch input to interact with content. This adds significant complexity and cost. Importantly, neither uses an augmented lamp as a form of direct manipulation on its own.

We use a pro-cam mounted in a standard architect lamp for explicit and direct manipulation of a SAR interaction space to extend and augment desktop computing (Figure 1). Our new design space, with proof-of-concept LuxAR system, demonstrates a novel form of direct manipulation leveraging unique characteristics of physical lamp movement. With a single button (mounted on the lamp hood like a standard switch), we show how this can be used to reposition displays, interact with applications, and adapt content to various surfaces, devices, and objects. To maintain the flow of desktop computing, the approach can integrate standard input from the mouse, as well as mobile devices with touch and pen.

Our work makes three main contributions: (1) a new interaction design space for physical direct manipulation using an articulated pro-cam lamp; (2) the LuxAR proof-of-concept system demonstrating usage applications in various scenarios; and (3) results of a two-phase user study evaluating the potential of such prototype and design implications for explicit and direct manipulation systems to extend and augment desktop computing. Together, these contributions establish a promising way to connect current desktop computing with the surrounding physical desktop space.

2 RELATED WORK

We discuss SAR projection systems for extending desktop computing and past examples of lamps as input or output devices. Table 1 summarizes our work compared to previous systems.

2.1 Extending Desktop Interaction with SAR

Conventional SAR systems (*e.g.* [10, 14, 30, 40]) typically use fixed pro-cams mounted on ceilings or tripods, but these require multiple pro-cam setups to cover different surfaces with reasonable resolution [10, 14, 30]. Prior research has investigated ways to minimize the number of pro-cams while still supporting interactive experiences across multiple surfaces. The Everywhere Display [26] and the Escritoire system [1] employ a manually adjusted mirror to redirect the projection display on different surfaces. Maeda et al. [24] use a fixed fisheye pro-cam for omnidirectional displays, but its configuration limits projected information to areas with fiducial markers. To relax such constraints, researchers explored *implicit* and *explicit* methods for controlling the projected display. Implicit control leverages contextual cues to interpret interactions [33]. For instance, Project LFX [27] and Beamatron [44] track a user in a room and implicitly repositions projected content onto the nearest projectable surface. Joshi et al. [15] mount a pro-cam on a chair, using the chair's pose to implicitly control the projected display content and location in the environment. However, implicit control relies on behaviour and context analysis which may not always accurately reflect the user's intentions and needs. In contrast, explicit control means the user decides exactly where a projected display is positioned and how it is used.

Explicit user control over SAR systems can be *indirect* and *direct*, depending on whether the projection display is manipulated physically by users. For instance, LuminAR [22]uses hand gestures to move an actuated pro-cam lamp to position a projection display on a single desktop surface; Beamatron [44] combines body poses and voice commands from users to guide the display to a target position in a room. These indirect control systems enable users to control where the projected display is located, but users have limited interactions with the content when adjusting the systems' pose. Direct manipulation increases input expressiveness and enables users to fully control all aspects of a projection display. A straightforward approach is to hold a pro-cam in the air and aim it at surfaces. Cao and Balakrishnan [7] developed a handheld SAR system to annotate virtual and physical space. By aiming a handheld pro-cam in space, it can reveal stories hidden in a storybook [43] or enable collaborative interactions between multiple users on a wall [41]. MotionBeam [42] engages users with projected anime characters through its manipulation and awareness of the physical environments. As a variation of the handheld approach, Blasko et al. [5] developed a wrist-worn SAR system for manipulating projected content, like selecting, panning, and scrolling information on a wall. A handheld system can also project information onto physical objects to interact with them, such as controlling lamps and televisions [34].

However, a notable limitation with most of these SAR systems is that interactions are *attached* to the direct manipulation. Users need to continuously carry and operate them, making it impossible to interact without holding the system in hand and easily leading to fatigue. In contrast, user interactions in implicit or indirect systems are usually *detached* to their pose and manipulation, *e.g.* users can control a virtual car across various surfaces using a joystick controller [44], or type on a virtual keyboard with both hands on the table [22], without the need to hold and manipulate the system. Besides, most direct manipulation SAR systems emphasize on interacting with content, which may lead to a lack of space awareness and adaptive interactions based on the environments the system references, including surfaces, objects, and devices.

Our approach extends desktop computing through explicit and direct manipulation. It enables the user to interact with the projected augmentation by adjusting the pose of an architect lamp (*attached to the manipulation*) and when the lamp is stationary and maintains the augmentation, it allows users to interact with the augmentation through other inputs (*detached from the manipulation*). This bridges desktop and physical space, incorporates multiple inputs, and facilitates adaptive adjustments of virtual content to environmental and contextual changes.

2.2 Lamps as Interactive Systems

Lamps are ubiquitous objects in many interactive systems and offer various opportunities for interactions based on their forms. When not manipulated, they serve as strong supports to hold a system and allow users to interact with the content using different inputs. Using only a camera mounted inside a lamp, HuddleLamp [28] tracks devices and recognizes configurations so users can annotate

	Form	Implicit	Explicit Indirect	Explicit Direct	Display	Reposition	Interact	Track	Wall	Table	Ceiling	Objects	Devices	Env. Awareness	Handheld	Mouse	Pen	Touch	Device	Gestures	Voice	Attached	Detached	User Eval
Our Work	**Architect**			✓	✓	✓	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓			✓	✓	✓
Lantern Demo [19]	Architect			✓	✓	✓			✓	✓		✓							✓				✓	
LuminAR [22]	Architect		✓		✓	✓			✓	✓								✓		✓			✓	
MotionBeam [42]	Torch			✓	✓	✓	✓		✓					✓	✓							✓		✓
Cao et al. [7]	Torch			✓	✓	✓	✓		✓	✓		✓			✓		✓					✓	✓	✓
Beardsley et al.[2]	Torch			✓	✓	✓	✓		✓			✓			✓							✓		
PICOntrol [34]	Torch			✓	✓	✓	✓		✓			✓		✓	✓							✓		
SideBySide [41]	Torch			✓	✓	✓	✓		✓						✓							✓		
AR Magic Lantern [12]	Torch			✓	✓	✓			✓			✓		✓	✓							✓		
Omnilantern [46]	Lantern			✓	✓	✓			✓	✓		✓			✓							✓		
HideOut [43]	Torch			✓	✓				✓	✓		✓			✓							✓		
Project LFX [27]	Ceiling	✓			✓	✓			✓	✓		✓		✓					✓		✓		✓	
HuddleLamp [28]	Architect							✓					✓						✓				✓	
IllumiShare [16]	Architect				✓					✓							✓	✓					✓	✓
AR Lamp [18]	Architect				✓					✓							✓	✓					✓	✓
FACT [21]	Tabletop				✓					✓							✓						✓	
Xiao et al. [45]	Ceiling				✓					✓								✓					✓	
Lamposcope [39]	Architect				✓					✓							✓						✓	
		User Control			**Functions**				**Surfaces**						**Inputs**							**Modality**		

Table 1: Previous interactive lamp systems compared to our work (sorted by relevance, see text for comparison criteria).

the same information across different devices. But this is only a computer vision input platform. Some systems use a fixed display, for example integrating a pro-cam into a ceiling light [45] or a stationary lamp on the desk [16, 18, 21, 39]. These systems do not focus on input, limiting interactions to specific regions with hands or a pen. The form factor of a torch (i.e. a "flashlight") [2, 7, 12, 41–43] and lanterns [46] have inspired handheld display devices for users to interact with virtual information directly. Of course the handheld nature means at least one hand is occupied during active usage to maintain the position and to perform direct manipulation for input. This makes it challenging to establish sustainable augmented physical environments for desktop computing spaces.

In contrast, articulated lamps, especially architect lamps, featuring unique mechanical structures, provide flexible manipulation and maintain their pose after manipulation. This creates an opportunity for users to interact with information using interactions and inputs that are *attached* or *detached* to the direct manipulation of an architect lamp. In addition to hold a pro-cam unit in space [16, 18, 39], for example, the Lantern demo [19] allows users to rotate the lamp's head to adjust projected displays on a desired surface, and then use phones to change the content. LuminAR [22] integrates architect lamps with robotic systems and enables users to control the projected display through mid-air gestures, and interact with content through direct touch. However, unlike torch or lantern inspired systems, both fail to consider designing interactions based on their direct manipulation and poses, and are unaware of changes of physical spaces to adapt content accordingly. Critically, while various types of lamps have been proposed as interactive systems, there are surprisingly few user evaluations of these kinds of systems.

To the best of our knowledge, the direct manipulation potential of architect lamps has not been investigated yet in order to interact with virtual content and extend and augment desktop computing. We examine potential design possibilities in this work and introduce LuxAR as a system to elevate user experiences from the desktop space to various physical surfaces.

3 LuxAR DESIGN GOALS

We want to fully leverage architect lamp properties for direct manipulation and ability to remain stationary after control. The goal is to design an explicit and direct SAR system to extend and augment a desktop computing space onto nearby physical environments with an awareness of the immediate environment. In contrast to most torch-based systems using direct manipulation input, it should also support interaction with content using other inputs. We summarize with the following designs goals:

- DG1 Content interactions can be attached to lamp's direct manipulation, like rotating for repositioning or adjusting distance to change displayed information.
- DG2 When adjusting the lamp, content can appear on various desktop surfaces like tables, nearby walls, and the ceiling. Content could also be projected onto physical objects and mobile devices on these surfaces. The floor is excluded due to clutter and user movement.
- DG3 The system should detect changes in physical surfaces, objects, and devices and adapt the displayed content accordingly.
- DG4 Content interactions can also be detached to the manipulation of the lamp, allowing users to use other inputs on the desktop.

By implementing these design goals, our approach serves as both an input and output device to connect desktop computing with the nearby desktop environment through explicit and direct manipulation. A key aspect is how it leverages physical characteristics of an architect lamp, occupying a unique position relative to previous work (Table 1).

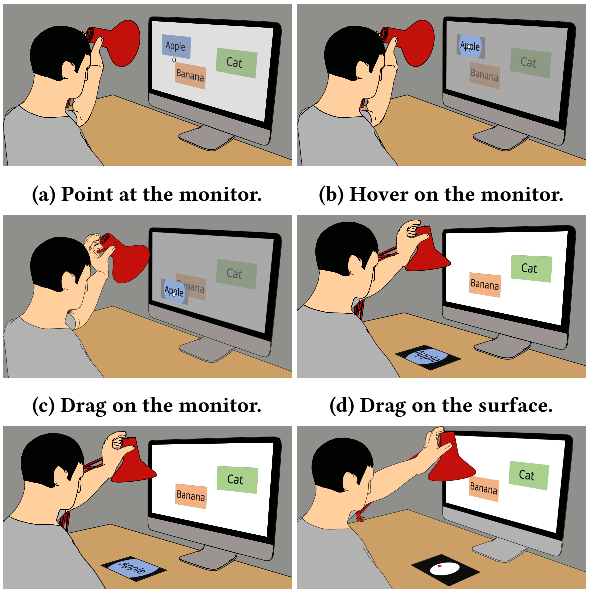

(a) Point at the monitor. (b) Hover on the monitor.

(c) Drag on the monitor. (d) Drag on the surface.

(e) Focus on the surface. (f) Radar on the surface.

Figure 2: From (a) to (e), users manipulate the LuxAR to move an application from the monitor onto a nearby surface.

3.1 Content Repositioning [DG1, DG2]

Content can be moved beyond monitor boundaries onto surrounding surfaces such as desks, walls, and ceilings. Ray-casting is used to interact with virtual content [2, 5, 7]. The lamp head is the ray origin with the lamp head angle determines the ray direction. A button on the lamp shade can be pressed with an index finger or thumb while manipulating the lamp. A spotlight metaphor visualization [29] highlights the location where the lamp points. Two variations of this spotlight visualization are used.

When pointing at a monitor, the visualization of the ray intersection point is a small white dot, essentially forming a "lamp cursor" (Figure 2a). To emphasize the dot, the rest of the screen dims. Hovering the lamp cursor over an application window further dims the screen and highlights the window in an oval shape (Figure 2b), indicating it is select-able by the lamp. To move the window to a surface outside the monitor, the user presses and holds the lamp button to drag, much like dragging a window with a mouse (Figure 2c). Releasing the button drops the window onto the surface. The window remains anchored in place until it is picked up again.

When pointing at surfaces outside the monitor, the spotlight visualization changes according to four modes based on the visibility of the window and the interaction status (Figure 3). In radar mode, a circular display appears when the lamp is not directed at any window. Coloured icons within indicate the direction and location of an anchored but not visible window (Figure 2d). Hover mode activates when the lamp hovers above a virtual window, displaying a large oval shape for detailed examination (Figure 2e). Clicking or pressing a button on the window transitions to focus mode, allowing the window to receive additional input events. Clicking on an empty area reverts to hover mode, and moving the lamp away from the window exits hover or focus mode, returning to radar mode. Users press and hold the button to enter drag mode. The lamp shows a smaller oval display to maintain contextual visibility during dragging (Figure 2f). The window maintains orientation invariant and aligns parallel to the user's seated location.

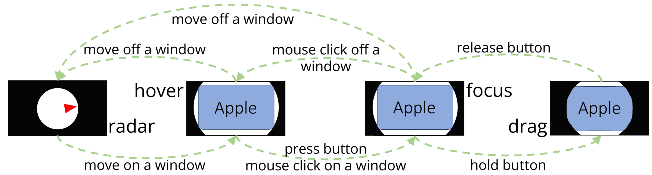

Figure 3: Interaction state machine with lamp display styles.

3.2 Surface Adaptation [DG1, DG2, DG3]

When a user moves a virtual window across different surfaces, like from the desktop to a desk, walls, and the ceiling, both the lamp display size and window visibility change, impacting user behaviors around the desk: A window on the desk is easily visible to a single user, while on the wall, it becomes noticeable to multiple users; users may find it less convenient to tilt their heads upward when seated to view information on the ceiling; yet, when users move away from the desk, information displayed on the ceiling serves as an ambient cue, easily noticeable from a distance.

We employ design principles inspired by Vogel and Balakrishnan's work [38] to facilitate interaction transitions. Our approach allows explicit manipulation of the lamp to reposition content on different surfaces, adapting the virtual window to distinct interaction spaces: the desk for personal interactions, the wall for public interactions, and the ceiling for ambient interactions, as shown in Figure 4. With all the information located on the desk, this space serves as a decision-making platform for users to determine the optimal placement and presentation of the virtual window based on their preferences. When the window moves to the wall, it becomes visible to a larger audience, fostering collective information sharing. On the ceiling, it may escape the seated user's immediate notice but serves as an ambient cue for others.

3.3 Transition Granularity [DG1, DG2]

Building on prior work that used proximity to adjust the detail and granularity of information in virtual content, such as PenLight [35], PaperLens [36], and Cao and Balakrishnan's handheld system [7], we embrace a comparable approach. Proximity, in our context, is defined as the distance between the lamp and the projected surface, categorized into three levels: High, Medium, and Low. When our lamp is oriented to the desk surface, its height can be adjusted between approximately 7 cm and 80 cm, maintaining a stable position. Given that the size of the lamp's projected display is influenced by its height adjustment, ensuring optimal visibility and interactivity of the projected content within the lamp display becomes essential. Based on our lamp's physical shape, we empirically established threshold values (25 cm and 50 cm) to determine transitions between different proximity levels for the desk surface. We adopted these same values for the wall for simplicity, and they are ignored for the ceiling because of the absence of a manipulation axis.

At each proximity level, as the lamp's distance changes, the content and size of the virtual window remain consistent. This

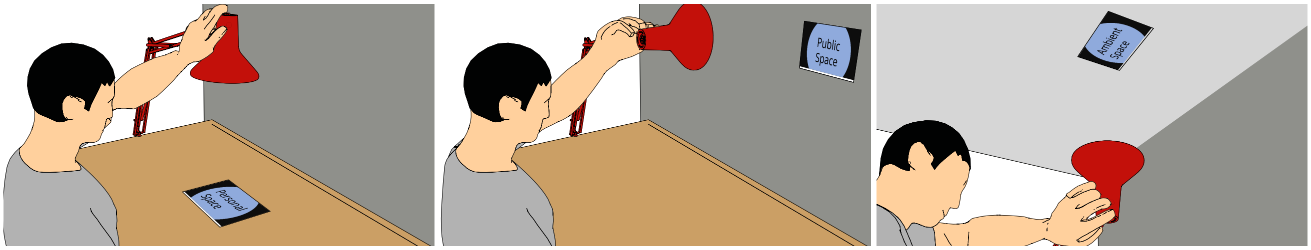

Figure 4: The content in the window adapts to the change of surface when users manipulate LuxAR.

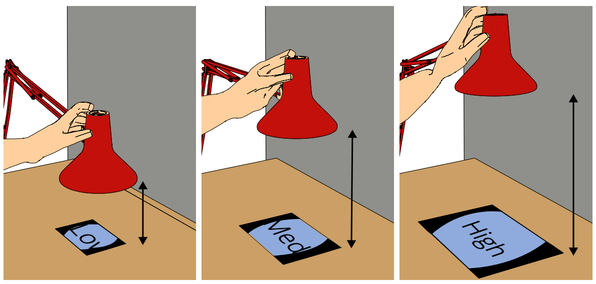

Figure 5: Adjusting the height alters the context of the window to accommodate the changing size of the lamp display.

enables the system to simulate a zooming effect by dynamically adjusting the display size of the lamp. When the height of the lamp crosses a proximity threshold and transitions to a different level, the targeted virtual window adjusts its content, similar to the PaperLens technique [36]. Specifically, when the lamp is raised and positioned farther from the surface, the size of the lamp display increases. This exhibits higher information density in the virtual window, providing a broader context. Conversely, as the lamp is gradually brought closer to the surface, the size of the lamp display decreases. The virtual window reduces information density while enhancing information quality, offering a focused view of the virtual window.

3.4 Object and Device Augmentation [DG2, DG3]

The desk space, with various displays, devices (*e.g.* smartphones), and tangible objects (*e.g.* keyboards, mice, and clocks), requires thoughtful consideration for the system 's awareness of its surroundings. Taking into account the timing (Temporary vs. Permanent) and presentation (Overlay vs. View) aspects of augmentations to physical objects and devices, we suggest three interactions: Temporary Overlay, Permanent Overlay, and Temporary View. We exclude Permanent View since LuxAR inherently transforms physical surfaces into permanent displays for presenting information.

Permanent Overlay (PO) persistently augments physical objects, anchored to their location, revealed by the lamp hovering. For example, a keyboard is always augmented and the shortcuts for an application are overlaid on it when hovering over it with the lamp.

In contrast, Temporary Overlay (TO) consists in temporarily enhancing a physical object based on the content displayed in a virtual window. A user picks up and drops a virtual window on a physical object. The window disappears and the content blends into the object, aligning with its context and shape. To retrieve information from the object, users replicate the same action on the object, transforming the blended content back into a window. An example is to augment a clock with virtual window information, such as reminders, calendars, and weather forecasts, allowing users to change the augmentation by dropping different windows.

For mobile devices on the desk, Temporary View (TV) enables users to expand the small display of the device using the lamp. The procedure is akin to Temporary Overlay: users pick up and drop a virtual window into the mobile device. Subsequently, the virtual window is transferred to the device, and additional views are displayed around it to enrich interactions.

3.5 Other Inputs [DG4]

Device Interaction. Temporary View enables users to interact with projected virtual content via touch screens on mobile devices, the modified content can then synchronize across various devices, surfaces, and desktops.

Mouse Interaction. Given LuxAR's expansion of existing desktop environments into physical spaces, maintaining mouse-based interaction is crucial, as they have demonstrated efficiency in SAR [10, 17]. The virtual cursor is present on both the monitor and the lamp display, seamlessly transitioning between them. Users can edit content on physical surfaces using the mouse when the cursor is in the virtual window. Additionally, the spatial relationship between the lamp and the physical display allows users to reposition the virtual cursor. This design also extends cursor functionality to augmented mobile devices, enabling interaction through both touch input and mouse control with Temporary View augmentation.

Pen and Touch Interactions. We enable touch interaction through a ring-mounted marker on the index finger and pen interaction using a registered pen, both tracked by the OptiTrack.

4 PROOF-OF-CONCEPT SYSTEM

4.1 Augmented Architect Lamp

Our proof-of-concept prototype was built in two simple steps. First, we created the LuxAR prototype by modifying a Ledu architect lamp. We integrated a Nebra AnyBeam laser projector[1] beneath the lamp, measuring 103mm (Length) × 50mm (Depth) × 19mm (Height). This projector supports auto-focus and provides a resolution of 1280 × 720p at 60 FPS. Additionally, to enable the input capability, we installed a button on top of the lampshade, using an Arduino MKR WiFi 1010 (Figure 6). This design allows users to use a finger,

[1]https://github.com/NebraLtd/AnyBeam

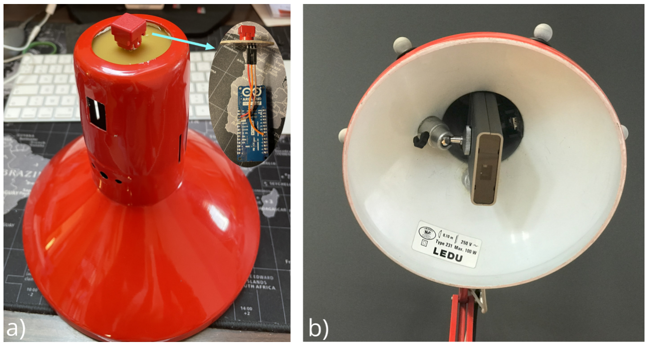

Figure 6: System hardware: (a) the standard lamp switch is replaced with an input button (b) a pico laser projector mounted inside the shade, and motion tracking markers attached on the lower bezel of the lamp shade

typically their index or thumb, to press the button while the other four fingers hold the lamp firmly. Next, we installed six reflective markers around the lower bezel of the lamp shade, which were tracked by an OptiTrack system[2]. These markers formed a rigid body used to update the position and orientation of the lamp in the system. Further details on environment reconstruction and object registration and tracking are provided below.

4.2 Environment Setup and Object Registration

We employed the OptiTrack system, featuring four cameras, to track and partially reconstruct the desktop workspace (Figure 7). Markers were placed on the wall, desk, and monitor to define their location and associate virtual objects. We built a custom pen equipped with three markers (see Figure 8f) and a unique OptiTrack ID to track its position and orientation. The tip of the pen was then used to register the pose and dimensions of static objects (*e.g.* a desk keyboard) by placing the tip at the outline of them and marking its positions. For objects with mobility, such as the lamp, smartphone, and ring, markers were attached and assigned unique IDs if possible, allowing continuous tracking of their position and orientation.

4.3 Touch Tracking and Detection

There is a challenge with the single marker on the fingertip, as it receives a different ID each time the cameras lose tracking due to occlusion or when dummy markers appear. To mitigate this issue, we implement a frame-by-frame exclusion of tracked points registered for moving objects (the lamp, the phone, and the pen). Then, when the ID for the index finger marker is missing, we reassign that ID to the closest individual marker that is below a 10 cm range of the previously known position of the index. If no point is found, we keep the previously known position of the index.

To accommodate different thicknesses and poses of the index finger, we perform a calibration phase for each participant. Users place their finger in contact with the desk to simulate a touch interaction and keep that pose for at least 5 seconds. The average distance between the marker and the desk surface is measured during that time to define a threshold distance used to consider whether there is a touch with the surface.

[2]https://optitrack.com/

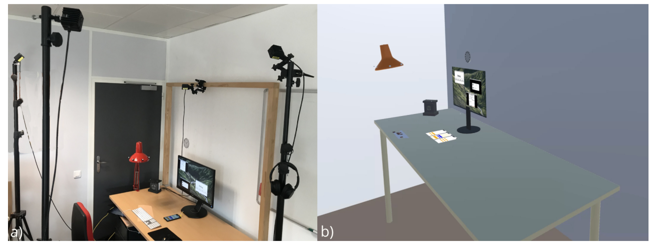

Figure 7: (a) desk space with four OptiTrack cameras and physical objects and devices. (b) reconstructed space in Unity.

4.4 Implementation

Our system was developed in Unity 2021.3.8f1, using Motive software 1.10.2 and the corresponding OptiTrack Unity Plugin 1.4.0. Camera settings in Motive included a frame rate of 180, exposure of 40, brightness threshold of 254, and LED illumination of 15. Cameras used low-range gain, infrared spectrum filter, and precision mode. While the physical lamp's position and orientation were tracked, the laser projector within remained hidden, with its exact orientation and position unknown. For consistency between virtual and physical environments, the laser projector and virtual camera were treated as identical. We fine-tuned the virtual camera's sensor sizes and offsets within Unity, aligning the calibration using the CS-200 calibration square of the OptiTrack system. This manual calibration process ended when the virtual and physical markers were approximately aligned in position, size, and orientation.

Our software, a virtual desktop that manages application windows on various surfaces, featured independent WebViews for each virtual window. These were supported by the Embedded Browser package [3]. Windows adapted their content based on factors such as current location (*e.g.* the monitor, desk, wall or ceiling), awareness of the physical environment (*e.g.* clock, speaker, phone), lamp height, and user input (*e.g.* lamp, mouse, pen, touch). Applications were implemented outside of Unity and hosted on a local server. Virtual window content was displayed based on the URLs provided and responded to events triggered by user input, and communication between the lamp system and the phone, including touch events, was established over a local network.

5 USAGE SCENARIOS AND APPLICATIONS

To demonstrate the supported interactions, we provide potential usage scenarios with three applications: a calendar, a music player, and an architectural design drawing tool. The accompanying video provides full demonstrations.

Calendar. Alice uses LuxAR to streamline her calendar management. By pointing at the calendar window on the monitor, she can easily drag it out, placing it on her desk with a press-and-hold button action. The height adjustment feature of the lamp allows Alice to access various levels of detail in the calendar window. At a medium height, she views the weekly schedule (Figure 8a). Bringing the lamp closer reveals the daily schedule (Figure 8b), while raising it higher displays the monthly view. Alice moves the calendar window onto the wall, showcasing schedules to her colleagues.

[3]https://assetstore.unity.com/packages/tools/gui/embedded-browser-55459

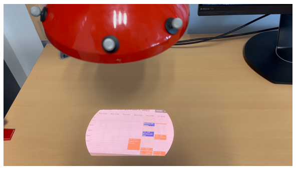

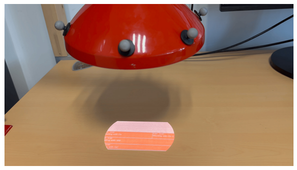

(a) Calendar (medium height) **(b) Calendar (low height)**

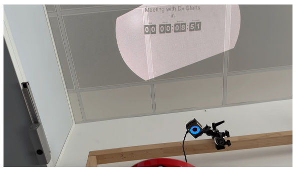

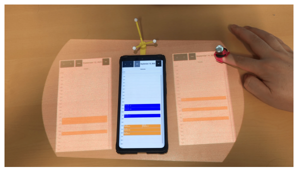

(c) Calendar on ceiling **(d) Calendar on phone**

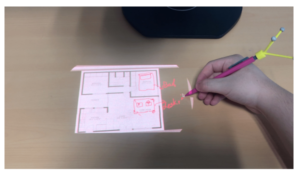

(e) Music player on speaker **(f) Pen-input under the lamp**

Figure 8: Applications used in demonstration scenarios.

Alternatively, placing the window on a clock creates a visual timeline. During breaks, she moves it to the ceiling, turning it into a countdown timer for the next event, reminding others and herself (Figure 8c). Alice can edit the calendar events on the lamp display using her mouse or finger. Hovering the lamp over the keyboard reveals the calendar application shortcuts. She can also transfer the calendar to a nearby mobile device, to better manage the events on her phone (Figure 8d).

Music Player. To relax, Alice plays music and moves the player to the surface using LuxAR. To choose a new song, she adjusts the lamp's height to explore different interfaces: at a medium height, she sees a list of songs with covers, or she can use buttons to loop through the list when the lamp is low. To share music publicly, she moves the player window to a nearby sound speaker, which offers music controls for volume and song selection (Figure 8e). To enhance the environment, Alice can also place the music player on the ceiling to create ambient visualizations that sync with the music.

Drawing Annotation Tool. Later, LuxAR assists Alice with annotation tasks, such as drawing on a floor plan (Figure 8f). She can move her digital workspace to a surface by pulling a floor-plan window from the monitor. Alice can switch between different floor plan views by changing the lamp's height: a medium height reveals the entire floor, lowering it unveils a bedroom-scale view, and raising it presents a 3D-rendered floor. Alice enhances her workflow by annotating directly on the surface with a pen, combining the benefits of digital and physical interaction.

6 USER STUDY

We conducted a two-phase user study to evaluate our system with the proposed interactions. In Phase 1, we evaluated participants' understanding and execution of these interactions. In Phase 2, we invited them to explore various scenarios. We further collected their feedback to identify ways to improve our system and enhance the current desktop experience.

6.1 Participants

We recruited 12 participants from a local institution, aged 20 to 33 (3 identify as women). 11 participants were right-handed and 1 was ambidextrous. 5 participants answered "yes" and 2 "maybe" on whether they have experience with architect lamps, and 2 reported using such lamps daily and weekly respectively. The study took approximately 60 minutes.

6.2 Procedure

The study consisted of three stages and employed a think-aloud protocol encouraging participants to provide immediate commentary. All comments and feedback were audio-recorded for subsequent analysis. The emphasis was on participant feedback and experience with LuxAR to extend desktop environments, rather than assessing task completion speed or efficiency.

6.2.1 Introduction. After reading the information letter and signing the consent form, participants were invited to adjust the seat and have a comfortable setting in the desktop space (Figure 7). Then, they were introduced to the study agenda, how our system was built, and how it could be used to manipulate virtual windows on the physical display or surfaces. Participants were allowed to manipulate the lamp and other devices to get a sense of the system.

6.2.2 Phase 1. In the first phase, participants engaged in application windows that display generic information and manipulated them within the desk space using the lamp. Each window featured a central icon, two navigation buttons to change icons, and a slider to adjust the background color. Additionally, windows provided information about the names of the object and the surface the lamp pointed at, and a distance indication categorized into three levels (small, medium, and high). Participants were assigned 15 tasks, which included activities such as changing icons and background colors on the display and surfaces using the mouse or touch and moving windows across surfaces and displays using the lamp. Following the completion of the task, the participants assessed physical and mental demands on a 10-point scale (1—very low, 10—very high). These two scales are a subset of the NASA TLX, using only two focused our investigation and reduced the burden for participants. Subsequently, open-ended interviews were conducted to collect feedback in depth on the manipulations and interactions. See also the accompanying materials.

6.2.3 Phase 2. In the second phase, participants used the interaction techniques acquired in the previous phase to explore the applications and scenarios described in section 5. Following the exploration, participants were interviewed to share their experiences with the applications, offering insights into potential system enhancements or scenario improvements.

6.3 Data Collection

We observed and noted how participants used our system to complete tasks during the study. We transcribed audio recordings and extracted comments from each question, as well as comments during each interaction. These notes were organized by question in a visual diagramming tool. To analyze the data, we applied a content analysis approach [8] to examine participant feedback on the direct manipulation of the lamp and associated interaction designs. Specifically, the main author analyzed the semantics of the data, grouped them by question, and then assigned them to each interaction design category. This analysis could help identify potential design considerations for designing explicit and direct manipulation interfaces for lamp-based systems that extend desktop computing.

6.4 Results

Participants (9 out of 12) found the system engaging and highlighted its ability to extend the current desktop space as *"...you are not trapped in your screen and have more spaces"* [P8] and its *"concepts are novel and practical"* [P10]. Participants also highlighted areas for improvement, particularly in the mechanical design of the lamp. They pinpointed specific challenges related to certain axes and recommended that enhancing degrees of freedom, coupled with increased lubrication in some joints, could further improve their experience. These observations were reflected in the self-evaluated mental demand (M = 3.33, SD = 2.01) and physical demand (M = 5.42, SD = 2.27).

6.4.1 Intuitive manipulation across surfaces, but angles matter. The reposition of windows on the table was considered straightforward, but challenges arose when moving them to the wall or ceiling. Participants noted that such manipulation *"requires a lot of body movement"* [P4] and that *"[manipulation] angles matter"* [P7] with awkward wrist positions for wall and ceiling interactions. As commented by P12, *"For some angles, it may be relatively easy to handle, but for others, it may not be so straightforward ... I feel that it may be related to how I grasp the lamp."* We also observed that in Phase 1, six participants accidentally dropped a window behind the monitor during table-to-ceiling manipulations, an issue not observed in other across-surface manipulations, and none occurred in Phase 2. This showed how direct manipulation angles potentially impacted content repositioning to extend desktop computing onto diverse physical surfaces.

Participants were also positive about using the lamp to transfer virtual information between the monitor, physical surfaces, and mobile devices, which was described as *"neat from a digital point of view"* [P8], and *"context-wise interesting"* [P7]. However, concerns arose regarding the repositioning from a different source to a phone. Some participants saw it as a means to *"transfer data from computers to phones"* [P2], while others found the window's shape inconsistent on the phone, prompting questions about its purpose (P5, P7). P10 believed it existed a discontinuity between surfaces and devices when the same information was splitted across them, as shown in Figure 8d and *"...the information should be contained in one space."*

Only four participants noticed the visual change of the style on the lamp display as *"area is different"* [P7] but *"I did not know why it is happening"* [P4, P12] and three participants used the radar mode to find missing windows during the interaction without noticing other modes (hover, focus, and drag). Participants' focus was mainly on the manipulated windows, as *"I don't notice it and focus on the app as long as I can see"* [P6] and *"...I think I perform the manipulation earlier than noticing the visual effect"* [P11].

6.4.2 Surface adaptation influenced by content and context. The adaptability to surface changes was prominently influenced by the content and contexts of manipulated applications, particularly exemplified in scenarios such as the calendar and music player. In the calendar example, P7 and P10 noted privacy concerns when asked to share the calendar with the public on the wall, but none raised this for the music player example. Besides, the ceiling was deemed unsuitable for the calendar by participants who rarely looked at it, as *"I really don't look at the ceiling at all and won't notice it"* [P3]. However, the adaptation was well received for the music player by participants: P1 expressed enthusiasm about the music player's ceiling visualization, envisioning it as *"super cool"* for wall visualizations while working, and P3 thought the visualization created a mood in an office, *"like in concert or club setting"*. Moreover, the adaptation should extend beyond the content within the window to encompass window sizes, particularly on larger surfaces, as P3 conveyed dissatisfaction with uniform window sizes on both a table and a wall.

6.4.3 Varied opinions on interactions based on height adjustment. Similarly to the results in §6.4.2, participants perceived the ability to change details and content through the manipulation of the lamp's height differently based on tasks and application contexts. Initially, eight participants were perplexed by this interaction in Phase 1, with P4 expressing *"this provided new possibility to interact with the lamp but did not provide clear use cases"*. However, in Phase 2, participants favored the floor plan example to illustrate this interaction, with positive comments such as *"Floor plan is cooler... it makes more sense to me"* [P8], and P2 adding, *"You can see different aspects, and this makes sense for the architect lamp.""* For the calendar example, while participants understood the concept, the limited lamp display size prevented them from reading the information, showing only a restricted part of the window when the lamp reached a lower level. Comments included *"It is more private, but not convenient. I don't like the view to be so limited ... have to go back and forth to see everything"* [P3], *"It is disturbing to move the lamp to explore the calendar because only parts of the calendar are shown"* [P2], and *"The closer the lamp is, the less information is shown. I have to move"* [P7]. Moreover, adjusting the height posed a challenge when pushing the lamp to a lower level and maintaining that pose. P8 highlighted this issue, stating, *"The concept is easy to understand. Just the mechanical parts... you could not keep it low."* This observation raises ergonomic concerns when designing interactions based on surface proximity for lamps.

6.4.4 Object augmentation enjoyable, but tricky to trigger. Participants positively responded to the augmentation of physical objects like the clock and speaker, finding it interesting and enjoyable, such as *"music player is funny with the speaker"* [P2]. While augmenting the phone with additional views was favored, as *"this is nice to have extended views for phones"* [P8], opinions about using the lamp to complete the information transmission from different surfaces to the phone varied, as mentioned in §6.4.1. While participants enjoyed the augmentations, they encountered challenges when triggering them and sought the help of the study facilitator. Although they grasped the concept of picking up a window and moving it to objects, they struggled with the drop action, specifically, where and when to release the button. For instance, some participants were observed to wait for the interaction to activate automatically or try to induce it by pressing the lamp downward. As the release action was non-intuitive, P4 and P8 suggested adding visual cues and feedback around the physical objects to guide the release action. P12 also suggested to increase the boundary for object detection.

6.4.5 Challenges with detached interactions. Participants were tasked with changing the icon and background color of the window on using the mouse, touch and mobile device. Although the use of the mouse on various surfaces was generally perceived as straightforward, several participants noted challenges in tracking the cursor as the lamp moved across surfaces (P2, P7, P11) or in low resolution within the lamp display (12) and P8 reported misaligned mouse coordinates between the lamp display and the application due to a tilted lamp angle. Only half participants were positive about touches on surfaces. P4 preferred to *"perform pinch with touches [on surfaces]"* but for *"simple actions, I would go back to the mouse and keyboard"*. P5, who was used to touchscreen interfaces, enjoyed interacting with the same content on physical surfaces. In contrast, P9 noticed a latency for touches on surfaces and P11 felt that *"...the sensitivity is not the same [for touches on screens and surfaces]"*. The introduction of the pen tool during the floor plan example did not elicit additional feedback from participants. They focused on how the height adjustment of the lamp influenced the content, implicitly reflecting the naturalness of using the pen with LuxAR.

6.4.6 Other Possible scenarios. Participants suggested various applications for LuxAR in both phases. In Phase 1, eight participants envisioned a collaborative environment, including multiplayer games and information sharing across devices. Two focused on information management with applications for maps, calendars, and drawings. One participant suggested using virtual objects for decoration, revealing them with the lamp and one believed that our system could be quite useful for parts assembly scenario when instructions can be superimposed onto objects and users can use both hands to achieve assembly tasks. In Phase 2, two participants proposed context-aware applications. P1 recommended a smart home control concept, suggesting, *"...put everything on the clock to show time-related events..."*. P4 suggested to *"...show the content depending on where the lamp is actually located... closer to the bed... closer to the desk or in the kitchen"*. P9 believed that it can be used for a football match such as *"pointing at a player and I can see the player's statistics for this match"* and P12 planned to *"show API documents when I point at the code in a window and then move the documents on other space"*. Two participants suggested replacing application windows with information windows to aggregate data from various sources.

7 DISCUSSION

Throughout our two-phase study, participants demonstrated a clear understanding of how to manipulate LuxAR to explore different scenarios and grasp the interaction designs centered on an architect lamp. They found it intuitive to use the lamp as both an input and output device to extend and augment desktop computing environments. This ease of use may be attributed to the design of the direct input, where our lamp is employed to move content on surfaces, akin to mouse input in desktop computing: button-based interactions closely resemble mouse clicks, while physical lamp manipulations mirror mouse movements on a surface. Although some participants faced challenges in operating the button in specific lamp orientations (*e.g.* moving the window from the table to the ceiling), mental demand was low overall, and none of the participants had difficulty learning the manipulations and interactions.

Our results also showed that participants had varied feedback on how projected content adapted to lamp heights. For instance, participants favored this interaction to draw on the floor plan compared to using it with a calendar. We also noticed that the ceiling was less favored compared to other surfaces. These observations suggest that the adaptation of virtual content to dynamic surfaces and the direct manipulation of a SAR system could be surface- and context-dependent or a combination of both. The limited use of the ceiling may also stem from the lack of a tool for displaying information on it. These open up a broader design space for directly manipulated SAR, offering potential avenues for future research.

Although understanding the interaction is straightforward, uncovering all available interactions can be a bit tricky. Our results revealed that while manipulating the system, participants primarily focused on the manipulated window, so the majority of participants neglected the visual changes in the lamp display, which help to signify the interaction state of the window. Meanwhile, they faced challenges triggering augmentations on physical objects, highlighting the necessity to amplify visual feedback around the physical objects when they were under the lamp. This indicated that rather than changing the projection display style, placing obvious interaction indicators around virtual windows or physical objects was more visually perceivable.

7.1 Design Implications

Our findings indicate that using an architect lamp as both input and output device is intuitive, but some aspects of its manipulation may be challenging. Additionally, designing attached and detached interactions for direct manipulation systems requires careful thought to accommodate different contexts. We revisit our design goals and suggest further considerations for explicit and direct manipulation of lamp-based systems, and SAR more broadly, when integrating desktop computing into physical spaces.

Mechanical Designs. Compared with previous efforts, such as the Lantern [19] and LuminAR [22], our system explores an architect lamp's direct manipulation potential to interact with the content across different spaces **[DG1, DG2]**. It also leverages the unique mechanical structure of the lamp to offer stable positioning and persistent augmentation without requiring continuous user engagement. This allows detached interactions and inputs from the lamp's manipulation, such as mouse, touch, pen and devices **[DG4]**. However, as the lamp is used as an direct manipulation input, participants encountered challenges in operating the lamp, desiring more degrees of freedom (DoFs) to address issues with specific axes. It is also noted that architect lamps are predominantly designed to face downward. Consequently, unlike handheld systems [2, 7, 12, 42, 46] employing the torch form, certain axes and placements of our system were challenging to execute during the study due to the lamp not being explicitly designed for such movements. Additionally, the fixed placement of the lamp on the desk limits mobility.

To address these concerns in architect lamp-based SAR systems, potential solutions include attaching wheels for easy desk movement and replacing current joints with spherical joints or 6-Axis force-torque sensors with locking capabilities for improved flexibility in positioning and orientation **[DG1, DG2]**. Future designs should balance the trade-offs between mobility and stability **[DG1, DG4]**, considering the specific form factors of SAR systems.

Attached Interactions with Spatial Awareness. Direct manipulation of the lamp requires some degree of spatial awareness to effectively interact with content. When moving a window across surfaces, we show how the content can adapt, offering a new style of interaction compared to other direct manipulation pro-cam systems systems [5, 7, 42]. Explicit control of augmentation also allows users to decide where to display information and what information to display. This makes interaction more focused and avoids being distracted by augmented information displayed elsewhere with a multiple pro-cam system simultaneously displaying content on many nearby surfaces. However, this may also limit people from browsing information outside the chosen augmented area (*e.g.* consider P3 when reading the calendar) and could slow down workflows spanning multiple sources of spatial information. Note many single pro-cam systems face the same challenges [5, 7, 15, 42].

Our results suggest that ceiling usage was less favored, possibly due to users finding it irrelevant when seated. Despite this, the ceiling could serve as ambient cues for conveying messages **[DG2, DG3]**, *e.g.* a countdown event could inform Alice's colleague. Moreover, while Tomitsch et al. [37] suggested ceiling use for information visualization and Satkowski et al. [32] recommended low visual complexity content on the ceiling in an HMD setting, most prior manipulation-based SAR systems did not consider the ceiling as an interaction space. This creates an opportunity to explore ceiling interactions, with LuxAR serving as a tool for such exploration. For example, parents and children can lie on the bed to use architect lamps around their table to control multiple characters, as in [41, 42] or tell bed stories, as in [43].

Our system employs three proximity levels for information granularity [7, 36]. Content remains consistently sized within each level, zooming with lamp height adjustments and crossing proximity levels alters the context displayed. Our findings show that the participants preferred to use this interaction to draw floor plans, while they encountered limitations in searching and reading information on the calendar. This suggests a potential context-dependent adaptation for height-adjusted interactions in the design of direct manipulation-based SAR systems — less suitable for content with high text density **[DG1, DG2]**. Additionally, our prototype employed consistent values for desks and walls, omitting this interaction for the ceiling due to the available manipulation axes. Future research could explore this area, broadening this interaction based on varied combinations of manipulations and surfaces with a more flexible manipulation-based lamp SAR system.

It is also important to effectively inform users about the interaction state and guide them on when, where, and how to initiate interactions, especially when interactions are attached to the direct manipulation of the system itself **[DG1]**. While we propose changing the lamp display style to indicate the interaction state of a content, future work should consider applying the visual signifiers to the virtual content and physical objects and devices directly and explore how different designs and visual significance could better guide users in direct manipulation-based interactions.

Detached Interactions with Stationary Lamp. When LuxAR left untouched, it maintains the augmentation and allow users to interact with the content with other inputs **[DG4]**. Although user can directly annotate information by touch and pen on the table, they are not suitable for the walls and ceiling, and the mouse is generally preferred but participants encountered tracking challenges. While cursor loss is a common issue [23], enhancing cursor visibility during manipulation, as previously discussed, is a potential solution for SAR systems using direct manipulation. Additionally, considering coordinate remapping between the manipulated display and the content can assist users in performing effective mouse interactions across different surfaces. While a mobile device support touch inputs, but its high resolution touchscreen make it a better holder for content. Therefore, future design on Temporary View could split up the information and control: placing information on a device and moving control and interactions onto surfaces lighted by the lamp.

8 CONCLUSION

Our work presents LuxAR, a directly manipulated SAR prototype that uses an architect desk lamp as an input and output tool, extending desktop computing to various surfaces in the desk space. By leveraging the lamp's manipulation capabilities, users can manage virtual content across physical displays, objects, surfaces, and mobile devices, interacting with it through the lamp and three other inputs. Our user study demonstrates that participants understand the design of our lamp system and use it to explore three demonstration examples, and showcases its potential to enhance the capabilities and accessibility of desktop computing environments.

Future work could focus on exploring bi-manual interactions using the lamp, alongside input modalities like the pen and lamp input [7], across diverse physical surfaces. Integrating sensors into each joint could further broaden the design space for physically manipulating architect lamp-based SAR systems. While our system employed the OptiTrack system for rapid prototyping, its cost and portability limitations in other environments suggest a potential future direction: developing a self-contained LuxAR using pro-cam and IMU systems, supported by SLAM algorithms [6] for pose, depth, and trajectory estimation. A self-contained system, complemented by a compatible toolkit, would empower practitioners to design and enhance personalized desktop experiences.

ACKNOWLEDGMENTS

Nikhita Joshi and Alessandra Luz for help improving readability, our participants for volunteering to do the study, and reviewers for valuable feedback. Made possible by the LAI Réapp, NSERC Discovery

Grant 2018-05187, and Canada Foundation for Innovation Infrastructure Fund 33151 “Facility for Fully Interactive Physio-digital Spaces”. Study protocol received clearance from the University of Waterloo Office of Research Ethics.

REFERENCES

[1] Mark Ashdown and Peter Robinson. 2004. A personal projected display. In *Proceedings of the 12th Annual ACM International Conference on Multimedia* (New York, NY, USA) *(MULTIMEDIA '04)*. Association for Computing Machinery, New York, NY, USA, 932–933. https://doi.org/10.1145/1027527.1027739

[2] P. Beardsley, J. van Baar, R. Raskar, and C. Forlines. 2005. Interaction using a handheld projector. *IEEE Computer Graphics and Applications* 25, 1 (2005), 39–43. https://doi.org/10.1109/MCG.2005.12

[3] Xiaojun Bi, Tovi Grossman, Justin Matejka, and George Fitzmaurice. 2011. Magic Desk: Bringing Multi-Touch Surfaces into Desktop Work. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems* (Vancouver, BC, Canada) *(CHI '11)*. ACM, New York, NY, USA, 2511–2520. https://doi.org/10.1145/1978942.1979309

[4] Verena Biener, Snehanjali Kalamkar, Negar Nouri, Eyal Ofek, Michel Pahud, John J. Dudley, Jinghui Hu, Per Ola Kristensson, Maheshya Weerasinghe, Klen Čopič Pucihar, Matjaž Kljun, Stephan Streuber, and Jens Grubert. 2022. Quantifying the Effects of Working in VR for One Week. *IEEE Transactions on Visualization and Computer Graphics* 28, 11 (2022), 3810–3820. https://doi.org/10.1109/TVCG.2022.3203103

[5] G. Blasko, F. Coriand, and S. Feiner. 2005. Exploring interaction with a simulated wrist-worn projection display. In *Ninth IEEE Int. Symposium on Wearable Computers (ISWC'05)*. 2–9. https://doi.org/10.1109/ISWC.2005.21

[6] Carlos Campos, Richard Elvira, Juan J. Gómez Rodríguez, José M. M. Montiel, and Juan D. Tardós. 2021. ORB-SLAM3: An Accurate Open-Source Library for Visual, Visual–Inertial, and Multimap SLAM. *IEEE Transactions on Robotics* 37, 6 (2021), 1874–1890. https://doi.org/10.1109/TRO.2021.3075644

[7] Xiang Cao and Ravin Balakrishnan. 2006. Interacting with Dynamically Defined Information Spaces Using a Handheld Projector and a Pen. In *Proc. of the 19th Annual ACM Symposium on User Interface Software and Technology* (Montreux, Switzerland) *(UIST '06)*. ACM, New York, NY, USA, 225–234. https://doi.org/10.1145/1166253.1166289

[8] Satu Elo and Helvi Kyngäs. 2008. The qualitative content analysis process. *Journal of advanced nursing* 62, 1 (2008), 107–115.

[9] Alex Endert, Lauren Bradel, Jessica Zeitz, Christopher Andrews, and Chris North. 2012. Designing Large High-Resolution Display Workspaces. In *Proc. of the Int. Working Conf. on Advanced Visual Interfaces* (Capri Island, Italy) *(AVI '12)*. ACM, New York, NY, USA, 58–65. https://doi.org/10.1145/2254556.2254570

[10] John Hardy. 2012. Experiences: A Year in the Life of an Interactive Desk. In *Proc. of the Designing Interactive Systems Conf.* (Newcastle Upon Tyne, United Kingdom) *(DIS '12)*. ACM, New York, NY, USA, 679–688. https://doi.org/10.1145/2317956.2318058

[11] Koike Hideki, Sato Yoichi, and Yoshinori Kobayashi. 2001. Integrating Paper and Digital Information on EnhancedDesk: A Method for Realtime Finger Tracking on an Augmented Desk System. *ACM Trans. Comput.-Hum. Interact.* 8, 4 (dec 2001), 307–322. https://doi.org/10.1145/504704.504706

[12] Paul Hine, Loza Tadesse Mamo, and Narcis Pares. 2022. AR Magic Lantern: Group-Based Co-Located Augmentation Based on the World-as-Support AR Paradigm. In *Extended Abstracts of the 2022 CHI Conf. on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI EA '22)*. ACM, New York, NY, USA, Article 207, 5 pages. https://doi.org/10.1145/3491101.3519918

[13] Steven Houben, Paolo Tell, and Jakob E. Bardram. 2014. ActivitySpace: Managing Device Ecologies in an Activity-Centric Configuration Space. In *Proc. of the Ninth ACM Int. Conf. on Interactive Tabletops and Surfaces* (Dresden, Germany) *(ITS '14)*. ACM, New York, NY, USA, 119–128. https://doi.org/10.1145/2669485.2669493

[14] Brett Jones, Rajinder Sodhi, Michael Murdock, Ravish Mehra, Hrvoje Benko, Andrew Wilson, Eyal Ofek, Blair MacIntyre, Nikunj Raghuvanshi, and Lior Shapira. 2014. RoomAlive: Magical Experiences Enabled by Scalable, Adaptive Projector-camera Units. In *Proc. of the 27th Annual ACM Symposium on User Interface Software and Technology* (Honolulu, Hawaii, USA) *(UIST '14)*. ACM, New York, NY, USA, 637–644. https://doi.org/10.1145/2642918.2647383

[15] Nikhita Joshi, Antony Albert Raj Irudayaraj, Jeremy Hartmann, and Daniel Vogel. 2022. An Instrumented Office Chair with a Steerable Projector for Personal Spatial Augmented Reality. In *Symposium on Spatial User Interaction* (Online, CA, USA) *(SUI '22)*. ACM, New York, NY, USA, Article 1, 12 pages. https://doi.org/10.1145/3565970.3567705

[16] Sasa Junuzovic, Kori Inkpen, Tom Blank, and Anoop Gupta. 2012. IllumiShare: Sharing Any Surface. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems* (Austin, Texas, USA) *(CHI '12)*. ACM, New York, NY, USA, 1919–1928. https://doi.org/10.1145/2207676.2208333

[17] Daekun Kim, Nikhita Joshi, and Daniel Vogel. 2023. Perspective and Geometry Approaches to Mouse Cursor Control in Spatial Augmented Reality. In *Proc. of the 2023 CHI Conf. on Human Factors in Computing Systems* (Hamburg, Germany) *(CHI '23)*. ACM, New York, NY, USA, Article 692, 19 pages. https://doi.org/10.1145/3544548.3580849

[18] Jeongyun Kim, Jonghoon Seo, and Tack-Don Han. 2014. AR Lamp: Interactions on Projection-Based Augmented Reality for Interactive Learning. In *Proc. of the 19th Int. Conf. on Intelligent User Interfaces* (Haifa, Israel) *(IUI '14)*. ACM, New York, NY, USA, 353–358. https://doi.org/10.1145/2557500.2557505

[19] Lantern 2022. *Lantern*. Retrieved Novemeber 22, 2022 from https://nordprojects.co/lantern/

[20] Zhen Li, Michelle Annett, Ken Hinckley, Karan Singh, and Daniel Wigdor. 2019. HoloDoc: Enabling Mixed Reality Workspaces That Harness Physical and Digital Content. In *Proc. of the 2019 CHI Conf. on Human Factors in Computing Systems* (Glasgow, Scotland Uk) *(CHI '19)*. ACM, New York, NY, USA, 1–14. https://doi.org/10.1145/3290605.3300917

[21] Chunyuan Liao, Hao Tang, Qiong Liu, Patrick Chiu, and Francine Chen. 2010. FACT: Fine-Grained Cross-Media Interaction with Documents via a Portable Hybrid Paper-Laptop Interface. In *Proc. of the 18th ACM Int. Conf. on Multimedia* (Firenze, Italy) *(MM '10)*. ACM, New York, NY, USA, 361–370. https://doi.org/10.1145/1873951.1874001

[22] Natan Linder and Pattie Maes. 2010. LuminAR: Portable Robotic Augmented Reality Interface Design and Prototype. In *Adjunct Proc. of the 23nd Annual ACM Symposium on User Interface Software and Technology* (New York, New York, USA) *(UIST '10)*. ACM, New York, NY, USA, 395–396. https://doi.org/10.1145/1866218.1866237

[23] Chuanyi Liu and Rui Zhao. 2018. Find the 'Lost' Cursor: A Comparative Experiment of Visually Enhanced Cursor Techniques. In *Intelligent Computing Theories and Application*, De-Shuang Huang, Kang-Hyun Jo, and Xiao-Long Zhang (Eds.). Springer Int. Publishing, Cham, 85–92.

[24] Kosuke Maeda, Mitski Piekenbrock, Toshiki Sato, and Hideki Koike. 2018. A Tabletop System Using an OmniDirectional Projector-Camera. In *Proc. of the 2018 ACM Int. Conf. on Interactive Surfaces and Spaces* (Tokyo, Japan) *(ISS '18)*. ACM, New York, NY, USA, 311–314. https://doi.org/10.1145/3279778.3279911

[25] Leonardo Pavanatto, Chris North, Doug A. Bowman, Carmen Badea, and Richard Stoakley. 2021. Do we still need physical monitors? An evaluation of the usability of AR virtual monitors for productivity work. In *2021 IEEE Virtual Reality and 3D User Interfaces (VR)*. 759–767. https://doi.org/10.1109/VR50410.2021.00103

[26] Claudio Pinhanez. 2001. The Everywhere Displays Projector: A Device to Create Ubiquitous Graphical Interfaces. In *Ubicomp 2001: Ubiquitous Computing*, Gregory D. Abowd, Barry Brumitt, and Steven Shafer (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 315–331.

[27] ProjectFLX 2021. *ProjectFLX*. Retrieved January 30, 2023 from https://lightform.com/lfx

[28] Roman Rädle, Hans-Christian Jetter, Nicolai Marquardt, Harald Reiterer, and Yvonne Rogers. 2014. HuddleLamp: Spatially-Aware Mobile Displays for Ad-Hoc Around-the-Table Collaboration. In *Proc. of the Ninth ACM Int. Conf. on Interactive Tabletops and Surfaces* (Dresden, Germany) *(ITS '14)*. ACM, New York, NY, USA, 45–54. https://doi.org/10.1145/2669485.2669500

[29] Stefan Rapp, G Michelitsch, M Osen, J Williams, M Barbisch, R Bohan, Z Valsan, and M Emele. 2004. SPOTLIGHT NAVIGATION: INTERACTION WITH A HANDHELD PROJECTION DEVICE.

[30] Ramesh Raskar, Greg Welch, Matt Cutts, Adam Lake, Lev Stesin, and Henry Fuchs. 1998. The Office of the Future: A Unified Approach to Image-Based Modeling and Spatially Immersive Displays. In *Proc. of the 25th Annual Conf. on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. ACM, New York, NY, USA, 179–188. https://doi.org/10.1145/280814.280861

[31] Jun Rekimoto and Masanori Saitoh. 1999. Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems* (Pittsburgh, Pennsylvania, USA) *(CHI '99)*. ACM, New York, NY, USA, 378–385. https://doi.org/10.1145/302979.303113

[32] Marc Satkowski, Rufat Rzayev, Eva Goebel, and Raimund Dachselt. 2022. ABOVE & BELOW: Investigating Ceiling and Floor for Augmented Reality Content Placement. In *2022 IEEE Int. Symposium on Mixed and Augmented Reality (ISMAR)*. 518–527. https://doi.org/10.1109/ISMAR55827.2022.00068

[33] Albrecht Schmidt. 2000. Implicit human computer interaction through context. *Personal technologies* 4 (2000), 191–199.

[34] Dominik Schmidt, David Molyneaux, and Xiang Cao. 2012. PICOntrol: Using a Handheld Projector for Direct Control of Physical Devices through Visible Light. In *Proc. of the 25th Annual ACM Symposium on User Interface Software and Technology* (Cambridge, Massachusetts, USA) *(UIST '12)*. ACM, New York, NY, USA, 379–388. https://doi.org/10.1145/2380116.2380166

[35] Hyunyoung Song, Tovi Grossman, George Fitzmaurice, François Guimbretiere, Azam Khan, Ramtin Attar, and Gordon Kurtenbach. 2009. PenLight: Combining a Mobile Projector and a Digital Pen for Dynamic Visual Overlay. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems* (Boston, MA, USA) *(CHI '09)*. ACM, New York, NY, USA, 143–152. https://doi.org/10.1145/1518701.1518726

[36] Martin Spindler, Sophie Stellmach, and Raimund Dachselt. 2009. PaperLens: Advanced Magic Lens Interaction above the Tabletop. In *Proc. of the ACM Int.*

Conf. on Interactive Tabletops and Surfaces (Banff, Alberta, Canada) *(ITS '09)*. ACM, New York, NY, USA, 69–76. https://doi.org/10.1145/1731903.1731920

[37] Martin Tomitsch, Thomas Grechenig, Andrew Vande Moere, and Sheldon Renan. 2008. Information Sky: Exploring the Visualization of Information on Architectural Ceilings. In *2008 12th Int. Conf. Information Visualisation*. 100–105. https://doi.org/10.1109/IV.2008.81

[38] Daniel Vogel and Ravin Balakrishnan. 2004. Interactive Public Ambient Displays: Transitioning from Implicit to Explicit, Public to Personal, Interaction with Multiple Users. In *Proc. of the 17th Annual ACM Symposium on User Interface Software and Technology* (Santa Fe, NM, USA) *(UIST '04)*. ACM, New York, NY, USA, 137–146. https://doi.org/10.1145/1029632.1029656

[39] Viveka Weiley and Matt Adcock. 2013. Drawing in the lamposcope. In *Proc. of the 9th ACM Conf. on Creativity & Cognition*. 382–383.

[40] Pierre Wellner. 1991. The DigitalDesk Calculator: Tangible Manipulation on a Desk Top Display. In *Proc. of the 4th Annual ACM Symposium on User Interface Software and Technology* (Hilton Head, South Carolina, USA) *(UIST '91)*. ACM, New York, NY, USA, 27–33. https://doi.org/10.1145/120782.120785

[41] Karl D.D. Willis, Ivan Poupyrev, Scott E. Hudson, and Moshe Mahler. 2011. SideBySide: Ad-Hoc Multi-User Interaction with Handheld Projectors. In *Proc. of the 24th Annual ACM Symposium on User Interface Software and Technology* (Santa Barbara, California, USA) *(UIST '11)*. ACM, New York, NY, USA, 431–440. https://doi.org/10.1145/2047196.2047254

[42] Karl D.D. Willis, Ivan Poupyrev, and Takaaki Shiratori. 2011. Motionbeam: A Metaphor for Character Interaction with Handheld Projectors. In *Proc. of the SIGCHI Conf. on Human Factors in Computing Systems* (Vancouver, BC, Canada) *(CHI '11)*. ACM, New York, NY, USA, 1031–1040. https://doi.org/10.1145/1978942.1979096

[43] Karl D. D. Willis, Takaaki Shiratori, and Moshe Mahler. 2013. HideOut: Mobile Projector Interaction with Tangible Objects and Surfaces. In *Proc. of the 7th Int. Conf. on Tangible, Embedded and Embodied Interaction* (Barcelona, Spain) *(TEI '13)*. ACM, New York, NY, USA, 331–338. https://doi.org/10.1145/2460625.2460682

[44] Andrew Wilson, Hrvoje Benko, Shahram Izadi, and Otmar Hilliges. 2012. Steerable Augmented Reality with the Beamatron. In *Proc. of the 25th Annual ACM Symposium on User Interface Software and Technology* (Cambridge, Massachusetts, USA) *(UIST '12)*. ACM, New York, NY, USA, 413–422. https://doi.org/10.1145/2380116.2380169

[45] Robert Xiao, Scott Hudson, and Chris Harrison. 2017. Supporting responsive cohabitation between virtual interfaces and physical objects on everyday surfaces. *Proc. of the ACM on Human-Computer Interaction* 1, EICS (2017), 1–17.

[46] Ruopeng Xu, Toshiki Sato, Shio Miyafuji, and Hideki Koike. 2022. OmniLantern: Design and Implementations of a Portable and Coaxial Omnidirectional Projector-Camera System. In *Proc. of the 2022 Int. Conf. on Advanced Visual Interfaces* (Frascati, Rome, Italy) *(AVI 2022)*. ACM, New York, NY, USA, Article 13, 5 pages. https://doi.org/10.1145/3531073.3531126

[47] Ye Yuan, Nathalie Riche, Nicolai Marquardt, Molly Jane Nicholas, Teddy Seyed, Hugo Romat, Bongshin Lee, Michel Pahud, Jonathan Goldstein, Rojin Vishkaie, Christian Holz, and Ken Hinckley. 2022. Understanding Multi-Device Usage Patterns: Physical Device Configurations and Fragmented Workflows. In *Proc. of the 2022 CHI Conf. on Human Factors in Computing Systems* (New Orleans, LA, USA) *(CHI '22)*. ACM, New York, NY, USA, Article 64, 22 pages. https://doi.org/10.1145/3491102.3517702