The Magic Lens: Refractive Steganography

Marios Papas[1,2] Thomas Houit[2] Derek Nowrouzezahrai[1,3] Markus Gross[1,2] Wojciech Jarosz[1]

[1]Disney Research Zürich [2]ETH Zürich [3]Université de Montréal

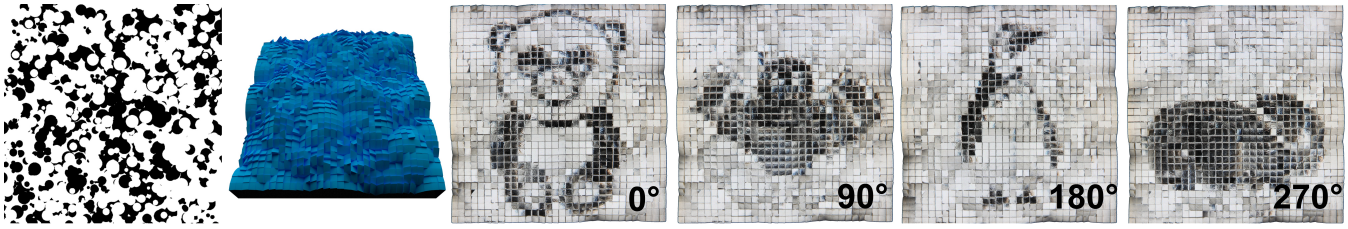

Figure 1: *We automatically design and manufacture magic lenses to warp source images into specified target images. Here we photograph a source image (far left) viewed through a manufactured lens with 32×32 facets (left), resulting in four images depending on the lens' orientation atop the source.*

Abstract

We present an automatic approach to design and manufacture passive display devices based on optical hidden image decoding. Motivated by classical steganography techniques we construct *Magic Lenses*, composed of refractive lenslet arrays, to reveal hidden images when placed over potentially unstructured printed or displayed source images. We determine the refractive geometry of these surfaces by formulating and efficiently solving an inverse light transport problem, taking into account additional constraints imposed by the physical manufacturing processes. We fabricate several variants on the basic magic lens idea including using a single source image to encode several hidden images which are only revealed when the lens is placed at prescribed orientations on the source image or viewed from different angles. We also present an important special case, the *universal lens*, that forms an injection mapping from the lens surface to the source image grid, allowing it to be used with arbitrary source images. We use this type of lens to generate hidden animation sequences. We validate our simulation results with many real-world manufactured magic lenses, and experiment with two separate manufacturing processes.

CR Categories: I.3.7 [Image and Video Processing]: Novel Display Technologies—Multi-View and 3D

Keywords: steganography, image morphing, lens fabrication

Links: DL PDF WEB VIDEO

1 Introduction

Steganographic techniques, from simple hidden message decoders to invisible inks and complex watermarking schemes, have led to active areas of research and have been applied in a wide variety of fields. Searching for and finding structure in unexpected places is also a fun and insightful process. Some common day examples of this expedition include the pursuits of a child armed with only a magnifying glass and their imagination, to a family huddled around a table, completing a jigsaw puzzle.

We leverage and incite this sense of wonder, encountered when inanimate objects suddenly convey a unexpected message or reveal surprising behavior, by combining the ideas of steganography, hands-on physical user manipulation, and structure from unstructured patterns. We design and construct several different types of *Magic Lenses*, using a custom computational procedure, capable of warping both structured and unstructured image sequences into unexpected target images. Our magic lenses are composed of lenslets that, when placed atop an image/video and viewed from prescribed locations, warp the image through refraction to form the desired images specified during lens generation.

We pose secret image encoding as an inverse light transport problem and present a fully-automatic approach for designing and manufacturing various types of magic lenses (see Figure 1). We experiment with various use-cases, for example enabling multiple target images to be warped from a single source image depending on the viewing angle between the user and the lens, or depending on the relative rotation or alignment of the lens and the source (see Section 6 for more results). In addition, while we experiment with two manufacturing processes to generate physical prototypes of hand-sized magic lenses, nothing about our technique precludes more exotic use-cases such as those depicted in Figure 2: e.g., replacing architectural fixtures with large-scale magic mirrors, revealing hidden messages for interactive and exploratory museum exhibitions, sending secret messages that can only be viewed with a user's magic lens, or embedding thin, flexible magic lenses in paper currency as an anti-counterfeiting and validation measure.

We are motivated by recent work on computationally embedding images into physical material properties, classic steganographic techniques such as the Cardan grille, as well as "magic decoder rings" which reveal secret messages already present in the source image using masking or subtractive transmission. In contrast, our lenses use optical refraction (or reflection), and we require little relation between the input and output images as long as the original

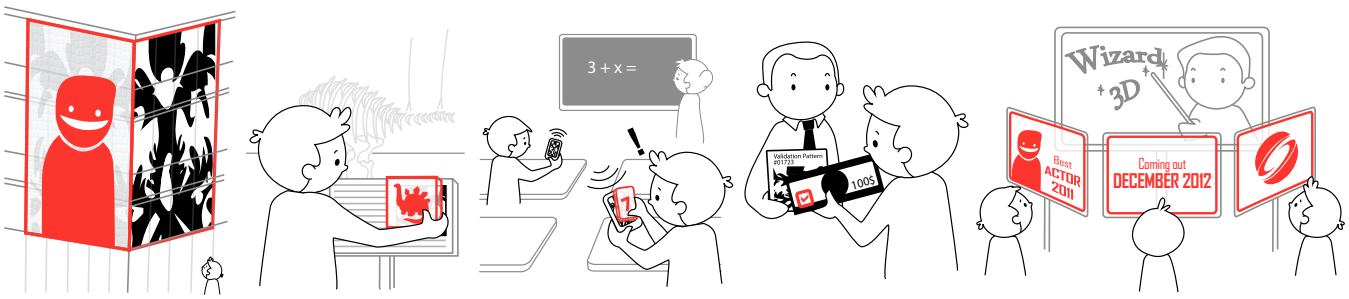

Figure 2: *A sketch of other potential applications. Left to right: architectural-scale magic lenses reflect hidden images to pedestrians as they walk past buildings; a teenager learns about dinosaurs by placing his magic lens over a museum exhibit pedestal; students find clever ways to cheat their professors, exchanging encoded SMS images which are decoded with magic lenses they manufacture at home and share amongst each other; a $100 bill has an embedded magic lens that, when placed over a scrambled verification pattern, confirms the validity of the bill; cinephiles view a movie billboard through specialized magic lens side-walk installations, revealing interesting facts about the movie.*

image contains all the colors of the target image. Furthermore, our approach can be passive, removing the need to carefully design or modify the source image (carrier signal) to encode the secret image. We also present an important special case of a magic lens called a *universal lens* (Section 5) that completely removes the dependence of the lens on the source image by generating an optical injection between the lens and an arbitrary source image grid.

Magic lenses are specialized passive display devices, related to light field displays, and we validate our simulation results with many real manufactured surfaces. Our approach is a first step towards realizing ideas such as those sketched in Figure 2. It has possible applications not just in art, education and entertainment (e.g. optical illusions, hidden message retrieval, optical decoder rings, holography), but also in banknote verification and security (see Section 8).

2 Related Work

Hiding Images. Recent works in computer graphics encode hidden visual information into 2D images or 3D objects. Different encoding and decoding schemes distinguish seemingly unrelated approaches that all share the same basic goal: encoding structured image information into potentially unstructured images or objects.

Autostereograms [Tyler and Clarke 1990] encode a (depth) image that can be decoded by a viewer by controlling vergence, while camouflage and emerging images [Chu et al. 2010; Mitra et al. 2009] hide images decoded by the temporal delays of the human visual system. Classical "invisible ink" writing using for instance lemon juice is only revealed when exposed to a catalyst such as heat. More recent computational approaches use metallic inks to embed images that only appear under specular reflection [Hersch et al. 2003], or print with specialized inks that are only visible under UV light [Hersch et al. 2007]. Band Moiré images [Hersch and Chosson 2004] encode a hidden image in interfering Moiré patterns caused by superimposing transparent sheets. We also manufacture transparent surfaces, but instead rely on a customized *refractive* lenslet-array to expose our hidden images and animations.

Recent techniques manufacture 3D objects which reveal hidden images when lit in a controlled manner. Niloy and Pauly [2009] design 3D objects which, when lit from different directions, cast distinct shadow images, while Baran et al. [2012] construct multilayer attenuators which cast several colored images depending on lighting conditions. Yue et al. [2012] allow users to manually arrange refractive pixel-like sticks in a grid arrangement to generate projective pixel art. Relief images [Alexa and Matusik 2010] encode two unique images in the diffuse shading of a height-field lit from two directions. Weyrich et al. [2009] design custom micro-facet BRDFs to embed images in the specular reflectance distribution of a material. Related works [Finckh et al. 2010; Papas et al. 2011] optimize the geometry of refractive surfaces to deform incident light into desired caustic patterns. Our work is, in a sense, the optical dual of such approaches. We apply a similar manufacturing process and optimization framework as Papas et al. [2011], but our optical system operates in reverse: we focus outgoing reflected or emitted light from a source image, using a lenslet array, to form an image (or stereo pair, or animation) directly on an observer's eye(s).

Patch Matching. The shape of our lenslet arrays are determined by matching regions between the source and target images. Many existing image feature and patch matching operations, commonly used in texture synthesis and image recognition [Lowe 1999; Pritchard and Heidrich 2003; Barnes et al. 2009; Barnes et al. 2010; Barnes et al. 2011; Barnes 2011], could be adapted to our task; however, our matching criteria is rather unique since we must constrain the smoothness of the resulting refractive lenslet array. As such, we use a simple matching procedure specialized to our task.

3D Displays. Simple autostereoscopic lenticular sheets [Lippmann 1908] "descramble" multiplexed left/right image pairs into distinct images seen from different viewpoints. Our approach can also be used to encode distinct images and, as such, can be seen as a special type of autostereoscopic display. We additionally encode target images completely unrelated to the structure of the source images by providing more degrees-of-freedom during lens optimization. With a directionally-dependent source image, our approach can also be interpreted as a way to optically warp one light field into another, opening up potential applications in light field display design [Gotoda 2010; Lanman et al. 2011; Wetzstein et al. 2011].

Hidden Watermarks. Another family of methods generate digital watermarks which can then be decoded using e.g. lenticular sheets [Alasia 1976; Alasia 1998; Brosh and Wright 1994]. These methods require that the parameters of the lenticular sheets are known prior to the generation of the watermarks [Renesse 2004].

3 Problem Statement and Goals

Our goal is to design and manufacture passive micro-lenslet arrays, which we call *magic lenses*, capable of generating one or many *target images* when placed over one or many physically displayed *source images* (e.g., printed on paper or displayed on a monitor). The manufactured magic lenses we present (see Section 6) are all hand-sized, however the approaches we present are not limited in

the size or form-factor of the lenses they can generate and one can imagine both larger-scale lens installations (in architectural settings) or miniaturized lenses (for banknote verification).

We will begin by describing the simplest case of warping a single source image to a single target image (i.e., viewed from a single position), and later address more interesting use-cases (Section 4.6).

Each magic lens is divided into an $N \times N$ grid of refractive *facets*[1] which we initially assume to be flat to simplify our exposition. Our target images are also divided into $N \times N$ tiles, which we wish to observe through the facets. We decompose source images, which are warped by our lenses, into a set of unstructured *image patches*, which are arbitrarily-sized rectangular regions in the source.

Each facet is constructed to refract light from the source image so that, when viewed from above, the image seen through each facet matches that of the corresponding image tile. In Section 4.5 we eliminate the flat-facet constraint using a multi-scale refinement, and so our finalized magic lenses are ultimately composed of a grid of $N \times N$ smooth lenslets. We will present two different types of magic lenses: *source-optimized* lenses and *universal* lenses.

Source-Optimized Lenses. In the simplest case, a source-optimized lens will "index" (through refraction) only the subset of source image patches necessary for it to reproduce its target image. As such, a source-optimized lens is "tied" to its input source image. This constraint allows us to better enforce smoothness and manufacturability, but at the cost of flexibility. This case is particularly useful when the source image must remain fixed (e.g. in the case of a well-known painting). We detail our approach for generating source-optimized magic lenses in Section 4.

Universal Lenses. These lenses are an important special case and are constructed so that each lens facet indexes a unique source image patch, forming an *optically injective mapping* between the lens and the source (i.e., each pixel in the source image maps to *at most* a single facet). As such, a single universal lens can be used with an arbitrary number of source images which can be scrambled, according to the optical bijection, in order to generate a target image. This case is most useful when there is flexibility in manipulating a source image or when a single lens must be used with an undetermined set of images. We discuss the modified approach we use for universal lenses in Section 5.

4 Generating Source-Optimized Lenses

Source-optimized lenses are most related to goal-based caustics [Papas et al. 2011], where a lenslet array warps incident light into a desired caustic image. Papas et al. form a bijection between facets and caustic image splats, however our search space is far larger as we are not constrained to the subset of solutions induced by one-to-one mappings. Instead, for source-optimized lenses, each facet has complete freedom to refract to an arbitrary source patch.

While having this larger space does expose a wider range of potential facet-patch pairings, special care must be taken to maintain a reasonable computation cost. We juggle these constraints using a multi-step and multi-pass approach: we ensure a baseline image quality using facet-patch matching (Section 4.2) *prior to* optimizing the geometric lens quality (Sections 4.3 and 4.5).

We decompose the procedure for generating source-optimized magic lenses into four mandatory (and one optional) steps (Sections 4.2 to 4.5): facet-patch matching, facet orientation, simulated annealing, and multi-scale height optimization are the mandatory steps. Optional facet-patch matching variants are detailed to handle special use-cases (Section 4.6) such as multi-source/multi-target warping, multi-view warping, and multi-rotation warping. These steps are repeated until a termination criterion is met. Figure 3 overviews the entire procedure.

[1] A similar principle could also be applied to *reflective* facets.

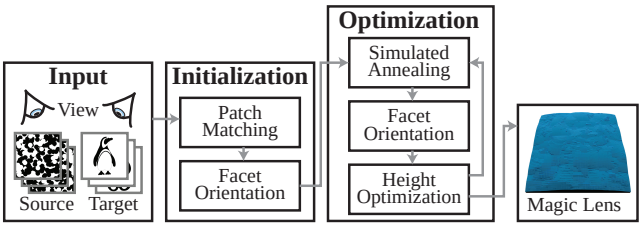

Figure 3: *Overview of source-optimized magic lens generation.*

4.1 Input

Source-optimized lenses require at least one input image and one or more target images to be specified. Focal distances from the lens to the source, and from the viewer to the lens must also be specified.

4.2 Facet-Patch Matching

Each facet of a source-optimized lens should refract a portion of the source image that matches the facet's associate target image tile. The first step of surface generation is to find and assign suitable facet-to-patch matches.

We divide *target images* into grids of $N \times N$ tiles and map each tile to its facet on the lens. Now, for each facet, we must find and rank regions in the *source image* which can serve as potential matches (according to pixel-wise differences) to the target image tiles.

We then search image regions within a neighborhood $\{R_x, R_y\} \in [-m/2, m/2]^2$ of $m \times m$ pixels in the source image, starting directly underneath the facet (assuming the lens is aligned atop the source). We compute a *matching score* for every eligible region, $M(R_x, R_y) = M_d(R_x, R_y) + \lambda M_s(R_x, R_y)$, where M_d is the image color matching term

$$M_d(R_x, R_y) = \sum_{i \in \mathbf{P}} \sum_{j \in \mathbf{P}} [\mathbb{T}(i, j) - \mathbb{S}(i + R_x, j + R_y)]^2 \quad (1)$$

and $M_s(R_x, R_y) = (m^2 - R_x^2 - R_y^2)^{-1} - m^{-2}$ is a smoothing term that penalizes distant matches which cause steep normals. Here, $\mathbf{P}$ is the target patch region, $\mathbb{S}$ indexes the source, $\mathbb{T}$ indexes the target, and $\lambda = 255^2\ \mathbf{P}^2$ is our regularization constant.

We compute M for every $R_x \times R_y$ offset and additionally compute matching scores at multiple scales S. We initially associate the best (lowest) match to the facet. Apart from a matching score, each potential match is assigned a *facet-patch transformation* (FPT) which consists of the match's (R_x, R_y) offset and the corresponding scale S, which map the facet to the associated region on the source. All matches with scores below an acceptance threshold are retained and sorted for future consideration, at each facet.

This marks the half-way point of the initialization phase for source-optimized lens generation.

4.3 Facet Orientation

Once each facet has been associated its top-ranked FPT, we proceed with the final initialization phase: assigning an orientation to each

facet in order to induce the refraction necessary to match its FPT's transformation. It is important to note that facet heights are not yet modified and thus, after these first two initialization steps, the surface is discontinuous and therefore not manufacturable.

To determine facet normals we solve an inverse lighting problem: the normal of each facet is computed so that an eye ray passing through the facet center will refract to the matching source image patch. We define this mapping as $g_0 : (x, y, \Delta x, \Delta y) \rightarrow n_{xy}$, where (x, y) index the lens facets in the plane of the lens, $(\Delta x, \Delta y)$ are the offsets onto the source image from the facet (in the source image's parallel plane), and n_{xy} is the facet normal. With $\theta_b = \arctan(d/D)$, we can solve for θ_r using Snell's Law and the fact that $\theta_i = \theta_r - \theta_b$:

$$\theta_r = [\arctan(\eta_o - \sin\theta_b)] \,/\, [\eta_o \cos\theta_b - \eta_i] \ . \tag{2}$$

We solve for the surface normal that induces the unique refractive light path between the source patch and facet with a tailored version of Walter et al.'s [2009] technique, which simplifies the problem to discrete root finding solved using a handful of Newton iterations.

At this point, the initialization of the source-optimized lens generation is complete and we enter the optimization phases.

4.4 Simulated Annealing

We find solutions in our large, non-linear search space using a stochastic simulated annealing search. This search will consider FPTs across neighboring facets while increasing surface smoothness and maintaining high quality reproduction of the target image.

We enforce smoothness between neighboring facet normals using stochastic search: we choose a facet at random and consider swapping its current FPT with another from its list (see Figure 5) if the new FPT assignment improves the lens smoothness of the lens (see below). We repeat this process at least 25 N^2 times.

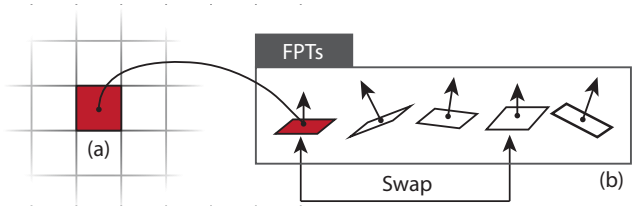

Figure 5: *During simulated annealing, a facet is randomly selected (left) and its current FPT is swapped with another from the list of FPTs generated during facet-patch matching (right).*

When considering the swap we compute the facet's new normal using facet orientation, and we compute a *facet energy* consisting of a geometric smoothness term E (see below). Prior to the first swap, the initial lens energy $\mathrm{E}_{\text{initial}}$ is the sum of E for every facet.

After computing the facet energy E, we perform an *energy validation* step to check whether the change in facet energy satisfies a threshold $T = \mathrm{E}_{\text{initial}} \cdot \mathbf{e}^{0.01t} / \left(100 \cdot N^2\right)$, where t is the iteration number corresponding to simulated annealing's temperature metaphor. After repeating the swapping and energy validation steps, accepting FPT changes only if the energy change is below the threshold T, we compute the *total lens energy* as the sum of all facet energies. If this energy is lower than $\mathrm{E}_{\text{initial}}$, we update the lens with the new facet settings and t is incremented.

This entire process is repeated until either the total energy of the lens does not decrease for 10 iterations in a row, or the reduction of energy is less than $10^{-1} \times \mathrm{E}_{\text{initial}}$. An example of a lens before and after simulated annealing is shown in Figure 4 with its corresponding simulated results.

Smoothness Term for Normals. The per-facet geometric smoothness energy term

$$E = \alpha||\Delta g_0||^2 + \beta||E_1||^2 + \gamma||E_2||^2 \tag{3}$$

is computed every iteration at each facet. Here Δg_0 is a discrete geometric boundary gradient term

$$\begin{aligned} A &= ||g_0(X+1,Y) - g_0(X,Y)||^2 = ||(a_x, a_y)||^2 \\ B &= ||g_0(X-1,Y) - g_0(X,Y)||^2 = ||(b_x, b_y)||^2 \\ C &= ||g_0(X,Y+1) - g_0(X,Y)||^2 = ||(c_x, c_y)||^2 \\ D &= ||g_0(X,Y-1) - g_0(X,Y)||^2 = ||(d_x, d_y)||^2 \\ \Delta g_0 &= \sqrt{(A+B+C+D)/4}\,, \end{aligned} \tag{4}$$

measuring the deviation of the two normals in the x (A and B) and y (C and D) directions, and E_1 and E_2 are component-wise geometric error terms

$$E_1 = (a_y + b_y + c_x + d_y)/4, \quad E_2 = ((a_y - b_y) - (c_x - d_x))/8,$$

which correspond to deviations of the normal's y components in the x direction, and the normal's x components in the y direction (the change of the z component is corrected during height optimization; see Section 4.5) for E_1, and a constraint on the integrability of the surface for E_2. We normalize each term and set $\alpha = 0.125$, $\beta = 0.375$ and $\gamma = 0.5$ so that $E \in [0, 1]$.

4.5 Multi-Scale Height Optimization

Once every facet has been assigned an FPT (and hence an orientation), we formulate the problem of optimizing heights for surface smoothness as a global optimization in the heights of the facets. This stage serves two important purposes. Firstly, facet heights are adjusted to maximize coincidence with neighboring facets along facet boundaries (see Figure 6), which increases manufacturability[2]. Secondly, after facet heights are optimized, each facet is decomposed at a finer-scale into a grid of micro-facets which are optimized to approximate a smooth lenslet shape over the facet. Height optimization is equivalent to integrating a normal field into a height field and we use the over-constrained linear system of Papas et al. [2011] to solve this problem.

We repeat height optimization with facet orientation until convergence since: any height change may invalidate the selected FPT, which would require facet re-orientation; and, any facet re-orientation may reduce facet boundary smoothness, requiring further height optimization. After facet height and orientation are set, we dice facets into grids of micro-facets, performing a nested iteration of facet-orientation and height optimization on the micro-facets grids of each facet in order to generate smooth lenslets instead of planar facets for our final surface. Micro-facets are generated once, at the end of the entire procedure.

4.6 Facet-Patch Matching [optional]

Several variants of the standard single-source/target magic lens can be realized by adjusting the facet-patch matching process. We detail three important examples below, but note that the application of these steps is optional depending on the desired use of the lens.

[2]Smoothness is necessary for manufacturing with an engraving milling machine, but can be relaxed when using a 3D printer (Section 6).

Figure 4: *Source/target images (left pair) are mapped with a source-optimized lens without (middle set) and with (right set) simulated annealing. Simulated results are indistinguishable but the lens generated with simulated annealing is smoother and manufacturable. The lens (128×128 facets with 11×11 micro-facets) is designed to be placed at $10cm$ from the source and 40 cm from the viewer.*

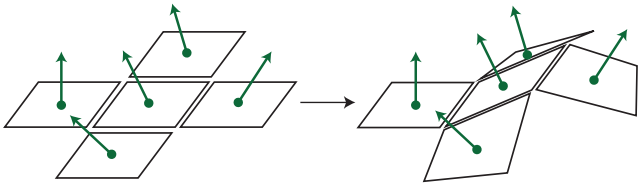

Figure 6: *Height optimization displaces facets with (non-geometric) normals (left) along z, enforcing smoothness along facet boundaries and inducing the correct geometric normal (right).*

Rotating Lenses. One interesting use-case for source-optimized lenses is when a single lens produces several different target images depending on its relative rotation about the source image. To realize this type of lens, we can extend our matching scores to include scores (and FPTs) for rotated versions of the source image. We do so by augmenting the source sampling function $\mathbb{S}$ with a rotation parameter, and, for our square-shaped lenses, we have experimented with $0°, 90°, 180°$, and $270°$ rotations.

Multiple Source-to-Target Pairings. Another use-case is when a single source-optimized lens is used to warp several source images to different target images. In this case the matching score is computed as the sum of matching scores for each source/target pair.

Multi-view Lenses. Multi-view source-optimized lenses are a more complex use-case: a lens that warps a single source image into different target images depending on the viewing angle of the user with the lens. This is a more complicated case since changes in the FPT for a given view direction induce changes in other view directions. We alternatively set each viewing direction as a "master" direction and optimized separately according to its FPT. The final matching score, as in the case of multiple source-to-target pairings, is the sum of the matching scores over all viewing direction conditions given the facet orientation for the master direction. After each direction has been selected, multiple matching scores corresponding to each master direction are sorted and only the top score corresponding to the best master direction is kept for each FPT.

For each multi-target use-case, we require an additional optimization step after height optimization to re-introduce new FPTs that take the updated lens geometry into account. This occurs since height changes affect the regions each facet can "index" from the source (and thus invalidates the facet's FPTs).

4.7 Output

The result of this process is a mesh of a magic lens surface which we then forward to one of two manufacturing pipelines for physical construction (see Section 6).

5 Generating Universal Magic Lenses

While source-optimized lenses are tied to the input source images, we can also devise a generalized patch-matching process that does *not* depend on the source image, allowing us to construct a magic lens that can be re-used with an arbitrary number of different sources. The dependence on the source image in source-optimized lenses arises from the fact that many lens facets may target the same region on the source image. Figure 7 visualizes the difference between the facet-patch mappings of a source-optimized lens and a universal lens, where we can clearly see that each patch on a source image is sampled by *at most* one facet on the universal lens.

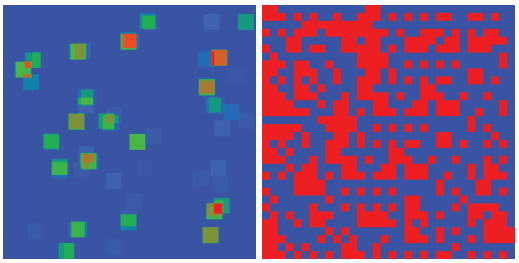

Figure 7: *Heatmap for a source-optimized lens (left) with values from 0 (blue) to 18 (red) overlaps. Clearly, several facets are matched to the same source patch. The heatmap for a universal lens (right) illustrates the injective mapping: red corresponds to a single match and blue to no match.*

Such a *universal lens* could be thought of as a "key" that can "unlock" hidden target images embedded in seemingly arbitrary (but specially authored) sources. An example use-case where this could be useful is the generation of encoded animations (see Figure 9).

The fundamental difference between universal and source-optimized lenses is that, once constructed, source images must abide by the optical injective mapping induced by the lens in order to be decoded by it. Figure 8 overviews the different steps of the universal lens generation process: assuming a single-view use-case again for simplicity, we substitute the facet-patch matching and simulated annealing steps of the source-optimized procedure with a customized dart-throwing approach that forms the injective mapping between our lens and the source. In addition, after repeating iterations of height optimization and facet orientation, as in the source-optimized lens procedure, we follow with an inpainting stage that embeds the target image into an arbitrary source image.

5.1 Orientation Initialization with Dart-Throwing

We devise a modified dart-throwing approach which replaces facet-patch matching and facet orientation with: choosing a patch on the source image to match to facets on the lens, and determining the orientation of the facet in order to target the matched patch. For universal lenses, however, each source patch can only be matched to one facet at most.

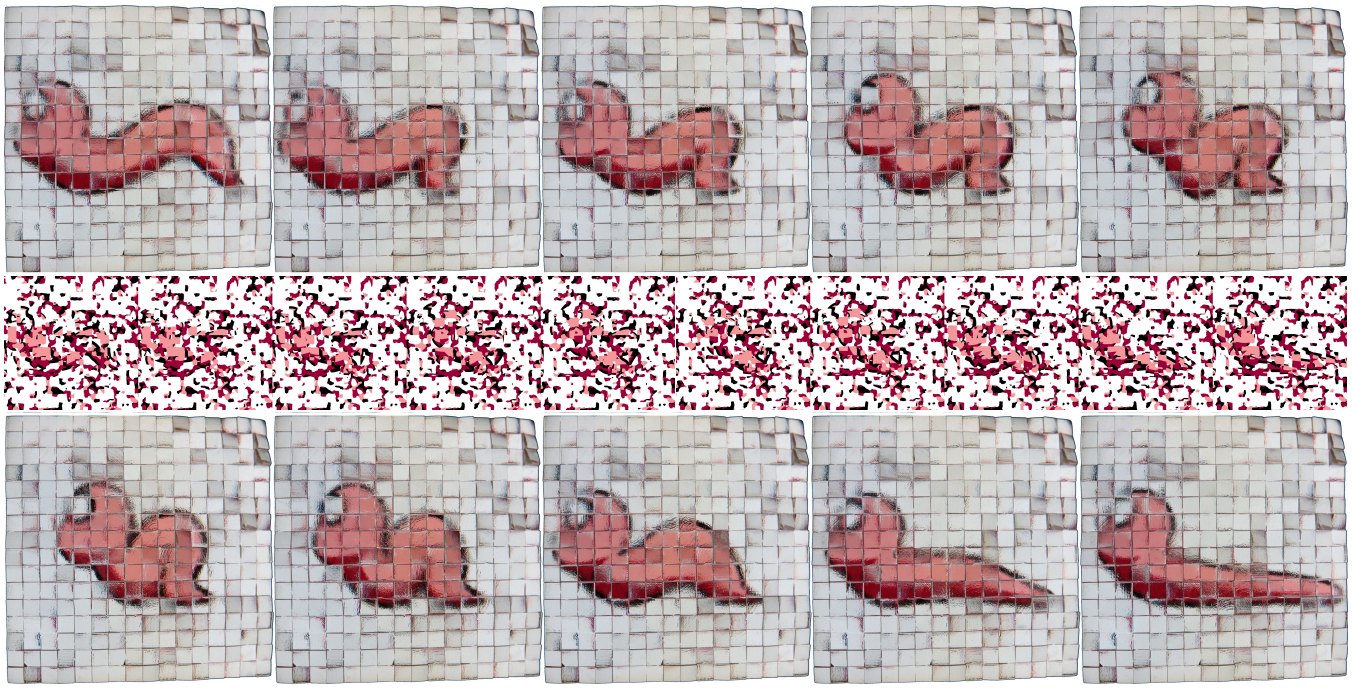

Figure 9: *Warping an unstructured video sequence with a universal lens with* 16×16 *facets. The middle row has the different (shrunken) source images, and the top and bottom rows are photographs of a manufactured lens placed over the source images as they animate.*

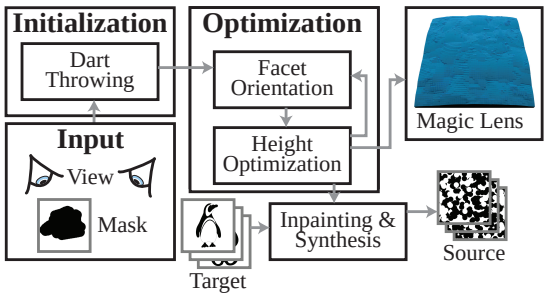

Figure 8: *Overview of universal magic lens generation.*

We employ a multi-scale approach, dividing the source into $N \times N$ square patches, aligned with the facets of the lens. Then, each patch is divided into smaller square sub-regions in order to promote flexibility in finding a match without resulting in an identity mapping. At each facet, we then select one of the physically-reachable (via refraction) regions at random and mark it as used so it cannot be selected by any other facet. After repeating for all the facets we are left with an injective mapping between the facets on the universal lens and the sub-regions (and patches) on the source.

We can optionally supply a source *mask* image to invalidate regions on the source from use in the dart-throwing selection procedure. The mask can be useful to reserve regions on the source images that are guaranteed to not perturb the output target image.

5.2 Inpainting and Texture Synthesis

Given a universal lens, specified by the output of the dart-throwing process, and a target image, we take an empty (or arbitrary) source image and inpaint the sub-regions in order to match the target image, given the injective mapping.

Once in-painting is complete, some pixels of the source image (which are never indexed by any facets) are uninitialized (white). We optional perform an additional texture synthesis step in order to fill the uninitialized regions with camouflage: we isolate unmapped sub-regions in the immediate neighborhood surrounding mapped sub-regions, and synthesize an input pattern or a background image in these regions. This approach was employed in our *worm animation* example in Figure 9.

5.3 Iterative Adjustment and Convergence

After the execution of the dart-throwing injective initialization and the inpainting steps, the final universal lens is obtained by similarly repeating the height optimization and facet orientation stages until our stopping criterion is met. As with source-optimized lenses, this results in a mesh surface which we forward to our manufacturing pipeline for physical construction (Section 6).

6 Fabrication Details and Results

We evaluate our lenses with simulated and fabricated results.

Fabricated Results. We physically manufacture our lenses using two separate processes: milling of acrylic blocks using a computer controlled engraving machine (a Roland EGX-600), and 3D printing with a clear material on a fast 3D prototyping machine (an Object Connex 350). These two manufacturing modes have different physical constraints, which we can account for in our optimization. All lenses are 10×10cm, with thickness varying between 2 to 4 cm. All lenses are placed directly on the source image, unless stated otherwise in the caption.

Figure 11 shows a source-optimized magic lens that transforms an image of random circles into a panda. We visualize the magic lens surface and show a simulated result as seen through the lens. We

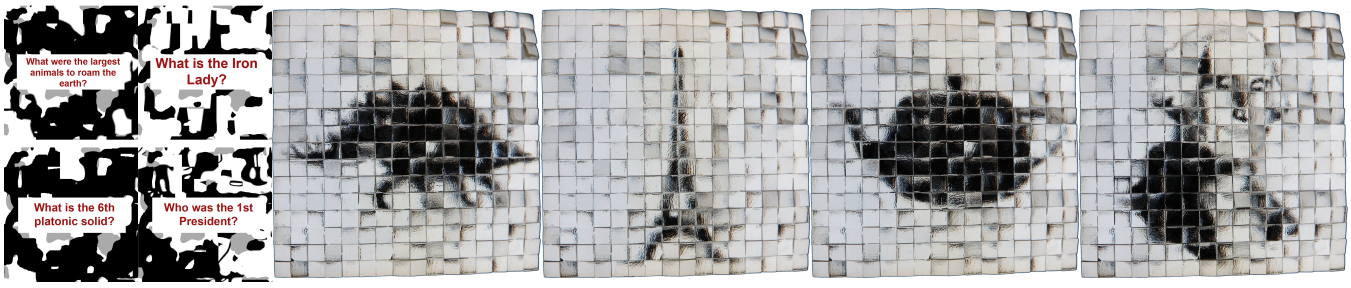

Figure 10: *Four source images with question (far left; center zoom-ins) are warped to reveal pictographic answers with a single manufactured magic lens (16×16 facets). Note that the region containing the question is not targeted by any of the facets. More examples and the full-sized sources are included in our supplemental document.*

manufactured this lens using our milling approach — a photograph of the fabricated lens placed over the source image is also shown.

The precision of the milling machine is higher than our 3D printer, but the surfaces it produces have a frosted finish and need to be manually polished. This can be quite time consuming, and can also alter the shape of the lenslets, reducing quality. Because of this, and because 3D printed lenses do not need to be as continuous, we produce all our remaining results with the 3D printing method.

In Figure 12 we show another source-optimized magic lens with the panda as a target image but, when viewed at a different angle, the lens reveals a penguin image, all using the same random circle pattern as a source. Figure 1 also uses the same random circle source, but the source-optimized lens is capable of displaying four different targets (panda, bat, penguin, and whale) at four 90 degree rotation increments.

We also manufactured a universal lens that we place over a scrambled image sequence (Figure 9). When viewed through the physical lens, an animated sequence of a worm is revealed (see our supplemental video).

Magic lenses can also be used for more creative interactions. For instance, question-answer lenses can be designed to reveal pictorial answers when placed over source images posing questions. Figure 10 shows four different question images being answered using the same universal lens as in Figure 9. Since we use a universal lens, there is actually no limit on the number of randomized question-answer image pairs that can be used with this single lens (see the supplemental document for more examples).

Simulated Results. In Figure 4 we illustrate, using simulation, the theoretical output that could be obtained with a higher quality manufacturing process. We simulate a high resolution lens (128×128 facets) that warps two images with fine-scale details. Our simulated annealing process is clearly capable of generating smooth, continuous lenses, but these lenses are beyond the manufacturing capabilities of our current milling or printing processes (especially at hand-held output sizes). In Figure 16 we show another example of a question-answer lens, this time source-optimized to the four question images, and simulated at a higher resolution (128×128 facets) than our current manufacturing processes allow.

7 Discussion

We have presented the idea of magic lenses as well as a general approach for creating various special case lenses. This is a significant first step towards realizing additional use-cases in the future, and we will discuss lessons learned and some considerations for extending our work in this section, before itemizing concrete ideas for future work in Section 8.

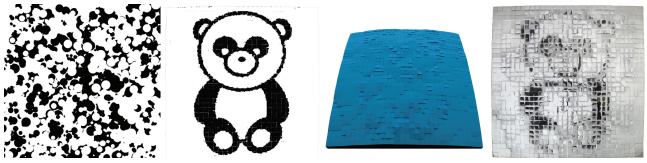

Figure 11: *A physically milled result. Simulation (middle left) closely matches the target, however the smooth lens (middle right; 32×32 facets with 11×11 micro-facets) is degraded during milling and manual polishing, resulting in a lower quality physical result (right) than that of a 3D printer (see Figure 12).*

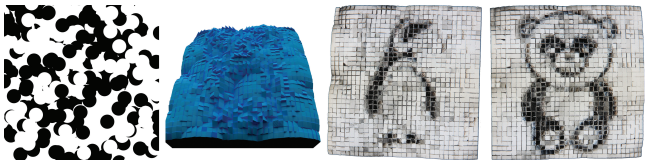

Figure 12: *A physically printed multi-view result. Left to right: source image, lens (32×32 facets), and photographs of a manufactured lens viewed from two locations.*

Manufacturing Quality and Considerations. Our proof-of-concept manufacturing process cannot yet yield lenses which match the high quality of our simulation results. In general, these discrepancies arise due to the differences between our idealized model of specular refraction used in simulation, and the additional surface roughness introduced during the manufacturing process (for both milling and 3D printing processes). In the case of milling, surface roughness is further diminished through a manual polishing post-process, which is necessary to render the lens smooth enough for refraction but which also degrades the surface quality and accuracy (when compared to the prescribed mesh output). Milling also imposes additional surface smoothness constraints, limiting the space of image patches each facet can index. To alleviate this constraint, we place milled lenses higher above the source and not directly atop it (as is the case for 3D printed lenses). Figure 11 is the only milled result we illustrate, and it is clearly of lower quality than the results generated using the 3D printer, due to the aforementioned issues. This lens is placed 10cm from the source and is intended to be viewed at 40cm above the lens.

As can be seen in our video, our fabricated lenses are currently thick and likely too bulky and expensive for immediate public use. Improved manufacturing processes in the future can permit us to

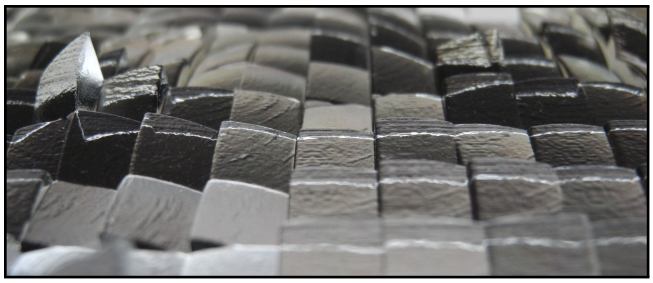

Figure 13: *A photographic zoom-in of one of our 3D printed magic lenses. Notice the easily visible z layers deposited by the printer and the semi-transparent nature of the printing material.*

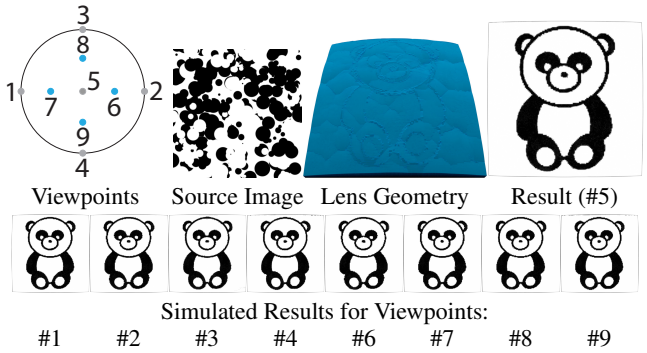

Figure 14: *We induce a stable viewing zone by exploiting multi-view lens generation in a novel way: by using the **same** target images at each of the nine views, arranged in the layout on the left. The lens consists of 128×128 facets.*

shrink the thickness of the final lens object without compromising the warping capabilities of the lens, especially in the case of 3D printed surfaces where smoothness is less of an issue. In general, we do not see any reason why a higher quality (e.g., commercial) manufacturing process could not rectify many of the discrepancies between our simulation and manufactured results.

Physical Limitations. We consider the physical manufacturing limitations that must be taken into account when designing our lens optimization. We constrain the search range of each facet to 10% of the source image width to reduce the maximum steepness of facets and increase the lens surface regularity. When milling, lens thickness is limited to 4cm by the milling machine and vertical transitions are limited in range due to the V-groove shape of the milling bit. On the other hand, two pragmatic limitations of the 3D printing process are its price and limited resolution. Transparent 3D printing material costs several orders of magnitude more than an equivalent volume of acrylic, and the 3D printer's precision is limited for our purposes, with significant surface roughness added as a side-effect of printing (Figure 13).

We found that, despite the additional undesired surface roughness introduced by 3D printing process, we can actually place the final lens objects *directly* on the source image mainly due to fewer constraints on inter-lenslet smoothness. By moving the lens closer to the source image we reduce blurriness and calibration sensitivity, resulting in increased usability. Reducing the distance between the lens and source often requires an increase in the facets' slopes. This increase introduces a non-negligible Fresnel reflection effect, especially when the angle between the average facet normal and the view direction is greater than 60^o. We account for these effects during lens construction using the search distance threshold and using the physical lens in low (external) light settings.

Source Image Constraints and Manipulation Currently, the two types of lenses we produce place different constraints on the input. Source-optimized lens assumes a very rigid constraint on the source image (namely, that it is provided by the user as input and cannot be edited), whereas the universal lens allows complete freedom when generating the source image(s). However, one could imagine a continuum of approaches between these two extremes: for instance, where a source image is provided, but the optimization is allowed to modify it for increased fidelity or smoothness. This type of approach could be very useful for our source-optimized examples which use an unstructured random pattern as the source. We have indeed experimented with applying variants of inpainting and texture synthesis procedure *during* the simulated annealing component of source-optimized lens generation. In doing so, we allow (un-masked) regions of the source image to be manipulated during lens optimization. Our initial results have proven promising and have permitted the generation of smoother and more accurate source-optimized lenses.

Use-case Customization. We have demonstrated that our approach can be readily and easily extended to handle more complex use-cases (see Section 8 for examples, as well as Figure 2), and it is our hope that this will promote future work in not only investigating interesting use-cases that we have not conceived, but also in using our approach for larger-scale applications.

Another example of this type of customization is presented in Figure 14, where we modify our multi-view lens generation to enhance the viewing stability of our manufactured lens: we generate a multi-view lens, with views arranged in a circular-cross layout, and with all the target images set to the same panda image. The apex of the inner viewing cone is 1.15^o. The resulting multi-view lens affords a much larger "viewing zone", in case a user slightly shifts their viewing position.

Cryptographic and Steganographic Strength. For certain applications steganographic or cryptographic strength is an important consideration. For a useful discussion, "strength" needs to be better defined in our context and depends on the usage scenario. If an "attacker" gains access to the lens, the mapping can be easily inverted (i.e., it is weak). However, for image-specific lenses, if the attacker obtains the source image it is impossible to deduce the target without the lens (i.e., it is strong).

For our physically-manufactured universal lenses the mapping is not cryptographically strong since parts of the target could be deduced from the source. However, this is a limitation of the manufacturing quality (using fewer/large facets), and not of our overall approach. With improved manufacturing, or in simulation, our universal lenses can be "cryptographically strong" in the limit of pixel-wide facets: e.g., by randomly scrambling the image pixels plus injecting arbitrary noise pixels to further evade detection.

A useful analogy of our universal lenses are paper shredders (cutting images into small slices) but our method shreds and scrambles in 2D while also allowing for zooming effects. A shredder that produces only a few slices (large facets) is less secure. Security is improved by increasing the number of slices (smaller facets) and mixing slices with other random shredded documents (inpainting with random pixels). We illustrate the mapping obtained from a universal lens that satisfies manufacturing constraints in Figure 15.

Visual Cryptography [Naor and Shamir 1994] can provide stronger

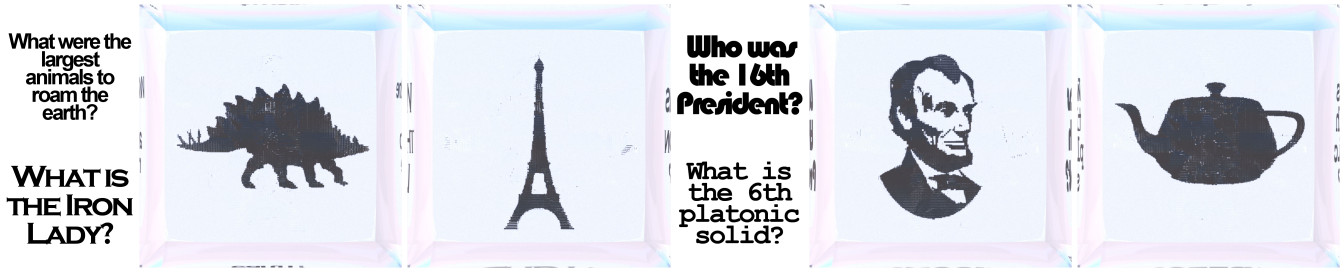

Figure 16: *A single source-optimized lens (128×128 facets) encodes four hidden pictographic answers to four textual questions (simulated).*

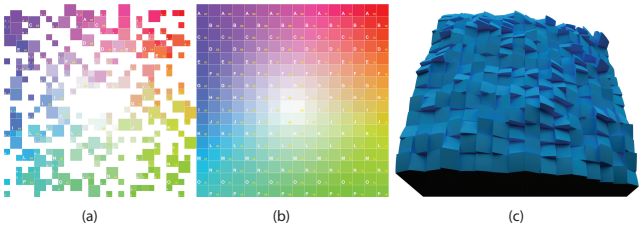

Figure 15: *Visualization of the warping behavior of a universal lens with* 16×16 *facets: Given the target image (b) and lens geometry (c), we compute the source image (a) using our inpainting method.*

cryptographic guarantees than our method. Our method does not explicitly optimize for cryptographic strength, though we believe this is an interesting avenue for future work.

3D Displays. Mass-produced 3D technology (e.g., lenticular sheets, parallax barriers) are less expensive to manufacture than our lenses, but these have significant limitations such as low signal to noise ratio, aliasing and crosstalk. Customizing 3D display optics to the expected displayed content is one possible direction which could overcoming these issues. Our lenses are an example of such customized optics and, although they are not yet cost-effective, fabrication quality and cost will only improve in time, opening the door to content-optimized display techniques, like ours, in this domain.

8 Conclusion and Future Work

We introduced *magic lenses*, a form of passive display device for exposing hidden messages and images from both seemingly random as well as structured source images. We are able to efficiently navigate a large search space of viable solutions using stochastic search and multi-scale height optimization and our process also addresses manufacturability constraints of the final output lens surfaces. We have manufactured physical prototypes using two processes: commodity 3D printing and engraving with a milling machine. We obtained high quality simulation results and also demonstrated promising manufactured results, despite using a proof-of-concept manufacturing process.

Future Work. In the case of colored source and target images, the source *must* contain all the colors desired in the target image. One can imagine using back-lit source projectors and modeling, or even inducing, *dispersion* effects in order to side-step this constraint.

As we currently assume perfect refractive surfaces, a better modeling of the surface roughness resulting from our manufacturing processes can be achieved by using a physically based BSDF such as the one introduced by Walter et al. [2007]. Renderings of Magic Lenses with the aforementioned rough glass BSDF can be found in the supplemental document. This may lead to better agreement between simulated and manufactured results, however we believe that a higher-quality manufacturing process is a more important first step of investigation towards this end.

We similarly also assume perfect diffuse reflectance of the source image(s), and an interesting direction of future work would be to model the effects of glossy reflection or directional emission off of the source. In this case, the observed radiance of each facet will change not only as a function of location but also as a function of the facet's "viewing" direction. As discussed in Section 7, investigating methods for optimizing or manipulating a set of random source images to best reproduce a desired set of target images is also an interesting area of future work.

Our multi-view magic lenses may be suitable as a form of stereoscopic or light field display. We have experimented with the stereoscopic use case (using a 3D stereo pair as target images, and specifying view conditions that line up with inter-eye distance), however with mixed results. We have yet to find a solution that consistently generates stereo image pairs that are comfortable to fuse.

As we feel that the general idea of magic lenses can be exploited for many different specialized applications, the most exciting areas of future work may in fact be the ones we have not yet conceived of. It is quite possible that tailored magic lens solutions will find their way into unique optical illusions (e.g., imagine an enhanced version of a magic mirrors room at an amusement park), cheap personal encryption, holography, banknote security, or architectural art (Figure 2).

Acknowledgments. We thank Olga Sorkine, Paul Beardsley, Simon Heinzle and Marc Pollefeys for their feedback on an early draft of the paper, as well as the anonymous reviewers for their helpful feedback and recommendations. We thank Maurizio Nitti for the worm animation in Figure 9 and Alessia Marra for the sketches in Figure 2.

The Panda and Bat images are available from `webdesignhot.com` and the Penguin, Whale, Lincoln, Washington, Eiffel Tower and Dino images from `clker.com` as public domain. The Floral and Truck images in Figure 4 were provided as public domain by `keepdesigning.com` and `chadlonius.com` respectively.

References

ALASIA, A. V., 1976. Process of coding indicia and product produced thereby. U.S. Patent Number 3937565. Filed Jun 3, 1974.

ALASIA, A. V., 1998. Digital anti-counterfeiting software method and apparatus. U.S. Patent Number 5708717. Filed Nov 29, 1995.

ALEXA, M., AND MATUSIK, W. 2010. Reliefs as images. *ACM Transactions on Graphics 29*, 4 (July), 60:1–60:7.

BARAN, I., KELLER, P., BRADLEY, D., COROS, S., JAROSZ, W., NOWROUZEZAHRAI, D., AND GROSS, M. 2012. Manufacturing layered attenuators for multiple prescribed shadow images. *Computer Graphics Forum (Proceedings of Eurographics) 31*, 2.

BARNES, C., SHECHTMAN, E., FINKELSTEIN, A., AND GOLDMAN, D. B. 2009. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH) 28*, 3 (Aug.).

BARNES, C., SHECHTMAN, E., GOLDMAN, D. B., AND FINKELSTEIN, A. 2010. The generalized PatchMatch correspondence algorithm. In *European Conference on Computer Vision*, Springer, K. Daniilidis, P. Maragos, and N. Paragios, Eds., vol. 6313 of *Lecture Notes in Computer Science*.

BARNES, C., GOLDMAN, D. B., SHECHTMAN, E., AND FINKELSTEIN, A. 2011. The PatchMatch randomized matching algorithm for image manipulation. *Communications of the ACM 54*, 11 (Nov.), 103–110.

BARNES, C. 2011. *PatchMatch: A Fast Randomized Matching Algorithm with Application to Image and Video*. PhD thesis, Princeton University.

BROSH, S., AND WRIGHT, T., 1994. Anti-counterfeiting process using lenticular optics and color masking. U.S. Patent Number 5303370. Filed Nov 13, 1992.

CHU, H.-K., HSU, W.-H., MITRA, N. J., COHEN-OR, D., WONG, T.-T., AND LEE, T.-Y. 2010. Camouflage images. *ACM Transactions on Graphics 29*, 4 (July), 51:1–51:8.

FINCKH, M., DAMMERTZ, H., AND LENSCH, H. 2010. Geometry construction from caustic images. In *Proceedings of the European Conference on Computer Vision (ECCV)*.

GOTODA, H. 2010. A multilayer liquid crystal display for autostereoscopic 3d viewing. *Stereoscopic Displays and Applications XXI 7524*, 1–8.

HERSCH, R. D., AND CHOSSON, S. 2004. Band moiré images. *ACM Transactions on Graphics (Proceedings of SIGGRAPH) 23*, 3 (Aug.), 239–247.

HERSCH, R. D., COLLAUD, F., AND EMMEL, P. 2003. Reproducing color images with embedded metallic patterns. *ACM Transactions on Graphics (Proceedings of SIGGRAPH) 22*, 3 (July).

HERSCH, R. D., DONZÉ, P., AND CHOSSON, S. 2007. Color images visible under uv light. *ACM Transactions on Graphics 26*, 3 (July), 75:1–75:9.

LANMAN, D., WETZSTEIN, G., HIRSCH, M., RASKAR, R., AND HEIDRICH, W. 2011. Polarization fields: Dynamic light field display using multi-layer LCDs. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia) 30*, 6 (Dec.).

LIPPMANN, J. F. G. 1908. Epreuves reversibles donnant la sensation du relief. *Journal of Physics*.

LOWE, D. G. 1999. Object recognition from local scale-invariant features. In *The Proceedings of the Seventh IEEE International Conference on Computer Vision*, vol. 2, 1150–1157.

MITRA, N. J., AND PAULY, M. 2009. Shadow art. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia) 28*, 5.

MITRA, N. J., CHU, H.-K., LEE, T.-Y., WOLF, L., YESHURUN, H., AND COHEN-OR, D. 2009. Emerging images. *ACM Transactions on Graphics 28*, 5 (Dec.), 163:1–163:8.

NAOR, M., AND SHAMIR, A. 1994. Visual cryptography. In *Advances in Cryptology EUROCRYPT*, Springer Berlin / Heidelberg, A. De Santis, Ed., vol. 950 of *Lecture Notes in Computer Science*, 1–12.

PAPAS, M., JAROSZ, W., JAKOB, W., RUSINKIEWICZ, S., MATUSIK, W., AND WEYRICH, T. 2011. Goal-based caustics. *Computer Graphics Forum (Proceedings of Eurographics) 30*, 2.

PRITCHARD, D., AND HEIDRICH, W. 2003. Cloth motion capture. *Computer Graphics Forum 22*, 3 (Sept.), 263–271.

RENESSE, R. L. 2004. *Optical Document Security*, 3 ed. Artech House, 161–164.

TYLER, C. W., AND CLARKE, M. B. 1990. The autostereogram. In *Proceedings of SPIE*, vol. 1256 of *Stereoscopic Displays and Applications*, 182–196.

WALTER, B., MARSCHNER, S. R., LI, H., AND TORRANCE, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Rendering Techniques 2007: 18th Eurographics Workshop on Rendering*, 195–206.

WALTER, B., ZHAO, S., HOLZSCHUCH, N., AND BALA, K. 2009. Single scattering in refractive media with triangle mesh boundaries. *ACM Transactions on Graphics 28*, 3 (aug).

WETZSTEIN, G., LANMAN, D., HEIDRICH, W., AND RASKAR, R. 2011. Layered 3D: Tomographic image synthesis for attenuation-based light field and high dynamic range displays. *ACM Transactions on Graphics (Proceedings of SIGGRAPH) 30*, 4.

WEYRICH, T., PEERS, P., MATUSIK, W., AND RUSINKIEWICZ, S. 2009. Fabricating microgeometry for custom surface reflectance. *ACM Transactions on Graphics 28*, 3 (July).

YUE, Y., IWASAKI, K., CHEN, B.-Y., DOBASHI, Y., , AND NISHITA, T. 2012. Pixel art with refracted light by rearrangeable sticks. *Computer Graphics Forum (Proceedings of Eurographics) 31*, 2.