

A Perspective on Distortions *

Rahul Swaminathan, Michael D. Grossberg and Shree K. Nayar

Department of Computer Science, Columbia University
New York, New York 10027

Email: {srahul, mdog, nayar}@cs.columbia.edu

Abstract

A framework for analyzing distortions in non-single viewpoint imaging systems is presented. Such systems possess loci of viewpoints called caustics. In general, perspective (or undistorted) views cannot be computed from images acquired with such systems without knowing scene structure. Views computed without scene structure will exhibit distortions which we call caustic distortions.

We first introduce a taxonomy of distortions based on the geometry of imaging systems. Then, we derive a metric to quantify caustic distortions. We present an algorithm to compute minimally distorted views using simple priors on scene structure. These priors are defined as parameterized primitives such as spheres, planes and cylinders with simple uncertainty models for the parameters. To validate our method, we conducted extensive experiments on rendered and real images. In all cases our method produces nearly undistorted views even though the acquired images were strongly distorted. We also provide an approximation of the above method that warps the entire captured image into a quasi single viewpoint representation that can be used by any “viewer” to compute near-perspective views in real-time.

1 What Are Distortions?

An image is essentially a projection of a three dimensional scene onto a two dimensional surface. The nature of this projection depends on the geometry of the imaging system used to sample the light rays, and the scene structure. For instance, the human eye forms an image on the retina, which is perceived by our brain as a *perspective image*. As a result, straight lines in the scene appear as straight lines in this perspective image. If the imaging system is not perspective, straight lines in the scene appear curved in the image. This deviation from perspective provides the notion of *image-distortions* and is different from *perceptual image distortions* [27].

Distortions are clearly of great importance in vision and graphics. Interestingly, a general treatment of the subject has not yet been conducted. Before analyzing distortions, it behooves us to develop a taxonomy of distortions. Fig. 1 shows a classification of distortions based on the geometry of the imaging system. In particular, we

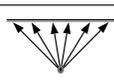
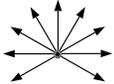
| | Imaging Model | Imaging Optics | Type of Distortion | Information to Undistort |
|---|----------------------|--|--|----------------------------------|
|  | Perspective | Pinhole cameras and lens-based cameras | None | N/A |
|  | Single Viewpoint | Wide-angle cameras, fish-eye cameras, catadioptric cameras and rotating cameras | Radial, tangential distortions and image warps | Camera model |
|  | Non-Single Viewpoint | Wide-angle cameras, fish-eye cameras, catadioptric cameras, camera clusters and moving cameras (mosaics) | Caustic distortions | Camera model and scene structure |

Figure 1: Taxonomy of Distortions: We classify distortions based on the geometry of the imaging model into three main categories: (a) perspective, (b) single-viewpoint, and (c) non-single viewpoint. The table lists typical realizations of each imaging model, their inherent distortions and the information needed to undo the distortions.

consider: perspective, single viewpoint and non-single viewpoint imaging systems. In perspective imaging systems, all the incoming light-rays intersect at a single point before being imaged on a planar detector (CCD or film). Conventional cameras typically adhere to the this imaging model and therefore produce no distortions.

A drawback of perspective systems is their limited field of view. In order to capture a wider field of view, cameras have been designed to deviate from perspective projection, while still maintaining a single viewpoint. Lens based (*dioptric*) non-perspective cameras include wide-angle and fish-eye¹ cameras [12]. Fig. 2(a) shows an image acquired with such a wide-angle camera, exhibiting *radial* and *tangential* distortions[1]. However, because the camera maintains a single viewpoint, a simple image warp can undistort the acquired images [2, 23, 20]. Single viewpoint *catadioptric* cameras (employing lenses and mirrors) have also been suggested [14, 17, 26] that provide omnidirectional (hemispherical or panoramic) fields of view. Fig. 2(b), shows an image acquired with such a camera. The image appears distorted but can be undistorted into perspective views using only an image warp [5, 17].

*This work was supported by the DARPA HumanID Program (Contract No. N00014-00-1-0929) and an NSF ITR Award IIS-00-85864

¹Such lenses actually have a locus of viewpoints called a *diacaustic* which is typically small and may be approximated by a point.

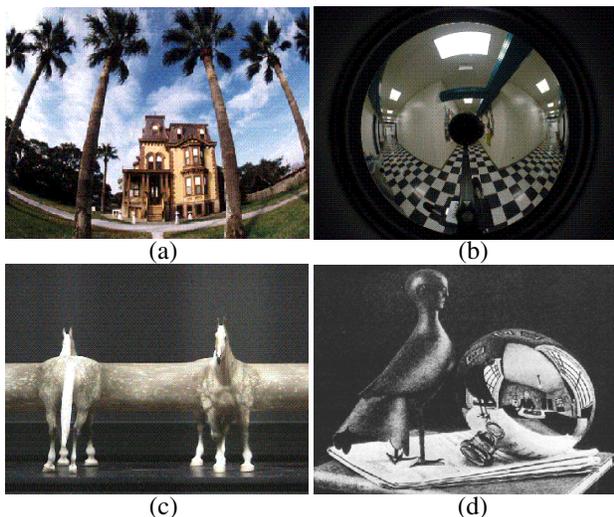


Figure 2: Examples of distorted images. Single viewpoint cameras (row 2 of Fig. 1) produce radial and tangential distortions as seen in (a) an image acquired with a wide-angle camera and (b) an image acquired with a single viewpoint catadioptric camera. In both cases, an image warp can be used to undistort the images into a one or more perspective views. Non-single viewpoint cameras (row 3 of Fig. 1) produce distortions we call *caustic distortions* as seen in (c) a cyclograph of a horse (adapted from [18]) and (d) M.C Escher's rendering of reflections on a spherical object acting as a non-single viewpoint catadioptric camera.

Non-single viewpoint imaging systems, shown in row three of Fig. 1, possess a locus of viewpoints called a *caustic* [21]. We refer to an image acquired with such a sensor as a *Multi-Viewpoint Image*²(MVI). Figs. 2 (c,d) show examples of MVIs. Due to the multiple viewpoints, perspective views cannot be computed from MVIs unless scene structure is known. Any view computed without using scene structure is guaranteed to exhibit distortions, which we refer to as *caustic distortions*.

Recently, there has been growing interest in designing imaging systems that possess specific resolution characteristics. Examples of such systems include equi-areal [10] and equi-resolution [7] cameras, amongst others [25, 3, 6, 9]. These sensors do not maintain a single viewpoint. Also, traditional mosaics computed either from camera motion [19, 16, 15] or from camera clusters [13, 11, 22], are another type of multi-viewpoint images. Furthermore, MVIs are often used as an efficient and compact representation for stereo [18] as well as image-based rendering [8]. Almost always, MVIs are used to create views of the scenes they represent. It is therefore highly desirable that these views have minimal distortion. Clearly, the functionality of MVIs would be greatly enhanced by adding the capability to compute such undistorted views.

²Such images are also referred to as multi-perspective images[24] in the graphics and vision literature.

Towards this goal, our paper makes the following key contributions:

- A metric to quantify distortions in a view is presented. Distortions are measured with respect to a perspective view that is computed assuming knowledge of scene structure. Note that this assumption is made for the sole purpose of deriving the metric.
- We develop a method to compute minimally distorted views from an MVI using simple scene priors defined by primitives (spheres, planes, or cylinders) with simple uncertainty models for the parameters.
- We also present an approximation to the above method, which morphs an entire MVI into a quasi-single viewpoint representation. This facilitates fast near-perspective view synthesis.

We demonstrate the effectiveness of our techniques using a variety of synthetic and real multi-viewpoint images. The synthetic images were rendered by modeling non-single viewpoint catadioptric sensors with parabolic, spherical and elliptical reflectors. Experiments were also conducted on MVIs, such as pushbroom, center-strip and concentric mosaics, of real scenes acquired with a moving camera. In all cases, the views computed using our method appear nearly undistorted.

2 Caustic Distortions

Consider an algorithm that computes views from a multi-viewpoint image without knowing the true scene structure. We refer to this algorithm as a *view creation map*. Since scene structure is not known, the computed view is guaranteed to be distorted. We now present a metric to quantify distortions in such views. Note that, to formulate the distortion metric, we assume knowledge of scene structure. Later we will show how to eliminate this assumption.

2.1 Quantifying Caustic Distortions.

Let I_a be an MVI of a scene acquired with a non-single viewpoint imaging system. Let I_v be a view computed from a region within I_a using the view creation map \mathcal{M} . As shown in Fig. 3, a point \mathbf{q}_a in the MVI is mapped by \mathcal{M} to a point \mathbf{q}_v in the computed view I_v :

$$\mathbf{q}_v = \mathcal{M}(\mathbf{q}_a). \quad (1)$$

Given scene structure \mathcal{D} , we can project the point \mathbf{q}_a in the MVI onto the scene point \mathbf{Q}_a as:

$$\mathbf{Q}_a = \mathbf{S}_c(\mathbf{q}_a) + \mathcal{D}(\mathbf{q}_a) \cdot \mathbf{V}_r(\mathbf{q}_a), \quad (2)$$

where, $\mathbf{S}_c(\mathbf{q}_a)$ and $\mathbf{V}_r(\mathbf{q}_a)$ denote the viewpoint and viewing direction, and $\mathcal{D}(\mathbf{q}_a)$ denotes the depth of the scene point.

Based on our taxonomy in Fig. 1, a view is considered undistorted if it adheres to perspective projection. Therefore, we insert a hypothetical camera in this scene and acquire a perspective image I_π (see Fig. 3). If we let the hypothetical camera parameters be \mathbf{K} (intrinsic) and $\{\mathbf{R}, \mathbf{t}\}$ (pose), the perspective projection of the scene point \mathbf{Q}_a in I_π can be written as:

$$\mathbf{q}_\pi = \Pi(\mathbf{q}_a, \mathbf{S}_c, \mathbf{V}_r, \mathcal{D}, \mathbf{K}, \mathbf{R}, \mathbf{t}) \quad (3)$$

The disparity between the computed view I_v and the hypothetical perspective view I_π of the scene point \mathbf{Q}_a is $|\mathbf{q}_\pi - \mathcal{M}(\mathbf{q}_a)|^2$.

Over the entire field of view ($\mathbf{q}_a \in FOV$) of the computed view, the total disparity is given by:

$$\delta(\mathcal{M}, \mathcal{D}) = \sum_{FOV} \|\mathbf{q}_\pi - \mathcal{M}(\mathbf{q}_a)\|_2. \quad (4)$$

Note that this disparity depends not only on the view creation map \mathcal{M} but also on the viewpoint locus ($\mathbf{S}_c, \mathbf{V}_r$), pose (\mathbf{R}, \mathbf{t}), and intrinsic parameters \mathbf{K} of the hypothetical perspective camera.

We define caustic distortions as the total disparity in projections of scene points between I_v and the most similar perspective view. In other words, caustic distortions are defined as the minimum of Eq.(4) over all possible parameters ($\mathbf{R}, \mathbf{t}, \mathbf{K}$) of the hypothetical camera:

$$\Delta(\mathcal{M}, \mathcal{D}) = \min_{\{\mathbf{K}, \mathbf{R}, \mathbf{t}\}} (\delta(\mathcal{M}, \mathcal{D})). \quad (5)$$

We now have a simple metric to measure caustic distortions in a view computed using any view creation map \mathcal{M} given scene structure \mathcal{D} . Note that, knowledge of scene structure was assumed purely for purposes of defining the metric. As we shall show in section 3, the metric can be used without scene structure as well and serves as a valuable tool to evaluate different view creation maps.

2.2 Estimating Distortions: A Linear Solution

Using Eq. 5 to estimate caustic distortions involves finding the hypothetical camera parameters ($\mathbf{R}, \mathbf{t}, \mathbf{K}$) that minimize the disparity between I_v and I_π . Ordinarily, this would require a laborious search over the parameter space. Here, we present an alternative approach that results in a simple linear solution.

Instead of explicitly parameterizing the camera parameters, we define them as a 3×4 projective matrix \mathbf{P} . Now, given a view creation map \mathcal{M} we can compute the location of the point \mathbf{q}_a in the computed view \mathbf{q}_v using Eq. 1. Also, given scene structure \mathcal{D} , Eq.(2) yields the three-dimensional location \mathbf{Q}_a of the point \mathbf{q}_a . The hypothetical camera parameters \mathbf{P} that yield the closest perspective view to I_v are constrained as:

$$\tilde{\mathbf{q}}_v = \mathbf{P} \cdot \tilde{\mathbf{Q}}_a, \quad (6)$$

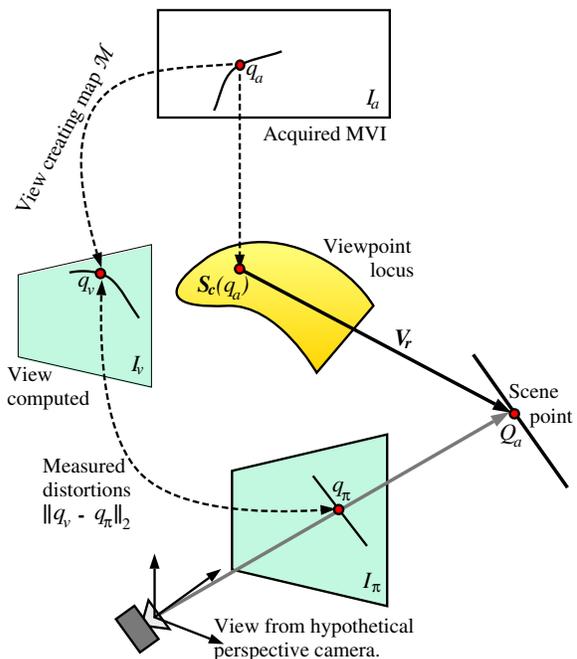


Figure 3: Measuring caustic distortions: A region within the MVI I_a is mapped into a view I_v using a view creation map \mathcal{M} . Each point \mathbf{q}_a in I_a maps to the point \mathbf{q}_v in I_v . Given scene structure \mathcal{D} , we can project \mathbf{q}_a onto the scene point \mathbf{Q}_a . A perspective view I_π can then be synthesized by a hypothetical camera. Distortions are defined as the projection errors between the perspective view I_π and the computed view I_v .

where, $\tilde{\mathbf{q}}_v$ and $\tilde{\mathbf{Q}}_a$ are the homogeneous point coordinates, such that: $\tilde{\mathbf{q}}_v = [\mathbf{q}_v \ 1]^T$ and $\tilde{\mathbf{Q}}_a = [\mathbf{Q}_a \ 1]^T$.

A minimum of 11 image points are required to solve for \mathbf{P} , since it is defined up to a scale factor. In practice, \mathbf{P} is found from an over-determined system of equations obtained by sampling several (> 11) points within the desired FOV in I_a . The residue obtained by substituting \mathbf{P} into Eq. 6 quantifies caustic distortions.

This linear solution is very fast to compute, but not always robust. When high accuracy is desired, it is preferable to estimate the camera parameters ($\mathbf{R}, \mathbf{t}, \mathbf{K}$) using non-linear search techniques.

3 Reducing Distortions in Views

To compute undistorted (perspective) views for regions within an MVI we need to know scene structure. When scene structure is not known, the view creation map must make assumptions about the scene structure. Most previous approaches model the scene as being distant (at infinity)[9] and use only light-ray directions to compute views. We refer to this mapping as the *infinity-map*. An exception being the column scaling method proposed by in [19] for the special case of concentric mosaics.

For systems with relatively large viewpoint loci, the infinity-map causes severe distortions, warranting bet-

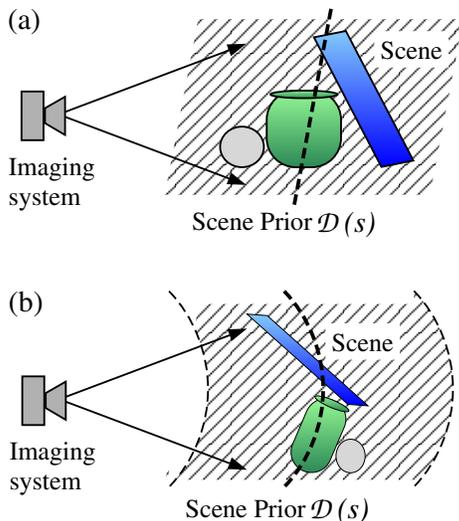


Figure 4: A non-single viewpoint imaging system is used to capture an arbitrary scene. The scene structure is not known and is modeled using simple priors. (a) The prior is defined as a plane $\mathcal{D}(s)$ at depth s . The uncertainty in depth s is modeled with a simple probability distribution function $\mathcal{P}(s)$. (b) The prior used in this scene is a sphere $\mathcal{D}(s)$ of radius s . Again, the uncertainty in the radius is modeled probabilistically by $\mathcal{P}(s)$.

ter models of scene structure. We model the scene using simple priors in the form of parameterized primitives $\mathcal{D}(s)$ such as *planes*, *spheres*, or *cylinders* and simple probability distribution functions $\mathcal{P}(s)$ associated with the parameters of the primitives. For example, the scenes shown in Figs. 4(a,b) are modeled by a parameterized plane and sphere, respectively. The parameters of these primitives can be modeled with simple distributions such as *uniform* or *gaussian*. The priors are similar to *imposters* [8, 4] commonly used in graphics, but differ in that no fixed primitive is specified. Furthermore as will be shown in Section 5, the prior need not be a parameterized shape.

In general, view synthesis is achieved by: (1) projecting a region within the MVI onto a chosen primitive and then (2) synthesizing a view from a virtual perspective camera. If the primitive used to project the scene onto is $\mathcal{D}(s)$, the view creation map corresponding to this view is denoted by \mathcal{M}_s . Given a family of primitives, we must find the *optimal* primitive parameter s^* which yields the least distortion. Using the distortion metric in Eq. 5, s^* is given by:

$$s^* = \operatorname{argmin}_s \left(\int \Delta(\mathcal{M}_s, \mathcal{D}(k)) \cdot \mathcal{P}(k) dk \right). (7)$$

Estimation of the optimal view creation map parameters can be posed as a non-linear optimization problem. The optimal view computed depends on the desired viewing direction and *FOV* of the view. This makes the method very computationally expensive for real-time ap-

lications. We therefore present an approximation of this approach in section 5. However, first we will present results of applying the above method to various MVIs.

4 Experimental Results

The MVIs used in our experiments include rendered scenes as well as real image mosaics.

4.1 Synthetic Scenes

The synthetic scenes were rendered by modeling a non-single viewpoint catadioptric sensor within a cubic enclosure of dimensions $5 \times 5 \times 5 \text{cm}^3$, whose walls have a checker-board pattern. In Figs. 5(a,b,c) we show the viewpoint loci for the simulated catadioptric cameras with: (1) a spherical reflector and telecentric lens, (2) an ellipsoidal reflector and telecentric lens, and (3) a parabolic reflector and perspective lens. The corresponding rendered MVIs are shown in Figs. 5(d,e,f).

Figs. 5(g,h,i) illustrate the severe caustic distortions inherent in views computed using the infinity-map. In contrast, Figs. 5(j,k,l) shows views computed using our method exhibiting negligible distortions. For comparison, we also present the ground truth views synthesized using perspective projection of the scene in Figs. 5(m,n,o).

As specified in the taxonomy of Fig. 1, distortion correction requires the viewpoint locus and scene structure. The viewpoint locus was computed using the model proposed in [21]. The scene prior used was a sphere with uniform distribution on its radius within the interval [1.0, 5.0]cm. The optimal sphere radii for the spherical, ellipsoidal and parabolic reflector based experiments were estimated to be 2.0, 2.5 and 2.2cm, respectively.

4.2 Real Panoramas

We also conducted extensive experimentation on real MVIs such as a pushbroom, center-strip, and concentric-mosaics acquired using camera motion. The camera trajectory was precisely controlled using a robotic arm. The scene prior was defined as a plane whose depth was assumed to be uniformly distributed between two planes. Figs. 6(a,d,g) show both the camera trajectory as well as the scene prior used in each experiment.

Pushbroom Mosaic: Fig. 6 (b) shows the view computed using the infinity-map for a pushbroom mosaic acquired under camera motion. Again, the view computed using our method (shown in Fig. 6(c)) exhibits almost no distortions. The optimum primitive (plane) depth assuming uniform distribution between 5 and 20cm was estimated to be 13cm.

Centre-strip Mosaic: Figure 6(e) shows a view computed using the infinity map from a centre-strip mosaic. In contrast, the view shown in Fig. 6(f), computed using our method, is nearly undistorted. The optimal primitive

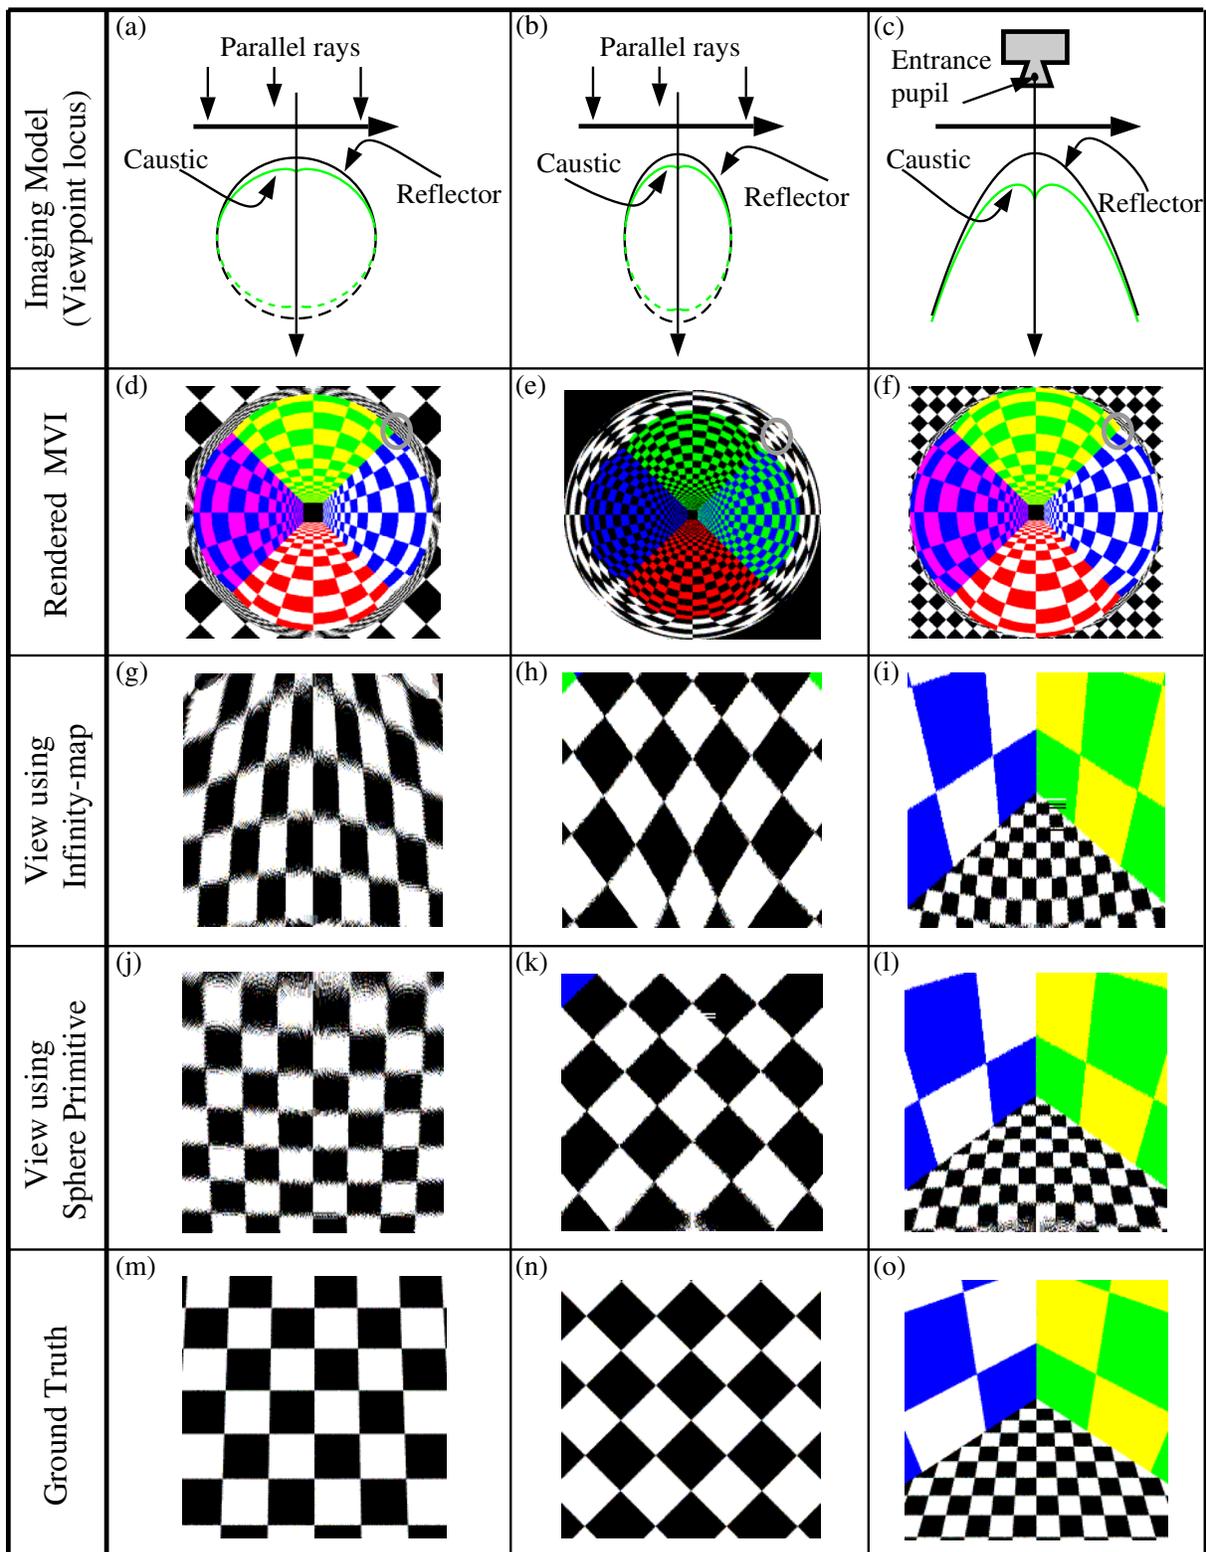


Figure 5: Viewpoint loci for catadioptric imaging systems consisting of : (a) a spherical reflector and telecentric lens, (b) an elliptic reflector and telecentric lens and (c) a parabolic reflector and perspective lens. (d,e,f) Images rendered for these imaging geometries of a cubic enclosure of size $5 \times 5 \times 5 \text{cm}^3$ having a checker board pattern on its walls. (g,h,i) Views computed for the region marked in (d,e,f) using the infinity-map. (j,k,l) Near-perspective views computed using our method. The prior used was a sphere and its radius is assumed to be uniformly distributed between 1 and 5cm. In all cases, the computed views are nearly perspective and exhibit virtually no distortion. Note that the aliasing effects noticeable in some views is due to inaccurate resampling of the MVI and not an artifact of our method.

| | Imaging Model (Caustic) | Acquired MVI | Computed View |
|---------------------|----------------------------|--------------|---------------|
| Pushbroom Mosaic | <p>(a)</p> | <p>(b)</p> | <p>(c)</p> |
| Centre-strip Mosaic | <p>(d)</p> | <p>(e)</p> | <p>(f)</p> |
| Concentric Mosaic | <p>(g)</p> | <p>(h)</p> | <p>(i)</p> |

Figure 6: (a,d,g) Viewpoint loci for MVIs (mosaics) computed for a pushbroom mosaic, center-strip mosaic and concentric mosaic respectively. (b,e,h) Views computed using the infinity-map from the mosaics clearly exhibiting strong caustic distortions. (c,f,i) Near-perspective views computed using the proposed method using a plane for the scene primitive. The shaded region in (a,d,g) represents the space of probable plane depths. Notice how the highly curved (distorted) regions in (b,e,h) have been undistorted to nearly straight lines in (c,f,i).

(plane) depth was estimated to be 1.3 assuming uniform distribution of depths within the interval [1.0, 2.0]cm.

Concentric Mosaic: The third experiment involves a concentric mosaic. Figure 6(h) shows the view computed using the infinity-map, exhibiting caustic distortions. In contrast, Fig. 6(i) was computed using our method and exhibits almost no distortions. The optimal plane primitive was estimated at a depth of 2cm assuming uniform

distribution within the interval [1.5, 3.5]cm.

We have used very simple primitives, such as planes and spheres, to represent otherwise complex scenes. In spite of using such simple primitives, the computed views appear almost perspective (undistorted).

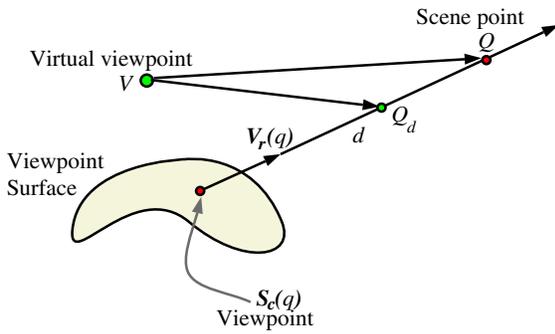


Figure 7: A point $S_c(q)$ on the viewpoint surface observes the scene point Q along the ray direction $V_r(q)$. If the scene point were assumed to lie at depth d , the resulting angular distortion from the virtual viewpoint V is $\angle QVQ_d$.

5 Real-Time Undistorted View Synthesis

The technique presented in section 3 to compute minimally distorted views is optimal but depends on the desired viewing direction and *FOV*. Therefore, every new view computed with a different *FOV* or viewing direction requires estimating the parameters of the view creation map. This is computationally expensive and precludes usage in real-time applications. It is therefore desirable to transform the entire MVI into a single new representation, from which undistorted views can be created in real-time.

5.1 Angle-based Distortion Metric

The distortion metric in section 2 measures distortions in the image plane. This made the metric dependent on the viewing direction and *FOV*. Instead, we may approximate distortion as the angle subtended by the disparity at the virtual viewpoint. This angle-based distortion metric is invariant to the viewing direction (orientation of the imaging plane).

Figure 7 shows a viewpoint surface of a non-single viewpoint imaging system. A scene point Q is visible from the viewpoint point $S_c(q)$ along the viewing direction $V_r(q)$. The angular distortion with respect to the point Q_d at distance d from a virtual viewpoint V is given by:

$$\angle QVQ_d = \arccos(\overrightarrow{VQ} \cdot \overrightarrow{VQ_d}).$$

Similar to section 3, we model the scene using priors defined as a probability distribution $\mathcal{P}(s)$ over scene depth s . The effective angle-based distortion assuming a scene point Q_d to lie at depth d is:

$$\xi(d) = \int \angle Q_s V Q_d P(s). \quad (8)$$

5.2 The Global View Map

We wish to estimate a single global transformation for the MVI so as to warp it into a nearly-single viewpoint



(a)



(b)



(c)

Figure 8: (a) A concentric mosaic constructed by concatenating the center columns of images acquired by moving a camera along a circular trajectory shown in Fig. 6(g). (b) A view computed using the infinity-map from the MVI in (a) clearly exhibiting *caustic distortions*. (c) A near-perspective view computed using the approximate method exhibits significantly lower distortions. The view was computed by estimating the optimal depth at each point in the mosaic. The prior used (two bounding planes) restricted the depths to lie uniformly between 1.5 and 3.0 cm.

representation. At every point in the MVI we wish to find a transformation that minimizes the distortion given by Eq. 8. Therefore, we may define the transformation in terms of the optimal depth d^* at every scene point that minimizes Eq. 8 as:

$$d^* = \operatorname{argmin}_d (\xi(d)). \quad (9)$$

Using Eq. 9 a single *global view map* such as a spherical panorama [5] can be computed. Once this is done, perspective views can be computed in real-time.

5.3 Experimental Verification

We now present results of applying the angle-based distortion metric to the concentric mosaic shown in Fig. 8(a). Figure 8(b) shows the caustic distortions visible in the view computed using the infinity-map. In contrast, Fig. 8(c) was computed using the *global view map* technique and appears nearly undistorted.

In this experiment, the scene prior was modeled as a plane with depth bounded uniformly between 1.5 and 3.0cm from the center of the viewpoint locus. Using Eq. 9, the optimal depth at every point in the MVI was estimated. Since the image is very large, we estimated the optimal depths for a coarse grid of points in the MVI and then interpolated these depths across the MVI. The resulting depth map has no parameterized form and is used as imposters are commonly used in graphics applications.

This technique, although approximate, provides a simple method to compute a quasi-single viewpoint transformation for the entire MVI. The transformation needs to be applied only once and facilitates real-time view creation.

Summary

In this paper, we introduced a taxonomy of distortions based on the geometry of the underlying imaging system. Using this taxonomy, we defined the notion of caustic distortions in views computed from multi-viewpoint images (MVI). In general, perspective views cannot be computed from MVIs unless scene structure is known. Any view computed without knowing scene structure will exhibit caustic distortions.

We derived a metric to quantify distortions in views computed from an MVI, given a view creation algorithm. Using this metric we presented a method to compute views with minimal distortions. To do so, we assumed scene priors defined in terms of very simple parametric primitives such as spheres, planes and cylinders, with probability distributions for their parameters. We demonstrated the effectiveness of our approach on several rendered and real MVIs. In all cases, our method computed views that appear nearly perspective (undistorted), in spite of the MVIs being extremely distorted.

We also presented an approximate approach that warps the entire MVI into a quasi-single viewpoint representation. This representation makes it possible to reduce distortions globally in an MVI and use the resulting image with any off-the-shelf “viewer” to render near-perspective views in real-time.

References

- [1] M. Born and E. Wolf. *Principles of Optics*. Pergamon Press, 1965.
- [2] D.C. Brown. Close range camera calibration. *Photogrammetric Engineering*, 37(8):855–866, August 1971.
- [3] J. Chahl and M. Srinivasan. Reflective surfaces for panoramic imaging. *Applied Optics*, 36(31):8275–8285, 1997.
- [4] J. Chai, X. Tong, S. Chan, and H. Shum. Plenoptic sampling. In *SIGGRAPH*, pages 307–318, 2000.
- [5] S.E. Chen. QuickTime VR — An Image-Based Approach to Virtual Environment Navigation. *Computer Graphics*, 29:29–38, 1995.
- [6] S. Derrien. and K. Konolige. Approximating a single viewpoint in panoramic imaging devices. In *International*

- Conference on Robotics and Automation*, pages 3932–3939, 2000.
- [7] J Gaspar, C. Decco, J. Okamoto Jr, and J. Santos-Victor. Constant resolution omnidirectional cameras. In *Third Workshop on Omnidirectional Vision*, pages 13–18, 2002.
- [8] S.J. Gortler, R. Grzeszczuk, R. Szeliski, and M.F. Cohen. The Lumigraph. In *SIGGRAPH 96*, pages 43–54, 1996.
- [9] R.A. Hicks and R. Bajcsy. Catadioptric Sensors that Approximate Wide-Angle Perspective Projections. In *Proc. CVPR*, pages I:545–551, 2000.
- [10] R.A Hicks and R.K Perline. Equi-areal catadioptric sensors. In *Third Workshop on Omnidirectional Vision*, pages 13–18, 2002.
- [11] A. Majumder, W.B. Seales, M. Gopi, and H. Fuchs. Immersive teleconferencing: a new algorithm to generate seamless panoramic video imagery. In *ACM Multimedia (1)*, pages 169–178, 1999.
- [12] K. Miyamoto. Fish Eye Lens, Fish Eye Lens, *JOSA*, 54(8):1060-1061, 1994.
- [13] V. Nalwa. A True Omnidirectional Viewer. Technical report, Bell Laboratories, Holmdel, NJ 07733, U.S.A., February 1996.
- [14] S. K. Nayar. Catadioptric Omnidirectional Cameras. In *Proc. CVPR*, pages 482–488, 1997.
- [15] S. K. Nayar and A. D. Karmarkar. 360 x 360 Mosaics. In *Proc. CVPR*, pages I:388–395, 2000.
- [16] S. Peleg, B. Rousso, A. Rav-Acha, and A. Zomet. Mosaicing on adaptive manifolds. *PAMI*, 22(10):1144–1154, October 2000.
- [17] V. N. Peri and S. K. Nayar. Generation of Perspective and Panoramic Video from Omnidirectional Video. *DARPA-IUW*, pages I:243–245, December 1997.
- [18] Seitz S and J. Kim. The Space of All Stereo Images. In *Proc. ICCV*, pages I:26–33, July 2001.
- [19] H. Shum and L. He. Rendering with concentric mosaics. *Computer Graphics*, 33:299–306, 1999.
- [20] G.P. Stein. Lens distortion calibration using point correspondences. In *Proc. CVPR*, pages 143–148, June 1997.
- [21] R. Swaminathan, M. D. Grossberg, and S. K. Nayar. Caustics of Catadioptric Cameras. In *Proc. ICCV*, pages II:2–9, 2001.
- [22] R. Swaminathan and S.K. Nayar. Nonmetric calibration of wide-angle lenses and polycameras. *PAMI*, 22(10):1172–1178, October 2000.
- [23] R. Y. Tsai. A versatile camera calibration technique for high-accuracy 3d machine vision. *International Journal of Robotics and Automation*, 3(4):323–344, August 1987.
- [24] D.N. Wood, A. Finkelstein, J.F. Hughes, C.E. Thayer, and D.H. Salesin. Multiperspective Panoramas for Cel Animation. *Computer Graphics*, 31:243–250, 1997.
- [25] Y. Yagi and M. Yachida. Real-Time Generation of Environmental Map and Obstacle Avoidance Using Omnidirectional Image Sensor with Conic Mirror. In *Proc. CVPR*, pages 160–165, 1991.
- [26] K. Yamazawa, Y. Yagi, and M. Yachida. Omnidirectional imaging with hyperboloidal projection. In *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1029–1034, July 1993.
- [27] D. Zorin and A. H. Barr. Correction of geometric perceptual distortions in pictures. *Computer Graphics*, 29:257–264, 1995.