

Compressive multi-mode superresolution display

Felix Heide,^{1,*} James Gregson,¹ Gordon Wetzstein,² Ramesh Raskar,²
and Wolfgang Heidrich^{3,1}

¹University of British Columbia, 2329 W Mall, Vancouver, BC V6T 1Z4, Canada

²MIT Media Lab, 75 Amherst St, Cambridge, MA 02139, United States

³King Abdullah University of Science and Technology, Thuwal 23955-6900, Saudi Arabia

*fheide@cs.ubc.ca

Abstract: Compressive displays are an emerging technology exploring the co-design of new optical device configurations and compressive computation. Previously, research has shown how to improve the dynamic range of displays and facilitate high-quality light field or glasses-free 3D image synthesis. In this paper, we introduce a new multi-mode compressive display architecture that supports switching between 3D and high dynamic range (HDR) modes as well as a new super-resolution mode. The proposed hardware consists of readily-available components and is driven by a novel splitting algorithm that computes the pixel states from a target high-resolution image. In effect, the display pixels present a compressed representation of the target image that is perceived as a single, high resolution image.

© 2014 Optical Society of America

OCIS codes: (120.2040) Displays; (100.3190) Inverse problems; (070.6120) Spatial light modulators.

References and links

1. H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcovs, "High dynamic range display systems," *ACM Trans. Graph. (SIGGRAPH)* **23**, 760–768 (2004).
2. M. Grosse, G. Wetzstein, A. Grundhöfer, and O. Bimber, "Coded Aperture Projection," *ACM Trans. Graph.* **29**, 22:1–22:12 (2010).
3. D. Lanman, M. Hirsch, Y. Kim, and R. Raskar, "Content-adaptive Parallax Barriers: Optimizing Dual-layer 3D Displays using Low-rank Light Field Factorization," in "*ACM Trans. Graph. (SIGGRAPH Asia)*," **29** (2010), p. 163.
4. 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 Trans. Graph. (SIGGRAPH)* **30**, 1–12 (2011).
5. D. Lanman, G. Wetzstein, M. Hirsch, W. Heidrich, and R. Raskar, "Polarization Fields: Dynamic Light Field Display using Multi-layer LCDs," *ACM Trans. Graph. (SIGGRAPH Asia)* **30**, 186 (2011).
6. G. Wetzstein, D. Lanman, M. Hirsch, and R. Raskar, "Tensor Displays: Compressive Light Field Synthesis using Multilayer Displays with Directional Backlighting," *ACM Trans. Graph. (SIGGRAPH)* **31**, 80 (2012).
7. J. Platt, "Optimal Filtering for Patterned Displays," *IEEE Signal Processing Letters* **7**, 179–180 (2002).
8. N. Damera-Venkata and N. L. Chang, "Display Supersampling," *ACM Trans. Graph.* **28**, 9:1–9:19 (2009).
9. C. Jaynes and D. Ramakrishnan, "Super-resolution composition in multi-projector displays," in *Proc. of IEEE International Workshop on Projector-Camera Systems*, 354–356 (IEEE, 2003).
10. W. Allen and R. Ulichney, "Wobulation: Doubling the addressed Resolution of Projection Displays," in "*Proc. SID 47*," **36**, 1514–1517 (2005).
11. F. Berthouzoz and R. Fattal, "Resolution Enhancement by Vibrating Displays," *ACM Trans. Graph.* **31**, 15:1–15:14 (2012).

12. P. Didyk, E. Eisman, T. Ritschel, K. Myszkowski, and H.-H. Seidel, "Apparent Resolution Display Enhancement for Moving Images," *ACM Trans. Graph. (SIGGRAPH)* **29** (2010).
 13. B. Sajadi, M. Gopi, and A. Majumder, "Edge-guided Resolution Enhancement in Projectors via Optical Pixel Sharing," *ACM Trans. Graph. (SIGGRAPH)* **31**, 79:1–79:122 (2012).
 14. F. Durand, N. Holzschuch, C. Soler, E. Chan, and F. X. Sillion, "A Frequency Analysis of Light Transport," *ACM Trans. Graph. (SIGGRAPH)* **24**, 1115–1126 (2005).
 15. S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning* **3**, 1–122 (2011).
 16. M. Yan, "Convergence Analysis of SART by Bregman Iteration and Dual Gradient Descent," *UCLA CAM report* pp. 10–27 (2010).
 17. M. Gross, S. Würmlin, M. Naef, E. Lamboray, C. Spagno, A. Kunz, E. Koller-Meier, T. Svoboda, L. Van Gool, S. Lang *et al.*, "blue-c: a spatially immersive display and 3d video portal for telepresence," in *ACM Trans. Graph. (Proc. SIGGRAPH)*, **22** 819–827 (2003).
 18. S. Izadi, S. Hodges, S. Taylor, D. Rosenfeld, N. Villar, A. Butler, and J. Westhues, "Going beyond the display: a surface technology with an electronically switchable diffuser," in *Proc. UIST* (2008), pp. 269–278.
 19. S. Baker and T. Kanade, "Limits on Super-Resolution and How to Break Them," in *Proc. IEEE CVPR*, **24**:9, 1167–1183(2000).
 20. P. D. Burns, "Slanted-edge mtf for digital camera and scanner analysis," in "Is and Ts Pics Conference," (Society for Imaging Science & Technology, 2000), pp. 135–138.
-

1. Introduction and related work

Throughout the last few years, display technology has undergone a major transformation. Whereas improvements of display characteristics, such as resolution and contrast, have traditionally relied exclusively on advances in optical and electrical fabrication, computation has now become an integral part of the image formation. Through the co-design of display optics and computational processing, computational and compressive displays have the potential to overcome fundamental limitations of purely optical designs. Characteristics that can be improved by a co-design of display optics and computation include dynamic range [1] and depth of field of projectors [2]. A significant amount of research has recently been conducted on compressive light field display for glasses-free 3D image presentation [3–6].

There has also been considerable work on computational superresolution displays in recent years [7]. Many recently proposed approaches integrate different images optically by superimposing multiple projections on the same screen [8, 9]. This can also be achieved in single display/projector designs using mirrors [10] or by displaying different patterns on a quickly moving device [11, 12]. These approaches, however, have the need for multiple projection devices and mechanically moving parts. One solution that overcomes these problems is *Optical Pixel Sharing* (OPS) [13], which uses two LCD panels and a 'jumbling' lens array in projectors to overlay a high-resolution edge image on a coarse image to adaptively increase resolution. This way only the most strongest edges are improved adaptively, while smooth areas are unchanged. This approach brings the problem of the decision of when an edge should be superresolved or not, causing noticeable artifacts due to the global edge threshold used in their method. It is not immediately obvious how such approaches could be adapted from projectors to flatpanel displays.

In this paper, we explore a new compressive display design that can be switched between light field, high dynamic range, and superresolution modes. The display design is inspired by previously proposed light field displays [3], and is comprised of two high-speed liquid crystal displays (LCDs) that are mounted in front of each other with a slight offset (Fig. 1). To support a super-resolution display mode, we introduce an additional diffuser covering the LCD closest to the observer. The two stacked LCDs synthesize an intermediate light field inside the device; the diffuser then integrates the different views of that light field such that an observer perceives a superresolved, two-dimensional image. Making the diffuser electronically switchable allows for the display to be used in 3D or high dynamic range mode.

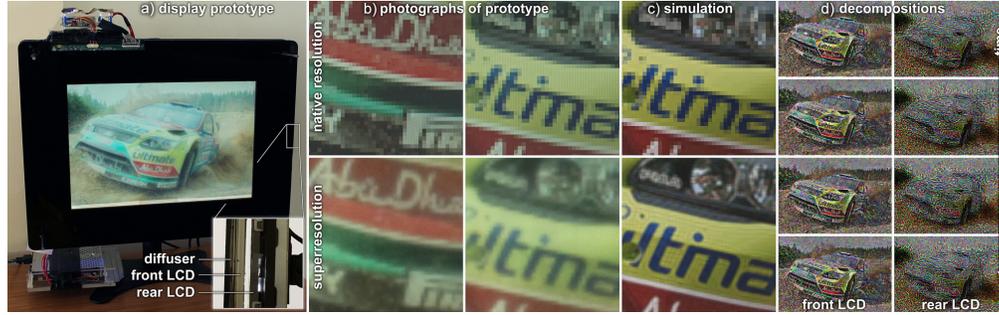


Fig. 1. Compressive superresolution display. The proposed display architecture comprises two stacked high-speed liquid crystal displays (LCDs) covered by a diffuser in a). A target high-resolution image is then decomposed into a set of patterns that are shown on the front and rear LCD in quick succession in d). Compared to the native resolution of each panel in top c), the proposed compressive display approach achieves significant improvements in resolution bottom c) without any mechanically moving parts. Display content under CC license: <http://goo.gl/KYR3Tp>

2. Superresolution mode

In this section, we derive models for superresolved image formation, based on compressive light field displays. We also introduce inverse methods to compute optimal pixel states for a target high-resolution image.

2.1. Image formation

The proposed optical display configuration comprises a light field display behind a diffuser. The image $i(\mathbf{x})$ observed on the diffuser is the integration of the incident light field $l(\mathbf{x}, \mathbf{v})$ over the angular domain $\Omega_{\mathbf{v}}$:

$$i(\mathbf{x}) = \int_{\Omega_{\mathbf{v}}} l(\mathbf{x}, \mathbf{v}) d\mathbf{v}. \quad (1)$$

Here, \mathbf{x} is the 2D spatial coordinate on the diffuser and \mathbf{v} denotes the angle. The light field is weighted with angle-dependent integration weights of the diffuser. We employ a relative two-plane parameterization of the light field [14] (see Fig. 2 in a)). Conceptually, any light field display can be placed behind the diffuser; we follow Lanman et al. [3] and use two stacked liquid crystal displays (LCDs). Driven at a speed beyond the critical flicker frequency of the human visual system (HVS), an observer perceives the temporal integral of the sets of patterns shown on the display. The light field that is synthesized inside the display and incident on the diffuser is

$$\tilde{l}(\mathbf{x}, \mathbf{v}) = \frac{1}{K} \sum_{k=1}^K f^{(k)}(\mathbf{x} - d \cdot \mathbf{v}) \cdot g^{(k)}(\mathbf{x} - (d + d_l) \cdot \mathbf{v}), \quad (2)$$

where d is the distance between diffuser and front panel and d_l is the distance between front and rear panel (Fig. 2 in a). The spatial coordinates on the panels are denoted by ξ whereas the functions $f(\xi_1)$ and $g(\xi_2)$ give the transmittance of front and rear panel at each position.

In this model, the panels run at a frame rate that is K times faster than the HVS. As observed by Lanman et al. [3], the emitted light field of any pair of LCD patterns corresponds to their outer product and is therefore rank-1. The light field observed from the high-speed panels $\tilde{l}(\mathbf{x}, \mathbf{v})$ is rank- K due to the retinal integration of K rank-1 light fields. Combining Eqs. (1), (2)

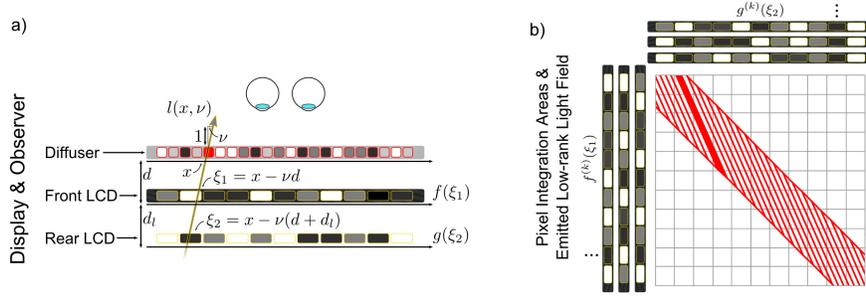


Fig. 2. Schematic of display components and parameters in a). A diffuser is directly observed by the viewer and optically projects a 4D light field into a superresolved 2D image. The light field is emitted by two high-speed LCD panels. Optically, their combined effect is a multiplication allowing the light field to be represented by the outer product of their respective patterns $f(\xi_1)$ and $g(\xi_2)$ in b). The pixels on the diffuser have a resolution exceeding that of either LCD panel, their integration areas are illustrated in red in b).

results in the following expression for the image observed on the diffuser:

$$\begin{aligned} \tilde{i}(\mathbf{x}) &= \int_{\Omega_\nu} \frac{1}{K} \sum_{k=1}^K \left(f^{(k)}(\mathbf{x} - d \cdot \nu) \cdot g^{(k)}(\mathbf{x} - (d + d_l) \cdot \nu) \right) d\nu \\ &= \frac{1}{K} \sum_{k=1}^K \iint \phi(\mathbf{x} - \xi_1, \mathbf{x} - \xi_2) \left(f^{(k)}(\xi_1) \cdot g^{(k)}(\xi_2) \right) d\xi_{1,2}. \end{aligned} \quad (3)$$

Equation (3) shows that each location on the diffuser integrates over some area on the front and rear LCDs. This integration is modeled as a convolution with a 4D kernel ϕ . For an infinitely small point \mathbf{x} on the diffuser, the kernel is

$$\phi(\xi_1, \xi_2) = \text{rect}(\xi_1/\mathbf{s}_1) \text{rect}(\xi_2/\mathbf{s}_2) \delta\left(\xi_2 - \frac{\mathbf{s}_2}{\mathbf{s}_1} \xi_1\right), \quad (4)$$

where $\mathbf{s}_{1,2}$ represent the spatial extent of the diffused point on the front and rear panel, respectively, and $\text{rect}(\cdot)$ is the rectangular function. These sizes depend on the distance between the panels d_l , that between panel and diffuser d , and the angular diffusion profile of the diffuser (see Fig. 2). In practice, the integration areas of each superresolved pixel are calibrated for a particular display configuration (see Sec. 5.1). Discretizing Eq. (3) results in

$$\mathbf{i} = \mathbf{P} \text{vec}(\mathbf{F}\mathbf{G}^T). \quad (5)$$

Here, the K time-varying patterns of front and rear LCD panels are encoded in the columns of matrices $\mathbf{F} \in \mathbb{R}^{M \times K}$ and $\mathbf{G} \in \mathbb{R}^{M \times K}$, respectively (see Fig. 2) in b). The resolution of the observed image $\mathbf{i} \in \mathbb{R}^N$ is larger than that of either panel, i.e. $N \geq M$. $\text{vec}(\cdot)$ is a linear operator that reshapes a matrix into a vector by stacking up its rows.

The matrix $\mathbf{P} \in \mathbb{R}^{N \times M^2}$ is a discretized version of the projection operator $\phi(\cdot, \cdot)$, which maps the product space of all possible combinations of pixel values for the front and rear LCD to the space of superresolution pixels. \mathbf{P} is a Toeplitz matrix, with each row consisting of a shifted copy of the point spread function of the display, represented as $\phi(\cdot, \cdot)$. We also note that \mathbf{P} is sparse since each superresolution pixel only depends on a small number of pixels in each of the front and rear LCD panel.

Figure 2 in b) illustrates the low-rank light field matrix emitted by the two display layers along with the integration areas of the superresolved pixels. Although each of these is smaller than the regular grid cells of light rays spanned by the display, the superresolved pixels are not aligned with the grid and each pixel receives contributions from multiple different rays, allowing for superresolution image synthesis.

2.2. Superresolution image synthesis

Given a target high-resolution image \mathbf{i} and the image formation derived in the last subsection, we can formulate an objective function that minimizes the ℓ_2 -norm between the target and emitted images given physical constraints of the pixel states

$$\begin{aligned} & \underset{\{\mathbf{F}, \mathbf{G}\}}{\text{minimize}} && \|\mathbf{i} - \mathbf{P} \text{vec}(\mathbf{F}\mathbf{G}^T)\|_2^2 \\ & \text{s.t.} && 0 \leq \mathbf{F}, \mathbf{G} \leq 1 \end{aligned} \quad (6)$$

This objective is difficult to deal with, as it involves a large matrix factorization embedded within a deconvolution problem. To make the problem manageable, we split the objective using the intermediate light field \mathbf{I} produced by the display as a splitting variable

$$\begin{aligned} & \underset{\{\mathbf{F}, \mathbf{G}\}}{\text{minimize}} && \|\mathbf{F}\mathbf{G}^T - \text{ivec}(\mathbf{I})\|_F^2 \\ & \text{s.t.} && \mathbf{P}\mathbf{I} = \mathbf{i}, 0 \leq \mathbf{I}, 0 \leq \mathbf{F}, \mathbf{G} \leq 1 \end{aligned} \quad (7)$$

Here, $\text{ivec}(\cdot)$ is a linear operator reshaping the vector into a matrix, and the Frobenius norm $\|\cdot\|_F^2$ measures the sum of squared differences of all matrix elements. Although the objective function is non-convex, it is convex with respect to each individual variable $\mathbf{F}, \mathbf{G}, \mathbf{I}$ with the other two fixed. The first constraint is affine in \mathbf{I} , an additional slack variable that splits the matrix factorization from the linear operator, while both are still coupled via the added consensus constraint. We solve Eq. (7) using the alternating direction method of multipliers (ADMM [15]).

After deriving the augmented Lagrangian for (7), the minimization of the augmented Lagrangian in each step leads to the following algorithm specific for our problem (see [15] again for more details of the derivation of ADMM in general):

$$\begin{aligned} & \mathbf{I} \leftarrow \underset{\{\mathbf{I}\}}{\text{arg min}} \quad \mathcal{L}_\rho(\mathbf{F}, \mathbf{G}, \mathbf{I}, \lambda) = \underset{\{\mathbf{I}\}}{\text{arg min}} \|\mathbf{F}\mathbf{G}^T - \text{ivec}(\mathbf{I})\|_F^2 + \rho \|\mathbf{P}\mathbf{I} - \mathbf{i} + \mathbf{u}\|_2^2 \\ & \text{s.t.} \quad 0 \leq \mathbf{I} \\ & \{\mathbf{F}, \mathbf{G}\} \leftarrow \underset{\{\mathbf{F}, \mathbf{G}\}}{\text{arg min}} \quad \mathcal{L}_\rho(\mathbf{F}, \mathbf{G}, \mathbf{I}, \lambda) = \underset{\{\mathbf{F}, \mathbf{G}\}}{\text{arg min}} \|\mathbf{F}\mathbf{G}^T - \text{ivec}(\mathbf{I})\|_F^2 \\ & \text{s.t.} \quad 0 \leq \mathbf{F}, \mathbf{G} \leq 1 \\ & \mathbf{u} \leftarrow \mathbf{u} + (\mathbf{P}\mathbf{I} - \mathbf{i}) \end{aligned} \quad (8)$$

where $\mathbf{u} = 1/\rho * \lambda$ is a standard substitution that simplifies the notation [15].

Using ADMM allows Eq. (7) to be transformed into a sequence of simpler subproblems. The first step of Eq. (8) is a deconvolution problem, which we solve using SART iterations [16], while the second step is a matrix factorization problem similar to the one in the work by Lanman et al. [3]. We note that by splitting the objective in this manner, different light field display technologies could be employed and would only require the second term in the objective function to be replaced with the appropriate image formation and inversion model.

3. Three-dimensional and high dynamic range modes

As mentioned, we aim for a display architecture that can be operated in both a superresolution mode as well as 3D and HDR modes. The difference between these modes is the presence

of the diffuser, which should therefore be made electronically switchable for a multi-mode display [17, 18].



Fig. 3. **a) - c): High dynamic range display mode.** The diffuser is switched off and a light field with no angular variation, but a higher contrast than that of the LCD panels is emitted. With a simulated black level of 15%, a conventional display only achieves a low dynamic range in a) whereas the proposed dual layer display (4-frame time-multiplexed here) significantly increases the dynamic range for 2D image display in b). **d) - e): Compressive light field display mode.** We can use the algorithm by Lanman et al. [3] to generate low-rank glasses-free 3D content with the diffuser switched off. Simulated here is a rank-8 light field that has 5×3 views. In e) we show a single time frame for both the front and the rear liquid crystal panel.

With this diffuser switched off, the display hardware turns into a duallayer light field display that is functionally equivalent to the hardware configuration used by Lanman et al. [3]. We can therefore use the same algorithms proposed in that work to produce glasses-free 3D content. Figure 3 shows a simulation of the 3D mode with the same parameters (LCD size, resolution, and spacing) as in the superresolution mode.

As first pointed out by Wetzstein et al. [4], a multi-layer display can also be used to represent a 2D image with an increased dynamic range compared to the maximum available contrast on either the front or the back panel. Whereas Wetzstein et al. rely on a tomographic image formation model to generate HDR imagery, we directly apply a low-rank approximation using the proposed mathematical framework.

Figure 3 in a) - c) illustrates the approach by simulating a black level of 15% of the peak intensity for both of the two LCD panels. When presenting a conventional 2D image, its contrast is severely reduced in a). Using low-rank factorization, the contrast of the displayed 2D image is significantly increased in b). The rank-4 decompositions for both LCDs are shown in the bottom c).

4. Analysis and evaluation

In the following we analyze the different parameters of the display design in simulation.

4.1. Analyzing device parameters

Diffusion Kernel and Spacing In superresolution image acquisition techniques, the quality of an optical setup is often analyzed using the condition number of the projection matrix (e.g., [19]). We follow this approach: assuming that the proposed duallayer setup can synthesize any light field, we plot the condition number of the projection matrix \mathbf{P} for a varying distance between front LCD panel and diffuser as well as for a varying diffuser spread. Other

display parameters match those of the prototype described in the Method section of our paper. Figure 4 shows the results for a target resolution increase of $2\times$.

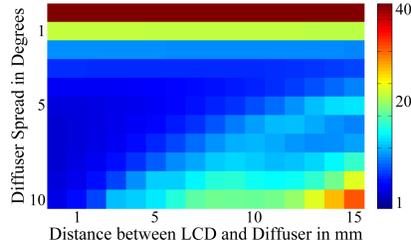


Fig. 4. Conditioning analysis for a target $2\times$ superresolution. We evaluate the condition number of the projection matrix for a varying distance between front LCD and diffuser as well as varying diffuser spread. A lower condition number corresponds to optical setups that are better suited for superresolution display.

A lower condition number corresponds to an optical configuration that is better suited for inversion or, similarly, for superresolution display. We observe that a very small distance with a diffusion spread between 5 and 10 degrees results in the best expected image quality.

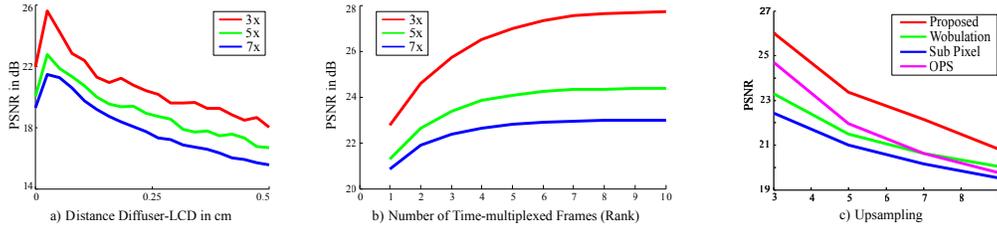


Fig. 5. Quantitative analysis of the proposed display architecture. a): we simulate reconstructions of a test scene for a varying distance between front LCD and diffuser. The resulting image quality is best for a small distance. b): reconstruction quality of a test scene on the prototype device is simulated for an increasing rank or number of subframes. While a higher rank allows for more degrees of freedom in the light field synthesis, only minor improvements in image quality are observed for ranks higher than six. Therefore, readily-available LCD panels with 120 or 240 Hz are well-suited for computational superresolution display. c): we compare the proposed method to Optical Pixel Sharing, wobulation, and subpixel rendering for a varying superresolution factor.

In Fig. 5 in a), we analyse the parameters diffuser-spread, rank (time-multiplexing) and the upsampling factor w.r.t. the superresolution performance using for simulations using a set of natural images, which confirm the findings from the conditioning analysis. Using the prototype device parameters, we simulate reconstructions of the scene from Fig. 1 for different superresolution factors. The number of subframes (rank) has been fixed to 4 for this experiment to match the frame rate capabilities of our prototype hardware. Both the PSNR analysis and the scene-independent conditioning analysis are consistent: we observe that a small distance between diffuser and front LCD is best suited for different superresolution factors.

Rank of the Light Field Factorization Faster LCD panels support more subframes, which in our system equates to a higher rank approximation of the inverse problem. We therefore expect the image quality to increase with the speed of the panels. Assuming a critical flicker frequency of about 30 Hz, for instance, allows two 120 Hz LCD panels to emit a rank-4 light field because a human observer perceptually averages over four time-multiplexed frames. We

analyze the quality of superresolution image generation for a varying light field rank in Fig. 5 in b). In this experiment, the diffuser spacing is fixed at 0.03 cm. For this example, we simulate reconstructions of the scene from Fig. 1 on the prototype device. As seen in the plots, reconstruction quality using our method asymptotically reaches a maximum for a given image as a function of the light field rank. Only minor improvements are observed for ranks higher than six.

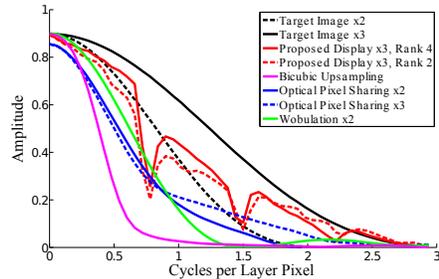


Fig. 6. Quantitative resolution analysis for several different super-resolution displays based on simulated images for each system. We show the modulation transfer function (MTF) for each system, as measured with Burns' slanted edge method [20]. The curves are normalized such that the Nyquist limit for the native hardware resolution is at 1. The black curves show slanted edge results for $2\times$ and $3\times$ superresolution target images. Our approach (red curves) generally preserves the most high frequencies, compared to other methods such as Optical Pixel Sharing and wobulation.

Resolution Analysis Next, we perform a quantitative analysis of the resolution achievable with our approach. Due to the compressive nature of our display design, the target image resolution used in the optimization procedure is a goal that is not, in general, achieved in the actual result. In order to obtain a quantitative estimate of the actual resolution increase, we measure the Modulation Transfer Function (MTF) of the display using the slanted edge procedure Burns [20].

The MTF curves for our display and a number of comparison methods are shown in Fig. 6. All MTF curves are normalized such that the Nyquist limit for the pixel size of one of the LCD panels is at 1. The black curves show MTFs for 2×2 (dashed) and 3×3 (solid) superresolution target images. For comparison we also include a cubic upsampling filter (magenta), which performs a sharpening operation and therefore contains slightly higher frequencies than the Nyquist limit of the original pixel grid. Our approach for a target 3×3 superresolution is shown in red for a rank 4 display (solid) and for a rank 2 display (dashed). We note that both variants a resolution increase between a factor of 2 and 3, with a diminishing resolution gain for higher ranks.

Comparison to Other Superresolution Displays Fig. 6 also provides a comparison with other proposed superresolution display technologies. The blue curves refer to Optical Pixel Sharing (OPS) with 2×2 (dashed) and 3×3 superresolution (solid). We note that in both cases, the OPS results are dominated by our approach, indicating that our design is overall better at reproducing higher frequencies than OPS for the same resolution LCD panels. Another way to compare the two methods is by considering parameter selections for both methods that exhibit the same compressibility requirements, i.e. the same ratio between degrees of freedom in the display hardware vs. in the target image. The two dashed curves provide that comparison, as both methods use the same number of subframes (2) for the same superresolution target (3×3).

Figure 6 shows that our method also outperforms the wobulation approach (shown in green). With an equal number or half the degrees of freedom (solid and dashed red lines, respectively)

available to wobulation, our method is able to successfully preserve much of the high-frequency content up to the target 3 cycles per LCD pixel.

Figure 5 c) shows a different way to compare with wobulation [10], [7] and OPS.

The target image for this experiment was an image showing a chirp that contains a range of frequencies. We see graceful degradation of image quality as the superresolution factor is increased. Simulated results from our method at a superresolution factor of three exceed the reconstruction quality of the other two methods by a factor of two.

5. Prototype

5.1. Methods

Display A superresolution display prototype was constructed from two Viewsonic VX2268wm 120 Hz LCD panels. All diffusing and polarizing films were removed from the front panel. This includes the light shaping diffusers that are responsible for angular dependency of the LCD panels in conventional displays. The front-most (diffusing) polarizer was replaced by a clear linear polarizer. Mounted on a rail system, the panels were adjusted to have a spacing of 19 mm in between. The rear panel had an unmodified backlight that illuminated the dual-layer device. Fixed to a frame that was adjustable on the rail system, the diffuser was mounted at a distance of 6 mm to the front LCD.

The prototype was controlled by a 3.4 GHz Intel Core i7 workstation with 4 GB of RAM. A four-head NVIDIA Quadro NVS 450 graphics card synchronizes the two displays and an additional external monitor. With the diffuser in place, the display functioned in superresolution mode using content generated by the algorithm discussed in Section 2. With the diffuser removed, the display functions in glasses-free 3D or HDR modes. While electronically-switchable diffusers are commercially available, we did not use one in our experiments.

Calibration For the prototype display, we calibrated display gamma curves, geometric alignment of the LCD panels, and diffuser point spread function (PSF).

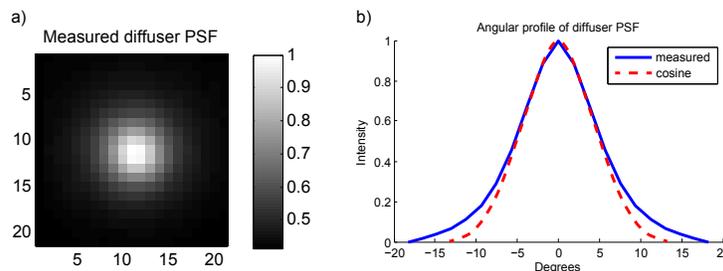


Fig. 7. Measured diffusion point spread function of the prototype in a). The angular profile of the measured PSF is well-modeled by a rotationally-symmetric cosine function in b).

Gamma curves are calibrated using standard techniques: uniform images with varying intensities are shown on the display and captured with a linearized camera in RAW format. The acquired curves are inverted in real-time when displaying decomposed patterns. The display black level is incorporated as a constraint into the nonnegative matrix factorization routine.

A second calibration step requires geometric registration of front and rear LCDs. For this purpose, we aligned the panels mechanically as well as possible and fine-tuned possible misalignments in software. With the diffuser removed, we showed crossbars on both screens that were aligned for the perspective of a calibration camera.

Finally, we measured the point spread function (PSF) of the diffuser by displaying white pixels on a uniform grid on the front panel, with the rear panel fully illuminated. The PSFs

were then extracted from linearized RAW photographs of the prototype by extracting areas around the according grid positions. The PSFs measured on the prototype were approximately uniform over the display surface, hence we averaged all PSFs and used a spatially-invariant PSF in the computational routines. Figure 7 shows the calibrated PSF captured from the device; this PSF is well modeled as a rotationally-symmetric angular cosine function with a field of view of 15 degrees (shown in b).

5.2. Results

Simulated results Figure 8 in a) shows a simulated result for the proposed superresolution display mode for a target upsampling factor of $5\times$ and rank 4.

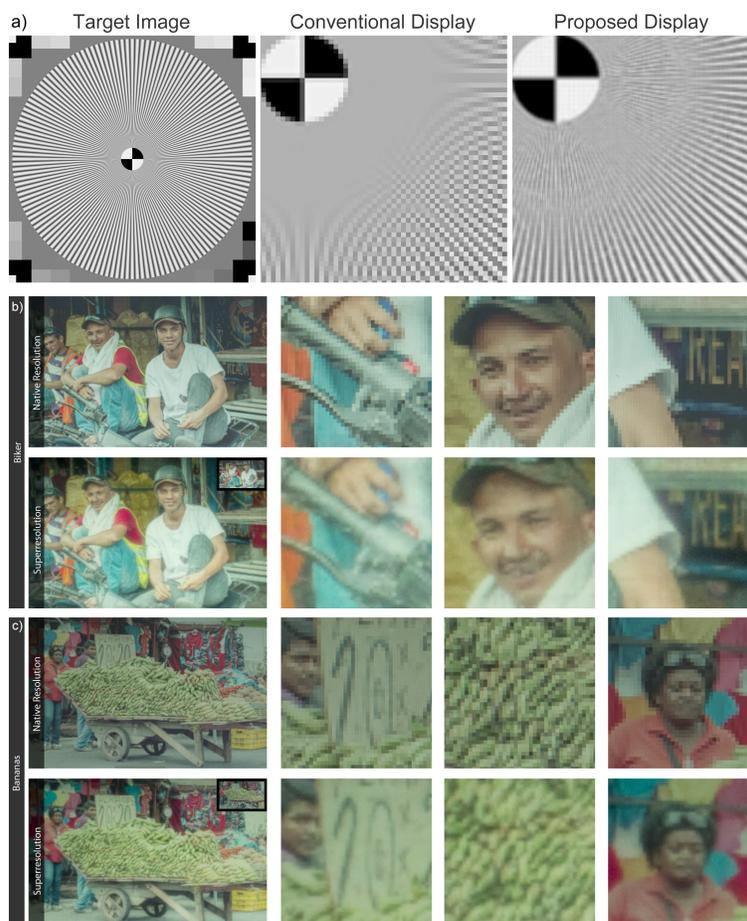


Fig. 8. **a)**: Showing a high resolution target image (left) on a lower-resolution display results in a loss of image features (center). The proposed method is capable of preserving such features by displaying superresolved image content (right). **b) - c)**: Photographs of prototype computational superresolution display. Each of the example scenes is superresolved at $5\times$. The top row of each example shows photographs captured at the native display resolution while the bottom shows our method. Isolated and thin high-frequency details such as text on signs, butterfly antennae, and sharp off-axis edges are significantly enhanced by the proposed method. Faces are also dramatically improved as are textured regions such as the bananas, plant, and clothing. Display content under CC license: <http://goo.gl/RwXvVV> <http://goo.gl/DdNka3>

Results of prototype display The two bottom rows of Fig. 8 show a result captured using our prototype hardware. Images were captured by photographing the display using an Canon T3 SLR camera. The bottom row shows a result obtained at the nominal image resolution while the bottom row shows superresolution results with a nominal upsampling of $5\times$.

The example images have several faces and off-axis edges as well as sparse text and textured regions. All are improved on the prototype, including increased detail on the brake handle, fine-scaled features around the eyes and mouth of the face and dramatic improvements to the pixel-scaled text.

Notice that the Siemens star chart result in Fig. 8 is a simulation. The according capture of the prototype you can find in Figure 9. The Siemens star result shows that we achieve a significant resolution increase on the prototype. The bunny result in b) shows that this can lead to drastic increases in perceived contrast around edges. Although in this case we suffer from slight color artifacts due to errors in the calibrated response curves of the panels, you can notice a clear improvement in image quality.

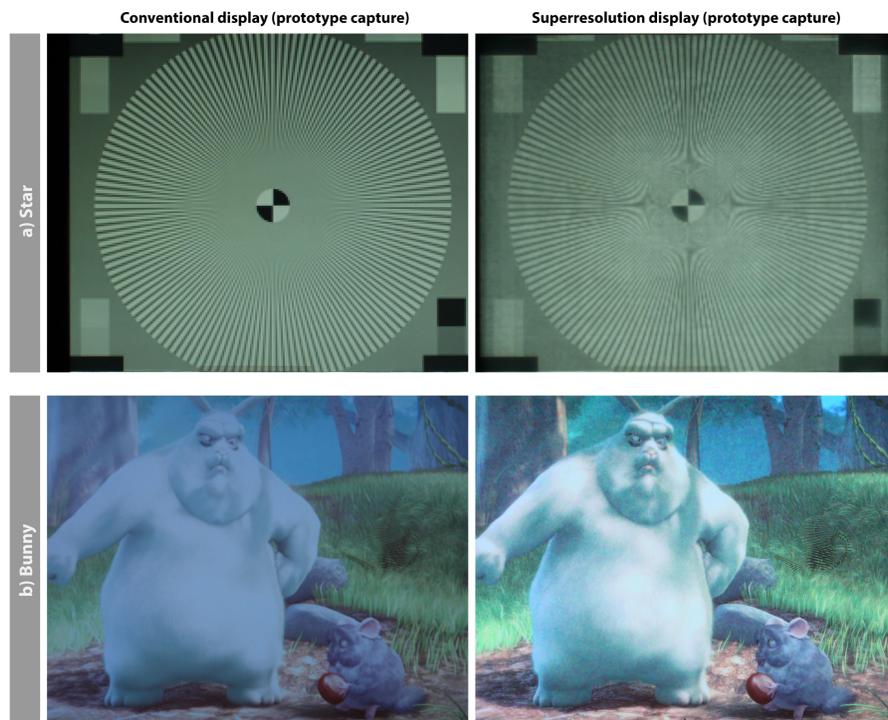


Fig. 9. **Additional prototype results:** a) and b) are showing a lower-resolution conventional display results in a loss of image features (on the left). The proposed method reconstructs high frequencies leading to increases contrast and sharper edges (right). Display content under CC license: <http://www.sintel.org>.

6. Discussion

In summary, we present a computational display approach to superresolution image synthesis. Through the co-design of display optics and computational processing, the proposed architecture is the first to facilitate superresolution with a single device and without the need for mechanically moving parts in a form factor suitable for televisions and computer monitors.

Moreover, it is the first design that can be reconfigured on the fly between three radically different display modes: superresolution, glasses-free 3D, and 2D high dynamic range.

The off-the-shelf hardware components used in our prototype exhibit limited contrast, color cross-talk, and interreflections between LCD panels that result in slight deviations of observed results from simulations; these deviations are perceived as ringing and loss of contrast. While the employed 120 Hz LCD panels allow for rank-3 to rank-4 decompositions, depending on the viewing conditions, higher speed panels are desirable. Our simulations indicate that a refresh rate of 240 Hz would be ideal for the application of superresolution image display. Currently, we process each color channel separately and do not take the panel-specific subpixel structures into account. Diffraction has not been an issue in our experiments, but significantly smaller pixel sizes would require such effects to be included in the image formation. Finally, we envision the proposed display to operate with an electronically-switchable diffuser but leave this to future engineering efforts.