

# Reprocessing Anaglyph Images

Henry G Dietz

Electrical and Computer Engineering, University of Kentucky, Lexington, KY USA

## ABSTRACT

In related work, we have shown that conventional digital cameras easily can be modified to directly capture anaglyphs. Anaglyph images commonly have been used to encode stereo image pairs for viewing, but anaglyphs also can be treated as an efficient encoding of two-view image data for reprocessing. Each of the two views encoded within an anaglyph has only partial color information, but our preliminary results demonstrate that the "lost" information can be approximately recovered with any of a variety of reasonably efficient algorithms. This not only allows credible full-color stereo pairs be computationally extracted, but also enables more sophisticated computational photography transformations such as creation of depthmaps and various types of point-spread-function (PSF) substitutions.

**Keywords:** Anaglyph, Stereo Pair, Color Analysis, PSF, Deconvolution

## 1. INTRODUCTION

An anaglyph, as shown in Figure 1, is a color image that creates the illusion of depth when viewed through color filters that separate the left and right views. The concept is usually credited to Wilhelm Rollman for experiments he performed in 1853, but may be much older.<sup>1</sup>

The most common photographic process for creating an anaglyph has generally been to obtain two separate images from left and right viewpoints, superimpose and align the images, and then remove (approximately) complementary colors from each of the views. Most often, the green and blue color channels are removed from the left image and the red color channel is removed from the right image. Less commonly, but with potentially

---

Further author information: E-mail: [hankd@engr.uky.edu](mailto:hankd@engr.uky.edu), Homepage: <http://aggregate.org/hankd>



Figure 1. Sample Green-Left/Magenta-Right Anaglyph



Figure 2. Sample Anaglyph Channels: Red, Green, and Blue

better quality, the red and blue channels can be removed from the left and green from the right. We will refer to these two encodings as red/cyan and green/magenta, which both happen to have the property that two of the three color channels are copied from the right side view. Reconstructing a stereo image pair from either type of anaglyph requires estimation of the missing color channel values at each pixel location.

The complexity of finding credible values for the missing color channel data is highly scene dependent. For example, if the stereo baseline between the views is negligible compared to the minimum distance to an object in the scene, both views are identical copies of the anaglyph. Recovery also is simple if the scene is known to consist entirely of neutral gray shades – the missing channels can be trivially reconstructed by simply replicating the present color channel(s) in each view for those that were missing. As the anaglyph color channel extractions shown in Figure 2 demonstrate, it can be very difficult to “match up” portions of scenes for which the colors used to encode the anaglyph vary dramatically in sensitivity (here, also in object transparency – one can see things through the red stripes). As an extreme case, missing colors in the reconstructed stereo pair for a scene consisting of small, randomly colored, dots cannot be recovered.

The goal of this paper is to suggest a variety of new ways in which anaglyphs can and should be reprocessed to produce full-color stereo pairs and various other types of data and images.

## 2. INPUT ANAGLYPH PROPERTIES

Although all anaglyphs use color to encode left and right views, the encoding and other qualities of a particular anaglyph can have a huge impact on the effectiveness of reprocessing. Some of these properties of an anaglyph can be corrected before reprocessing the anaglyph, while others render the anaglyph unsuitable as input. These issues can be naturally divided into relevant problems that exist in conventional 2D images and those that are specific to anaglyphs.

Key problems that also occur in conventional 2D photography include:

- Many cameras use encodings that apply a non-linear gamma, most often JPEGs<sup>2</sup> with a gamma of approximately 2.2. The anaglyph image data should be converted to a linear encoding prior to performing any of the other processing described in this paper. If desired, the resulting linear gamma stereo pair can be converted back to the original gamma as a final step.
- Transverse (lateral) *chromatic aberration (CA)* typically appears in the form of different color channels having slightly different magnifications. It is not uncommon for modern lenses to cause a color channel to be offset by as much as several pixels by the edge of the frame. Given that the color misalignment that gives depth to anaglyphs is often on the same scale, correction of this CA always should be performed by scaling the color channels appropriately.
- Lens projection distortions should be corrected. For example, uncorrected barrel and pincushion distortion both cause minor alignment errors in anaglyphs. When using lenses with very wide view angles, even the “stretching” of off-center objects by a perfect rectilinear projection can cause alignment errors.

Other quality issues are specific to anaglyphs:

- JPEG<sup>2</sup> perceptual compression was not designed for scenes with the color misalignments found in anaglyphs. Empirically, the YUV color space and discrete cosine transform (DCT) appear to be particularly damaging to red/cyan anaglyphs. Image formats like PNG<sup>3</sup> are far superior for preserving the quality of anaglyph images. Thus, although individual views may be captured as JPEGs, anaglyph images constructed from them should not be.
- Yaw, pitch, or roll of the left and right frames relative to each other corrupts the stereo baseline. The left and right views in an anaglyph should differ only by a horizontal displacement, but the majority of anaglyphs sampled from the WWW as part of this research were found to have a measurable amount of this type of misalignment. Roll is most visible, and can be corrected by simple inverse rotation of the 2D view image.

- *Retinal rivalry* is a phenomenon in anaglyphs in which large differences in the color and brightness of the left and right views of a portion of a scene prevent that region from being seen with a consistent color. Instead, the region is seen as randomly oscillating between the colors in the two views. Much of the recent commercial interest in anaglyphs has centered on methods by which a small loss in color fidelity can be traded for elimination of rivalry. Independent of the color set used (although our discussion with assume red-left), The various different color codings have significantly different implications for reprocessing:

**Monochromatic** Monochromatic left and right views are simply treated as the red and blue channels of the anaglyph. The green channel is either dark or copied from the right image. Reprocessing to recover the monochromatic stereo views is trivial, but no information is present about scene colors.

**Color** The red channel from the left view is combined with green and blue channels from the right view. Perhaps because of its simplicity, this is the most common scheme used for anaglyph coding. Although retinal rivalry can be severe when the anaglyph is viewed, this channel-separable encoding is ideal for color reprocessing.

**Half-Color** The left view is converted to a monochrome image from which the red channel is combined with the green and blue channels from the right view – reducing both retinal rivalry and color quality. This type of coding is amenable to reprocessing, and the fact that the anaglyph red channel is largely derived from the left view’s green and blue information implies that computational matching between views will be superior. However, this type of coding is rare, and thus is not discussed further in the current work.

**Optimized/“Leaky” Color** In any of various ways, the left and right channels are forced to be more similar than they would be for a conventional color anaglyph. Most methods involve polluting one or more of the anaglyph color channels with the values “leaked” from other channels in the same view. The leakage can be linear or can be computationally optimized by replacing problematic colors in full-color stereo pairs before performing conventional color or half-color conversion into anaglyphs. Commercial anaglyph systems using “leaky” colorings include Anachrome,<sup>4</sup> InfiniColor 3D,<sup>5</sup> and ColorCode 3D.<sup>6</sup> Anaglyphs created with arbitrary “leaky” coloring can be reprocessed as ordinary color anaglyphs, but generally will suffer ghosting artifacts and color errors.

The remainder of this paper will assume that all input anaglyphs are appropriately pre-processed color anaglyphs using pure red/cyan, green/magenta (also known as TrioScopic<sup>7</sup>), or blue/yellow (which is very rare) encoding. These three color sets share the important property that the left and right views can be separated by extracting color channels from an RGB representation of the anaglyph. To simplify the descriptions in this paper, it is further assumed that the left view is encoded by a single color channel and the right by two.

It is perhaps surprising that the commonly-discussed anaglyph issues of focus plane position, length of the stereo baseline, etc., are not directly relevant. These choices only have secondary impact on reprocessing. For example, very short stereo baselines with very distant scenes may virtually eliminate view misalignment, but all processing works the same way that it would have if a longer baseline and nearer scene had been used.

### 3. DECONVOLUTION

In computational photography, deconvolution is commonly used to sharpen a “blurry” image. The key concept is that each pixel in the original image,  $d_i$ , has a convolved value:

$$d_i = \epsilon_i + \sum_j p_{ij} u_j \quad (1)$$

in addition to a small level of noise,  $\epsilon_i$ , each point in the scene,  $u_j$ , contributes a fraction of its photons to  $d_i$ . The fraction contributed by each point,  $p_{ij}$ , is determined by the (normalized) point spread function (PSF), which is essentially the image created for a single point of light at position  $j$ . A wide range of techniques, largely involving analysis in the frequency domain, deconvolve using a given PSF to determine  $u_i$  values. Typically, it

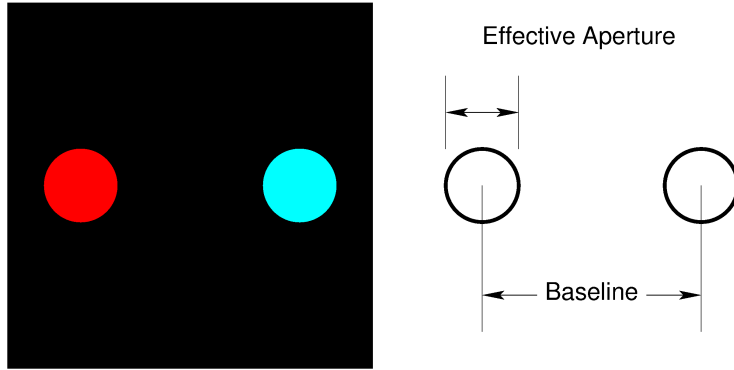


Figure 3. Red/Cyan (Virtual) Coded Aperture and Parameters

is assumed that the PSF is given as a known input; where it is not specified, it is determined automatically as part of blind deconvolution.<sup>8</sup>

Perhaps surprisingly, the problem of mapping an arbitrary anaglyph into a full-color stereo pair also can be treated as a deconvolution problem. In fact, “deblurring” and other PSF-substitution transformations in general can be applied to anaglyphs... but what is the PSF of an anaglyph?

### 3.1 Anaglyph PSFs

Traditionally, anaglyphs are produced by any of a variety of methods:

- sequential capture shifting the camera position between exposures
- simultaneous capture using a mirror assembly to produce side-by-side views
- simultaneous capture through two separate lenses

All of these methods actually capture a full-color stereo pair. The anaglyph is then produced as a color-filtered overlay of the two images in each stereo pair. In order to minimize retinal rivalry while retaining natural colors in viewing, the color filtering process may be surprisingly complex. However, in the simplest case, the filters merely extract disjoint color channels from the left and right images. For example, a red/cyan anaglyph can be produced by adding the red channel from the left image to the green and blue channels from the right image.

However, it also is possible to directly capture an anaglyph. To do this, the filters are applied during exposure and the images are combined within the camera. Obviously, the sequential capture method can accomplish this by using a two-shot multiple exposure with the appropriate filter for each shot. Less obviously, it is actually even easier to capture an anaglyph in a single shot using a conventional camera and lens with a coded aperture. This idea was discussed in Songer’s 1973 patent,<sup>9</sup> applied commercially circa 1991 in the famous but unsuccessful “Vivitar Series 1 Qdos 70-210mm  $f/2.8$ - $f/4.0$  3-Dimensional Lens System Macro 1:2.5x,” and heavily refined and tuned for digital still and video cameras by us over the past two years.<sup>10</sup>

The details of our single-shot anaglyph capture method are not the topic of this paper; the important insight is that **any anaglyph is equivalent to a single capture using a special type of color-coded aperture**, like that shown in Figure 3. The PSF image of an out-of-focus (OOF) point of light is generally clipped by the aperture of the lens, thus taking on the aperture’s shape. The reason it takes that shape is simple ray tracing: the aperture blocks rays coming from viewpoints outside the aperture’s opening. The shocking realization is that it doesn’t matter if the two apertures used to isolate the stereo viewpoints are taken from within one lens or two well-matched separate lenses. All that matters is the coding of the dual-aperture.

Let’s be more precise about this. Each ray of light passing through a single, circular, aperture, from an OOF source lands at a different position on the sensor or film. Contrary to intuition, nothing OOF is really “blurred”;





Figure 4. PSF Images to Substitute to Restore Left and Right Images

it is merely convolved with light coming from multiple viewpoints. Thus, no matter how the anaglyph was captured or created, an anaglyph is simply an image convolved by a scalable PSF that for a red/cyan anaglyph looks like that shown in Figure 3. Although the size of the PSF changes with how OOF the point source is, the *ratio of the Effective Aperture to the Baseline is constant* for all points imaged. Even if each aperture has an arbitrary shape, rather than being circular, the same ratio invariance property still holds.

Using deconvolution based on a scalable PSF like that shown in Figure 3 directly produces estimates of the  $u_i$  values as per Equation 1. Others have explored use of multiple apertures for computational photography,<sup>11-14</sup> but we believe we are the first to exploit the relationship to anaglyphs and the simplicity of imposing an appropriately coded aperture as a filter in front of an unmodified lens.<sup>10, 15</sup>

Beyond the awkwardness of processing a multi-channel image, our method must deal with the fact that the PSF scaling includes flipping of the PSF on opposing sides of the plane focused upon: if distant objects have red-then-cyan edges, close ones will have cyan-then-red. This actually can be considered as an advantage in recovering depth, because, unlike depth recovery using more conventional coded apertures,<sup>16</sup> identically-sized PSFs before and after the focus plane are trivially disambiguated. An alternative deconvolution procedure would treat individual color channels of the PSF and anaglyph separately, using an additional algorithm to match points across the channels based on the preservation of the ratio of the Effective Aperture to the Baseline.

Once deconvolution has produced the  $u_i$  estimates, it also has produced a (signed, before/after the focal plane) estimate of the PSF scale. Thus, creating the left and right images of a full-color stereo pair is simply a matter of substituting the appropriately scaled PSFs shown in Figure 4.

### 3.2 Problems With Deconvolution

Given that any anaglyph can be treated as a coded aperture image, it seems natural to use any of the standard known-PSF deconvolution methods to create the full-color stereo pair. However, there are significant problems.

One issue with deconvolution in general is that it is not particularly fast. A problem specific to anaglyphs is that they often are deliberately created to have very large depth of field (DOF). This implies that the Effective Apertures will be small enough to make measurements of the ratio between Effective Aperture and Baseline imprecise.

Conventional deconvolution, in general, also suffers from the assumption that convolution is additive. Clearly, as the apertures become small enough to have any structures (e.g., corners) approaching the diffraction limit, wave properties of light can become evident. Avoiding fine structures in the aperture design can minimize this effect. Unfortunately, there is another way in which the assumption that light adds is wrong.

Consider Figure 5. The first image shows a paper cut-out of the Electronic Imaging logo shot at about twice the focus distance – significantly OOF. The second is a point light source several times further away. One might expect that placing the paper between the point source and the camera would have added the two images, but the actual result is quite different. The third image shows that objects in front of an OOF point remove



Figure 5. OOF Cut-Out Pattern, OOF Point Light Source, OOF Cut-Out In Front Of OOF Light

rays, thus yielding a sharp image despite neither the cut-out nor the point source being in focus. Actually, this is a variant of the phenomenon well-known as *occlusion* in the stereo vision community. This sensitivity to depth order actually reveals more information about the scene, but clearly violates the assumption that light intensities simply add. Put another way, deconvolution assuming addition does not create credible data for occluded portions of a scene.

In summary, the concept of reprocessing anaglyph images using deconvolution shows great promise. Our ongoing research in deconvolution of anaglyph images centers on methods which respect the non-additive “depth layering” properties of the PSFs.

#### 4. STEREO MATCHING AND REGISTRATION

In the computer vision community, there are many algorithms designed to register multiple views<sup>17</sup> or more specifically to recover 3D data about a scene by matching corresponding points in a stereo image pair.<sup>18–20</sup> Clearly, given a correspondence between coordinates in the left and right color channels of an anaglyph, it is trivial to create full-color left and right images by for each position copying the missing color(s) from the corresponding position in the other side’s image. As for deconvolution processing, occlusions are problematic. However, many stereo matching algorithms at least explicitly identify the regions that are occluded as such, so they could be handled by applying a second algorithm – such as the color analysis discussed in Section 6. Another possibility is to use something like resynthesizer<sup>21</sup> to create credible data for occluded areas.

The primary problem is that stereo matching depends strongly on having features or patterns that can be identified in both images. Once again, consider the anaglyph color channel extractions shown in Figure 2. Although some features certainly do appear clearly in both images, the red stripes in the flags are actually transparent in the red channel image, letting various structures behind them show through. For this scene, the problem would be even worse using a red/cyan anaglyph instead of green/magenta, because the stripes are transparent in neither green nor blue. The wavelength-dependence of *albedo*, or reflection coefficient, can make the correspondence between misaligned different-channel images ambiguous.

When we attempted to use traditional dense stereo matching to create full-color stereo pairs, the results were very poor. The problem lies in the fact that tonal properties of different color channels differ greatly, so attempting to match tonal features from either channel in the right view to the channel in the left view is not effective.

This observation led to the idea of trying to map the existing channels into computed channels that would have more consistent features for stereo matching. An obvious transformation to try was creation and matching of gradient channels, which worked significantly better than using the anaglyph channels directly, but left large areas of the views between features ambiguous. An alternative approach was to use the color analysis described in Section 6 to create an estimate of the missing channel of the right view and to match that against the left

view’s channel. This was more effective, but even minor errors in the shading of the synthesized right channel had major impact on the quality of the matching obtained.

The problem is essentially that the matching process needs to be more capable of finding credible colorings for all pixels. This realization led to a variety of novel alignment schemes that computed weighted average color values for each pixel rather than finding the best matches. Some of these schemes are computationally very efficient, and credibility of the reconstructed stereo pair is usually good, but large regions of the images often have their colors “contaminated” by small numbers of incorrectly weighted pixel values. A color histogram matching algorithm was able to reduce, but not remove, these tints.

## 5. SHAPE-BASED (SUPERPIXEL) PROCESSING

Grouping of pixels into “superpixel”<sup>22</sup> shapes can produce structures that are relatively independent of tonal differences between channels. Initially, a simple recursive seed fill algorithm was adopted to group pixels that had similar values in the selected channel. That shape-grouping algorithm was replaced by either of two different non-recursive algorithms, the second of which is very “cache friendly” and typically runs in approximately  $O(n)$  time for an  $n$ -pixel image. Shapes are not just collected once, but created for a range of different shape-inclusion thresholds. Thus, most pixels are part of multiple shapes, not just one.

The shape information has been used in two ways. First, shapes were used to smooth colors such that pixels within each shape have the values in the reconstructed channel(s) shifted toward the average color or average hue for the containing shape. Alternatively, shape-based stereo matching was used to copy missing channel data from corresponding pixels in a shape matched from the other view. Testing included both accepting the best shape match for each pixel and using a weighted average over all approximately matching shapes.

## 6. COLOR ANALYSIS

Another family of approaches to recovery of the missing color channels is based entirely on the colors within a view rather than spatial relationships between the views. Because more color information is missing from the left side view in our input anaglyphs, this type of approach is less likely to produce satisfactory results for the left channel than for the right.

As Edwin Land observed in 1959,<sup>23</sup> and later attempted to codify in the *Retinex Theory of Color Vision*,<sup>24</sup> real scenes commonly allow all colors to be recovered despite having fewer than 3 channels recorded. Directly applying Retinex to produce the missing color channels in an anaglyph was ineffective. However, the work of Land and others in 2-color systems has demonstrated that colors that were not explicitly recorded often can be recovered.

In addition to the known color channels in each view, the approximate statistical properties of all color channels over both views also is known. The complete color data in an anaglyph suffers typically small, but unspecified, spatial mis-registration between left and right channels. This imperfect information is far more than a simple “gray world” assumption would provide.

### 6.1 The Color-Guessing Game

Given that an anaglyph actually does contain all three color channels, the question becomes how much of that misaligned color information can be extracted?

In a typical anaglyph, some fraction of the scene is naturally aligned to within a pixel. For example, a baseline of a few inches is negligible relative to a background that is miles away. If the naturally aligned regions can be recognized as such, the color information they provide is known accurate. Not only is it accurately describing the full color of each aligned pixel, but it also is establishing that those specific RGB triples actually occur in the luminant-colored scene. For example, if RGB value 0x802a1c is known to occur frequently in the aligned portion of a scene, it seems likely that a pixel in a not-aligned region with R=0x80 and B=0x1c should have G=0x2a.

Suppose we are examining a region of an anaglyph which has a relatively constant color for a horizontal distance which exceeds the maximum possible stereo misalignment between left and right views. The object in this region might not be fully aligned – there may be color fringes on either or both sides of it. However, the

colors that appear on the wide overlapping portion of the object are approximately correct. In fact, if there are fringe colors on either side of this region, the actual alignment is most likely a shift corresponding to the width of the fringe, so even more precise color samples can be obtained.

The identification of color fringe areas is easy for humans, and approximately identifying them can be fairly easy for a computer. A simple approach is to look for areas with colors that are heavily biased toward either of the anaglyph coding colors; for example, in a red/cyan anaglyph, strongly red or strongly cyan pixels are good candidates for fringe areas. More accurately, a region with an RGB color tuple in the anaglyph that is relatively far from any RGB tuple known to occur in the image is probably fringe. The accuracy can be further improved by imposing the rule that the pixels can only be fringe if they are in a horizontal group less wide than the maximum shift. We do not know the true colors of fringe areas, but it is likely that they do not introduce color tuples that are dramatically different from any seen in other portions of the anaglyph.

Spatially biasing color guesses further improves accuracy. For example, if both RGB value 0x802a1c and 0x80741c occur frequently in the aligned portion of a scene, the choice of G=0x2a or G=0x74 for a pixel in a not-aligned region with R=0x80 and B=0x1c can be biased based on which is more common in that locality of the image. Further, using approximate color matching in a color space like HSV empirically worked better than RGB.

## 6.2 From Right To Left

The above color guessing logic alone often produces acceptable right-view reconstructions. Left-view reconstructions generally lack detail, especially in the case of a red/cyan anaglyph. Red/cyan anaglyphs that have been “optimized” to minimize retinal rivalry (as discussed in Section 2 yielded particularly poor results because the optimized coloring introduces “ghost” (print-through) edges in color channels that should not have had edges, thus introducing errors in the fringe region and shape recognition processes.

Additional improvement has been obtained by using the colors of the reconstructed right view, rather than the anaglyph, to provide a color histogram that colors for the left view could be shifted to match. A more recent scheme iteratively uses the right view to estimate the left, then the left to refine the estimate for the right, etc.

Despite various imperfections, at this time, color analysis is clearly the fastest and usually most effective basis for reprocessing anaglyphs into full-color stereo pairs.

## 7. PRELIMINARY RESULTS AND CONCLUSION

Anaglyph images are one of the oldest representations of images for stereo viewing, but we suggest that they also are a viable representation for image data to be reprocessed by computational photography means. In fact, we suggest that anaglyphs may be a better form for capture and reprocessing than plenoptic captures, conventional structured aperture images, etc.

While anaglyphs might not be the simplest form to work with, they have two huge advantages. First, there is a relatively large set of anaglyph images already available. Second, as we have shown in other work, **nearly any conventional still or video camera can be reversibly modified for less than \$1 to be able to capture anaglyph images in a single exposure. All of the anaglyph images in this paper were captured using this single-exposure trick.**<sup>10</sup> In fact, to give a realistic worst-case scenario, the corrections described in Section 2 were not even applied.

Figure 6 shows six green/magenta anaglyphs and the full-color stereo pairs produced by reprocessing them using an algorithm based on color analysis augmented by shape-based smoothing. More conventional methods for anaglyph capture can produce better-quality anaglyphs, and thus superior stereo pairs. Even the flags image – with which this paper begins and ends – is converted into a reasonably credible stereo pair.

These preliminary results are far from perfect, but these are the first reprocessing results we have formally published. The algorithm used for the examples in Figure 6 was developed in February 2011, and since that time our reprocessing algorithms have continued to improve. For example, we are now using both neural network and genetic algorithm (GA) techniques to develop better reprocessing algorithms.






Figure 6. Single-Shot Green/Magenta Anaglyphs Reprocessed Into Left/Right Stereo Pairs

In conclusion, the point of this first paper on anaglyph reprocessing is not the details of any particular algorithm, but the simple proof of concept validation that anaglyphs – which can be captured with nearly any camera in a single shot – are a viable representation for reprocessing.

## REFERENCES

- [1] Zone, R., [*Stereoscopic Cinema and the Origins of 3-D Film 1838-1952*], The University Press of Kentucky (2007).
- [2] Pennebaker, W. B. and Mitchell, J. L., [*JPEG Still Image Data Compression Standard*], Van Nostrand Reinhold, New York (1992).
- [3] Hammel, M. J., “Png: The definitive guide,” *Linux J.* **2000** (Jan. 2000).
- [4] Silliphant, A., “Anachrome,” <http://www.anachrome.com/> (2010).
- [5] “TrioViz,” <http://www.trioviz.com/> (2010).
- [6] “ColorCode 3-D,” <http://www.colorcode3d.com/> (2010).
- [7] “TrioScopics 3D,” <http://www.trioscopics.com/> (2010).
- [8] Levin, A., Weiss, Y., Durand, F., and Freeman, W., “Understanding and evaluating blind deconvolution algorithms,” in [*Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*], 1964–1971 (june 2009).
- [9] Songer, J. D., “Three-dimensional color photographic process, apparatus and product,” *United States Patent #3,712,199* (January 1973).
- [10] Dietz, H., “Use Your Camera To Capture ”3D” Anaglyphs,” <http://www.instructables.com/id/Use-Your-Camera-To-Capture-3D-Anaglyphs/> (Nov 2010).
- [11] Green, P., Sun, W., Matusik, W., and Durand, F., “Multi-aperture photography,” *ACM Trans. Graph.* **26** (July 2007).
- [12] Bando, Y., Chen, B.-Y., and Nishita, T., “Extracting depth and matte using a color-filtered aperture,” *ACM Trans. Graph.* **27**, 134:1–134:9 (December 2008).
- [13] Baek, J., “Multi-channel coded-aperture photography,” *Ph.D. Dissertation* (2008).
- [14] Zhou, C., Lin, S., and Nayar, S. K., “Coded Aperture Pairs for Depth from Defocus,” in [*IEEE International Conference on Computer Vision (ICCV)*], (Oct 2009).
- [15] Dietz, H., “Single-Lens Anaglyph Aperture Tool,” <http://aggregate.org/anaperture> (Aug 2010).
- [16] Levin, A., Fergus, R., Durand, F., and Freeman, W. T., “Image and depth from a conventional camera with a coded aperture,” in [*SIGGRAPH ’07: ACM SIGGRAPH 2007 papers*], 70, ACM, New York, NY, USA (2007).
- [17] Brown, L. G., “A survey of image registration techniques,” *ACM Comput. Surv.* **24**, 325–376 (December 1992).
- [18] Medioni, G. and Nevatia, R., “Segment-based stereo matching,” *Computer Vision, Graphics, and Image Processing* **31**(1), 2 – 18 (1985).
- [19] Kanade, T. and Okutomi, M., “A stereo matching algorithm with an adaptive window: theory and experiment,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **16**, 920 –932 (sep 1994).
- [20] Sun, J., Zheng, N.-N., and Shum, H.-Y., “Stereo matching using belief propagation,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **25**, 787 – 800 (july 2003).
- [21] Harrison, P. F., “Image Texture Tools: Texture Synthesis, Texture Transfer, and Plausible Restoration,” *Ph.D. Dissertation* (2005).
- [22] Achanta, R., Shaji, A., Smith, K., Lucchi, A., Fua, P., and Susstrunk, S., “SLIC superpixels,” *EPFL Technical Report* (June 2010).
- [23] Land, E. H., “Experiments in color vision,” *Scientific American* **200**, 84–99 (May 1959).
- [24] Land, E. H., “The retinex theory of color vision,” *Scientific American* **237**(6), 108–128 (1977).