

Image Enhancement in Projectors via Optical Pixel Shift and Overlay

Behzad Sajadi, Duy Qoc-Lai, Alex T. Ihler, M. Gopi, Aditi Majumder
Department of Computer Science
University of California, Irvine

{bsajadi, laid, ihler, gopi, majumder}@ics.uci.edu

Abstract

Earlier work has explored enhancing the perceived resolution of a display by shifting multiple different low-resolution images by fractions of a pixel and overlaying them in a temporally multiplexed fashion. This increases the manufacturing cost and also sacrifices the temporal resolution that can compromise other capabilities like 3D active stereo. In this paper we propose a method to achieve the same goal in projectors by performing the pixel shift and superposition optically by introducing a simple and inexpensive optical ensemble of a set of lenses on the projector light path. This does not sacrifice the temporal resolution and is extremely easy to implement in practice.

However, instead of overlaying different images, we overlay an image with one or more sub-pixel shifted copies of itself. Therefore, we seek a single $n \times n$ image which when shifted and overlaid with itself creates a perceptually closer to a higher resolution $2n \times 2n$ target image. This changes the optimization formulation significantly and requires solving a system of sparse linear equations. We take advantage of this sparsity and design a parallel implementation of this optimization in GPUs for real-time computation of the input image critical for its practical implementation. But, since this system is more constrained than using multiple overlaid images, the enhancement of resolution is compromised. However, since the optical design is very simple and inexpensive, it can be deployed on a variety of low-cost projectors and still offer a significant image quality benefit.

1. Introduction

Several earlier works have proposed overlaying multiple different low-resolution $n \times n$ images shifted by a fraction of a pixel to enhance the resolution of a display to be closer to a higher resolution $2n \times 2n$ target image [6, 1, 2, 4, 7]. When using a single display [2, 4, 7], this is achieved by temporal multiplexing of multiple images shifted differently. The shifts are achieved either by wobulating the image [2] or by vibrating the entire display [4, 7]. While wobulation is too cumbersome and expensive to implement in practice, the latter is not practical in scenarios like mobile devices.

Further, since these methods compromise the temporal resolution of the display, the quality of other capabilities like 3D active stereo can be compromised significantly. Alternatively, such resolution enhancement can also be achieved from a more complex multi-projector setup [6, 1] driven by one or more machines.

In this paper, we focus on the single display scenario and propose a technique where the shift and overlay can be achieved in projectors by introducing a very simple and inexpensive ensemble of a set of lenses in the optical path of the projector (Section 3.2). However, in this case, the enhancement needs to happen by shifting and overlaying the *same* image, as opposed to prior works which used different images. Hence, the problem formulation changes significantly. Formally, let us consider a target image I_T of $2n \times 2n$ resolution. Traditionally, this image is reproduced on an $n \times n$ projector by low pass filtering I_T to create I_L and hence the loss in resolution. In this paper, we seek a different low-resolution ($n \times n$) image I such that when overlaid with an *identical* copy of itself which is shifted by a fraction of a pixel, the resulting image, I_R , provides a closer perceptual match to I_T than I_L (Section 3). We formulate this as a linear optimization problem with boundary conditions that constrain the minimum and maximum possible value of the pixels. We present an efficient solution to this problem and demonstrate its parallel implementation on GPUs to achieve interactive performance (Section 3). The GPU implementation of the optimization process to generate I is a relatively small cost to pay to achieve the enhancement in the image quality. But, since we are constrained to shift and overlay a *single* image I , as opposed to using multiple different images as in [2, 4], the enhancement in resolution is compromised (Section 5). However, the simplicity of our design makes it extremely suitable to be deployed to a variety of low-cost projectors still allowing them to achieve a significant resolution enhancement. We use a structural similarity metric that provides a quantitative measure of the degree of enhancement achieved. We evaluate and compare the improvement in visual quality resulting from our technique with existing methods using this quantitative analysis. Finally, using our optical design we build a working proto-

type at a very low cost (Section 5).

Interestingly, from a different perspective, overlaying a pixel with a shifted copy of itself, essentially achieves a non-square (often non-rectangular) point spread function (PSF) that overlaps with adjacent pixels. Though this may seem counter productive at first sight, earlier work has identified its ability to enhance the perceived resolution – [23] used a defocused projector with overlapping Gaussian PSFs to achieve enhancement in the perceived image quality by computing an image that, when projected through an out-of-focus projector, results in a sharp appearance with less aliasing artifacts compared to a low-resolution image projected through a traditional projector. However, unlike this work, we modify the PSF to have a higher frequency bandwidth. Hence, our method can produce superior resolution enhancement (Section 5).

Finally, the proposed concept of overlaying a pixel with a shifted copy of itself has the potential to create different pixel shapes using different kinds of shifts. It even opens up the possibility of shifting different channels differently. We present an empirical cost and benefit analysis of some of these design options in Section 4.

2. Related Work

Computational displays have been explored in several directions, such as using stacked LCD panels to create parallax barrier displays [13, 14, 21, 22], bi-directional displays for capturing the gestures and displaying the content using the same device [11], and enhancing the resolution of the displays by modifying the optics of the projectors [2, 18]. Creating high resolution imagery by overlaying multiple jittered lower resolution images was inspired by the analogous concept explored in cameras – an excellent survey available in [3]. However, early theoretical analysis of the problem [15] showed that achieving super-resolution in displays is only possible if smaller pixels (i.e. a wider band PSF) can be generated; something that exists in a very limited manner when projecting on a planar surface. [12, 6] showed that this can nevertheless provide high-frequency boosting yielding to lesser aliasing artifacts and hence enhanced perceived resolution. However, [1] showed that the enhancement is more significant (1.3 times from three superimposed projectors) on non-planar surfaces since the variation in pixel PSFs can be much larger, allowing the projectors to have much wider bandwidth PSFs at some regions of the surface. Similar resolution enhancement was explored for *single displays*. Pixels were ‘wobulated’ to temporally multiplex multiple different low-resolution frames, each shifted by a fraction (usually half) of a pixel [2]. [4, 7] proposed to achieve the shifting by vibrating a display in a predefined trajectory and dividing the vibration cycle to multiple sub-frames. [4] showed that, unlike projectors, LCD panels are pixelated (can be seen when very close to the panel) and

hence offer the desired wider band pixel PSF. Hence, resolution enhancement of almost 2 times is possible. However, in both these works, enhanced resolution is achieved by sacrificing temporal resolution compromising other capabilities of the display. Our optical design results in a lesser resolution enhancement but does not sacrifice the frame rate and can be a useful for low-cost projectors.

Since, the same pixel is overlaid with a shifted copy of itself, our work creates a display with overlapping pixels. *Image refocussing* methods have explored the benefits of overlapping pixels in the past. Most of these works [23, 5, 16], unlike our work, do not have much control over the pixel PSFs. When a projector projects on surfaces with a large range of depth, the image is not perfectly focused and each pixel is blurred by a Gaussian PSF. This blur is compensated via deconvolution. Like our method, [23] model the compensation problem using a linear system with boundary constraints and show that their defocus compensation method can be used to reduce the pixelization artifacts. However, unlike a Gaussian blur, our overlapping pixel shape have sharp corners when in focus. Thus, our PSF can reinforce certain directional gradients (Section 4), providing higher resolution enhancement (Section 5). More importantly, the PSF of our reshaped pixels results in a linear system that is considerably sparser compared to the linear system resulting from a Gaussian PSF. Hence, we use a fast solver based on message passing using Gaussian belief propagation, resulting in further improvement of the image correction performance (Section 3.1). These differences enables computation of the appropriate input images at interactive rates of more than 60 frames per second.

From an alternate perspective, our work is similar to earlier efforts that *modify the pixel PSF* using hardware solutions. [10] proposed coded aperture projection (with both static and adaptive codes) that provides a pre-designed broadband PSF to alleviate the ringing artifacts due to deconvolution during image refocussing. However, this work only deals with the problem of compensating for the image blur without trying to better represent an image of higher resolution. Therefore, the defocus compensated image is computed using a regularized inverse filter assuming that the resolution of the target image is the same as the resolution of the projector. Further, unlike our technique, the coded aperture reduces the light efficiency (Section 3.2). Our work is closer to edge-guided resolution enhancement in projectors using an optical pixel sharing [18] unit instrumented via an array of $c \times c$ lenses between two $n \times n$ light modulator panels in a projector. This unit increases the density of pixels by c^2 at specific regions guided by edges. This higher density is achieved by reducing the size of the pixels in these regions by a factor of c^2 thereby scaling the resolution to $cn \times cn$, where c is a small integer (usually 2 or 3). However, the higher resolution achieved by optical pixel sharing comes at




Figure 1: The basic concept of our image enhancement. The high resolution pixels of I_T are shown using a blue grid. Left: The low resolution image I with pixels twice the size pixels in I_T . Four pixels are shown with four colors. Middle: I shifted by $s_x = s_y = 0.5$. Right: I and its shifted copy overlaid on each other creating I_R . The pixels with dark boundaries get contributions from the same pixel in I while the others get contributions from different pixels in I .

the cost of almost 50% reduction in light efficiency, sacrificed temporal resolution and doubling the number of light modulator panels. In contrast to optical pixel sharing, in this paper we use a set of lenses to create multiple shifted and overlaid copies of the pixels without modifying their sizes and hence cannot achieve resolution enhancement of the same kind. But, we do not sacrifice temporal resolution or light efficient and do not use any additional light modulators in the projector, resulting in a very inexpensive hardware design.

3. Shifted Overlaid Pixels

Let I_T be the $2n \times 2n$ grayscale target image. Our goal is to compute an $n \times n$ image I , which when shifted and overlaid with itself creates an enhanced image I_R that is perceptually close to I_T . Every pixel in I_T is denoted by (i, j) , $1 \leq i, j \leq 2n$ and $0 \leq I_T(i, j) \leq 1$ is normalized. Hence, every pixel in I corresponds to a block of 2×2 pixels in I_T . Now, consider overlaying I with a copy of itself shifted diagonally by half a pixel in the coordinate system of I . We consider this first due to its practicality and ease of explanation. More general shifts are discussed in Section 4. Let the horizontal and vertical shifts be s_x and s_y where $s_x, s_y \in \{-0.5, 0, 0.5\}$. Figure 1 shows how different pixels of I line up with smaller pixels of I_T when thus shifted and overlaid. We approximate pixel $I_T(i, j)$ by the sum of the intensities of the underlying pixels of the two copies of I , as given by

$$I_T(i, j) \approx \frac{1}{2}I\left(\left\lceil \frac{i}{2} \right\rceil, \left\lceil \frac{j}{2} \right\rceil\right) + \frac{1}{2}I\left(\left\lceil \frac{i}{2} + s_y \right\rceil, \left\lceil \frac{j}{2} + s_x \right\rceil\right) \quad (1)$$

Assembling the equations for every pixel (i, j) in I_T we get a linear system of $2n \times 2n = 4n^2$ equations. The goal of finding the appropriate $n \times n = n^2$ input, I , is given by

$$\mathbf{A}\mathbf{I} \approx \mathbf{I}_T \quad (2)$$

where \mathbf{I}_T is a known column vector of size $4n^2 \times 1$, \mathbf{A} is a known matrix of size $4n^2 \times n^2$, and \mathbf{I} is the unknown column vector of size $n^2 \times 1$ we seek.

Algorithm 1 Constrained Gaussian BP

Input: Matrix J , vector h , initial step size α
Output: Image \mathbf{I}
Initialize $M_{ij} = 0, V_{ij} = 0$
while not converged **do**
 for $i \leftarrow 1$ to n , compute beliefs and estimate **do**
 $V_i = J_{ii} + \sum_{j \in N_i} V_{ji}$
 $M_i = h_i + \sum_{j \in N_i} M_{ji} + \gamma_i - \lambda_i$
 $\hat{\mathbf{i}}_i = M_i / V_i$
 end for
 for $i \leftarrow 1$ to $n, j \in N_i$, compute messages **do**
 $V_{ij} = -J_{ji}^2 / (V_i - V_{ji})$
 $M_{ij} = -J_{ji}(M_i - M_{ji}) / (V_i - V_{ji})$
 end for
 for $i \leftarrow 1$ to n , update Lagrange multipliers **do**
 $\gamma_i = \max[0, \gamma_i - \alpha \hat{\mathbf{i}}_i]$
 $\lambda_i = \max[0, \lambda_i + \alpha(\hat{\mathbf{i}}_i - 1)]$
 end for
 Decrease α if desired
end while

In order to find \mathbf{I} we solve the constrained linear least square problem. Note from Equation 1, each pixel in I_T has a contribution from the copies of the same pixel of I or copies of two different pixels in I that are adjacent to each other (Figure 1). Hence, \mathbf{A} is a very sparse matrix. In particular, each row of \mathbf{A} will have at most two non-zero values separated by at most n zeros in between. We exploit this structure to design an efficient parallelizable algorithm to solve for \mathbf{I} (Section 3.1). Figure 2 shows examples of the target image I_T , the corresponding low-resolution image I_L , the projector input image I , and the enhanced images I_R formed by overlaying the two copies of I with $s_x = s_y = 0.5$. Note that I_R provides a much closer match to I_T than the standard I_L projected from a traditional projector, particularly evident in the zoomed-in views of the pictures.

Evaluation: In order to provide an objective quality measure for comparing the images from our projector and a traditional projector, we use the average CIELAB ΔE difference of I_T from I_L and I_R (Figure 2). This demonstrates an improvement in the results of our technique when compared to a traditional projector of the same resolution. However, this metric does not consider the spatial relationships among the pixels. Therefore, we also use the more sophisticated metric of SSIM index [20] (Figure 2) that demonstrate a closer perceptual match for our technique. We also verify this improvement with a more extensive set of experiments in Section 5 on a large collection of images. However, in order to achieve such results in practice, we need to solve the constrained linear system in real-time and also incorporate the pixel shifting with minimal alteration of the projector design, described in details in the next two sections.

3.1. Efficient Constrained Linear Solver

Solving $\mathbf{A}\mathbf{I} \approx \mathbf{I}_T$ s.t. $0 \leq \mathbf{I} \leq 1$ is an over-constrained linear program; its solution can be re-expressed as the optimum of a constrained least-squares problem,

$$\min_{\mathbf{I}} \frac{1}{2} \mathbf{I}^T \mathbf{A}^T \mathbf{A} \mathbf{I} - \mathbf{I}^T \mathbf{A}^T \mathbf{I}_T \quad \text{s.t.} \quad 0 \leq \mathbf{I} \leq 1. \quad (3)$$

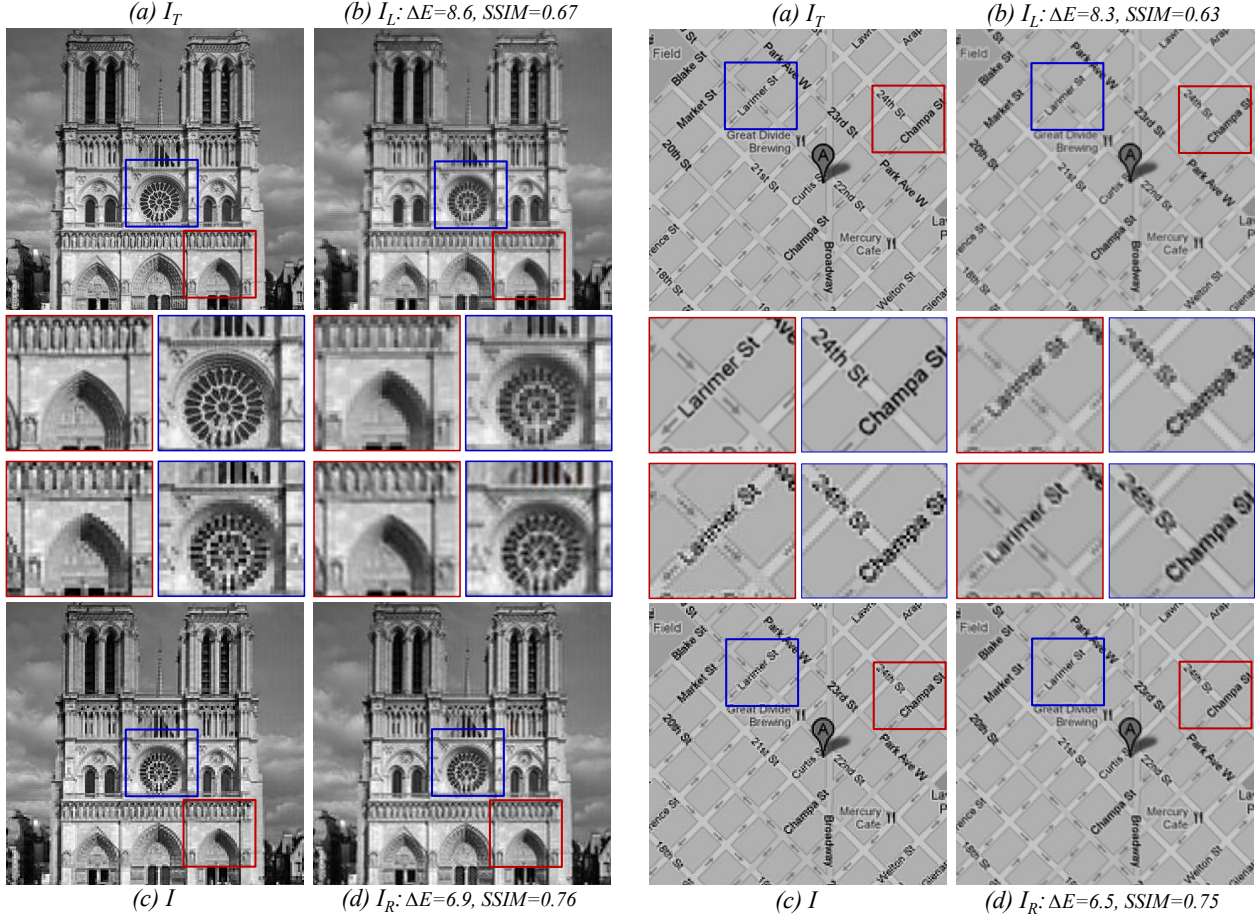


Figure 2: Two simulated examples of pixel shift and overlay. (a) high resolution target image (I_T); (b) down-sampled low resolution $n \times n$ image (I_L); (c) optimal input to our shift and overlay projector (I); (d) result from our projector (I_R). Please use the zoomed-in views to compare the quality of the images. We have also provided the CIE LAB ΔE and the Structural SIMilarity (SSIM) index between I_L and I_T and between I_R and I_T . A lower CIE LAB ΔE and a higher SSIM demonstrate a closer match.

Defining $J = \mathbf{A}^T \mathbf{A}$ and $h = \mathbf{A}^T \mathbf{I}_T$, the quadratic matrix J is extremely sparse and has significant local structure: $J_{pq} \neq 0$ only for pixels p and q in the input image \mathbf{I} that will both be projected onto the same pixel in the target image \mathbf{I}_T .

We exploit the local structure of J to provide an easily parallelized, iterative algorithm for GPU implementation. We use a Gaussian Belief Propagation (BP) solver [8] to minimize the quadratic component; Gaussian BP is an iterative fixed-point algorithm that, while not guaranteed to converge, is often faster than other similar iterative linear solvers such as the Jacobi or Gauss-Seidel methods [19]. In our experiments on more than 500 randomly chosen images, our Gaussian BP solver always converged to an image with unnoticeable difference, i.e., less than 3 CIE LAB ΔE , and was more than three times faster than the Gauss-Seidel solver. We enforce the linear constraints using Lagrange multipliers updated using gradient ascent at each iteration. Our final objective is thus given by the quadratic form

$$\max_{\lambda, \gamma \geq 0} \min_{\mathbf{I}} \frac{1}{2} \mathbf{I}^T \mathbf{J} \mathbf{I} - \mathbf{I}^T h - \mathbf{I}^T \gamma + (\mathbf{I} - 1)^T \lambda \quad (4)$$




Figure 3: A 2D schematic of our optical design. The two copies of the pixels are created using the two lenses with a small offset with respect to each other. The red and green pixels represent these two copies.

and the resulting algorithm is shown in Algorithm 1. We choose this solver since it is highly parallel, consisting of several simple computational kernels that are executed for each pixel in the projector input image \mathbf{I} , and provides approximate solutions if stopped early due to the real-time execution deadline. We implemented this algorithm in CUDA on a GPU yielding a frame rate of 60 fps or more.

3.2. Optical Design

A traditional projector consists of a projection lamp placed behind a light modulator panel, which modulates the light spatially. The pixels in the light modulator panel, e.g., transmissive LCD or a reflective DMD, form a regular 2D array and each pixel attenuates the light that passes through it to achieve a brightness between 0 (black) and 1 (white). The light that passes through the pixels then goes to a projection lens that focuses the light on a screen. The different colors are created using temporal multiplexing of the three channels in single-chip projectors, or by combining the light from three different optical paths in 3-chip projectors.

In order to create two shifted copies of each pixel, we modify this design by placing two lenses, L_1 and L_2 , with the same focal length, f , between the light modulator panel and the projection lens. We need a very small distance between the optical axes of these lenses to achieve a sub-pixel shift. Therefore, we use two split lenses for this purpose (Figure 3). When placed at distance l from the light modulator panel on a plane perpendicular to the optical axis of the projection lens, these lenses will have the same magnification factor of $m = \frac{f}{l-f}$.

The light from each pixel passes through both lenses, creating two copies of each pixel. Consequently, the shift between the two copies depends on the distance between the optical axes of the two lenses. Let us indicate this by vector $\mathbf{d} = (d_x, d_y)$ where the coordinate system is aligned with the pixel grid of the light modulator panel. Let us consider a pixel of the light modulator panel at an arbitrary distance \mathbf{d}' from the optical axis of L_1 . Consequently, it is located at distance $\mathbf{d} + \mathbf{d}'$ from the optical axis of L_2 . Therefore, the two shifted copies of this pixel are formed at distances $m\mathbf{d}'$ and $m\mathbf{d} + m\mathbf{d}'$ from the optical axes of L_1 and L_2 respectively. Since the distance between the optical axes of L_1 and L_2 is \mathbf{d} , the distance between the two shifted copies is $\mathbf{d} + m\mathbf{d}$. If the size of the pixel is a units in each dimension, the size of the pixel after magnification is ma units. Consequently, the shift between the two copies is

$$(s_x, s_y) = \frac{(1+m)\mathbf{d}}{ma} = \frac{(1+\frac{f}{l-f})\mathbf{d}}{\frac{f}{l-f}a} = \frac{\frac{l}{l-f}\mathbf{d}}{\frac{f}{l-f}a} = \frac{l\mathbf{d}}{fa}. \quad (5)$$

Interestingly, by choosing a large f , we can use a large \mathbf{d} to create a small shift. Further, the magnification factor m will be very close to 1 and therefore, addition of the two lenses will not considerably affect the FOV of the projection lens. Finally, the shift is the same for all the color channels and is independent of the placement of the projection lens since the projection lens magnifies the shift and the pixel size similarly resulting in the same shift in unit of pixels. This allows us to change the zoom or focus settings of the projection lens without changing the shift between the two copies of the pixels. We use this optical design to incorporate the pixel shift and overlay in a traditional projector.

4. Extensions

In Section 3, we consider only two overlays and shifts of only half a pixel for grayscale images. This allows the pixel boundaries in I and its shifted overlay to line up with the underlying smaller pixels of I_T making the problem formulation easier. In this section we consider the more general problem of overlaying m frames where $m \geq 2$, fractional shifts of $\frac{1}{k}$ where k is an integer, and extend our design to RGB color images.

Multiple Overlays: When a pixel is overlapped with a shifted copy of itself, the combination provides non-square *reshaped pixels*. When $s_x = s_y = 0.5$ (Figure 4(c)) the PSF has a larger frequency bandwidth in certain directions and hence favors certain gradients. In other words, gradients along the direction of the shift are represented better than gradients in the perpendicular direction. We explore the possibility of creating more varied shapes by overlaying more shifted frames (Figure 4). For example, if an image is overlaid with two shifted images, one with $s_x = s_y = 0.5$ and another with $s_x = -0.5, s_y = 0.5$ (Figure 4(e)), a close perceptual match to two gradient directions can be achieved. In order to verify this, we form a Histogram of Gradients (HoG) with 9 directions for every 4×4 pixels in the projected images and in the target image. We then measure, for each gradient direction, the correlation between the number of gradients detected in that direction in different regions of the projected images and the target image (Figure 5). This analysis verifies that pixel shapes created by only two shift and overlays cannot represent all the directional gradients well, resulting in low correlation values for some of the directions. On the other hand, the pixel shape in Figure 4(e) performs well for all the gradients.

Assuming m copies of I are overlaid with each other and the shifts are still multiples of half a pixel and p th copy, $1 \leq p \leq m$, is shifted by $s_x(p)$ and $s_y(p)$, Equation 1 becomes,

$$I_T(i, j) \approx \frac{1}{m} \sum_{p=1}^m I\left(\left\lceil \frac{i}{2} + s_y(p) \right\rceil, \left\lceil \frac{j}{2} + s_x(p) \right\rceil\right) \quad (6)$$

Note that size of \mathbf{A} in Equation 2 does not change due to multiple overlays and it is still sparse and can be solved using the same method. In practice, m shifts can be achieved by placing m small lenses before the projection lens.

General Shifts: Half pixel shifts let the pixel boundaries of I line up with the smaller pixels of I_T . For the more general case of fractional shifts, let $s_x = \frac{k_x}{k}$ and $s_y = \frac{k_y}{k}$ where k is a small integer and $0 < k_x, k_y < k$ (Figure 6). To make the boundaries of the larger and smaller pixels coincide, I_T now needs to be resampled and the dimensions of \mathbf{I}_T and \mathbf{A} in Equation 2 will be $k^2 n^2$ and $k^2 n^2 \times n^2$ respectively. Further, Equation 1 becomes,

$$I_T(i, j) \approx \frac{1}{2} I\left(\left\lceil \frac{i}{k} \right\rceil, \left\lceil \frac{j}{k} \right\rceil\right) + \frac{1}{2} I\left(\left\lceil \frac{i}{k} + s_y \right\rceil, \left\lceil \frac{j}{k} + s_x \right\rceil\right) \quad (7)$$

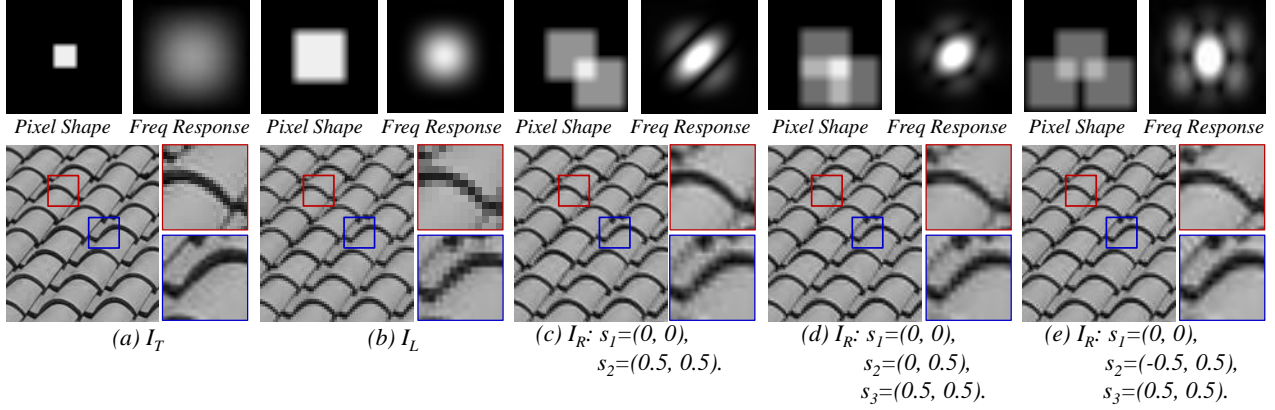


Figure 4: From left to right: (a) high resolution target image; (b) image from a low resolution projector; (c,d,e) images from our design with a diagonal shift (c), with a vertical and a diagonal shift (d), and with two diagonal shifts (e). For each image, the PSF and the frequency response are shown in the top row. The frequency responses suggests that using two diagonal shifts in different directions provides the most broadband frequency attenuation and is not biased towards a particular direction. This is also evident from the zoomed-in views of I_R ; the gradients are well preserved only in one diagonal direction in (c) and (d), while the gradients are preserved in a symmetric manner in (e).

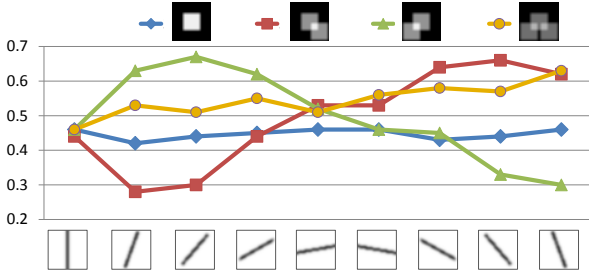


Figure 5: To measure the ability of different pixel shapes at preserving different gradients, we compute a HOG with 9 directions for every 4×4 pixels in the target and projected images. The graphs show for each gradient direction the correlation between the gradients detected in that direction in the projected and target images. The red and green graphs are higher than the blue graph for certain directions and lower in other demonstrating that only certain gradients are better represented. However, the orange graph is consistently higher than the blue graph, suggesting that all the directions are represented better using this pixel shape.



Figure 6: The different pixel shapes for different values of m (number of overlays), s_x (horizontal shift) and s_y (vertical shift). The overlaid pixels are shown in gray and the reshaped pixel is outlined in bold black. From left to right:

- (a) $m = 2$, $s_x(1) = s_y(1) = 0$, $s_x(2) = s_y(2) = \frac{1}{2}$;
- (b) $m = 2$, $s_x(1) = s_y(1) = 0$, $s_x(2) = \frac{1}{3}$, $s_y(2) = \frac{2}{3}$;
- (c) $m = 3$, $s_x(1) = s_y(1) = 0$, $s_x(2) = s_y(2) = \frac{1}{2}$, $s_x(3) = -\frac{1}{2}$, $s_y(3) = \frac{1}{2}$;
- (d) $m = 3$, $s_x(1) = s_y(1) = 0$, $s_x(2) = s_y(2) = \frac{1}{2}$, $s_x(3) = \frac{1}{4}$, $s_y(3) = \frac{1}{4}$.

Note that achieving general shifts does not require any change in the optical design and only needs the adjustment of the shift between the lenses (Section 3.2).

We evaluated the quality of the resulting images in Figure 7, which do not show any significant improvement compared to half pixel shifts. Further, upsampling of the target im-




Figure 7: Comparing the quality of images with different fractional shifts given in Figure 7 with respect to the high resolution target image, I_T , and the low resolution image, I_L . The images demonstrate considerable improvement compared to I_T . However, the difference between the images using half pixel shifts, i.e., (c) and (e), and the images using other fractional shifts, i.e., (d) and (f), is small.

age and the cost of converting the matrix to the quadratic form in Equation 4 increased the computation time of each frame. In our experiments, the total computation time for each frame increased by 20% after upsampling the target image to three times the resolution of the projector. This confirms that practically half pixel shifts are sufficient to achieve a significant resolution enhancement.




Figure 8: Comparison of the quality of the target high-resolution colored image with a low-resolution version of itself and the image created using our shift and overlay technique. The pixel shape used for this experiment is the same as the one in Figure 4(e). In the zoomed-in views, note the reduced jaggy artifacts and the extra sharpness in the results of our technique compared to a low resolution image.

Multiple Channels: We denote the images for the three different color channels with superscripts of r , g and b for the red, green and blue channels respectively. In order to extend our method to multiple color channels, we simply apply Equation 6 to each channel independently as follows.

$$I_l^i(i, j) \approx \frac{1}{m} \sum_{p=1}^m I^l \left(\left[\frac{i}{2} + s_y(p) \right], \left[\frac{j}{2} + s_x(p) \right] \right), l \in \{r, g, b\} \quad (8)$$

The results are demonstrated in Figure 8. Further, the optical design does not require any modification since multiple color channels can be shifted similar to a single channel.

However, optimizing the three channels separately does not guarantee that the combined effect of the three will preserve the luminance. Since the human visual system is highly sensitive to luminance changes [9], one can use additional constraints on the luminance. Different weights can be used to assign higher importance to the luminance or to any other color channel as demanded by the application using techniques similar to [17].

5. Implementation and Results

We used the GTX 560 GPU and CUDA for our real-time implementation resulting in more than 60 FPS for target color images of size 2048×1536 . We also implemented our algorithm in MATLAB and C++, both taking around 2-3 seconds to find I . We now describe our evaluation using a large set of simulated images. Next, we describe the details of our prototype projector and another set of experiments using this prototype.

Comparison with Other Techniques: We compare the quality of the images produced by our technique against the images from a traditional projector with the same $n \times n$ pixels. We also compare our technique with [23], which uses




Figure 10: A top view of the modified LCD projector. The LCD panels, the two half lenses, and the projection lens are shown in the picture. Note that the reason for the large size of the setup is that we were not able to place the half lenses inside the projector to keep them close to the LCD panels.

a slightly out of focus projector and image pre-sharpening. We also compare our method with wobulation [2] using two different diagonally shifted frames. Finally, we compare our results against the optical pixel sharing technique [18]. One example of this comparison is shown in Figure 9. We have also summarized the comparison results for a large database of different types of images in Table 1. Our method demonstrates superior quality compared to the low-resolution images and also compared to [23]. This turns out to be significant when considering only a small modification in the projector hardware without sacrificing frame rate. On the other hand, wobulation and pixel sharing techniques produce superior results compared to our method since they are not constrained by the use of a single input. However, wobulation comes at the cost of sacrificed frame rate and a considerably higher production cost. Similarly, optical pixel sharing uses an additional light modulator panel and comes at the cost of sacrificed frame rate, brightness and contrast and a more complex optical design.

Our Prototype: We modified an Epson EMP-74 3-chip LCD projector of resolution 1024×768 to create a projector with two shifted and overlaid copies with $s_x = 0.5$, $s_y = 0.5$ (Figure 10). First, we put two lenses with a small offset between them before the projection lens. In the original projector the projection lens is placed at $50mm$ distance from the LCD panels. However, it is hard to place the half lenses at this distance, which is inside the projector. Therefore, we decided to place the lenses at distance $l = 100mm$ from the LCD panels.

In order to create a small shift using a larger offset between the lenses (Section 3.2), we used lenses with focal lengths much larger than l ($f = 1000$) from several off-the-shelf-lens options ¹. Considering the size of the pixels in the

¹http://search.newport.com/?q=*&x2=sku&q2=KPX124

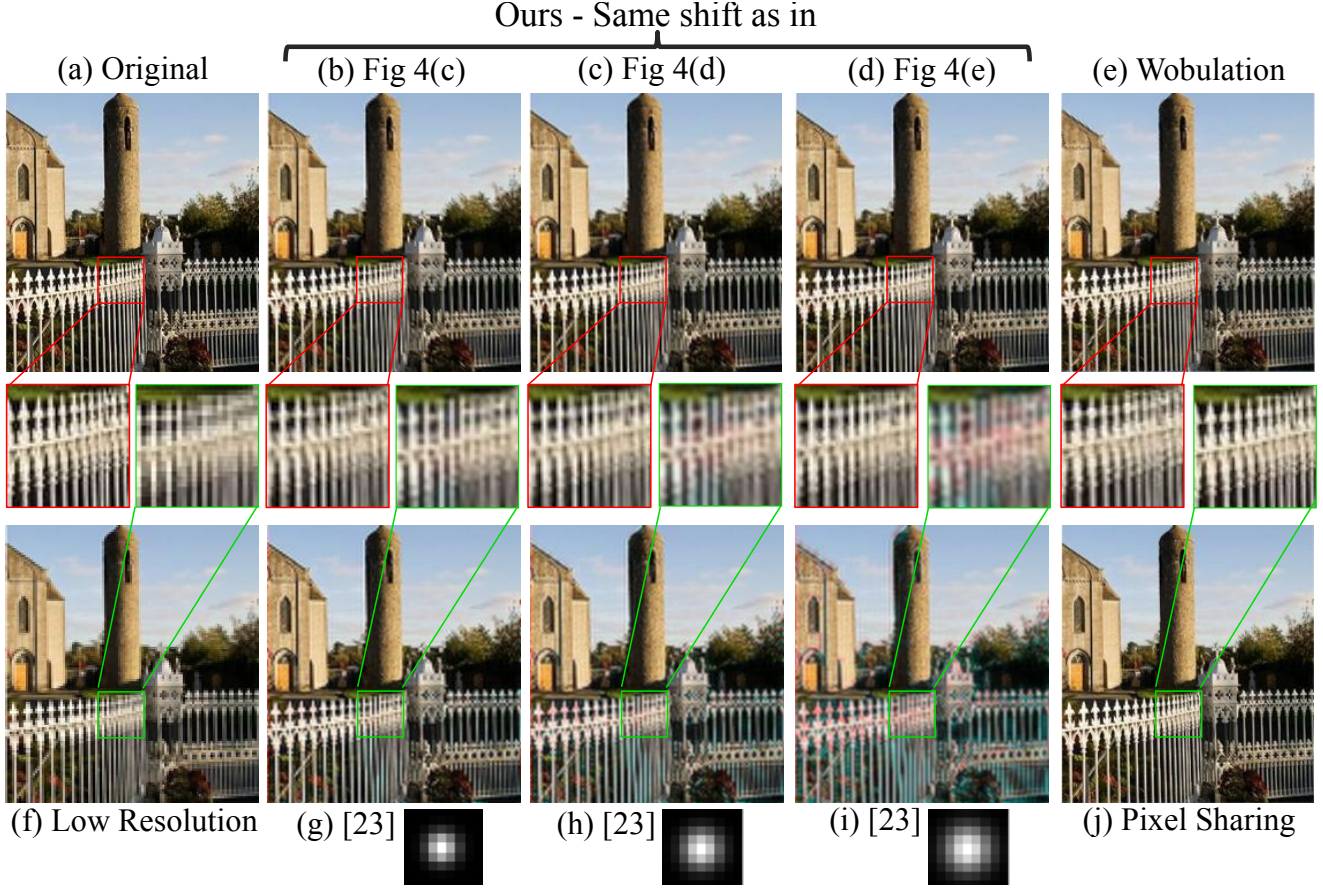


Figure 9: In this figure we visually compare the high-resolution input image (a) with the images from our technique (b-d), a wobulated image with two diagonally shifted frames [2] (e), a low-resolution image from a traditional projector (f), a sharpened image using [23] from an out-of-focus projector with different bluer sizes (g-i), and an image created using optical pixel sharing [18] with four copies for each pixel and 3JNDs threshold for choosing the edge pixels (j). Note that as we increase the blur radius for [23] more color artifacts are introduced and the optimization process takes longer to converge.

	Architecture		People		Nature		Oil Paintings		Line Drawings	
	CIELAB ΔE	SSIM	CIELAB ΔE	SSIM	CIELAB ΔE	SSIM	CIELAB ΔE	SSIM	CIELAB ΔE	SSIM
Ours (Same shifts as Fig 4(c))	7.2	0.75	6.9	0.80	7.4	0.75	6.6	0.8	7.0	0.74
Ours (Same shifts as Fig 4(d))	7.2	0.76	7.0	0.79	7.6	0.73	6.7	0.82	6.9	0.74
Ours (Same shifts as Fig 4(e))	7.1	0.78	6.8	0.81	7.5	0.76	6.3	0.82	6.7	0.78
Low Resolution	8.4	0.64	8.2	0.66	9.1	0.61	8.0	0.68	8.8	0.65
[23] (Same PSF as Fig 9(g))	8.2	0.67	7.9	0.70	8.5	0.64	7.7	0.71	8.7	0.65
[23] (Same PSF as Fig 9(h))	8.6	0.68	8.1	0.73	8.9	0.64	8.0	0.73	9.1	0.64
[23] (Same PSF as Fig 9(i))	10.1	0.62	9.3	0.70	10.5	0.60	9.4	0.70	9.9	0.59
Wobulation	6.0	0.84	5.6	0.86	6.1	0.82	5.4	0.89	5.5	0.83
Optical Pixel Sharing	2.1	0.96	2.2	0.97	2.4	0.95	1.9	0.97	1.3	0.98

Table 1: We perform our experiments on five categories of images (each 100 images of resolution 2048×1536): Architecture, People, Nature, Oil Paintings and Line Drawings. We compare the quality of the images from our technique with the images from a traditional projector, from an out-of-focus projector using [23], from the wobulation technique with two diagonally shifted frames [2], and from optical pixel sharing [18] with four copies for each pixel and 3JNDs threshold for choosing the edge pixels. For evaluation, we use the average CIELAB ΔE and SSIM of each image with respect to the target image.

LCD panels, $a = 14\mu\text{m}$ in each dimension, the required shift between the lenses based on Equation 5 is $d_x = d_y = fa/2l = 70\mu\text{m}$. This requires a shift of $|\mathbf{d}| = 70\sqrt{(2)}\mu\text{m}$ in the diagonal direction. Since putting two lenses with

circular boundaries with such a small offset between their optical axes is impossible, we cut a single lens in half and placed the resulting half lenses together. To create the small shift, we placed a thin piece of tape between them where

the thickness of the tape is almost equal to $|d|$. The two half lenses are placed tightly together inside a metric bar-type holder². Since the half lenses were placed at distance 100mm from the LCD panels, we had to use our own projection lens to accommodate for this change of location. This lens is placed in front of the two half lenses and is formed by a convex lens with focal length 75.6mm ³ followed by a concave lens with focal length -150mm ⁴ at distance 140mm from the first lens providing a field-of-view of 50° .

In order to compare our results against a traditional projector, we replaced the two half lenses with an identical single lens. Further, in order to compare our results against [23], we used the same setup, but with different degrees of out-of-focus. We then used their method to compute the proper input images. The results of this comparison are summarized in Figure 11. Since the lenses were chosen from COTs components, the images suffered from considerable chromatic aberration. Therefore, we have only used the green channel of the projector in our experiments and then converted the captured image to grayscale. The images demonstrate that our setup provides considerable improvement both compared to a traditional projector and compared to [23].

6. Discussion

Diffraction Limit: In our optical design, due to the use of a set of lenses, the effective aperture area is reduced by the number of the shifts. Fortunately, most projectors use high F-numbers (F/2 or higher) to improve the light efficiency. This is possible since depth of focus is not very critical in many projection applications compared to photography. Further, the size of the lens is often large, except for the new generation of mini and pico projectors. Therefore, the resolution of the projector is often limited by the number of pixels instead of the aperture size, making it feasible to divide the aperture to two or three smaller sub-apertures without introducing considerable blurriness.

Projection of Low Resolution Images: When the resolution of the target image and the projector are both equal, i.e., $n \times n$, a low resolution projector can accurately represent the target image. In this case, our design does not allow us to exactly represent the target image. In order to measure the effect of our shift and overlay design on such low resolution images we performed an experiment on a collection of 100 images of resolution 1024×768 . We first upsample the images to $2n \times 2n$ and then run the same optimization as before. The average CIELAB ΔE difference between the target image and the images from our projector was 2.1 units and the average SSIM index was 0.96, showing only a hardly noticeable degradation compared to the use of a tra-

ditional low resolution projector. However, when projected through our prototype, the images have a more pleasing appearance because of the reduction of the screen door effect.

7. Conclusion and Future Work

In conclusion, we proposed a simple modification in the physical construction of the projectors to create pixels of different shapes with small overlaps with the adjacent pixels. We demonstrated how the problem of producing a close match to a higher resolution target using such overlapping pixels can be modeled as a constrained linear system. We then use the sparsity of this linear system to efficiently solve it with a Gaussian belief propagation solver in real-time via a GPU implementation. Using an extensive set of experiments, both in simulation and in practice, we showed that our technique provides a closer match to a high resolution target image when compared to a traditional projector. This improvement involves a very small production cost and the image computation can use the existing GPU capabilities.

We believe that the concept of shifted and overlaid pixels is a powerful tool that provides more control over spatial areas that fall within the overlap of multiple pixels. Though using overlapping pixels has been explored to some extent in prior work [23], our method shows superior results compared to such efforts. However, we would like to analyze the different pixel shapes in depth both analytically and empirically to get more insights into an optimal pixel shape. We would also like to explore the possibility of extending this design to other display devices such as LCD panels and capture devices. Finally, we would like to explore the possibility of combining our design with other techniques such as wobulation and optical pixel sharing.

References

- [1] D. Aliaga, Y. H. Yeung, A. J. Law, B. Sajadi, and A. Majumder. Fast high-resolution appearance editing using super-imposed projections. *ACM TOG*, 2011.
- [2] W. Allen and R. Ulichney. Wobulation: Doubling the addressed resolution of projection displays. *SID*, 2005.
- [3] R. S. Babu and K. E. S. Murthy. A survey on the methods of super-resolution image reconstruction. *IJCV*, 15(2), 2011.
- [4] F. Berthouzoz and R. Fattal. Resolution enhancement by vibrating displays. *ACM Trans. Graph.*, 31(2):15:1–15:14, Apr. 2012.
- [5] M. S. Brown, P. Song, and T.-J. Cham. Image preconditioning for out-of-focus projector blur. *Computer Vision and Pattern Recognition, IEEE Computer Society Conference on*, 2:1956–1963, 2006.
- [6] N. Damera-Venkata and N. L. Chang. Display supersampling. *ACM TOG*, 28(1), 2009.
- [7] P. Didyk, E. Eisemann, T. Ritschel, K. Myszkowski, and H. Seidel. Apparent display resolution enhancement for moving images. 2010.

²<http://www.edmundoptics.com/onlinecatalog/displayproduct.cfm?productID=2190/>

³http://search.newport.com/?q=* &x2=sku&q2=KPX055

⁴http://search.newport.com/?q=* &x2=sku&q2=KPC073

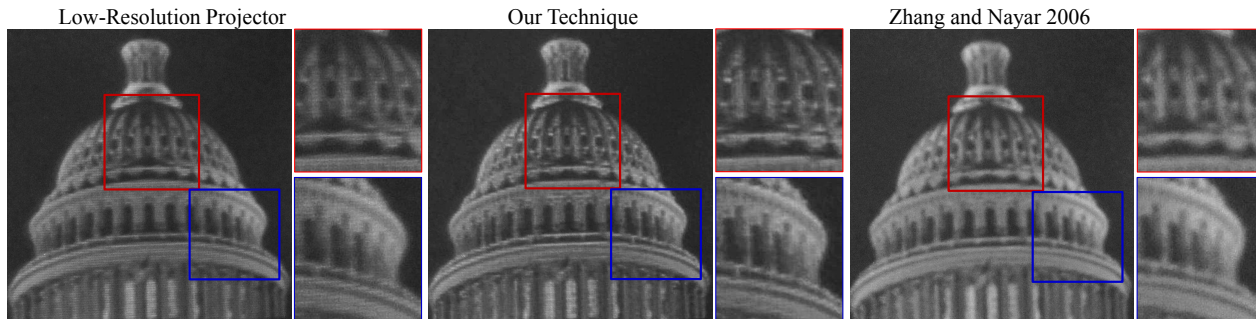


Figure 11: Comparison of images from a low-resolution projector, our prototype, and [23]. The zoomed-in views show that in comparison to the low-resolution image, the result of our technique looks sharper and shows less pixelization artifacts. [23] also removes the pixelization artifacts but the image looks less sharp compared to the result of our technique.

- [8] D. Dolev, D. Bickson, and J. K. Johnson. Fixing convergence of gaussian belief propagation. In *ISIT*, pages 1674–1678, 2009.
- [9] E. B. Goldstein. Sensation and perception, 2000.
- [10] M. Grosse and O. Bimber. Coded aperture projection. In *Proceedings of the 2008 workshop on Immersive projection technologies/Emerging display technologies*, IPT/EDT '08, pages 13:1–13:4, 2008.
- [11] M. Hirsch, D. Lanman, H. Holtzman, and R. Raskar. BiDi screen: A thin, depth-sensing LCD for 3D interaction using lights fields. *ACM TOG*, 28(5), 2009.
- [12] C. Jaynes and D. Ramakrishnan. Super-resolution composition in multi-projector displays. *PROCAMS*, 2003.
- [13] D. Lanman, M. Hirsch, Y. Kim, and R. Raskar. Content-adaptive parallax barriers: optimizing dual-layer 3d displays using low-rank light field factorization. *ACM TOG*, 2010.
- [14] D. Lanman, G. Wetzstein, M. Hirsch, W. Heidrich, and R. Raskar. Polarization fields: Dynamic light field display using multi-layer LCDs. *ACM TOG*, 30(6), 2011.
- [15] A. Majumder and G. Welch. Computer graphics optique : Optical superposition of projected computer graphics. *Proceedings of Joint Eurographics Workshop on Virtual Environments and Immersive Projection Technology*, 2001.
- [16] Y. Oyamada and H. Saito. Focal pre-correction of projected image for deblurring screen image. In *IEEE CVPR*, 2007.
- [17] J. Platt, B. Keely, B. Hill, B. Dresevic, C. Betrisey, D. P. Mitchell, G. Hitchcock, J. F. Blinn, and T. Whitted. Displaced filtering for patterned displays. *Proceedings of Society for Information Display*, 2000.
- [18] B. Sajadi, M. Gopi, and A. Majumder. Edge-guided resolution enhancement in projectors using optical pixel sharing. *ACM TOG*, year = 2012,.
- [19] O. Shental, P. H. Siegel, J. K. Wolf, D. Bickson, and D. Dolev. Gaussian belief propagation solver for systems of linear equations. *CoRR*, abs/0810.1736, 2008.
- [20] Z. Wang, A. C. Bovik, H. R. Sheikh, S. Member, and E. P. Simoncelli. Image quality assessment: From error measurement to structural similarity. *IEEE TIP*, 13:600–612, 2004.
- [21] G. Wetzstein, D. Lanman, W. Heidrich, and R. Raskar. Layered 3D: Tomographic image synthesis for attenuation-based light field and high dynamic range displays. *ACM TOG*, 30(4), 2011.
- [22] G. Wetzstein, D. Lanman, M. Hirsch, and R. Raskar. Tensor displays: compressive light field synthesis using multilayer displays with directional backlighting. *ACM Trans. Graph.*, 31(4):80:1–80:11, July 2012.
- [23] L. Zhang and S. K. Nayar. Projection defocus analysis for scene capture and image display. *ACM TOG*, 2006.