Privacy-Enabled Parallax Display

Category: Technical Papers

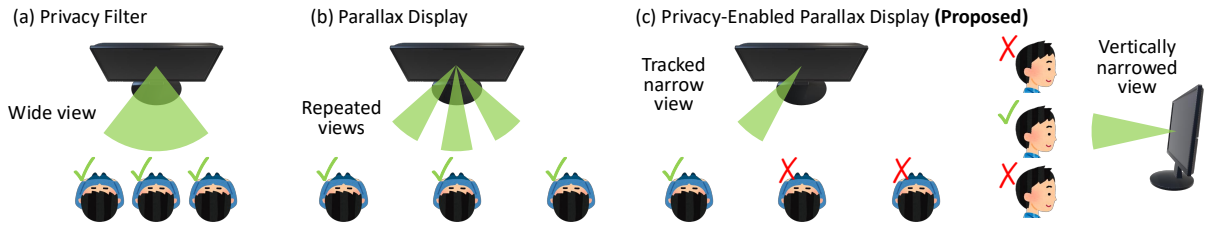

Figure 1: **Privacy-enabled parallax display. (a)** Existing privacy filters have wide viewing fields or limit user positioning. **(b)** One way to address this is to track the user for a dynamically adjusted, narrow viewing field centered around them. This real-time modulation of light fields can be achieved with a parallax display, which, however, suffers from view repetition. **(c)** Our *privacy-enabled* parallax display achieves high selectivity, ensuring *only the* tracked user sees the correct view. It also restricts the vertical viewing angle, preventing *shoulder surfing*, unlike conventional privacy filters.

Abstract

Privacy filters for displays are designed to obfuscate or hide visual content from unintended users, while making the displayed information visible only to selected viewers. Existing privacy filters suffer from either wide viewing field (low selectivity) or limited user positioning. To solve this dilemma, we propose a display technology that allows a narrow but adaptive viewing field that can be directed to arbitrary user location. While conventional parallax barriers provide such capability of modulating the light field according to the user location, it suffers from repeated views. Our key observation is that this view repetition originates from the periodicity of barrier patterns, and propose a privacy-enabled parallax display based on randomized barrier design. In addition to randomizing the locations of 1D slits, we also propose breaking down the slits into pinholes and randomizing their 2D locations, which results in privacy-preservation along the vertical direction as well. We build a hardware prototype using two off-the-shelf liquid-crystal displays. Experiments show that the proposed randomized parallax barrier can direct to the user a narrow viewing field of about $\pm 6^\circ$, providing a significantly improved privacy protection as compared to traditional privacy screens.

Index Terms: Privacy display, parallax display, parallax barrier.

1 Introduction

Privacy filters for displays, designed to protect confidential information from prying eyes, typically consist of micro-scale parallel slats that reduce light transmission as the viewing angle increases. Despite their widespread commercial use, these filters often provide a wide viewing field and limited selectivity, with the maximum viewing angle ranging from ± 30 to ± 45 degrees, leaving room for potential compromise of privacy as shown in Fig. 1(a). A critical challenge in their design lies in the delicate balance between privacy and usability: While it is feasible to create privacy filter with narrower viewing angles, it would confine the user to a limited range of positions, diminishing overall usability and user experience.

One potential solution to this tradeoff involves tracking the user and dynamically adjusting the viewing field accordingly, presenting a narrow viewing field centered around the user. To achieve this, real-time modulation of the emitted light field by the display is required. A cost-effective implementation of such a light field display is via parallax barriers or lenticular lenses. While these technologies can deliver the correct view to the user, they also simultaneously exhibit repeated views outside the intended viewing field (sometimes referred to as repeating "virtual lobes" [14]), compromising selectivity and privacy preservation, as illustrated in Fig. 1(b).

We propose a *privacy-enabled* parallax display, which achieves high levels of selectivity such that *only the* tracked user can access the correct view. Our key insight addresses the repetitive nature of views inherent in the periodic design of parallax barriers. Typically, a parallax barrier display utilizes an array of vertically placed slits to spatially multiplex multiple views periodically. As the user traverses from left to right by one period, they encounter the same image repeatedly, resulting in duplicated views. To disrupt this periodicity, our method involves *randomizing the locations of these slits*. Consequently, all locations, except those of the intended user, display a *scrambled* view, where each pixel originates from a random intended view. This effectively eliminates high-frequency content in the view, making it challenging or impossible to discern confidential text or intricate details from a peeking view. Furthermore, we observe that the parallax barrier need not be composed of slits; alternatively, the slits can be vertically replaced with randomly positioned *pinholes*. This further restricts the vertical viewing angle, preventing *shoulder surfing*–an aspect not achievable with existing privacy filters. Since the proposed technique is based on parallax barrier, it provides a privacy-enabled *3D viewing* experience as well.

To demonstrate this idea, we build a hardware prototype using two layers of off-the-shelf liquid-crystal displays (LCD), where the front layer is used as programmable pinholes. Due to this programmable capability, it is possible to apply time-multiplexing to the random pinhole patterns. This allows flexibly trading off the display's effective frame rate for achieving full-resolution of the displayed images (same as the native resolution of the monitors), addressing a key limitation of light-field displays which need to significantly sacrifice spatial resolution of the displayed content. Finally, the programmable design allows switching to a "normal" 2D display mode when privacy protection is not needed. Experiments show that the proposed randomized parallax barrier exclusively provides the correct view to the intended user within a narrow viewing field of only $\pm 6^\circ$. Since the method can direct the viewing field to be centered at arbitrary viewing angles, when combined with face tracking or eye tracking technologies, it can provide a significantly improved level of privacy without sacrificing usability.

Scope and contribution. We emphasize that our primary objective in this work is to protect privacy from curious bystanders

rather than prepared attackers. While it is conceivable that capturing a few images from different locations might empower skilled attackers to reconstruct the correct view, we assert that our approach is well-suited to meet the privacy needs of ordinary users in public places such as offices, cafes, and similar environments.

We also would like to highlight that, while various technologies have been developed over the years to achieve a steerable viewing cone, each technology has its own merits and limitations. Instead of trying to outperform all existing ones in all different aspects, we raised the following question: Can privacy be achieved as an "add-on" to existing commercial displays, with minimal hardware modifications and computational overhead? To our knowledge, among all autostereoscopic display technologies, only parallax displays, such as layered LCDs (e.g., Nintendo 3DS) or lenticular arrays (e.g., Asus Spatial Vision), have been successfully commercialized, which suggests its practicality particularly in terms of cost and ease of manufacturing. Given the recent growing interest in eye-tracked 3D parallax displays [1, 8], we believe it is worth exploring whether we can use such displays with enhanced privacy.

2 Related Work

Privacy displays. Privacy display technologies aim to prevent bystanders from witnessing confidential information displayed on the screen. Commercially-available privacy filters dim the display brightness at oblique horizontal viewing angles, but still allow a reasonably wide viewing angle range. [10] combines multiple sensors to detect the existence of screen peeping and adjust the brightness and contrast accordingly such that only the intended user can see the screen clearly. [2] utilizes the human vision and optical system properties to transform the display image into a discretized grid image that is hard to discern when viewed outside the desired angle range, which has the undesirable side-effect of low image quality for the primary user.

[19] combines polarized 3D glasses and polarized light source such that only the viewer wearing the glass can see a clear image. [16] combines a directional backlight and a spatial light modulator such that the angular viewing range can be adjusted dynamically. [15] proposes an electronically switchable privacy film that switches between a transparent filter and a privacy filter. Privacy displays based on parallax barriers have been studied before. To prevent screen peeping from different directions, [9] proposes a parallax barrier that can switch between vertical and horizontal directions, depending on the current orientation of the mobile device. To our knowledge, the proposed method is the only method that can achieve narrow viewing field and high selectivity (both horizontally and vertically), which, combined with recent eye-tracked parallax display technology, can support a user at any location in front of the display.

Parallax barriers for displays. The invention of parallax barrier can be dated back to over 100 years ago [4]. While the parallax barrier enables glass-free 3D (autostereoscopic) displays, it significantly reduces the spatial resolution and brightness of the display. Later, lenticular arrays were proposed to improve the brightness. [3] first proposes using an LCD as a parallax barrier to allow easy switching between 2D and 3D, or a combination of both. Sharp manufactured the first electronically switchable LCD-based parallax barrier [5]. Such technology led to the success of Nintendo 3DS, a commercial game console. [6] proposes electronically movable pinhole array which improves spatial resolution of the display via temporal multiplexing. More recently, parallax displays have been combined with a face/eye tracker to track the user location such that the spatial resolution can be increased. [14] proposes a system that can provide high-resolution imagery to two tracked users at the same time. Recent lenticular displays [8, 1] can support single-user high-resolution 3D views. Compressive light field displays [7, 17] combine layers of LCDs with directional backlight

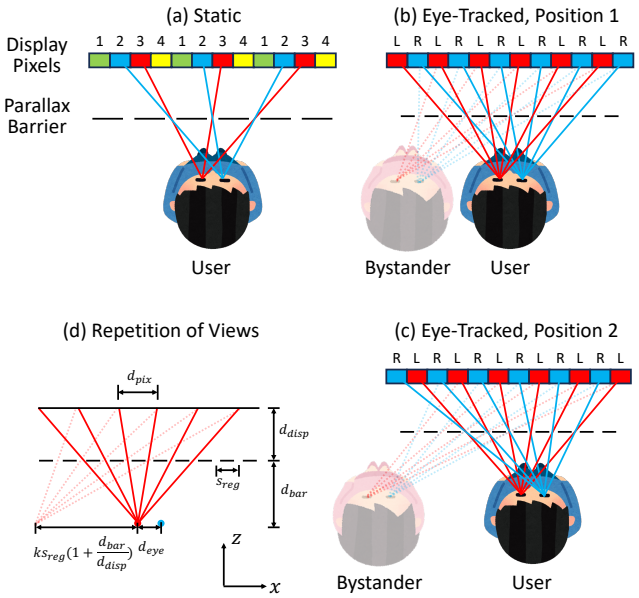

Figure 2: **Preliminary: parallax display. (a)** Static parallax displays show images to a fixed, predefined set of views, which usually has limited spatial resolution. **(b)** Eye-tracked parallax displays dynamically assign the display pixels to the left and right views, achieving improved resolution. **(c)** Alternative assignment of display pixels to the two views. **(d)** Eye-tracked parallax displays suffer from repetition of views: Bystanders standing at multiples of $s_{reg}(1+\frac{d_{bar}}{d_{disp}})$ from the user position can clearly see the user images.

to achieve high resolution and brightness, which rely on solving nonnegative matrix factorization iteratively and may not be suitable for computation-heavy applications such as 3D design and video games. [11, 18] proposes using random pinholes to enable the coexistence of multiple viewers. This results in multiple views with equally relatively low quality (grainy image with color errors). As compared to different variations of parallax barriers developed over a long period of time, we focus on delivering a private view to the user at arbitrary positions, with minimal computation overhead and latency, while maintaining high image quality for the user only.

3 Preliminary: Parallax Display

In this section, we review the basics of parallax displays. A conventional parallax barrier consists of an array of vertical slits placed in front of a display, such as an LCD monitor. Since light can only pass through the slits, each eye of the user sees a different set of pixels on the display. This is illustrated in Fig. 2(a) via a top view, where the display pixels are divided into n disjoint sets ($n=4$ in this example). At the current user position, the left eye sees all pixels from Set 3, and the right eye sees Set 2. By showing a stereo pair of images for each set of pixels, correct parallax is observed and 3D perception is achieved. One limitation of parallax barriers is that a fraction of the emitted light is blocked and the observed brightness is sacrificed. Lenticular displays mitigate this problem by replacing the slits with minuscule lenses which focus light and increase the overall brightness. Since they are based on the same principle to create parallax, in this paper we refer to both as *parallax displays*.

The static parallax display described above splits the display pixels into n *predefined* sets, with each set corresponding to a user position. In Fig. 2, at each specific location, the two eyes of the user see two different views at a time: (1,2), (2,3) or (3,4). In practice, to allow the user to move freely in front of the display, n is chosen to be a large number (e.g., ¿10). This means that the native resolu-

tion of the display is used to multiplex n views, which significantly reduces the spatial resolution of the image observed by the user.

To address this limitation, eye-tracked parallax displays have been proposed [14, 8]. Instead of using a fixed set of predefined views, a camera is used to track the face or eyes of the user, and the display pixels are *dynamically* assigned to the left and right views. Fig. 2(b) and (c) shows an example: Depending on the user position, a pixel can be assigned to either eye. Since there is no need to split views in advance, eye-tracked parallax displays significantly improve the effective resolution.

Repetition of views. Eye-tracked parallax displays offer a cost-effective way to dynamically modulate the light field presented by the display such that a correct view can be delivered to the user in a narrow angular range, irrespective of the user location. However, it suffers from the problem of repetition of views [14]. As shown in Fig. 2(b, c), in addition to the user, the bystander in pink can also see the correct left and right views. This effect is due to the periodicity of the parallax barrier: By moving horizontally by a multiple of the period, the correct views can again be observed, which can reveal confidential information to the bystander.

Fig. 2(d) shows a detailed analysis of this effect. Let us denote the distance between the two eyes by d_{eye} (inter-pupillary distance or IPD), the separation between neighboring slits by s_{reg}, the distance between the eye and the parallax barrier by d_{bar}, and the distance between the barrier and the display by d_{disp}. We choose a coordinate system where the x-axis is horizontally aligned with the barrier, and the z-axis is perpendicular to the barrier. Then, every point in this figure can be represented by its 2D (x,z) coordinates. Let the coordinates of the i-th slit be $(c_i,0)=(i\cdot s_{reg},0)$, and the coordinates of the left eye be $(e_{left},-d_{bar})$. Define the i-th *observed pixel* by the left eye as the pixel corresponding to the ray that passes through the i-th slit. Its coordinates (l_i,d_{disp}) can be computed as:

$$\frac{e_{left}-c_i}{-d_{bar}}=\frac{c_i-l_i}{-d_{disp}}, \tag{1}$$

$$l_i=i\cdot s_{reg}(1+\frac{d_{disp}}{d_{bar}})-e_{left}\frac{d_{disp}}{d_{bar}}. \tag{2}$$

This suggests that, if a bystander has their left eye positioned at $(e_{left}+ks_{reg}(1+\frac{d_{bar}}{d_{disp}}),-d_{bar}),k\in\mathbb{N}$, the corresponding observed pixel is given by:

$$l'_i=(i-k)\cdot s_{reg}(1+\frac{d_{disp}}{d_{bar}})-e_{left}\frac{d_{disp}}{d_{bar}}=l_{i-k}. \tag{3}$$

This proves that the bystander sees a shifted version of the user view. The same derivation can be applied to the right eye as well.

4 Privacy-Enabling Barrier Randomization

In this section, we propose methods to break the periodicity of the parallax barriers, thereby enabling it to provide private views in a narrow range of viewing angles. Periodic patterns, though intuitive and often easy to implement, do not always guarantee the optimal performance in many graphics problems. For example, pseudorandom patterns have been proposed for color filter arrays (CFAs), which avoid structured moiré artifacts that occur when demosaicing raw images captured using a periodic CFA. We draw inspiration from a wide range of methods that employ randomization, and propose *randomizing the locations of the view-blocking slits* in a parallax barrier to disrupt its periodicity.

4.1 Parallax Barrier with Randomized 1D Slits

To enable randomizing the slit locations, we consider a programmable parallax display, with the slit locations programmatically changeable in real-time. Specifically, we assume the parallax barrier is implemented using an amplitude-only spatial light modulator that consists of a 2D array of pixels, the light transmission of which can be independently controlled. In this work, we use an off-the-shelf LCD monitor with the backlight removed [7, 17]. As shown in Fig. 3(a), when the light transmission of a vertical column of pixels is set to one (fully transparent), it works as an open "slit". Therefore, by opening the slits at a regular spatial period, a conventional parallax display can be implemented.

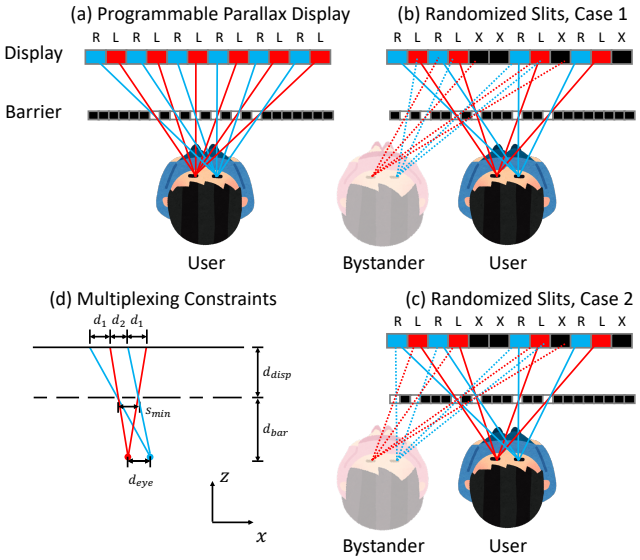

Figure 3: **Slit randomization. (a)** A programmable parallax display can be implemented by placing a amplitude-only spatial light modulator (SLM) in front of a display. Notice that this figure is not to-scale: In practice the SLM has the same pixel size as the display, and is placed much closer to the display than to the user. **(b)** By randomizing the locations of the slits, periodicity of the parallax barrier is disrupted and the bystander no longer sees a repeated view. **(c)** Another case of randomized slits. The bystander is able to see more correct pixels (first three). **(d)** The parameter choice for the random slits must conform to the multiplexing constraints (Eqs. (6) and (7)).

Randomizing slit locations. Let us denote the column indices of the open slits as $\mathbf{C}=\{c_1,c_2,\ldots,c_m\}$. $\mathbf{C}$ is sequentially generated by randomizing the distance between neighboring c_i:

$$\Delta c_i=c_{i+1}-c_i=\text{random\_integer}(s_{\min},s_{\max}), \tag{4}$$

where $s_{\min}$ is the minimum separation between two slits, which is determined such that pixels corresponding to the left view and right view cannot intersect each other. $s_{\max}$ is the maximum separation between two slits which is a tunable parameter. We will discuss the effect of both parameters later in this section.

Fig. 3(b) illustrates the slit randomization approach. The column indices of the slits are $\mathbf{C}=\{1,3,6,8,12,15\}$, with the distances between neighboring slits being $(2,3,2,4,3)$. The user sees the pixels on the display through each slit, where rays corresponding to the left and right eyes are denoted in red and blue, respectively. Similarly, the pixels seen by the left eye are shown in red, and the pixels seen by the right eye are shown in blue. Pixels that are seen by neither eye are kept black, and denoted as X in the figure. As is evident, this randomized barrier is still capable of multiplexing two views onto the display pixels.

Next, let us consider what a bystander (person in pink in the figure) may see. Since the slits are no longer periodic, the bystander does not see the same image as the user. The red rays do not always intersect with red pixels on the display; instead they can intersect

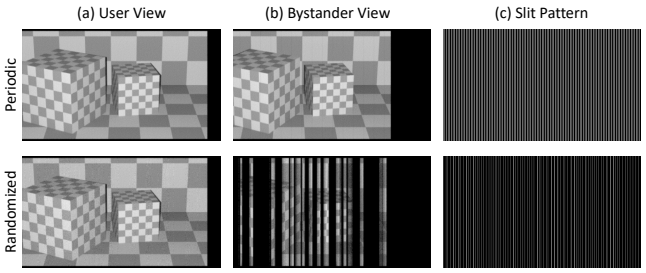

Figure 4: **Random slits example. (a)** Both patterns can show a clear image to the user. **(b)** The periodic pattern shows a repeated view to the bystander, while the randomized pattern shows scrambled columns with high frequency texture obfuscated. **(c)** Visualization of the slit patterns.

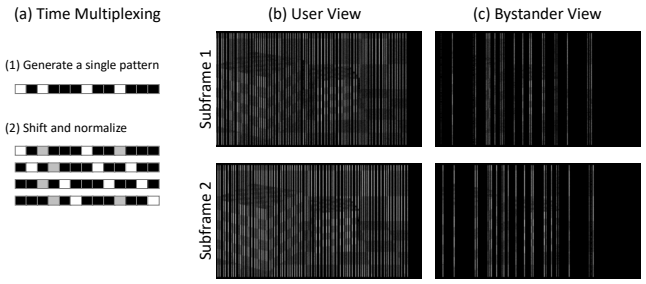

Figure 5: **Time multiplexing. (a)** Time multiplexing can be achieved by shifting and normalizing the initial pattern. **(b)** Visualization of subframes for the user view. **(c)** Visualization of subframes for the bystander view.

either red, blue or black (unused) pixels. Therefore, the left eye sees a *scrambled image*, where each pixel can come from either the left image, the right image or a black image. Fig. 4 shows a simulated example. With a regular slit pattern, the bystander can see an identical image as the user (cropped due to perspective difference). With a randomized slit pattern, the bystander sees a scrambled image, with a large fraction of black pixels. This effectively prevents the bystander from seeing confidential high-frequency information on the screen.

Time multiplexing. In the method explained above based on random patterns, the distance between neighboring slits is varying across the viewing field. This may create unnatural viewing experience for the intended users such as spatially-varying brightness and effective resolution across the image. To overcome this challenge, we propose a *time multiplexing* method: After generating the original slit pattern C, we shift it by one pixel horizontally every time-instant (called a subframe) to generate s_{max} patterns, and then normalize the intensity of each slit such that each pixel sums up to 1 if we add the shifted patterns together. This is shown in Fig. 5(a). The same image is displayed while the slit patterns are shifting across the parallax barrier. This approach lowers the effective temporal resolution (frame-rate), but ensures that the slit pattern covers the entire resolution, and the brightness of each pixel remains uniform when the patterns are added together. In other words, this method trades the temporal resolution of the display for increased spatial resolution.

When a user or a camera observes the screen, they see a single *subframe* at each time, which corresponds to the display time of a single pattern. As the human eye or camera sensor automatically blends the s_{max} subframes, a complete image is observed. Fig. 5(b) and (c) show two subframes as examples for the user and bystander view respectively. Adding all the subframes together gives the view images in Fig. 4.

4.2 Analysis: Why Does Randomization Work?

Let us take a closer look at the scrambling effect caused by the randomized parallax barrier. Fig. 3(c) shows a case where a *subsequence* of slits is repeated locally. The column indices of the slits are C = 0, 2, 6, 8, 12, 15, with the distance between neighboring slits being (2, 4, 2, 4, 3). The subsequence of (2, 4) is repeated once. As a result, when the bystander is at $k \cdot (2+4) p_{pix}(1+\frac{d_{bar}}{d_{disp}})$ (derived similarly to Eq. (3), where p_{pix} is the pixel pitch), they can see the same first three pixels as the user. This is demonstrated in Fig. 3, where the first three red rays from the bystander coincide with the red (L) pixels on the display, and the first three blue rays coincide with the blue (R) pixels. Therefore, it is possible that a local patch of the original image becomes visible to the bystander.

To quantify how often such an event can happen, consider the following problem: If we generate two random sequences of length m according to Eq. (4), what is the probability that they are identical? Since the slit distances are drawn from a uniform distribution, the probability can be easily computed as

$$P\{\text{identical subsequence}\} = \frac{1}{(s_{max} - s_{min})^m}. \tag{5}$$

For the images generated in Fig. 4 and Fig. 5, we choose $s_{min} = 4$ and $s_{max} = 8$. Consider a subsequence with $m = 8$ slits: Its length in the pixel space is given by $(4+8)/2 * 8 = 48$px, which corresponds to approximately six 16pt letters on a 27-inch 1080P display. The average probability of an identical subsequence is then $P = 1/(8-4)^8 = 1/65536$. Consider a normal display with horizontal resolution of 1920 pixels, such an identical subsequence is very unlikely to exist, which suggests that it is highly unlikely for a bystander to read this six-letter word. In other words, although small local patches may be visible to a bystander, they cannot discern sequences of meaningful text from the display.

4.3 Practical Consideration

Choice of parameters. The slit pattern is randomly generated according to the minimum separation s_{min} and the maximum separation s_{max}. The minimum separation s_{min} must be chosen such that the pixels corresponding to the left view and the right view do not intersect each other. Define d_1 and d_2 as the distances between the neighboring pixels that are seen by either eye, as shown in Fig. 3(d). d_1 and d_2 must be no smaller than a threshold d_{th}, *e.g.*, 2 pixels. We can write d_1 and d_2 as functions of the display hardware parameters:

$$d_1 = d_{eye} \cdot \frac{d_{disp}}{d_{bar}}, \tag{6}$$

$$d_2 = s_{min} \cdot (1 + \frac{d_{disp}}{d_{bar}}) - d_{eye} \cdot \frac{d_{disp}}{d_{bar}}. \tag{7}$$

Note that d_1 does not depend on s_{min}. Therefore, by setting $d_2 = d_{th}$, s_{min} can be computed.

s_{max} determines the maximum possible separation between adjacent slits. Eq. (5) shows that a larger s_{max} leads to lower probability of identical subsequences. However, it also increases the number of frames that need to be shown for time multiplexing, which reduces the frame rate. Furthermore, the duty cycle of each slit (column of pixels) becomes lower, which decreases the overall brightness. In practice, to balance these tradeoffs, we set s_{max} to be the maximum value that still allows a 30fps viewing experience.

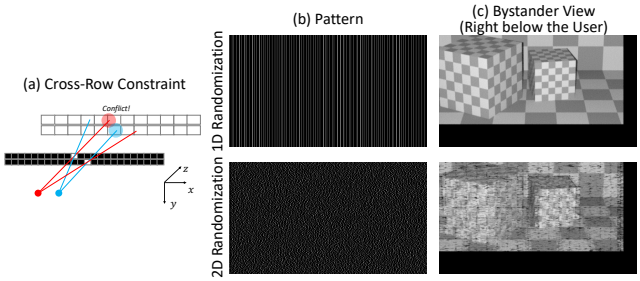

Figure 6: **Pinhole randomization. (a)** When replacing the slits with randomly located pinholes, cross-row constraint needs to be taken into consideration since each ray may depend on neighboring 2×2 pixels. **(b)** Visualization of 1D and 2D random patterns. **Zoom in to see details. (c)** Example bystander view right below the user. 2D randomization allows obfuscating viewpoints that only differ in the vertical coordinate.

Computing and rendering the display image. Once the random barrier is generated, we render the image to show on the display according to the user position, which we assume can be obtained through a face or eye tracker. We first generate the left and right view images that are supposed to be seen by user eyes, which can be (1) 2D planar image shown at a specific depth, such as a normal operating system interface, word processing software, *etc*., (2) pre-captured stereo pair, or (3) stereo images rendered from 3D models. Then we map the two images to pixels on the display.

One way to compute the display image is to formulate this as an inverse problem: It is possible to represent the image seen by the user as a linear function of the pixels shown on the display. Given two target images, what is the optimal display image that minimizes the difference between the target images and the actual observed images? However, this requires solving a large linear system that involves all pixels on the screen, which is not suitable for real-time applications. Furthermore, the optimized image may not be able to fully reproduce the target image, which sacrifices image quality. To solve these problems, we adopt a simple solution: We enforce a threshold d_{th} that is sufficiently large such that any two rays will not interfere with each other significantly. For example, suppose a ray passing through a slit intersects the display at a circle with a radius of 1 pixel. Then, its appearance will depend on the neighboring 2×2 pixels since the intersection is usually at a fractional pixel position. Therefore, we choose $d_{th} = 2$ to mitigate the interference between pixels. We then use a forward splatting method: We "paint" the corresponding pixels with the color of the ray that it is supposed to look like in the target images. This method only involves tracing a collection of rays and finding their intersections with a plane, which is lightweight and can be easily parallelized.

4.4 Parallax Barrier with Randomized 2D Pinholes

One challenge with the randomized slits approach proposed above is that since pixels along a column all correspond to the same view (left, right or black), it makes vertical structures more easily visible to the bystanders. As a result, randomized slits does not prevent *shoulder surfing* (peeking over the shoulder), since moving the viewpoint up and down only shifts the viewing rays accordingly along the slits.

Our key observation is that the parallax barrier does not have to be in the form of vertical slits. We propose *randomized 2D pinhole patterns* which break the slits into a set of pinholes and allow them to be arranged randomly from row to row. Specifically, we randomize the column indices of pinholes at each row according to Eq. (4). One important difference from the slit-only randomization is that, since the pinholes at each row can have different locations, it is necessary to consider *cross-row constraints*. This is shown in Fig. 6(a). Since a ray may intersect with the screen at fractional coordinates and have an impact on 2×2 pixels, a ray from the previous row may cause a conflict with a ray from the current row. Therefore, when we generate the random patterns, we exclude column indices with a distance of $\pm s_{min}$ from any pinhole from the last row.

$$c_{j(i+1)} = \text{random\_choose}(\{n \mid c_{ji} + s_{min} \le n \le c_{ji} + s_{max}, \\ n \neq c_{(j-1)k} \pm s_{min}, \forall k\}), \tag{8}$$

where c_{ji} denotes the i-th pinhole on the j-th row.

Fig. 6(b, bottom) shows the generated 2D random pattern. When viewed from a bystander who is right above the user, the image is scrambled such that high frequency details become corrupted, while the 1D randomization still gives a clear image, as shown in Fig. 6(c).

To conclude this section, Algorithm 1 summarizes the complete algorithm for generating the pinhole pattern. Note that the patterns only need to be generated *offline once* and can be used for any user location. At runtime, only the display image needs to be rendered according to the user location, as described in Sec. 4.1.

ALGORITHM 1: Offline Generation of Randomized Pinhole Pattern

```
patterns = array of s_max patterns;
/* Generating the first pattern                    */
patterns[0] = zeros((num_rows, num_cols));
for j = 0...num_rows−1 do
    patterns[0][j,:]= a row of pinholes by Eq. (8);
end
/* Generating time multiplexing patterns           */
sum_pattern = patterns[0];
for j = 1...s_max − 1 do
    patterns[l]=circular_right_shift(patterns[0], l);
    sum_pattern += pattern[l];
end
for j = 0...s_max − 1 do
    patterns[l] /= sum_pattern;
end
```

5 Experiments

5.1 Hardware Prototype

We use an LG 27GP750-B 27-in LCD monitor as our display, which has a spatial resolution of 1920×1080 and a frame rate of 240fps. By using 8-pattern time multiplexing, a 30fps effective frame rate is achieved. For the parallax barrier, we use the same display model as a amplitude-only spatial light modulator, with the stock frame and backlight removed. We use a Macbook Pro (16-inch, 2021) to drive both LCD panels. It is worth noting that the chosen LCD monitor supports G-SYNC, which allows synchronizing the two LCDs to the refresh signal coming from the DisplayPort of the Macbook (with ProMotion enabled) instead of individual clocks on the display, which is necessary for time multiplexing.

Fig. 7 shows four pictures of the hardware prototype from different angles. Each LCD is rigidly fixed to a metal frame. The two metal frames are drawn towards each other via springs, with a number of knobs that fix the relative pose between them. Knobs A can be adjusted such that the two LCDs are parallel to each other. Knobs B can be adjusted such that each row in the barrier LCD is parallel to each row in the display LCD.

Calibration Although the knobs can be used to finely adjust the relative pose of the two LCDs, in reality the LCDs cannot be perfectly aligned. Furthermore, since each LCD is just a thin layer of glass (especially the barrier LCD since it is only fixed around the perimeter) which is as big as 27-inch diagonal, they are not guaranteed to be planar. To enable tracing rays and finding intersections

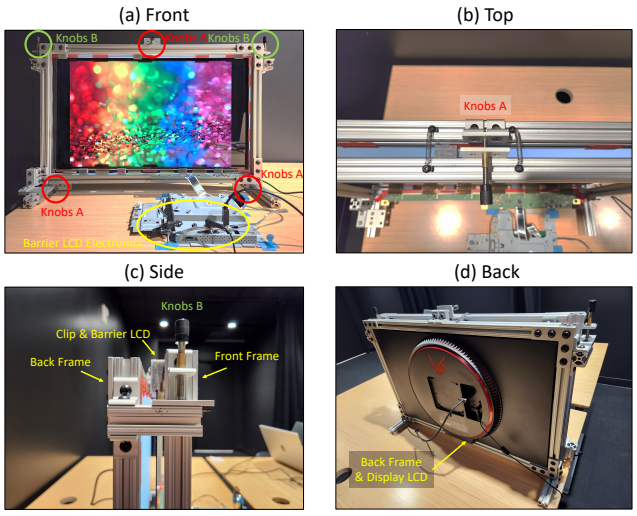

Figure 7: **Hardware prototype. (a)** Front view: The hardware is composed of two LCDs, with the front (barrier) LCD dissembled and backlight removed. The electronics of the barrier LCD is kept connected and fixed on the table. **(b)** Top view: The two LCDs are fixed to two separate metal frames, which are connected by springs. A set of three knobs (Knobs A) are used to make sure the two LCDs are approximately parallel. **(c)** Side view: The barrier LCD is fixed to the metal frame using 3D printed clips. A set of two knobs (Knobs B) are used to make sure the rows of two LCDs are approximately parallel. **(d)** Back view: The display LCD is glued to the back frame.

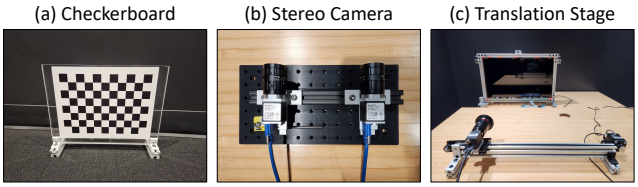

Figure 8: **Calibration and experiment setup. (a)** A printed checkerboard glued to a piece of acrylic for flatness. **(b)** A stereo pair fixed on an optical breadboard for rigidity. **(c)** A translation stage for accurate, repeatable control of camera positioning.

with the pixels, it is essential to build and calibrate a parametric geometry model of each LCD.

To calibrate the display, we first build a accurately calibrated stereo pair of cameras. To ensure that the stereo cameras stay accurately calibrated, we carefully glue a printed checkerboard on a piece of acrylic (Fig. 8(a)) to ensure its flatness. We also fix the two cameras on a breadboard (Fig. 8(b)). We then calibrate the shape of each LCD separately in a coordinate system centered at the left camera. Specifically, we show a sequence of graycode patterns on each LCD, as used in the structured light community. This establishes the correspondence between each LCD pixel, left camera pixel and right camera pixel. The depth of each LCD pixel can then be calculated through triangulation. To mitigate noisy measurements, we fit a fourth degree polynomial surface to each LCD, which we use as our geometry model. Ray tracing can then be accomplished by solving systems of fourth degree equations numerically.

5.2 Experiments

Obfuscating text. To demonstrate how the proposed privacy display can help hide confidential information from curious bystanders, we first conduct experiments on how text shown on the display are obfuscated. We displays four lines of text on a virtual plane. To quantitatively evaluate the performance, we place a camera on a translation stage (Fig. 8(c)) and take images at different viewing angles. For each captured image, we run a widely-used open-source OCR software [13]. We then use Python SequenceMatcher to compute a similarity metric between the detected text and the ground truth, which is defined as $R = \frac{2K_m}{|S_{detected}|+|S_{gt}|}$, where $|S_{detected}|$ and $|S_{gt}|$ are the lengths of the detected and the ground truth strings. K_m is the length of the longest common subsequence of the two strings. $0 \leq D \leq 1$.

We compare the privacy preserving performance of the proposed private parallax barrier and the conventional parallax barrier, which is implemented on the same hardware with a slit period equal to s_{max}. By shifting the slits by one pixel at a time, time multiplexing can be achieved at the same frame rate as the proposed private parallax barrier. Fig. 9(a) plots the computed similarity ratio R against the viewing angles. Both methods display the text clearly at the intended viewing angle (0 degrees). With the conventional parallax barrier, the display completely blacks out at 6 degrees. However, the image becomes visible again at 12 degrees, creating repeated views for bystanders. For the proposed private parallax barrier, the similarity ratio significantly drops at 8 degrees. It does not go zero since some individual letters can still be recognized, but it is extremely challenging for the OCR algorithm or a human bystander to discern a complete word, as shown in Fig. 9(c). Unlike the conventional parallax barrier, the image quality does not recover to recognizable level as the viewing angle increases, which proves the privacy-preserving capability of the proposed method. The images are clearly visible only within $\pm 6^\circ$ of the intended view, offering a viewing field much narrower than commodity privacy filters ($\pm 30 - 45^\circ$).

Obfuscation along the vertical direction. One important benefit of the proposed 2D randomized parallax barrier is that it prevents peeking bystanders at the same horizontal angles but different heights. To quantitatively evaluate this, we fix the camera on a height-adjustable standing desk and capture images at different vertical viewing angles. We show four different lines of text and compute the similarity ratio for each image. As shown in Fig. 10, the image quality of the conventional parallax barrier does not decrease as the vertical viewing angle changes. On the other hand, the similarity ratio of the proposed private parallax barrier significantly decreases and the text become indiscernible at 7°, proving the effectiveness of the proposed method at preventing shoulder-surfing.

Obfuscating QR codes. In addition to text, the scrambling of high-frequency information can lead to the obfuscation of other content. Fig. 11 shows that the success ratio of detecting QR codes decreases as the horizontal viewing angle increases. We display 10 different QR codes on the virtual plane, run a state-of-the-art QR code detector [12] on the captured image and record the ratio of codes that are successfully detected. Similarly, the success ratio rises again at 12 degrees for the conventional parallax barrier due to repeated views. The success ratio for the private parallax barrier keeps low for viewing angles larger than 4 degrees, proving that the proposed system can prevent bystanders from scanning sensitive QR codes such as friending or payment codes.

Showing general images. Besides structured content such as text and QR codes, Fig. 12 demonstrates the display's capabilities of showing general images. Fig. 12(a) and (b) show the images captured at the user's left and right views. To demonstrate the 3D effect, we also overlay the two images to build an anaglyph in Fig. 12(c), which visualizes the parallax between the two views. For readers who happen to have access to anaglyph glasses, a 3D view can be observed. We also show that for a bystander's view at 8°, the general images also get obfuscated. While a bystander

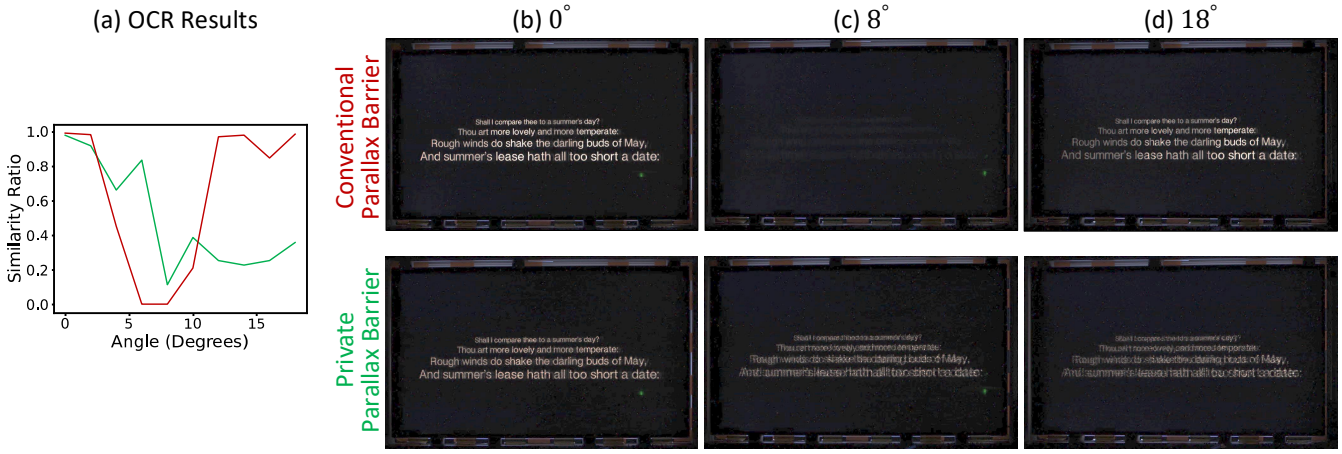

Figure 9: **Obfuscating text. (a)** We compare the similarity ratio between the detected text and the ground truth. **(b)** At the intended location, both barriers show the correct image. **(c)** At 8°, the conventional parallax barrier completely blacks out. The private parallax barrier shows a scrambled image. **(d)** At 18°, the conventional parallax barrier shows the correct image once again, while the private parallax barrier still obfuscates the image.

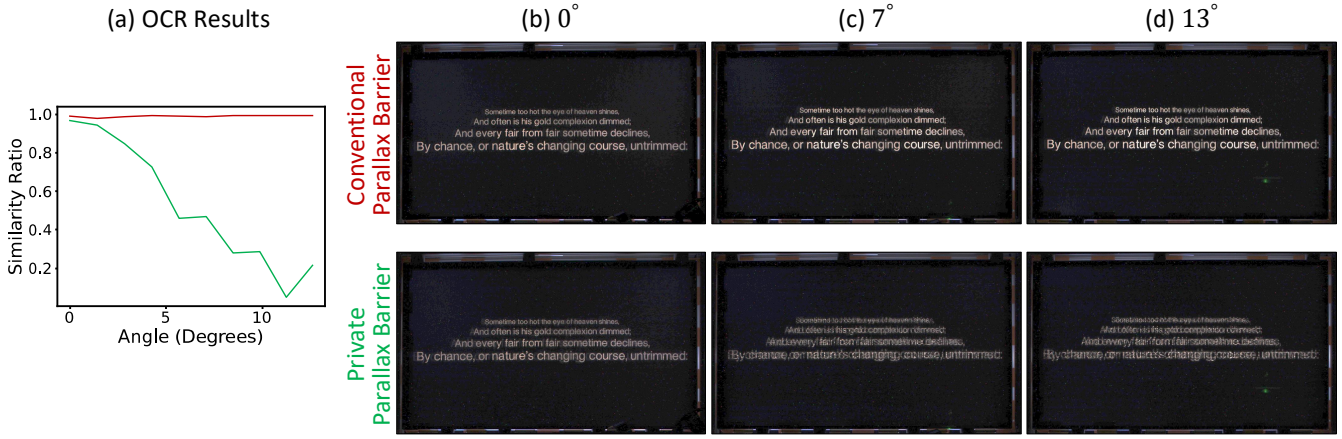

Figure 10: **Obfuscation along the vertical direction. (a)** Similarity ratio. **(b)** Clear images at 0°. **(c,d)** The conventional parallax barrier always show the correct image no matter what the vertical viewing angle is. The private parallax barrier shows a scrambled image.

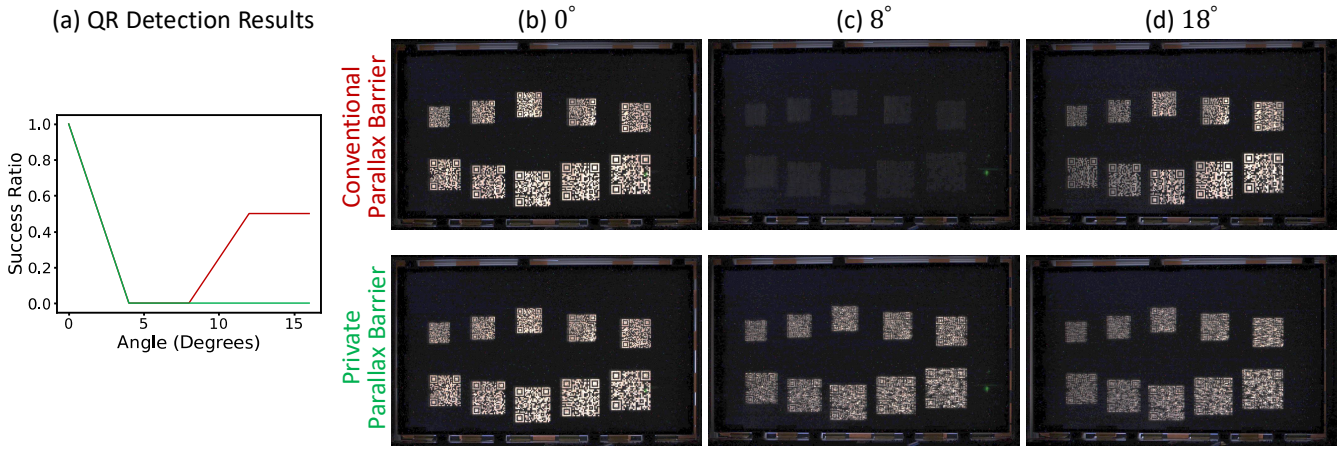

Figure 11: **Obfuscating QR codes. (a)** Ratio of successful QR code scans. **(b)** Successful scans at 0°. **(c)** Failed scans at 8°. **(d)** The conventional parallax barrier allows successful scans at 18° again. The private parallax barrier continues to obfuscates the codes.

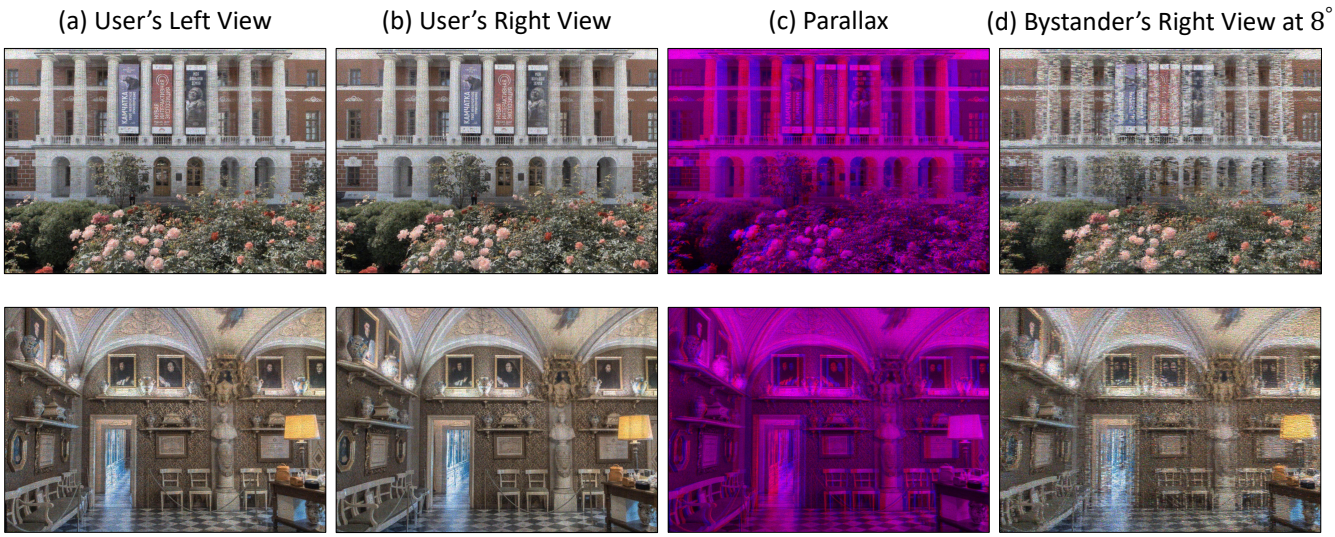

Figure 12: **General images.** The proposed private parallax barrier can show general images as well. **(a,b)** User's left/right view. **(c)** Anaglyph to visualize the parallax. **(d)** General images are also obfuscated in the bystander's view. The bystander can get the overall structure of the image, but fail to recognize details such as text (top) and portraits on the wall (bottom).

Figure 13: **More general images.**

can still get an overall idea of how the images look like, details in the images such as text (top) and portraits on the wall (bottom) becomes indiscernible. Fig. 13 shows a gallery of other captured general images. **See the supplementary video for the captured videos.**

6 Limitations and Future Work

Brightness. Since the proposed system is based on parallax barrier, it achieves a lower brightness level as compared to lenticular displays when the same input power is used. Although the overall brightness can be increased by using a backlight with higher power, the proposed method may not be suitable for mobile or portable devices with a limited power budget.

Alignment and calibration. In this work we present a hardware prototype where we build ad hoc frames to align the two displays and use a complicated calibration method. This has resulted in several inconveniences: First, the displayed images contain artifacts due to the error in alignment and calibration. Second, the computation of the display involves solving a system of fourth degree equations due to the warping of the LCDs. We emphasize that this is not a limitation of the proposed method, but rather the hardware prototype. With mature manufacture and assembly designs, such challenges are likely to be solved.

Eye tracker. In the hardware prototype we do not have a built-in face/eye tracker but assumes the user location is fixed and known a priori. Thankfully, recent face/eye tracking systems are becoming more and more mature. With the recent interest in eye-tracked parallax displays, we believe such technologies can be combined with the proposed parallax barrier randomization, realizing a privacy-enabled parallax display that has a higher level of privacy protection without sacrificing usability.

References

[1] ASUS. Asus spatial vision. https://www.asus.com/content/asus-spatial-vision-technology/, 2024. Accessed: 2024-01-24. 2

[2] C.-Y. Chen, B.-Y. Lin, J. Wang, and K. G. Shin. Keep others from peeking at your mobile device screen! In *The 25th Annual International Conference on Mobile Computing and Networking*, pp. 1–16, 2019. 2

[3] H. Isono, M. Yasuda, and H. Sasazawa. Autostereoscopic 3-d display using lcd-generated parallax barrier. *Electronics and Communications in Japan (Part II: Electronics)*, 76(7):77–84, 1993. 2

[4] F. E. Ives. A novel stereogram. *Journal of the Franklin Institute*, 153(1):51–52, 1902. 2

[5] A. Jacobs, J. Mather, R. Winlow, D. Montgomery, G. Jones, M. Willis, M. Tillin, L. Hill, M. Khazova, H. Stevenson, et al. 2d/3d switchable displays. *Sharp Technical Journal*, pp. 15–18, 2003. 2

[6] Y. Kim, J. Kim, J.-M. Kang, J.-H. Jung, H. Choi, and B. Lee. Point light source integral imaging with improved resolution and viewing angle by the use of electrically movable pinhole array. *Optics Express*, 15(26):18253–18267, 2007. 2

[7] D. Lanman, M. Hirsch, Y. Kim, and R. Raskar. Content-adaptive parallax barriers: optimizing dual-layer 3d displays using low-rank light field factorization. In *ACM SIGGRAPH Asia*, pp. 1–10. 2010. 2, 3

[8] J. Lawrence, D. Goldman, S. Achar, G. M. Blascovich, J. G. Desloge, T. Fortes, E. M. Gomez, S. Häberling, H. Hoppe, A. Huibers, C. Knaus, B. Kuschak, R. Martin-Brualla, H. Nover, A. I. Russell, S. M. Seitz, and K. Tong. Project starline: a high-fidelity telepresence system. *ACM Trans. Graph.*, 40(6), dec 2021. doi: 10.1145/3478513.3480490 2, 3

[9] G. Li. Parallax barrier, display device and display state control method thereof, Sept. 10 2019. US Patent 10,409,095. 2

[10] S. Lian, W. Hu, X. Song, and Z. Liu. Smart privacy-preserving screen based on multiple sensor fusion. *IEEE Transactions on Consumer Electronics*, 59(1):136–143, 2013. 2

[11] A. Nashel and H. Fuchs. Random hole display: A non-uniform barrier autostereoscopic display. In *2009 3DTV Conference: The True Vision-Capture, Transmission and Display of 3D Video*, pp. 1–4. IEEE, 2009. 2

[12] OpenCV. Wechat qr code detector. https://github.com/opencv/opencv_contrib/tree/4.x/modules/wechat_qrcode, 2020. 6

[13] PaddlePaddle. Paddleocr. https://github.com/PaddlePaddle/PaddleOCR, 2020. 6

[14] T. Peterka, R. L. Kooima, J. I. Girado, J. Ge, D. J. Sandin, A. Johnson, J. Leigh, J. Schulze, and T. A. DeFanti. Dynallax: Solid state dynamic parallax barrier autostereoscopic vr display. In *2007 IEEE Virtual Reality Conference*, pp. 155–162. IEEE, 2007. 1, 2, 3

[15] E. L. Schwartz, M. J. Pellerite, D. W. Hennen, M. D. Radcliffe, M. L. Steiner, M. A. Boulos, M. E. Lauters, G. T. Boyd, J. P. Baetzold, J. J. Stradinger, et al. Electronically switchable privacy film and display device having same, Jan. 5 2016. US Patent 9,229,253. 2

[16] E. Sommerlade, G. J. Woodgate, and H. Flynn. Intelligent privacy system, apparatus, and method thereof, Nov. 12 2019. US Patent 10,475,418. 2

[17] G. Wetzstein, D. R. Lanman, M. W. Hirsch, and R. Raskar. Tensor displays: compressive light field synthesis using multilayer displays with directional backlighting. 2012. 2, 3

[18] G. Ye, A. State, and H. Fuchs. A practical multi-viewer tabletop autostereoscopic display. In *IEEE International Symposium on Mixed and Augmented Reality*, pp. 147–156. IEEE, 2010. 2

[19] P. L. Yu, T. C. Lin, and S. W. Lin. Privacy filtering method and 3d display device, 2014. TW Patent TWI436099B. 2