

What's the best lens for stereo? An exploration of lens models and their stereoscopic depth performance

Leonid Keselman
Stanford University
leonidk@stanford.edu

Abstract

We introduce the use of various projection functions in the analysis of stereoscopic depth sensors. Through these methods, we are able to design stereo systems which are not bound by traditional quadratic depth error. Additionally, we demonstrate how existing correspondence algorithms can be modified to handle these lens designs. In addition, we can construct lens projection models which are more suited to natural lens designs, and build stereo systems with minimal re-sampling errors. This also allows us to construct wide-angler stereoscopic systems than previously possible, without significant re-sampling or sacrificing accuracy.

1. Introduction

Stereoscopic correspondence methods are at the heart of many practical depth sensors. This is because they're practical and simple to build, requiring only a rigid baseline and two sensors. Their use is widespread, from dental tooth scanners, to consumer electronics and advanced driver assist systems [7].

All of these systems, to the best of our knowledge, are built on an assumption that the lens used is a rectilinear lens. Rectilinear lenses create a projection model like that of a pinhole camera; characteristic of such a lens is that straight lines are always straight. This maps to humans own physical perception of the world. This is also known as perspective projection, and we will discuss some of its shortcomings, as applied to stereo, in this paper. The goal of this paper is to ask and analyze a simple question, *if one could design any lens system possible, how could you fully utilize it and what would the impact be?*

1.1. Stereoscopic depth systems

To generate valid stereo correspondence, we must first setup epipolar geometry. In this paper, we will use the convention of mapping the epipole at infinity, with the epipolar lines parallel to the horizontal axis of the imagers. While

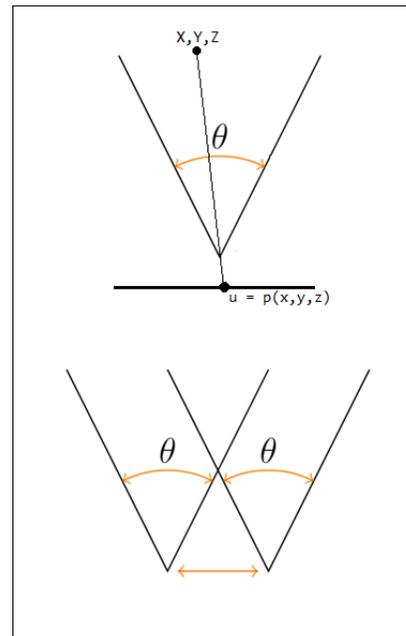


Figure 1. Two dimensional overview of our coordinate system. The field-of-view of each camera is identical and equal to θ . The spacing between the cameras, or baseline, is B . The projection function of the system is that which maps a 3D point (X, Y, Z) to a position on the sensor (using U and V as coordinate values)

exist methods exist for searching along other configurations of epipolar geometry, this paper will not explore them. Additionally, we assume that both camera use the same exactly lens model.

Additionally, the creation of valid depth correspondence requires that a lens projection equation be fully captured. this is because real lenses typically have some non-designed distortions. This is true for traditional rectilinear lenses, fisheye lenses, and likely any of the other lens models presented here. This distortion from a non-ideal lens is first captured and corrected for in a lens distortion step, before the lens intrinsics are estimated. While this is sometimes

a two-step process, many modern implementations [4] estimate both distortion parameters and the camera matrix simultaneously through the use of a Levenberg–Marquardt solver. This process, of capturing the full camera projection model, is required to create a consistent and known mapping from correspondences on the image plane to correspondences in the real world. And holds for all lens models presented here, although the camera model solver would have to be modified accordingly.

The model of projection most often used is that of a rectilinear lens, with a Brown-Conrady distortion model. While our paper discusses other lenses, the existing methods [24, 4] of using nonlinear optimization to capture a real lens distortion model can be trivially augmented to handle our proposed mechanisms and methods. For a general introduction to the basic geometries and methods, we recommend Szeliki’s book [21].

1.2. Lens Projections

This work proposes that lenses are analyzed strictly with their projection models. This allows us to perform straightforward analysis of the stereoscopic depth errors, along with building algorithms that perform directly depth matching directly on the projected images. Typical optical literature characterizes lenses via an angular specification, while this analysis is easier in Cartesian coordinates. A summary table of some common projections and their characteristics is available in table 1. A straightforward discussion of the benefits and types of common lens projections is readily available [3], from where the angular forms of these projections is obtained. The Cartesian forms were derived for the table on page 3.

1.2.1 Rectilinear lenses

A traditional lens is based on the pinhole model of light. It’s the projection that would be given with a pinhole lens. It is also known as rectilinear or perspective projection. It’s use is common because it preserves straight lines remaining straight. Additionally, it’s sampling is constant across a plane positioned in front of the camera, allowing a constant sized window to across the entire frame, as

$$x = \frac{1}{f} \cdot u \cdot z$$

$$u = f \cdot \frac{x}{z}$$

$$\frac{\partial u}{\partial x} = \frac{f}{z}$$

This is easy to see, as changes in real world position (x), have no dependence on position in the frame (u), only on the distance of the object and the focal length of the lens. Hence, for fronto-parallel objects, a constant sized window

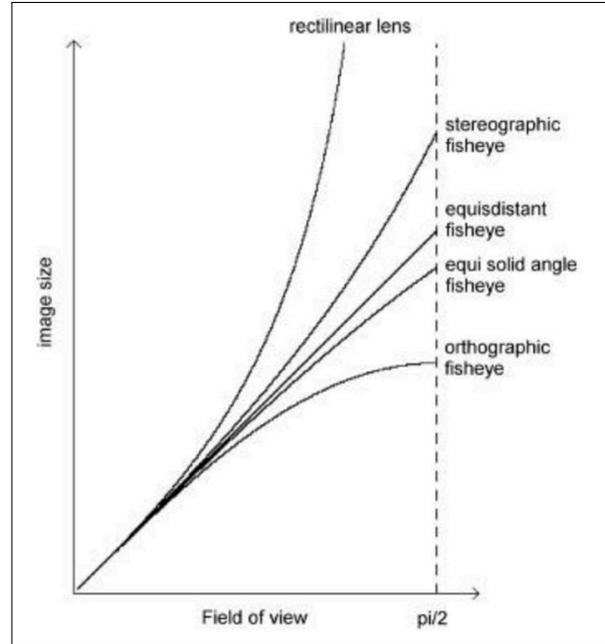


Figure 2. A figure taken from [3], showing the behavior of various lens projection types when given a fixed focal length.

would provide consistent matching over the entire frame, provided the focal lengths and distances of the images are matched. However, as described below, it has some very natural shortcomings.

1.2.2 Shortcomings of traditional lenses

Building very wide angle lenses is impractical under pin-hole projection. This is for two reasons: first, the number of pixels required for any object grows as $\tan(\theta)$, where θ is the off-axis angle of the optical ray. This trends to infinity as pinhole imagers try to cover a hemisphere (180 degrees)¹. Second, due to the physical properties of such a sampling function, there is strong natural vignetting. Specifically, perspective projection imagers exhibit $\cos^4(\theta)$ [2] light loss. In fact, this rule of thumb is often an underestimate due to the transmission properties of the optics involved; for details, see the cited paper.

1.2.3 Fish-eye Lenses

The common alternative lens model to perspective projection is a fish-eye projection. These have recently taken off for use in security, automotive, and consumer use. Unfortunately, this is a catch-all term for a variety of different lens projection models. We pick two to analyze here, the

¹Rectilinear lenses project as $u = \tan(\theta)$ which naturally has $x \rightarrow \infty$ as $\theta \rightarrow \frac{\pi}{2}$

Projection	Angular Formula	Cartesian Formula	Comments
Rectilinear	$u = f \cdot \tan(\theta)$	$u = f \cdot \frac{x}{z}$	Also known as perspective or pinhole projection
f- θ	$u = f \cdot \theta$	$u = f \cdot \tan(\frac{x}{z})$	A common type of fish-eye projection, also known as equidistant projection
Orthographic	$u = f \cdot \sin(\theta)$	$u = f \cdot \frac{x}{\sqrt{x^2+z^2}}$	A somewhat rarer type of fish-eye projection, but one that exhibits no natural vignetting.
Cubic	$u = f \cdot \tan(\theta)^3$	$u = f \cdot (\frac{x}{z})^3$	A lens model that has no practical examples that we know of, but has interesting theoretical properties described below.

Table 1. A table of the lens projections used in paper and their various forms



Figure 3. A picture of a real anamorphic lens with an oval aperture. On the right there is a demonstration of the stretching effect exhibited by the lens. These images and further information is available [22]

f- θ model and the orthographic model, which are both in common use. Additional information about the variations of fish-eye projection is available in [3], and their diagram is repeated in figure 2.

1.3. Anamorphic Lenses

Another interesting property of lens projections is that lenses often use spherical lens components, which create circular projections, and can be characterized by the off-axis angle of interest. While this simplifies lens design, there exist alternative lenses called anamorphic lenses [11]. These lenses have asymmetric behavior, typically stretching the image down only a single axis, while leaving the other axis undisturbed. This is used in the film industry to fit wide aspect ratio images onto a more square 35mm film. They're commonly used in cinema for these reasons and are readily available for purchase. This can be seen in figure 3.

2. Related Work

In some ways, this work is motivated by Fleck's article on the use of various lens projections in computer vision [5].

Although her work didn't focus on stereo systems, it discussed the concept that rectilinear projections have notable shortcomings for certain computer vision applications.

There have been many existing methods on expanding stereoscopic correspondence algorithms to handle non ideal lenses. Many authors build on wide angle lenses, distorting the imagers back to a linear projection for stereo matching [8] [6]. Other authors [18] [12] have pursued using fish eye lenses, which are a form of nonlinear lens, but did so mostly through the use of features which were invariant under their chosen lens projection. The shortcomings of using such invariant features are that these methods produced sparse depth maps.

Abraham and Förstner developed a method setting up valid epipolar geometry for using panoramic and wide angle lenses [1]. They analyze various model properties and generate a constraint under which epipoles at infinity can be preserved. However, while their work would correctly discover nonlinear lens distortion, they do not explore dense matching on their images. Without the restriction of rectilinear sampling down the matching axis, their model does not allow for use of standard correlation algorithms, instead requiring resampling at every disparity candidate. However, Roy et al. demonstrate the value of a cylindrical model like that proposed here [16]. However, they do not analyze its properties, use dense matching algorithms on such models, nor do they apply their model to anything but synthetic data.

Li [14] [15] developed methods of using nonlinear lens for generating depth, including calibration and search, but the method only applied to only spherical projections. Additionally, their use of nonlinearities was limited to only the axis perpendicular to the epipolar line.

However, since almost all of these methods undistort back to a rectilinear projection, there is no analysis of how to utilize and design lens projection geometries that provide favorable properties for stereoscopic imagers — including improved depth invariance at the expense non-uniform accuracy across the image.

3. Model for Lens Projections

3.1. Stereoscopic System

To setup the principles of a matching stereo system, we define the coordinate system defined in 1. There is a baseline, a focal length, and a projection function from the real world to the image.

When discussing various projection geometries, there is first a choice of parametrization. For an overview of various parametrization and models, please see [19]. While optics papers general prefer an angular parametrization (where fish-eye lenses are modeled as f - θ lenses, with a linear mapping between angle and sensor position), computer vision literature typically prefers a parametrization of Cartesian coordinates. While we've explored both parametrization, we've found that the Cartesian (on image plane) sampling of stereoscopic correspondence makes the Cartesian parametrization significantly simpler. For an example of various projection functions in both forms, see figure 1.

3.2. Stereo geometry with projection functions

To parametrize the equations of accuracy and correspondence, we can setup a general equation to define stereoscopic correspondence between two imagers as follows

$$x(u) = x(u - d) + B \quad (1)$$

where $x(u)$ is the projection function from the imager coordinate (u) to the real world (x), z is the distance from the plane of the imagers, B is the baseline, and d is the disparity. This is a two-dimensional parametrization, where the axis perpendicular to the epipolar lines is ignored. With this equation and the projection forms in table 1, it's easy to derive various properties of the stereo system, such as depth error, and mapping from disparity to depth.

3.3. Rectilinear Lens Derivation

For a traditional rectilinear projection, we'd see

$$x(u) = \frac{1}{f} \cdot u \cdot z \quad (2)$$

where f is the effective focal length. This equation can then be solved to generate the relationship between disparity and depth

$$\begin{aligned} \frac{1}{f} \cdot u \cdot z &= \frac{1}{f} \cdot (u - d) \cdot z + B \\ u &= (u - d) + \frac{Bf}{z} \\ z &= \frac{f \cdot B}{d} \end{aligned} \quad (3)$$

Additionally, we can take the derivative, $\frac{\partial z}{\partial d}$, and substitute in our relationship between z and d to generate an error relationship. Traditionally we take the error relative to a single pixel disparity error ($\partial d = 1$).

$$\begin{aligned} z &= \frac{f \cdot B}{d} \\ \frac{\partial z}{\partial d} &= -\frac{f \cdot B}{d^2} \cdot \partial d \\ \frac{\partial z}{\partial d} &= -\frac{f \cdot B}{\left(\frac{f \cdot B}{z}\right)^2} \cdot \partial d \\ \frac{\partial z}{\partial d} &= -\frac{z^2}{f \cdot B} \cdot \partial d \end{aligned} \quad (4)$$

3.4. Fish-eye Lens Derivation

The derivations in the rectilinear section can be repeated to generate results for stereoscopic depth systems with fish-eye lenses. The result for typical f - θ lenses is straightforward, although it's very complicated for orthographic lenses. However, both types of fish-eyes exhibit much less vignetting than rectilinear lenses and support imaging an entire hemisphere, so they're both adequate to address our issues with rectilinear lenses. Results are in table 2.

3.5. Cubic Lens Derivation

In the scope of this project, we also explored a variety of other lens projection functions, to see if we can solve an interesting issue in stereoscopic depth systems, their depth error function. Namely, in a traditional rectilinear system, depth error is quadratic with distance, as seen in equation 4. We looked for optics which used extreme distortion to counteract these properties. In our exploration, we found one, namely cubic projections. These create extremely barrel distorted images in the center of the image, but their depth error is

$$\frac{\partial z}{\partial d} = -\frac{3z^2}{Bf^3} \left(\frac{Bf^3}{z} - u^3 \right)^{\frac{2}{3}} \partial d$$

If we simplify the expression, by removing the dependence on the focal length and baseline, which work in tandem to set a coefficient of proportionality, we can better see that the errors are

$$\frac{\partial z}{\partial d} \propto \begin{cases} z^{\frac{4}{3}}, & \text{if } \frac{Bf}{z} \gg u^3 \\ z^2 \cdot u^2, & \text{if } u^3 \gg \frac{Bf}{z} \end{cases} \quad (5)$$

What this means is that for the center of the image, when u is small, we have sub quadratic depth error. Outside that, we have traditional quadratic depth error, which grows quadratically worse towards the edges. To the best of our knowledge, this is the first known result of a stereo system getting sub-quadratic depth error, even in theory.

Equation	World-to-Image	Image-to-World	Depth	Depth Error
Rectilinear	$u = f \cdot \tan(\theta)$	$z \frac{u}{f}$	$\frac{Bf}{d}$	$-\frac{z^2}{Bf}$
$f\theta$	$u = f \cdot \theta$	$z \tan\left(\frac{u}{f}\right)$	$\frac{B}{\tan\left(\frac{u}{f}\right) + \tan\left(\frac{1}{f}(d-u)\right)}$	$-\frac{1}{Bf} \left(z^2 + \left(B - z \tan\left(\frac{u}{f}\right) \right)^2 \right)$
Cubic	$u = (f \cdot \tan(\theta))^{\frac{1}{3}}$	$z \frac{u^3}{f^3}$	$\frac{Bf^3}{u^3 + (d-u)^3}$	$-\frac{3z^2}{Bf^3} \left(\frac{Bf^3}{z} - u^3 \right)^{\frac{2}{3}}$
Orthographic	$u = f \cdot \sin(\theta)$	$z \frac{u}{\sqrt{f^2 - u^2}}$	Omitted	Omitted

Table 2. Comparison of various projection functions. Due to the fact that spatial accuracy and depth depend on the derivatives of the projection function, the rectilinear projection is the only one with a constant derivative, and hence the only one with no dependence on where on the imager projection is taking place. These results were all verified using a computer algebra system [20]. The orthographic results are very challenging to derive as they’re not defined everywhere, and the results are omitted for clarity, although the included script can provide the results.

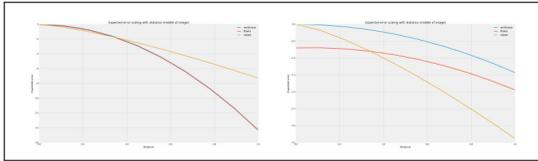


Figure 4. A simulation of the depth errors at given distances with a rectilinear lens, a fish-eye lens, and a cubic lens. The sub-quadratic performance of the cubic imager can be seen, along with the position invariance of the rectilinear lenses.

3.6. Conclusion: Projections of Interest

Of these various lens projections, there are two promising ones. First, the cubic lens projection is interesting because of its guarantee of superior depth accuracy in the middle of the image. However, as results in the next section show, these lenses require extreme magnification in the middle of the frame, and may not be feasible to build. Second, it seems that there are minimal accuracy advantages to using other lens projection. This is demonstrated in figure 4. This suggests that the reasons to use a different lens should come only from the angle-of-view and noise benefits, and not for any accuracy reason.

4. Simulation Results

In general, we will demonstrate the results of this work on anamorphic frames, with non-rectilinear distortion down only one axis. The methods for dealing with each case separately can then be combined for results that work for symmetric lenses. Constraining distortion to one axis allows for a simplification of the analysis. We will briefly discuss what it means to put the anamorphic axis in either direction.

When the axis of nonlinear distortion is perpendicular to the epipolar line, there is no change in the projection function as defined earlier. This allows for predictable depth accuracy. However, it requires that we change traditional

matching algorithms to search along a curved line instead of down a scan-line. We also need to take care to use the correct projection function when projecting a stereo disparity into a real-world coordinate (using the inverse of the projection equation chosen). While this causes variation in spatial resolution of image lens across the image, it is one simply way to create wide-angle stereo imagers. By mounting anamorphic lenses to both cameras, perpendicular to the epipolar line, we can create wide-angle stereoscopic systems which don’t suffer from heavy vignetting or perspective distortion. The limitation of this method is that the resulting images have very asymmetric aspect ratios (being much ”taller” than ”wide”).

When the axis of nonlinear distortion is parallel to the epipolar line, this affects the accuracy and matching algorithms required to generate stereoscopic depth. Matching algorithms then need to be modified to use variable sized windows. Additionally, algorithms must use the correction functions in table 2 to map back to consistent disparities.

4.1. Data source and algorithms

For this set of simulation, we will simply distort high-resolution stereo images from a standard dataset [17]. This allows us to have a ground truth and use the extra resolution from the image to compute arbitrary sampling functions. In addition, we implemented basic stereoscopic correspondence algorithms with sub-pixel resolution from scratch. Most of the results are with sum-of-absolute-differences block matching, but we also tried Sobel + BM [13] and Census [23]. These latter algorithms are important as they’re robust to noise and can be used to validate the best-case impact of photon noise. Additionally, we implement a standard left-right check [10] to remove noisy matches. All algorithms were implemented in C++, with the use of no external libraries (except for image loading), and tested on both Windows and Linux.

The modification of the algorithms is fairly straightforward. The easiest method is simply to adjust the location of

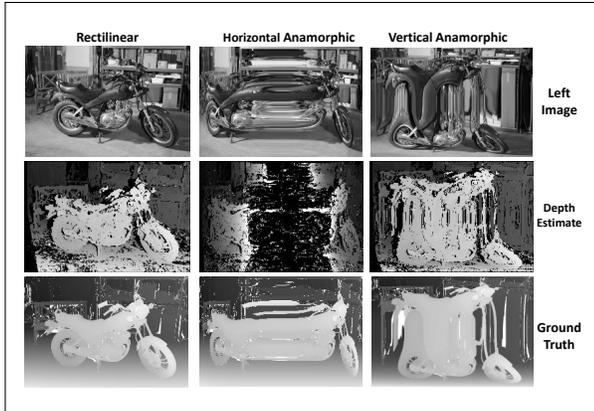


Figure 5. Example results for cubic anamorphic lenses from the Middlebury stereo dataset [17]. These images were distorted at the same time they were sub-sampled, to produce comparable results. The first row are left images from left/right pairs, while the middle row is results from an implemented depth matching algorithm. The last row are ground truth results.

searches based on where they'd be if the image were rectilinear. This can be done with equations from table 1. With the anamorphic assumption, this means that images from X (or horizontal) (or parallel to epipolar lines) distorted images require widening and shrinking the window. In our case, we simply use a fixed sized window on the smaller of the two windows and sampled the larger window according to the difference in projected patch sizes between the two u coordinates. For the vertical anamorphic distortion, one can change the window to keep consistent sampling, but keeping a fixed sized window for the whole image is acceptable in practice. We always used a fixed disparity search size for all the images.

4.2. Cubic Lens Performance

In our simulation of cubic lenses, we tried both anamorphic distortions in the horizontal and vertical axis. An example input image, the depth output, and the ground truth depth are seen in figure 5. As described earlier, the resulting algorithm causes a loss of data in the middle of the X anamorphic image. This is due both to the limited algorithm search window, and the fact that cubic distortion requires a huge magnification in the middle of the image (200 pixel wide image requires 30x more magnification in the center than the corresponding rectilinear) means that the dataset doesn't have any texture and the stereo algorithm fails. This math comes from the projection function gradient, where a rectilinear camera is constant, while a cubic camera is $\nabla u \propto 3 * x^2$.

The vertical anamorphic results work fine, but don't give the interesting depth properties that were discovered in the

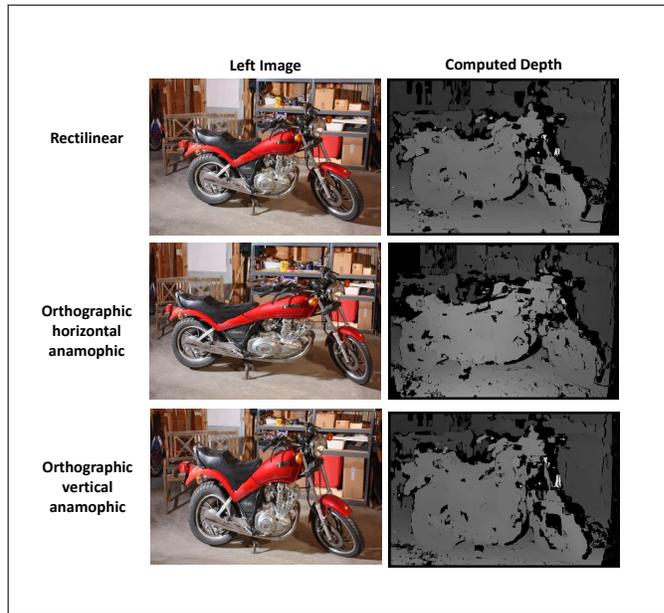


Figure 6. Example results for orthographic anamorphic lenses from the Middlebury stereo dataset [17]. These images were distorted at the same time they were sub-sampled, to produce comparable results.

analysis.

4.3. Orthographic Lens Performance

Since the cubic projection results were theoretically interesting, but the images were so extreme that it became obvious such lenses would be challenging to build, we decided to shift our focus to a more traditional projection function, the orthographic projection. Such lenses are interesting, practical to build [3] (e.g. Nikon 10mm f/5.6 OP), and exhibit no vignetting. That means, if such lenses can be used for stereo, they'd have better SNR performance due to less photon noise in the corners of the image. Our algorithms were validated on such images. The results are shown in figure 6. One caveat to this method is that

4.4. Proposal for a hardware system

The basic next steps for this component of the project would be to actually build a stereo system with either a rectilinear anamorphic lens, which are commonly available on the internet. Or build a stereo matching system with an orthographic fish-eye camera. Both would demonstrate the high fidelity results of a wide angle stereo depth system without any vignetting loss or a requirement for an infinite number of pixels (as would be needed in the traditional rectilinear lens case).

5. Synthetic stereoscopic matching

In order to further validate the contributions of these various lens models, we implemented a customized version of the tungsten C++ path-tracer which could handle arbitrary anamorphic projection functions. We used the well-known Sponza scene as our example data.

5.1. Synthetic GoPro Data

To select a camera design for these synthetic results, we decided to try and match the well-known cameras from GoPro. Specifically, we used a fish-eye lens, with an $f\text{-}\theta$ projection, and a horizontal field of view of 120 degrees. Our images were 2:1 aspect ratio (instead of the typical 1.7:1), and can be seen as slightly cropped versions of the images one would get from a GoPro.

5.2. Rectilinear Depth

As a baseline for these synthetic results, we also included a rectilinear projection with the same field of view and position. To better simulate the impact of a real $f\text{-}\theta$ lens, we also added synthetic vignetting to the rectilinear lens in order to account for its light loss properties. The model for vignetting was $\cos(\theta)^4$ [2], while our model for noise coming from vignetting was simply assuming photon noise [9]. We assumed 5,000 electron capacity wells and sampled the appropriate Poisson distribution at each location to emulate the noise.

5.3. Anamorphic Depth

Figure 7 includes an example of these results. It is visible that our proposed matching algorithms perform stereo matching correctly on these synthetic anamorphic picture, with little to no loss of performance. Although there's some variation in performance across the projections (due to having different image content), it's clear that the proposed methods provide comparable or better performance than a rectilinear projection, if real vignetting (+ photon noise) is taken into account, even using a robust matching function like Census [23] isn't enough to overcome the noise issues introduced. Our quantitative table of performance, normalized to the results of a rectilinear lens (without vignetting), is shown in figure 8.

We had originally planned to report the depth errors resulting from using a rectilinear un-distortion model instead of an anamorphic one. However, we found our generated datasets (which take 4 hours to render all example images) had scaled depth values which made it impossible to compute mean squared error or any other depth-based error metric. There was simply not enough time to fix the error in the tungsten renderer, regenerate these datasets, and compute depth-based error metrics. However, it is planned future work.

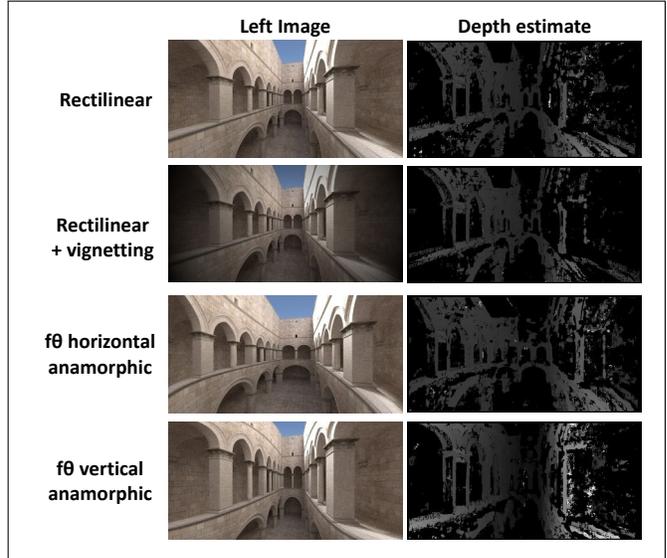


Figure 7. Example results for $f\text{-}\theta$ anamorphic lenses. The example scene is the commonly used Dubrovnik Sponza scene used in the graphics community. The renders were generated from a custom version of the Tungsten C++ ray-tracer.

Despite that, these results suggest that anamorphic software undistortion models would minimize resampling and make it easier to build stereo depth systems with fish-eye lenses.

6. Conclusions

We have demonstrated and derived a generalized form of stereoscopic correspondence, along the requisite generalized sampling algorithm for modifying existing algorithms. Additionally, we have shown synthetic renders supporting our mathematical conclusions, and demonstrated the intuition behind desiring non-linear lenses. This should open the path to making wide angle stereoscopic systems that naturally use the depth accuracy properties of the wide angle lenses they're designed with, instead of being distorted into rectilinear projections.

Of equal interest, we've shown that stereoscopic systems with properly designed optics are not limited to quadratic depth error, and instead have be built with a desired trade-off between spatial and depth invariant accuracy. This interesting mathematical and synthetic result paves the way for many interesting new stereoscopic systems. We'd like to pursue implementing a practical version of this concepts on top of consumer wide-angle cameras, to the aim of better computational efficiency and accuracy.

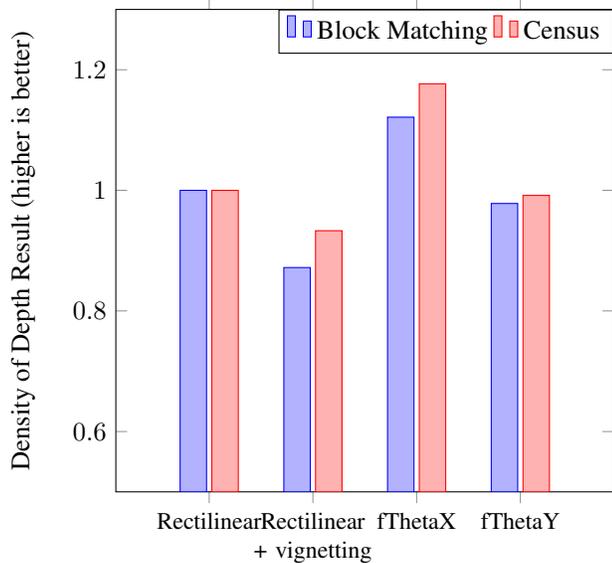


Figure 8. **Accuracy results.** Quantitative results from figure 7. The fThetaX and fThetaY examples are anamorphic distortions ($f\theta$) along a single axis. The metric is simply density of data after passing a standard left-right check [10]. Values are normalized to the scores from a rectilinear projection. There are two stereo cost correlation functions, a standard sum of absolute differences and Census [23]. The results from the proposed anamorphic methods are better than those from Rectilinear stereo, when realistic noise is taken into account.

7. Next Steps

We plan to validate these theoretical and synthetic results with real optics. We'd like to build stereo systems with actual anamorphic lenses to show the benefit of a new optical distortion model. In addition, we can use these results with traditional spherical lenses, in the form of an anamorphic computational model to minimize re-sampling, while continuing to use traditional correlation algorithms. This should make stereo systems with cameras like GoPros more efficient and practical.

References

- [1] S. Abraham and W. Förstner. Fish-eye-stereo calibration and epipolar rectification. *ISPRS J Photogramm*, 59(5):278–288, 2005.
- [2] M. Aggarwal, H. Hua, and N. Ahuja. On cosine-fourth and vignetting effects in real lenses. In *Computer Vision, 2001. ICCV 2001. Proceedings. Eighth IEEE International Conference on*, volume 1, pages 472–479. IEEE, 2001.
- [3] F. Bettonvil. Fisheye lenses. *WGN, Journal of the International Meteor Organization*, 33:9–14, 2005.
- [4] G. Bradski. Opencv library. *Dr. Dobb's Journal of Software Tools*, 2000.

- [5] M. M. Fleck. Perspective projection: the wrong imaging model. *Research report*, (Kingslake 1992):95–01, 1995.
- [6] S. Gehrig. Large-field-of-view stereo for automotive applications. *Omnivis 2005*, pages 1–8, 2005.
- [7] A. Geiger, P. Lenz, and R. Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [8] J. M. Gluckman, S. K. Nayar, and K. J. Thoresz. Real-Time Omnidirectional and Panoramic Stereo. In *DARPA Image Understanding Workshop (IUW)*, pages 299–303. Morgan Kaufmann, 1998.
- [9] S. W. Hasinoff. Photon, poisson noise. In *Computer Vision*, pages 608–610. Springer, 2014.
- [10] X. Hu and P. Mordohai. Evaluation of stereo confidence indoors and outdoors. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1466–1473. IEEE, 2010.
- [11] C. Iemmi and J. Campos. Anamorphic zoom system based on liquid crystal displays. *Journal of the European Optical Society-Rapid publications*, 4, 2009.
- [12] N. Kita. Dense 3D Measurement of the Near Surroundings by Fisheye Stereo. pages 148–151, 2011.
- [13] K. Konolige. Small vision systems: Hardware and implementation. In *Robotics Research*, pages 203–212. Springer, 1998.
- [14] S. Li. Real-Time Spherical Stereo. In *18th International Conference on Pattern Recognition (ICPR'06)*, pages 1046–1049, oct 2006.
- [15] S. Li. Binocular Spherical Stereo. *TITS*, 9(4):589–600, dec 2008.
- [16] S. Roy, J. Meunier, and I. J. Cox. Cylindrical rectification to minimize epipolar distortion. In *Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition (CVPR '97)*, CVPR '97, pages 393–, Washington, DC, USA, 1997. IEEE Computer Society.
- [17] D. Scharstein, H. Hirschmüller, Y. Kitajima, G. Krathwohl, N. Nešić, X. Wang, and P. Westling. High-resolution stereo datasets with subpixel-accurate ground truth. In *Pattern Recognition*, pages 31–42. Springer, 2014.
- [18] S. Shah. Depth Estimation Using Stereo Fish-Eye Lenses. In *Image Processing, 1994. Proceedings.*, volume 2, pages 740–744. IEEE, 2002.
- [19] P. Sturm. Camera Models and Fundamental Concepts Used in Geometric Computer Vision. *Foundations and Trends in Computer Graphics and Vision*, 6(1-2):1–183, 2010.
- [20] SymPy Development Team. *SymPy: Python library for symbolic mathematics*, 2014.
- [21] R. Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2011.
- [22] Wikipedia. Anamorphic format. *Wikipedia, The Free Encyclopedia*, 2016.
- [23] R. Zabih and J. Woodfill. Non-parametric local transforms for computing visual correspondence. In *Computer Vision ECCV'94*, pages 151–158. Springer, 1994.
- [24] Z. Zhang. A flexible new technique for camera calibration. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(11):1330–1334, 2000.