Seeing through Obscure Glass

Qi Shan, Brian Curless, and Tadayoshi Kohno

Department of Comptuer Science & Engineering, University of Washington
{shanqi, curless, yoshi}@cs.washington.edu

Abstract. Obscure glass is textured glass designed to separate spaces and "obscure" visibility between the spaces. Such glass is used to provide privacy while still allowing light to flow into a space, and is often found in homes and offices. We propose and explore the challenge of "seeing through" obscure glass, using both optical and digital techniques. In some cases – such as when the textured surface is on the side of the observer – we find that simple household substances and cameras with small apertures enable a surprising level of visibility through the obscure glass. In other cases, where optical techniques are not usable, we find that we can model the action of obscure glass as convolution of spatially varying kernels and reconstruct an image of the scene on the opposite side of the obscure glass with surprising detail.

1 Introduction

Obscure glass is a class of window glass used to separate spaces while allowing light and a limited amount of visual information to pass through. It is not uncommon to find such glass embedded, for instance, in office and conference room doors, in front doors of homes, in the exterior windows of bathrooms, and in the windows of close-proximity homes. Obscure glass is typically transparent, but has a surface texture that results in distorted and blurry images of a scene when photographing through the glass. The intent is to provide some degree of privacy, but how much privacy is actually provided?

In this paper, we explain the action of obscure glass and describe two complementary methods for seeing through it. First, we explore optical methods that include tuning the camera configuration, and show that, surprisingly, when the glass is textured only on the side facing the observer, household liquids may be used to cancel much of the obscuring effect of the glass. Indeed, it is possible to see through some visually impenetrable obscure glasses using a drop of honey and an inexpensive, compact video device, such as an iPod nano$^{\text{TM}}$. Second, we model the action of obscure glass as a spatially varying blur applied to a latent image. We then develop a calibration technique to recover this blur (assuming temporary access to both sides of the glass) which can be used to deblur a photo of a scene taken through the obscure glass at a later time, thus recovering the latent image of the scene.

Our contributions are threefold. First, to our knowledge, this is the first paper to define and address the problem of seeing through obscure glass. Second,

the optical methods, though building on related ideas in camera design and measurement of refractive substances, are novel in their application to minimizing optical degradation caused by obscure glass. Third, we have developed a new technique for recovering spatially varying blur kernels from small numbers of images by leveraging the sparsity of the kernel functions in this setting.

In the remainder of the paper, we review related work (Section 2), discuss the characteristics of obscure glass (Section 3), explore optical and calibrated deconvolution methods for seeing through it (Section 4), and conclude with results and discussion (Sections 5 and 6).

2 Related work

Refractive index matching. One of our approaches to seeing through obscure glass is to reduce its distorting and blurring effect with a substance (nearly) matching its index of refraction. The idea of matching index of refraction is well-known as a tool for measuring the index of refraction of an irregularly shaped, transparent solid by immersing it in liquids whose refractive indices are known or readily measured [1]. This idea has also been applied to the problem of multi-view 3D volumetric reconstruction of refractive objects by computed tomography [2,3].

Environment matting. We also attempt to undo the effects of obscure glass by measuring the blurring and distorting properties of the glass. This measurement step is known as "environment matting" [4]. The idea is to shoot photographs through a refractive object, behind which is a changing background. In the original formulation, the background was a monitor displaying horizontal and vertical hierarchical stripe patterns, and per pixel filtering rectangles were recovered [4]. Follow-on work used different monitor patterns and filter representations: a single, smooth color wash pattern to recover refractive distortion for smooth objects [5], many images of a Gaussian stripe sweeping in different directions to recover multi-modal, oriented, Gaussian filters [5], and wavelet patterns to recover per-pixel wavelet filters [6]. An alternative approach uses a single, large background image that is moved around behind the refractive object, recovering per pixel filters, typically for smooth objects [7]. Agarwal et al. [8] recover distortion maps for refractive objects by analyzing their effect on the optical flow of a video being played in the background. These methods generally require many images or they impose restrictive assumptions on the blurring properties of the glass. Many images is problematic in that access to both sides of the glass may require a relatively quick capture process so as not to be excessively intrusive. Peers and Dutre [9] reduce the number of images using wavelet noise functions and a non-linear wavelet approximation; we employ similar patterns but explicitly encourage sparsity in the kernels during recovery. Our work builds most directly on environment matting, but it is also related to work on inverse light transport [10,11].

Also relevant is work in computer vision on recovering the shape of refractive surfaces. A nice survey can be found in [12]. Our approach is to recover blur kernels that can arise from fine (sub-pixel) texture, rather than recovering an

explicit shape profile. The work of Murase [13] is particularly relevant in that an undistorted image is recovered as part of the shape estimation process, though the setting is different, as it depends on time-varying water motion.

Non-blind deconvolution. Recovering a latent image with a given image transform (blurring) matrix is known as non-blind image deconvolution. Recent research has yielded promising results using a natural image statistics prior and multi-scale techniques [14,15,16]. In this paper, we use a formulation similar to [15].

Security. The computer security community has several prior works studying "information leakage" via optical emanations. Kuhn found that the images on a CRT monitor can be reconstructed from reflections off a white wall [17]. Backes et al. extended this work to reconstruct images using reflections off diverse sets of objects, ranging from teapots to shirts [18,19]. Our work explores image reconstruction through refraction instead of reflections.

3 Characteristics of obscure glass

Obscure glass is plane glass with a surface texture – geometric perturbations – on at least one side of the glass. Light rays incident on the glass generally reflect and refract at the interface. The dominant visual effect when looking through obscure glass arises from refraction of viewing rays. If the glass is relatively smooth but wavy, the refraction results in a distorted image. If the glass has a finer texture, viewing rays can be significantly "scrambled" by refraction at the obscure glass surface, resulting in a blurred image of the scene.[1] Fig. 1(c) illustrates this effect and shows an image shot through an obscure glass sample; note the distortion and blur in this example.

Obscure glass varies in design and optical properties. The most noticeable variation is the geometric design of the surface texture, which can be locally smooth or rough, and exhibit various larger scale patterns for aesthetic appeal. Examples of obscure glass appear in Figs. 1 and 3-5. The surface perturbations are usually applied to one side of the glass, though glass with perturbations on both sides is not uncommon. When the texture is on one side, there is evidently no fixed rule as to which side will face outward when decorating offices and homes. Still, we note that for glass facing the outdoors, the flat side is often on the (dirtier) outside for ease of cleaning. Similarly, the flat side of a shower room door often faces toward the tub, for ease of cleaning mineral deposits, etc.

Finally, the refractive index of obscure glass, while nominally around 1.5-1.6, can vary according to the composition of the glass; manufacturers use different "recipes" for glass that results in this variation, sometimes using different recipes for different glasses in their own product lines.

[1] Taken to an extreme, the texture can be so rough as to completely scatter the rays, yielding no useful image; this kind of glass is called "frosted glass."

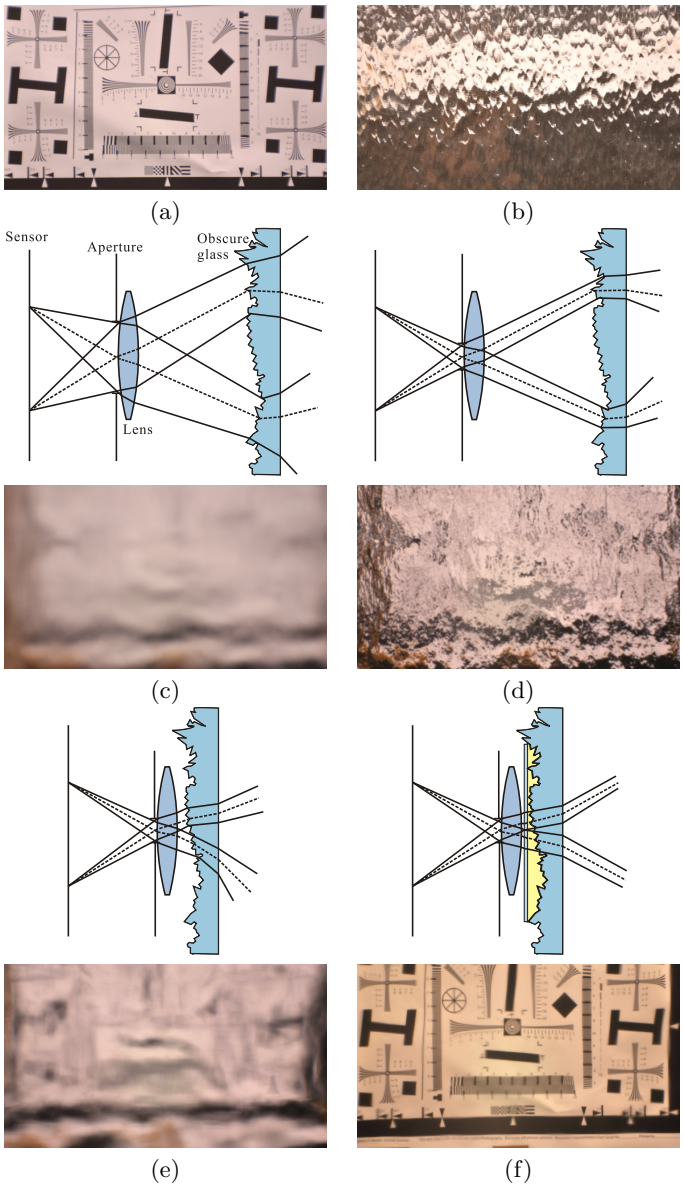

Fig. 1. A target scene (a) is photographed through a piece of obscure glass (b). A wide aperture (c) results in a blurry image, as the rays are broadly scattered. (The dotted lines represent chief rays; the solid lines are drawn to suggest the spread of refraction.) Narrowing the aperture (d) reduces the scatter, but severe distortion remains. Bringing the camera closer to the glass (e) yields a less distorted image. Interestingly, this image seems blurrier than the one in (d); this is because the distortion in (d) juxtaposes the blurred pixels at random, creating the false impression of higher frequencies. After applying a drop of liquid with refractive index close to that of the glass (pressed against the glass with a microscope slide cover slip), a nearly clear shot can be taken (f).

4 Seeing through obscure glass

In this section, we present two complementary approaches to the problem of trying to resolve a clear image of a scene observed through obscure glass: an optical approach to improve image clarity and a calibration approach for estimating and deconvolving blur kernels.

4.1 Optical approach

To improve the photographic quality when shooting through obscure glass, we propose three optical strategies: small aperture, close camera placement, and applying a substance to the glass to reduce the refractive effect of the near side of the glass before imaging.

Small aperture. To reduce the blurring effect of obscure glass, the aperture can be stopped down (reduced in diameter), in effect restricting the range of scene points that will contribute to the light recorded at a point on the sensor. Fig. 1(d) illustrates the effect of reducing the aperture.

It is well-known that most imperfections (defocus, spherical aberration, etc.) in lens systems themselves can be reduced by stopping down the aperture. What remains is geometric distortion. Lens manufacturers go to great lengths to correct for most of these imperfections across a range of settings in advance, so that the aperture need not be stopped down to take a clear image of the in-focus plane. However, with obscure glass, even when we take a photo with a well-corrected lens, the image is degraded because we have introduced a new optical element – the obscure glass – that introduces potentially severe aberrations. Stopping down the aperture reduces the blurring aberrations, but distortion remains.

Note that blurring occurs as long as the cone of rays from the scene arriving at the aperture must pass through a locally "rough" region. The scale of the roughness that matters is thus really a function of the size of that cone. An undulating surface may seem smooth at millimeter scale, but if the aperture is wide open so that the cone of rays passing through the surface is at centimeter scale, then the result will be blur.

Camera placement. Stopping down the aperture will reduce image blur introduced by obscure glass, but some amount of distortion will likely remain. The reason for the distortion is that points on the sensor are recording narrow cones of light arriving from very different parts of the obscure glass, which can have very different (uncorrelated) surface orientations. To minimize the spread of cones across the glass, we can simply place the aperture (more precisely, the entrance pupil) as close to the glass as possible, as illustrated in Fig. 1(e).

It is worth noting that a variety of commonly available camera and lens solutions may be applied at this point, e.g., an SLR with a conventional lens stopped down as far as possible or with a specialty lens capable of very narrow apertures (high F-numbers) [20]. Ideally, the optics and camera placement would be tuned, if possible, to place the entrance pupil at the surface of the obscure glass. There is, however, a surprisingly simple and effective alternative. Very small cameras

are becoming increasingly available in cell phones, webcams, and pocket video recorders. These cameras naturally have small apertures, and some of them can be placed very close to obscure glass, separated from the surface by a few millimeters or less. And, despite their small apertures, they have relatively low F-numbers, which means they can take (sometimes modest resolution) photos without requiring long exposures.

While placing the camera close to the glass can reduce the effect of distortion, a limitation is the fact that the camera itself becomes more noticeable from the opposite side of the glass, which makes the approach intrusive. However, this is not a concern when observing, e.g., a computer screen or documents in an office with no one inside. Further, a very small camera could be fairly unobstrusive; the camera (by itself) inside an iPod nano is already suggestively small.

Refractive index matching. Even with a closely positioned camera with a small aperture, if the surface has fine scale texture, then some amount of blurring and distortion will remain. When the texture is on the near side of the glass, then an unusual solution becomes possible. Recall Snell's law of refraction:

$$\eta_{\text{in}} \sin\theta_{\text{in}} = \eta_{\text{out}} \sin\theta_{\text{out}} \tag{1}$$

where θ_{in} and θ_{out} are the angles of the incident and refracted ray, respectively, taken with respect to the surface normal, and η_{in} and η_{out} are the indices of refraction on either side of the interface. The strength of refraction at the interface is controlled by how far the refractive index ratio $\eta_{\text{out}}/\eta_{\text{in}}$ is from unity. If we could smoothly "plaster over" the texture with a clear substance that had the exact same index of refraction as the glass, then the blurring, distorting effect of that texture would disappear. Fig. 1(f) illustrates this idea.

Precisely matching the index of refraction of obscure glass is challenging for several reasons. First, the index of refraction of glasses can vary depending on their compositions. Second, many of the standard, laboratory liquids used to cancel indices of refraction around $1.5 - 1.6$ are toxic. Finally, many of these liquids are very low viscosity; when applied to a vertical surface, they are difficult to contain, tending to run down the glass.

Instead, we propose to use non-toxic, high viscosity, household substances. For example, honey has proved to be very close in refractive index of glass and works well in experiments. The match is not exact, but combined with closely placed cameras with small apertures, visibility can improve significantly.

A critical limitation of this approach is the need for the textured surface to be on the near side of the glass. Still, for surfaces with texture on two sides, this approach can significantly reduce the degradation due to the obscure glass (cutting it "in half"). Further, if it is possible to deposit a (long lasting, ideally unnoticeable) substance on the far side of the glass in advance, then an optical portal is created.

4.2 Deconvolution approach

In this section, we formulate the action of obscure glass as convolution of spatially varying kernels with a latent image and describe calibration methods for

estimating these kernels and later recovering the latent image of an unknown scene. We note that this calibration scenario is plausible when access to both sides of the glass are available for a period of time, and then later, when only one-sided access is available, a degraded image of the scene is taken.

Image formation process. The image recorded at the sensor is a weighted sum of light rays scattering through the obscure glass and passing through the optical system of the camera. In principle, we need to recover a light field weighting function and latent light field within a small volume. To simplify the problem, we assume that the scene has constant depth and diffuse materials, or more loosely, minimal parallax and minimal view-dependent reflection with respect to the scattering of rays through a small portion of the obscure glass.

These assumptions allow us to model the scene as a latent image L and the formation of image I as a weighted sum of latent image pixels:

$$I = \mathbf{F}L + N, \tag{2}$$

where $\mathbf{F}$ is a degradation matrix whose rows correspond to spatially varying kernels, and N is sensor noise.[2] Note that $\mathbf{F}$ encodes both blur and distortion, and that L is generally larger in dimensions than I, i.e., $M_L > M_I$ where M_L and M_I are the number of pixels in L and I, respectively. Further, we assume the glass is not color tinted, and we do not model possible dispersion (wavelength dependent refraction), so that, while I and L are color-valued, $\mathbf{F}$ is not. Our goal is to recover $\mathbf{F}$ from measurements and then invert its effect to recover an image of a scene L restricted to L's overlap with I.

Recovering the degradation matrix F. We will recover $\mathbf{F}$ by recording a set of images $\mathcal{I} = [\mathbf{I}_1, \mathbf{I}_2, \mathbf{I}_3 \ldots]$ through obscure glass in front of a known, changing background $\mathcal{L} = [\mathbf{L}_1, \mathbf{L}_2, \mathbf{L}_3 \ldots]$. The image formation problem then becomes:

$$\mathcal{I} = \mathbf{F}\mathcal{L} + \mathcal{N}, \tag{3}$$

where $\mathcal{N}$ is sensor noise across the image set. Assuming independent and identically distributed Gaussian noise, the optimal $\mathbf{F}$ can be computed by minimizing:

$$E(\mathbf{F}) = \|\mathbf{F}\mathcal{L} - \mathcal{I}\|_2^2 \tag{4}$$

$\mathbf{F}$ is in general an $M_L \times M_I$ matrix; thus, in principle, M_L images are needed to estimate it, which could take a very long time to capture and process. By optimizing camera parameters, as described in the previous section, we have found that the filter support (and thus the number of free variables) in each row of $\mathbf{F}$ can be reduced significantly, e.g., to 100×100 for a 400×400 image. Still, an entirely brute force approach would require 10,000 images in this case.

To allow us to operate with fewer images (fewer observations than unknowns), we assume the blur kernels of the obscure glass are sparse in the spatial domain; i.e., we require most of the elements of each row of $\mathbf{F}$ to be zero. If the rows of

[2] We neglect lighting reflections from the camera side of the glass to the camera, which are minimized by placing the camera close to the glass.

F were instead dense, then the blur caused by them would be so severe as to make subsequent latent image recovery attempts essentially impossible. Thus, our assumption requires working with obscure glass and imaging setups where it actually *is* feasible to recover a useful image of the scene behind it.

We can encode a sparsity prior as a penalty on the L_1-norm of **F**, giving:

$$E'(\mathbf{F}) = \|\mathbf{F}\mathcal{L} - \mathcal{I}\|_2^2 + \gamma\|\mathbf{F}\|_1, \tag{5}$$

where γ is a weight to balance the data term and the prior term (set to 10^{-2} in all of our experiments). Thus, the problem is transformed into energy minimization with total variation regularization, solvable with existing techniques [15,21]. We note that this minimization can be performed independently for each row of **F**; i.e, the per pixel kernels can be estimated independently, and that each image pair provides three measurements (one per color channel) per pixel.

The obscure glass will not amplify or negate light, thus the elements $f_{i,j}$ of **F** must be in the range $[0, 1]$. Rather than impose these bounds directly, we take a simple (sub-optimal) approach of performing the optimization, then clamping all values $f_{i,j}$ to the range $[0, 1]$. Those values affected by the clamping step are then held constant and the optimization and clamping are repeated once more.

To reduce the number of variables and accelerate kernel estimation, we take a hierarchical approach that leverages the sparsity assumption. Specifically, we start with downsampled images and estimate kernels at the coarsest level where each kernel has at most 25 elements. After clamping, we then employ nearest neighbor upsampling to initialize the kernels at the next finer level. Any kernel variables at the finer level that are initialized to zero are held constant at zero. This process is repeated for each level. This approach reduces the required number of input images for two reasons. First, at the coarser levels, the size of the kernel is small enough to enable direct estimation with a small number of input image pairs. Second, at the finer levels, a large proportion of the pixels tend to be held constant at zero, thus reducing the number of unknowns, again requiring fewer images. See Fig. 2 for a visualization of the kernels.

We note that our solution method is tuned for speed, not optimality. Ignoring run time, we could explore slower methods that, e.g., exactly enforce constraints, and attempt to find the global optimum per pixel.

Image reconstruction process. After recovering **F**, we can now take a photograph of an unknown scene observed from the same viewpoint through the obscure glass and reconstruct the latent image L of the scene. The kernels in the rows of **F** generally perform some blurring, and it is well known that direct deconvolution of a blur filter (effectively inverting **F**) is highly sensitive to noise.

Recent deconvolution methods [14,15] have shown that regularizing priors based on natural image gradient statistics can significantly improve results. We follow the formulation of [15] which we summarize briefly here. In particular, we solve for the L and auxiliary function μ that minimizes

$$E(L, \mu) = \|I - \mathbf{F}L\|_2^2 + \lambda\Phi(\mu) + w\|\mu - \nabla L\|_2^2 \tag{6}$$

(a) (b)

Fig. 2. Each kernel is estimated in a multi-scale fashion, shown coarsest to finest (top to bottom) in (a). In (b), we visualize a set of computed spatially varying kernels in a local neighborhood.

where $\mu = (\mu_x, \mu_y)$ is encouraged to be similar to ∇L (the gradient of L) through the third energy term, and

$$\Phi(\mu) = \sum_{i,j} \phi(\mu_x(i,j)) + \phi(\mu_y(i,j)). \tag{7}$$

The function $\phi()$ encodes the natural gradient prior. We refer the reader to [15] for the exact form of the prior. The energy is minimized following the procedure (and constant parameters) in that prior work, generally alternating between optimizing μ with fixed L and optimizing L with fixed μ. In [15], $\mathbf{F}$ corresponds to a spatially invariant blur, and solving for L can be performed with the Fast Fourier Transform. In our problem, the kernels are spatially varying, so we must modify the algorithm; we solve for L using an iterative, least squares, conjugate gradient solver, terminating when the ℓ^2-norm of the difference between consecutively computed L's is less than a threshold, set to 10^{-3} in our experiments.

5 Results

In this section, we first show how the physical imaging configurations affect the quality of the images captured through obscure glass. Then we show image reconstruction results using the calibration-based approach.

5.1 Optical experiments

As discussed in Section 4.1, narrow apertures and close placement to the obscure glass gives the best imaging results. We experimented with four different camera configurations, listed here in order of decreasing aperture size and distance from the glass (due to physical limitations): a Nikon D90 SLR with an 18-105mm lens (set to 18mm and F/22), the same SLR with the LOREO Lens in a Cap [20] (35mm fixed focal length, F/64), an iPhone 3G camera, and an iPod nano video camera. Fig. 3 shows images captured with these optical devices through a challenging piece of obscure glass. The tiny aperture of the iPod nano and the ability

to place it closest to the glass leads to the clearest image. The nano's resolution is relatively low (640×480), but sufficient relative to the blurring caused by the glass in this example.

(a) (b) (c) (d)

Fig. 3. Images captured using different cameras placed as close as possible to the obscure glass: (a) a Nikon SLR with 18-105mm lens, 18mm, F/22, (b) Nikon SLR with Lens in a Cap (35mm, F/64), (c) an iPhone 3G, and (d) an iPod nano video camera.

We also experimented with different substances to cancel the refractive properties of the glass as described in Section 4.1. Fig. 4 illustrates the results for one piece of obscure glass using the Nikon SLR and the Lens in a Cap. The elastomer used in [22] would have been particularly convenient; it is flexible enough to fill in the crevices of the glass and does not leave a trace. Unfortunately, its refractive index did not match the glass very well. Karo Syrup, Grade A clover honey, and wintergreen oil performed better; we applied each to a microscope slide cover slip and placed it against the obscure glass. In our tests with various glasses, wintergreen oil was often the best match, but it has very low viscosity and thus quickly runs out from under the cover slip, making it extremely difficult to use. Honey was the best compromise, as it tended to match the glasses reasonably well and has very high viscosity, so would tend to stay in place long enough to shoot photos and short videos. We refer the reader to the supplementary material for an example of combining small camera imaging (iPod nano) through obscure glass with an applied substance.

5.2 Image deconvolution experiments

In this section, we show experimental results for calibrating spatially varying kernels and recovering a latent image. To be as realistic as possible, we took the following steps. First, the textured side of the glass was oriented away from the viewer, thus was not susceptible to applying a refractive index matching substance. Next, images were taken of a controlled background – an LCD monitor – placed roughly 30 cm behind the glass. Then, the camera was removed and then replaced near the obscure glass, to simulate the effect of performing calibration at one time and later having to reposition the camera to take a shot of a the scene. The repositioning was done manually, and multiple shots were taken in an attempt to collect an image similar in viewpoint to the original. In the experiments, a handful of images works surprisingly well for the glass we tested. Finally, the calibration background was photographed without the obscure glass, and these images were aligned to the calibration images taken through the obscure glass. Precise alignment of the calibration image sets was not necessary,

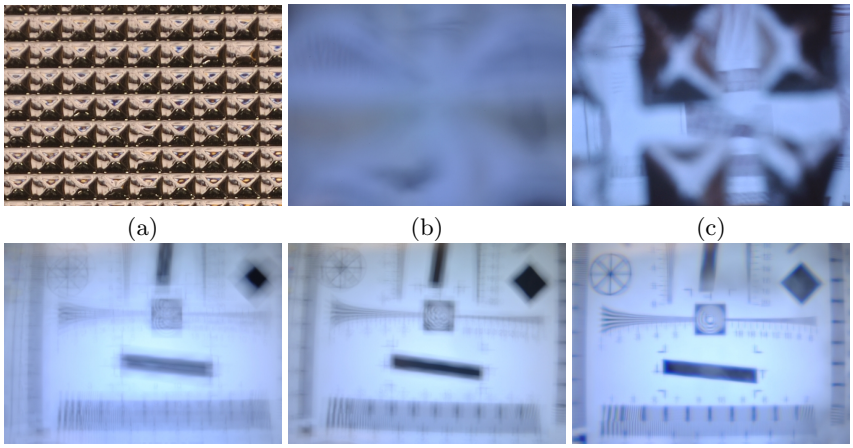

Fig. 4. Images captured with different kinds of liquid substances added to the obscure glass (a). Substances included: (b) no substance, (c) elastomer [22], (d) Karo Syrup, (e) Grade A clover honey, (f) wintergreen oil.

since the kernels being recovered can accommodate pixel offsets. The calibration pattern was 150 different instances of Perlin noise [23] (a different pattern per color channel, as well), which has both low and high frequency content. We used the Nikon SLR, as it was easier to maintain constant exposure settings for it than the iPhone 3G and iPod nano.[3] Images were taken with the 18-150mm lens at F/22 or the Lens in a Cap at F/64 using a tripod with exposures lasting several seconds. We note that, regardless of the imaging set-up, there will exist an obscure glass that introduces problematic degradations; our goal here is to demonstrate how much clearer an image can become given those degradations.

Fig. 5 shows the results for latent image recovery using several different obscure glasses. In each case, we show the best recovered image among the repositioned set, since these images were shot with the understanding that likely only one would be at the right position. The top three rows demonstrate the ability to recover latent images of non-planar scenes, with kernels restricted to 45×45 in size. Where distortion was large at the boundaries, we did not have sufficient calibration image coverage to recover the kernels, resulting in some artifacts.[4] The second row exhibits some artifacts due to repositioning error.

The bottom two rows correspond to a conference room scenario, where the photographs are taken of two consecutively projected slides on a presentation screen, as seen through obscure glass. This example illustrates the importance of using sufficiently large kernels. Severe artifacts are apparent for 45×45 kernels, with much better results obtained for 95×95 kernels.

We found that working with numbers of input images in the range of 100-200 generally worked well; not surprisingly, going much lower than this range

[3] The iPod nano had the added complication of recording compressed video, which resulted in severe blocking artifacts when compressing the noise patterns.

[4] We pad image boundaries of captured images using pixel replication as needed.

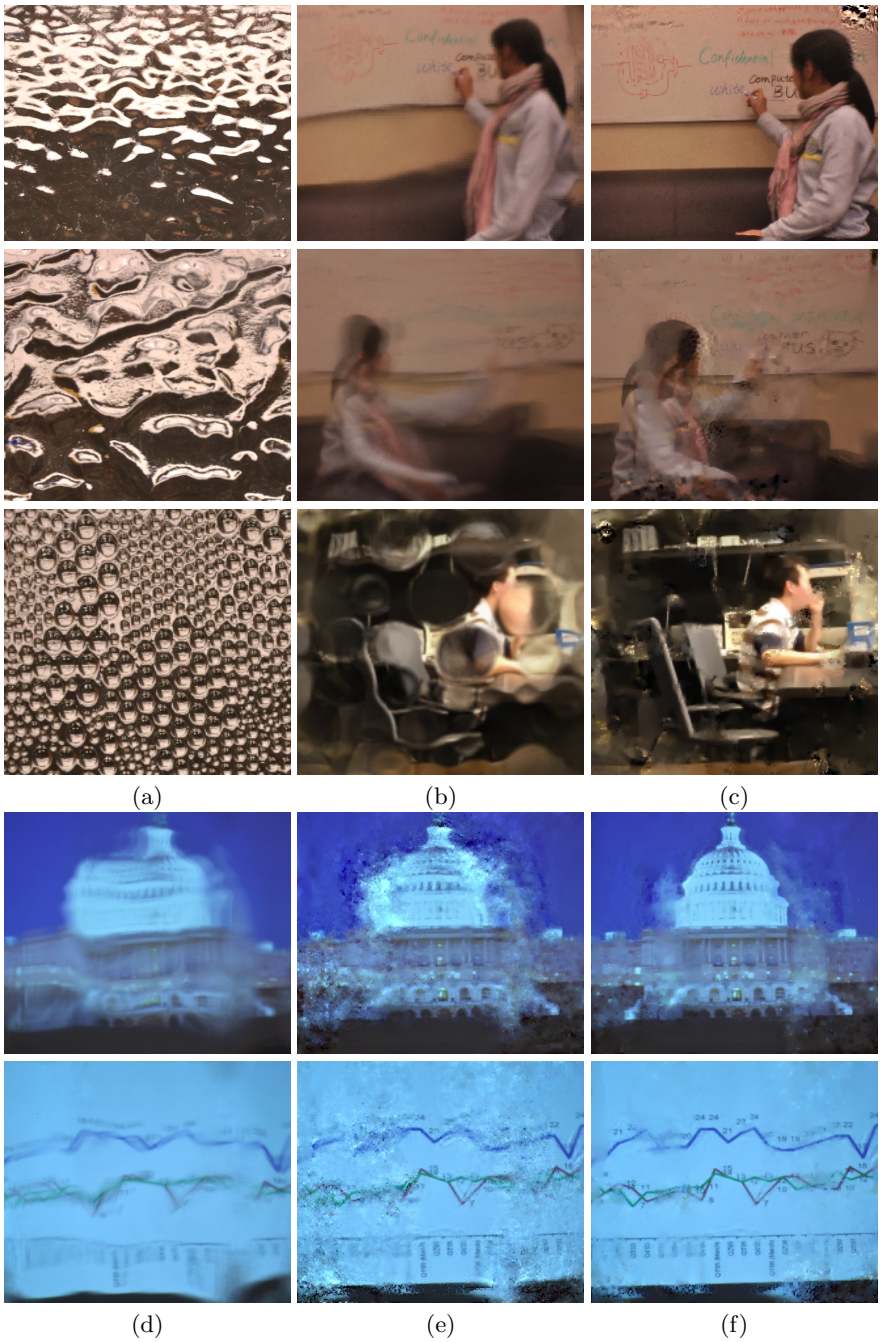

Fig. 5. Image reconstruction results with glass calibration. Column (a): three kinds of obscure glasses. Column (b): images captured through glasses in column (a). Column (c): reconstructed images with 45×45 kernels. Column (d) images captured through the obscure glass in Fig. 1 (b). Column (e): reconstructed images with 45×45 kernels. Column (f): reconstructed images with 95×95 kernels.

degraded the results. Solving for larger kernel sizes was generally preferable, at the expense of increased compute time. We also note that better results are possible without camera or monitor repositioning; we intentionally made the problem more challenging for the sake of realism.

Our kernel estimation procedure is fairly slow. For a 400×400 image, compute time is 40 CPU-hours for 45×45 kernels and over 200 CPU-hours for 95×95 kernels. The method is at least trivially parallelizable; we ran all of our computations on a cluster of 150 CPUs in under 2 hours in the worst case. High performance was obviously not our goal, though a notable area for future work. Deconvolution, by contrast, requires only one minute per 400×400 image.

6 Discussion and future work

We have posed the problem of trying to see through obscure glass, and developed both optical and calibration-based software techniques to do so. In our experience, the most effective solution is to apply a refractive index matching substance to the glass, essentially nullifying the effect of the glass when the match is exact. The match is, however, not always exact; further, when the textured surface is facing away from the viewer, it may not be feasible or desirable to leave a substance on that surface in real-world scenarios. Thus, it becomes important to undo the distorting and blurring effects of obscure glass. Our calibration approach is one way to accomplish this, and we have found it to be fairly effective when the blur and distortion are not extreme relative to the resolving capabilities of the optical setup. When the blur kernels are dense with large support, then it becomes difficult or impossible to recover a meaningful image. Thus, for privacy purposes, obscure glasses with dense blur kernels are preferable.

Our calibration-based approach is somewhat complex, requiring fairly careful re-positioning of the camera to recover a reasonably clear image, especially important for complex glass surface geometry. This step could be improved, e.g., by bracing the camera physically against the frame of the glass in a controlled manner. Another limitation of this approach is the required access to both sides of the glass at some point in time. An area of future work is to perform this step by taking a set of images of the natural scene on the other side and blindly recovering the latent images and the blur kernels. This scenario might be plausible if, for instance, one is attempting to recover images from a monitor, television, or projection screen on which slides or video are playing, providing a set of distinct images. Another future avenue would be to consider multiple viewpoints, moving the camera across the glass, and estimating the distorting and blurring structure of the glass, e.g., a height field and index of refraction.

Acknowledgements

We thank Steve Seitz and Sameer Agarwal for early discussions and ongoing feedback, and our sponsors, the University of Washington Animation Research Labs, Microsoft, Adobe, and Pixar.

References

1. Lide, D.R.: CRC Handbook of Chemistry and Physics, 90th Edition. (2009)
2. Sharpe, J., Ahlgren, U., Perry, P., Hill, B., Ross, A., Hecksher-Sørensen, J., Baldock, R., Davidson, D.: Optical projection tomography as a tool for 3d microscopy and gene expression studies. Science **296** (2002) 541–545
3. Trifonov, B., Bradley, D., Heidrich, W.: Tomographic reconstruction of transparent objects. In: Proc. Eurographics Symp. on Rendering. (2006) 51–60
4. Zongker, D.E., Werner, D.M., Curless, B., Salesin, D.H.: Environment matting and compositing. SIGGRAPH (1999) 205–214
5. Chuang, Y.Y., Zongker, D.E., Hindorff, J., Curless, B., Salesin, D., Szeliski, R.: Environment matting extensions: towards higher accuracy and real-time capture. SIGGRAPH (2000) 121–130
6. Peers, P., Dutré, P.: Wavelet environment matting. In: Rendering Techniques. (2003) 157–166
7. Wexler, Y., Fitzgibbon, A.W., Zisserman, A.: Image-based environment matting. In: Rendering Techniques. (2002)
8. Agarwal, S., Mallick, S.P., Kriegman, D.J., Belongie, S.: On refractive optical flow. In: ECCV. (2004) 483–494
9. Peers, P., Dutré, P.: Inferring reflectance functions from wavelet noise. In: Rendering Techniques. (2005) 173–182
10. Seitz, S.M., Matsushita, Y., Kutulakos, K.N.: A theory of inverse light transport. In: ICCV. (2005)
11. Sen, P., Chen, B., Garg, G., Marschner, S.R., Horowitz, M., Levoy, M., Lensch, H.P.A.: Dual photography. SIGGRAPH **24** (2005)
12. Ihrke, I., Kutulakos, K.N., Lensch, H.P.A., Magnor, M., Heidrich, W.: State of the art in transparent and specular object reconstruction. In: STAR Proceedings of Eurographics. (2008)
13. Murase, H.: Surface shape reconstruction of a nonrigid transparent object using refraction and motion. TPAMI (1992) 1045–1052
14. Levin, A., Fergus, R., Durand, F., Freeman, W.T.: Image and depth from a conventional camera with a coded aperture. SIGGRAPH (2007)
15. Shan, Q., Jia, J., Agarwala, A.: High-quality motion deblurring from a single image. SIGGRAPH (2008)
16. Yuan, L., Sun, J., Quan, L., Shum, H.Y.: Progressive inter-scale and intra-scale non-blind image deconvolution. SIGGRAPH (2008)
17. Kuhn, M.G.: Optical time-domain eavesdropping risks of CRT displays. In: IEEE Symp. on Security and Privacy. (2002)
18. Backes, M., Dürmuth, M., Unruh, D.: Compromising reflections – or – how to read LCD monitors around the corner. In: IEEE Symp. on Security and Privacy. (2008)
19. Backes, M., Chen, T., Duermuth, M., Lensch, H.P.A., Welk, M.: Tempest in a teapot: Compromising reflections revisited. In: IEEE Symp. on Security and Privacy. (2009)
20. LOREO: (Lens in cap) `http://www.loreo.com/pages/products/loreo_lenscap_spec.html`.
21. Kim, S.J., Koh, K., Lustig, M., Boyd, S.: An efficient method for compressed sensing. In: ICIP. (2007)
22. Johnson, M.K., Adelson, E.H.: Retrographic sensing for the measurement of surface texture and shape. In: CVPR. (2009) 1070–1077
23. Perlin, K.: An image synthesizer. SIGGRAPH Comput. Graph. **19** (1985) 287–296