

Generating Mosaics with Minimum Distortions*

Doron Feldman

Assaf Zomet[†]

*School of Computer Science and Engineering
The Hebrew University of Jerusalem
91904 Jerusalem, Israel*

doronf@cs.huji.ac.il zomet@cs.columbia.edu

Abstract

Manifold mosaicing is a fast and robust way to summarize video sequences captured by a moving camera. It is also useful for rendering compelling 3D visualizations from a video without estimating the 3D structure of the scene. However, since the result mosaics are not perspective images, their geometry is inherently distorted. These mosaics are commonly referred to as multi-perspective images, or multi-perspective mosaics.

In this paper we address the following question: Given a video captured by a moving camera, what is the best multi-perspective mosaic that can be generated from it? What is the mosaic with the best combination of large field-of-view and minimal geometric distortions? We define the necessary conditions for a good mosaic and a quantitative criterion for the geometric distortions, and derive analytically the optimal mosaic under this criterion. Results on video sequences confirm that indeed the optimal mosaic has significantly better quality than those generated by other techniques.

1. Introduction

Manifold mosaicing [8, 9] is a robust way to generate a panoramic image from a video sequence captured by a moving camera. The technique is simple, implemented by pasting thin strips from the video into the panoramic mosaic image. Since the camera is moving, every strip in the mosaic image is captured from a different viewing position. Therefore such a mosaic is called a *multi-perspective* [13, 15] image.

Multi-perspective images have been used for creating compact representations of video sequences [9, 16], for 3D reconstruction [13, 17], photogrammetry [2] and are recently

*This research was supported (in part) by the EU under the Presence Initiative through contract IST-2001-39184 BENOGO, and by the Israel Science Foundation grant no. 191/02.

[†]Current address: Computer Science Department, Columbia University, 500 West 120th Street, New York, NY 10027



(a)



(b)

Figure 1: **Typical distortions in multi-perspective mosaics. Image (a) is a part of a multi-perspective mosaic, and image (b) is one of the input images. Note the distortion of the building and the trees.**

becoming popular for 3D visualization and image based rendering [7, 10–12, 15, 18].

Figure 1 shows an example of a multi-perspective image (generated by the method in [8]). There is an apparent non-uniform distortion of scene objects, depending on their depth. The building and the bushes are stretched horizontally, while the cars are shrunk. Such distortions cannot be cancelled by any generic depth-independent transformation of the image.

In this paper we address the following question: Given a sequence taken by a calibrated camera moving on a known trajectory, and an *unknown scene*, what is the best multi-perspective image that can be generated from it? In other words, which strips should be copied from the images and

how should they be pasted into the mosaic, such that the result image will contain the maximal amount of visual information and minimal geometrical distortions?

We define the necessary conditions for a good multi-perspective mosaic and a criterion quantifying the geometric distortions, and derive the least-distorted mosaic under this criterion. The criterion is justified theoretically as well as empirically. It turns out that the mosaic with the minimal distortion also has the maximal field-of-view.

One may wonder why use multi-perspective mosaics at all, since they are distorted. An alternative to this would be to create a perspective panorama, by warping all input images into a common coordinate frame. The choice of multi-perspective mosaics comes from practical reasons. Multi-perspective mosaics are generated by a simple, fast and remarkably robust algorithm. On the other hand, in order to generate a perspective panorama, a depth representation of the scene should be available. The estimation of such a representation from a video sequence is computationally demanding and highly ill-posed (due to occlusions, reflections, transparencies etc.). By minimizing the distortions of multi-perspective images, it is possible to generate a visually satisfying image with minor geometric distortions. These distortions may in many cases be practically negligible, especially in comparison with artifacts in perspective panoramas due to errors in depth estimation.

The paper is organized as follows: In Section 2 we present a generalized formulation of manifold mosaicing, by allowing an arbitrary sampling of vertical strips from the images. Each sampling function defines a projection geometry relating a 3D point to a location in the mosaic image. In Section 3 we define the necessary conditions for a good mosaic, namely a large field of view and a continuous unique 3D to 2D projection. In Section 4 we define a quality criterion measuring the distortion in the image, and derive the optimal mosaic under this criterion. Examples show that mosaics generated by the optimal sampling are significantly less distorted than the ones generated by other sampling methods.

1.1. Manifold Mosaics

Manifold mosaicing [8, 16] was introduced for the purpose of generating a panoramic representation of a video. Vertical straight strips were copied from the centers of the images, and pasted into the multi-perspective mosaic. For a rectified camera moving sideways, the projection geometry of the mosaic can be described by the *linear pushbroom* projection model [2]: Parallel projection in the camera motion direction, and perspective projection in the orthogonal direction.

Manifold mosaicing was also used for 3D visualization by selecting different vertical strips for different views. A stereo pair of manifold mosaics, for example, can be gen-

erated by collecting different strips for the left-eye mosaic and for the right-eye mosaic. In [4, 7], one location was used for all the strips of the left-eye mosaic, and another location for all the strips of the right-eye mosaic. A different scheme for 3D visualization was proposed in [12, 18]. By a simple adjustment of the strip sampling function, the authors have generated realistic walkthrough sequences of multi-perspective images.

This work was inspired by [5, 11], which discuss all camera trajectories for generating stereo multi-perspective mosaics. We broaden this discussion, as we argue that the set of stereo mosaics should also be characterized by the strip sampling scheme.

There has been some work on manifold mosaicing with curved strips [9]. These are especially useful when the camera translation has a forward component. In this work, in order to simplify the analysis, we assume that the image plane is parallel to motion direction of the camera and limit the sampling accordingly to vertical lines.

2. Sampling Functions for Mosaicing

Consider a perspective video camera moving continuously on a curved segment with its image plane orientation tangent to the curve. Assume w.l.o.g. that the segment length is 1, and let $t \in [0, 1]$ be a parameter describing the location of the camera in the segment. A multi-perspective mosaic is generated by selecting a vertical line in each frame $I(t)$ according to a *sampling function* $\phi(t)$, and pasting it into the mosaic. $\phi(t)$ denotes the location of the line sampled from frame $I(t)$. Let $\pi(t)$ be the plane joining the camera center of projection at location t to the sampled line $\phi(t)$. The pasting location in the mosaic is defined by the intersection of $\pi(t)$ with the mosaic manifold. In case the camera moves on a linear trajectory, this manifold is a plane. Otherwise, the manifold is determined by the camera trajectory. As in [9], we set the distance of the manifold to be equal to the camera's focal length, in order to maintain the vertical resolution of the image.

It is assumed that the camera motion and internal calibration are known, or were estimated from the video sequences (For a review, see [3]), and that the horizontal field-of-view angle of the camera is θ .

We first analyze mosaics generated from linear camera trajectories, and find the optimal mosaic analytically. Two useful examples of sampling functions for linear camera trajectories, depicted in Figure 2, are the linear sampling function $\phi(t) = \alpha t + \beta$ and a special case of it, the constant sampling function $\phi(t) = \beta$ (where $\alpha = 0$). It was shown in [18] that in the former case, all rays pass through a vertical line in the plane $Z_2 = Z_1 + \frac{k}{\alpha}$, where Z_1 is the plane of the camera trajectory and $k = \cot(\frac{\theta}{2})$; in the latter case, this plane is at infinity.

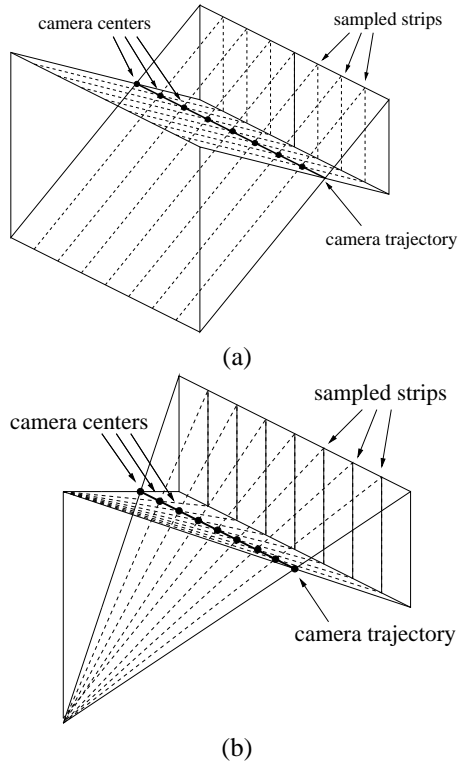


Figure 2: **Mosaicing by (a) a constant sampling function and by (b) a linear sampling function.**

General smooth trajectories are analyzed in Section 4.2 using local linear approximations.

3. Necessary Conditions for a Good Mosaic

Let $V \equiv \{(X, Y, Z) | Z > 0\}$ be the set of viewed scene points, i.e. the points in front of the camera.

We define the following necessary conditions for a good mosaic:

- **Unique Projection:** Every 3D point $P \in V$ is projected to a single point in the mosaic image.
- **Continuous Projection:** Connected sets of scene points are projected to connected sets of image points.
- **Data Utilization:** Strips are taken from all images.

A unique projection is important for avoiding duplicate images of an object in the mosaic image. In Section 3.1 we show that for linear camera trajectories, this condition holds if and only if the sampling function is monotonic non-decreasing¹. In Section 3.2 we relax the unique projection

¹Without loss of generality, it is assumed that the camera is moving from left to right



(a)



(b)

Figure 3: **3D object representation by mosaicing with a monotonic decreasing function. Mosaic (a) was generated by a linear sampling function, so every point on the object is associated with a single point on the image. Mosaic (b) was generated by a non-linear sampling function, and as can be seen, some scene points appear twice in the mosaic.**

condition by allowing a set G of points of measure 0 to violate the uniqueness condition. We show that in this case, G must be a line, and in case the camera moves on a linear trajectory, this corresponds to a linear sampling function. An almost-unique projection can be useful for constructing representations of convex objects. An example is shown in Figure 3.

The requirement for a continuous projection is obvious – to avoid discontinuities in the mosaic image. It follows that the sampling function must also be continuous.

The data utilization requirement is important for ensuring maximal field of view when minimizing the geometric distortion.

3.1. Projection Uniqueness

The projection is unique if every scene point is projected to a single point in the mosaic image. A key observation is

that the scene points V are *in front* of the camera. Hence the planes $\pi(t_1), \pi(t_2)$ must not intersect in front of the cameras for any $0 \leq t_1 < t_2 \leq 1$. For a camera moving on a linear trajectory, this implies that the sampling must be *non-descending monotonic*.

3.2. Uniqueness Excluding a Set of Measure 0

Another useful criterion relaxes the requirement by allowing some points to violate the uniqueness condition; This set of points G is required to be of measure 0 (e.g. a point or a curve). If G does not include any scene point, no scene point would appear multiple times in the mosaic. As we show below, this criterion implies that G is a line.

Theorem 1 *For any continuous sampling function $\phi(t)$, if the set of points that are not uniquely sampled is of measure 0, then this set is a line.*

Proof: The planes $\pi(0), \pi(1)$ intersect in a line \mathbf{l} , and all of its points are sampled by both cameras $t = 0, 1$. We will show that if there are points that are sampled by two cameras that are not on this line, then the set of all such points is of measure greater than 0: For every $s, t \in [0, 1]$, the intersection of plane $\pi(s)$ with plane $\pi(t)$ can be represented by the dual Plücker matrix:

$$\mathbf{L}^*(s, t) = \pi(s)\pi(t)^T - \pi(t)\pi(s)^T$$

(the planes are represented in homogeneous coordinates, see [3] - page 52). Since $\pi(t)$ is continuous, it follows that $\mathbf{L}^*(s, t)$ is continuous in s, t . If there exists a point not on \mathbf{l} that is sampled more than once, then it lies on a plane $\pi(a)$ for some $a \in (0, 1)$ such that $\mathbf{L}^*(0, 1) \not\propto \mathbf{L}^*(a, 1)$ (i.e., $\pi(0)$ and $\pi(a)$ intersect $\pi(1)$ in different lines). Refer to Figure 4 for illustrations. Consider the union of all lines of the form $\mathbf{L}^*(s, 1)$ for $s \in [0, a)$ (which are intersections of the planes $\pi(s)$ in this range with $\pi(1)$). Since $\mathbf{L}^*(s, t)$ is continuous, it follows that this union is a set of area greater than 0 on the plane $\pi(1)$. Let $A(t)$ denote the set of all points on the lines associated with $\mathbf{L}^*(s, t)$ for all $s \in [0, a)$. Then the above can be written as $|A(1)| > 0$. Due to the continuity of $\pi(t)$, there exists an interval $(b, 1]$ for which $|A(t)| > 0$ holds for every $t \in (b, 1]$. Therefore, since all planes $\pi(t)$ are distinct, it follows that the union $\cup A(t)$ is a set of volume greater than 0. Since it is contained in the set of all points that are sampled more than once, this set cannot be of measure 0. ■

Result 1 *In the case of linear camera motion, the sampling functions satisfying the uniqueness criterion up to measure 0 are either monotonic non-decreasing or linear (see Figure 3).*

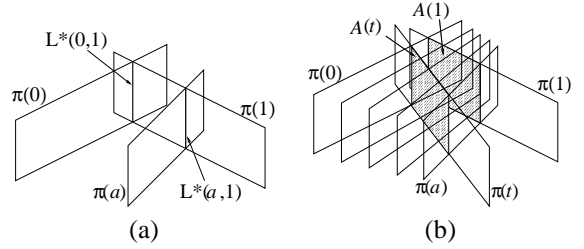


Figure 4: Illustrations for the proof of Theorem 1

4. Perspectivity: A Measure for Geometric Quality

We consider perspective images to be non-distorted. Hence the distortions in a mosaic image are measured with respect to the closest perspective image. In [14], a distortion was measured with respect to the closest perspective image, with the distance defined as the sum of distances of matching image points. Such a measure, while visually compelling, required knowledge of the scene depth. Since in our case the scene depth is unknown, we compare the *3D to 2D projections* rather than the images. That is, we would like the 3D to 2D projection induced by the mosaicing method to be as close as possible to a perspective projection. In a perspective projection, all rays intersect in a point. Hence, for a multi-perspective mosaic, the set of sampled rays should be as closely bundled as possible. We find a center point that has a minimal distance to all sampled rays, and we measure how small this distance is.

First, we define the *local perspectivity distortion*, which implements the idea above locally, for a neighborhood around an image point. We then define a *global perspectivity distortion* by integrating the local perspectivity distortion on the entire image. We selected an additive measure, so that the perspectivity of one region in the image is not influenced by other regions in the image.

We first analyze the case of linear camera motion. For this case, the least distorted mosaic is derived analytically, and it turns out that the least distorted mosaic also has the widest field of view. Non-linear camera trajectories are analyzed in Section 4.2 using local linear approximations. The global perspectivity is minimized numerically using standard optimization techniques.

4.1. Perspectivity: Linear Camera Trajectory

We consider only monotonic non-decreasing sampling functions satisfying the necessary conditions defined in Section 3.

Given a sampling function ϕ , each image point is associated with a single ray. Let us denote the intersection of the ray of image point p with the plane $Z = \hat{Z}$

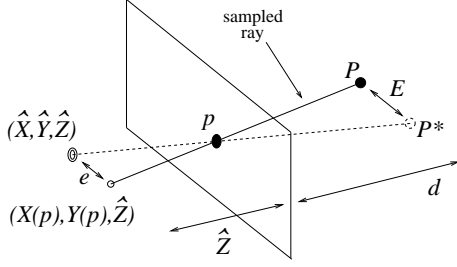


Figure 5: **The relation between the local perspective distortion and the error in estimation of the 3D scene. See Section 4.1 for details.**

by $(X(p), Y(p), \hat{Z})$. We define the distortion of the sampling function ϕ with respect to a candidate center point $(\hat{X}, \hat{Y}, \hat{Z})$ in a neighborhood ω of image point p as -

$$n(\phi, p, \hat{X}, \hat{Y}, \hat{Z}) = \int_{\hat{p} \in \omega} \frac{(X(\hat{p}) - \hat{X})^2 + (Y(\hat{p}) - \hat{Y})^2}{\hat{Z}^2} d\hat{p} \quad (1)$$

and the *local perspective distortion* at p as -

$$n_L(\phi, p) = \min_{\hat{X}, \hat{Y}, \hat{Z}} n(\phi, p, \hat{X}, \hat{Y}, \hat{Z}) \quad (2)$$

The expression given in (1) measures the distance e of the ray from the candidate center point $(\hat{X}, \hat{Y}, \hat{Z})$, on a plane, relative to the depth \hat{Z} of that plane (see Figure 5). The underlying idea is that an image is distorted if it is not consistent with a perspective image of a 3D scene. Consider a scene point P at depth d , which is projected by a ray whose error is e . Were the image perspective, this image point would seem to be the projection of a scene point P^* , and the error in the 3D scene would be E , such that $E = e \frac{d}{\hat{Z}}$.

The global measure of distortion is obtained by integrating a local perspective distortion on the image. To cancel the effect of the proportions of the neighborhood ω , we define the *global perspective distortion* of a given sampling function ϕ as follows:

$$n_G(\phi) = \int_{p \in I} \frac{n_L(\phi, p)}{n_L(\hat{\phi}, p)} dp \quad (3)$$

where $\hat{\phi}(t)$ is a reference sampling function which can be chosen arbitrarily, and p is integrated over the image domain $I = [x_{min}, x_{max}] \times [y_{min}, y_{max}]$. For simplicity, we choose the reference sampling function $\hat{\phi}(t) = 0$.

Theorem 2 *The global perspective distortion of a linear sampling function $\phi(t) = \alpha t + \beta$ is -*

$$n_G(\phi) = S \left(\frac{k}{k + \alpha Z_1} \right)^2 \quad (4)$$

where S is the image area (k and Z_1 are defined in Section 2).

The proof of the theorem above is omitted due to space limitations, and is given in [1].

A direct result of the theorem above is the following:

Result 2 *The global perspective distortion of a linear sampling function $\phi(t) = \alpha t + \beta$ with $\alpha \geq 0$ is monotonic decreasing in α . The most distorted linear sampling is the constant sampling.*

Note that the distortion of a linear sampling function depends only on the slope of the function. Now let us study the general case of continuous non-decreasing sampling functions:

Theorem 3 *Given a continuous non-decreasing sampling function $\phi(t)$, let us denote the linear sampling function which agrees with $\phi(t)$ at $t = 0$ and $t = 1$ as $\phi'(t)$. If $\phi \neq \phi'$, then $n_G(\phi) > n_G(\phi')$.*

In order to prove the above, we first prove it for a polygonal sampling function, i.e., a function $\phi(t)$ for which the interval $[0, 1]$ can be divided into segments $[0, t_1, t_2, \dots, 1]$ such that $\phi(t)$ is linear in each segment $[t_k, t_{k+1}]$.

Lemma 1 *Given a polygonal sampling function $\phi(t)$ and a linear sampling function $\phi'(t)$ such that $\phi(0) = \phi'(0)$ and $\phi(1) = \phi'(1)$, if $\phi \neq \phi'$ then $n_G(\phi) > n_G(\phi')$.*

Proof: The idea behind this proof is that by eliminating nodes in the polygon, the global perspective distortion does not increase. For any i , we eliminate the node i by defining a polygonal sampling function $\phi^*(t)$ which agrees with $\phi(t)$ everywhere except for the segment $[t_{i-1}, t_{i+1}]$, in which it is linear. As shown in (4), the distortion of a linear sampling function is proportional to the area of the image and to $\left(\frac{k}{\alpha Z_1 + k}\right)^2$. We denote the slopes of $\phi(t)$ in segments $[t_{i-1}, t_i]$ and $[t_i, t_{i+1}]$ by α_1 and α_2 , respectively, and the slope of $\phi^*(t)$ in $[t_{i-1}, t_{i+1}]$ as α_3 . The contribution of each segment to the global perspective distortion is proportional to its length, and therefore, $n_G(\phi) \geq n_G(\phi^*)$ if and only if

$$(t_i - t_{i-1}) \left(\frac{k}{\alpha_1 Z_1 + k} \right)^2 + (t_{i+1} - t_i) \left(\frac{k}{\alpha_2 Z_1 + k} \right)^2 \geq (t_{i+1} - t_{i-1}) \left(\frac{k}{\alpha_3 Z_1 + k} \right)^2$$

It can be shown that this inequality always holds, and that it becomes an equality if and only if $\alpha_1 = \alpha_2 = \alpha_3$, i.e., if $\phi = \phi^*$.

By repeatedly applying the result above on ϕ we obtain $n_G(\phi) \geq n_G(\phi')$, and $n_G(\phi) = n_G(\phi')$ only if $\phi = \phi'$. ■

Now we can proceed and prove the theorem:

Proof of Theorem 3: For any $\varepsilon > 0$, we divide the interval $[0, 1]$ into segments $[0, \varepsilon, 2\varepsilon, \dots, 1]$ and approximate $\phi(t)$ with a polygonal sampling function $\phi_\varepsilon(t)$ such

that $\phi(k\varepsilon) = \phi_\varepsilon(k\varepsilon)$ for all k , and $\phi_\varepsilon(t)$ is linear in each segment $[k\varepsilon, (k+1)\varepsilon]$. From Lemma 1 it follows that $n_G(\phi_\varepsilon) > n_G(\phi')$. Since this is true for all $\varepsilon > 0$, and since $\phi(t)$ is continuous, it follows that $n_G(\phi) > n_G(\phi')$. ■

Combining Theorem 3 with Result 2, we obtain:

Result 3 *For a camera moving sideways on a straight line, the sampling function with the minimal perspectivity distortion is $\phi_{opt}(t) = x_{min} + t(x_{max} - x_{min})$. This linear sampling function starts with the leftmost column of the first image and finishes with the rightmost column of the last image.*

4.2. Perspectivity: Non-Linear Trajectory

In order to handle non-linear camera trajectories, we define the local perspectivity (Equation 2) based on a local linear approximation of the camera trajectory and a local planar approximation of the manifold. For each frame I_f , where $1 \leq f \leq N-K+1$, we compute a discrete version $d_L(\phi, f)$ of the local perspectivity (equation 2) over a set of K neighboring frames I_f, \dots, I_{f+K-1} and minimize the sum of the discrete local perspectivities:

$$d_G(\phi) = \sum_f d_L(\phi, f) \quad (5)$$

To find the minimum of (5), we discretize the strip locations. Note that the local perspectivity $d_L(\phi, f)$ is defined by finding an optimal center of projection for each combination of rays. Computing these centers of projections for all possible sampling functions and for a large K is computationally intractable. This can be circumvented by selecting $K = 2$, in which case the local distortion $d_L(\phi, f)$ was derived analytically, as it corresponds to the linear perspectivity as defined in theorem 2. Once $d_L(\phi, f)$ is computed for all pairs of views, we use belief propagation [6] to find the optimum of equation 5. The complexity of this algorithm is linear in the number of frames, and quadratic in the number of possible strip locations in each frame.

5. Results

Figure 8 shows mosaicing results, using different sampling functions, from video sequences captured by a camera moving on a linear trajectory. We compare the optimal sampling function with the constant sampling function (linear pushbroom mosaicing), and with a non-linear monotonic sampling function $\tilde{\phi}(t)$ satisfying $\tilde{\phi}(0) = \phi_{opt}(0)$ and $\tilde{\phi}(1) = \phi_{opt}(1)$. This demonstrates two main results of this work. First, among all linear sampling functions $\phi(t) = \alpha t + \beta$, the least distorted results are achieved with the maximal α (compare Figure 8-b vs. 8-c). Second, among all monotonic functions aligned at the edge points $t = 0, 1$, the optimal sampling function is the linear one (compare Figure 8-b vs. 8-d).

Figure 6 compares a stereo mosaic generated by a constant sampling function (as done by [4,7,11,17]) to one generated by the optimal linear sampling function. Note that in addition to the distortions in the image, there is a distortions in the disparity which is larger with the constant sampling.

As for non-linear camera trajectories, we computed the least-distorted mosaics for various camera trajectories, some examples of which are shown in Figure 7. In all cases we tested, the least distorted mosaic was obtained when the projection rays intersect in a line, as in [18].

One practical case of a non-linear trajectory is when the camera moves on a circular arc. We examined visually the differences between the least-distorted mosaic and mosaics generated by constant sampling functions [7]. Various constant sampling functions were compared, each with a strip taken from a different offset from the center. The least distorted mosaic in this case is a Crossed-Slits mosaic [12, 18], as shown in Figure 7-a. We found that the differences in distortions with circular camera motion are not as significant as with linear camera motion, as the rays in this case are bundled together to begin with. Furthermore, in the case of non-linear trajectory, the distortion is also affected by the fact that the manifold is non-planar; this kind of distortion, which (unlike perspectivity distortion) can be treated with 2D warping, has not been discussed in this paper.

6. Summary

Multi-perspective mosaicing is a robust and efficient tool to summarize a video and to create 3D visualizations from a moving camera. In both applications, it is important to achieve least-distorted mosaics even when the 3D structure of the scene is unknown. We have quantified the image distortion and derived the optimal mosaic. When the camera moves on a linear trajectory, the least distorted mosaic is generated by the linear sampling function with the maximal slope. This mosaic also has the largest possible field of view. When the camera trajectory is not linear, the least-distorted mosaic can be derived numerically. We found that the distortions are especially significant when camera trajectory is close to linear.

While the scene depth and camera parameters were used in the analysis, the optimal sampling function is derived without knowing the scene depth. However, if some properties of the scene are known in advance, such as a rough depth estimate, they can be used to reduce the distortion, as was shown in [12, 18].

Acknowledgments

The authors would like to thank Daphna Weinshall, Shmuel Peleg and Michael Werman for their helpful comments.




Figure 6: **A comparison between different strip sampling methods for stereo mosaics. The images should be viewed in full color using anaglyphic 3D glasses. (a1) and (a2) are two rectified input images. Mosaic (b) was generated by the optimal sampling function, mosaic (c) by the constant sampling function (pushbroom mosaic).**

References

- [1] D. Feldman and A. Zomet. Least Distorted Mosaics. Hebrew University TR:2003-71, June 2003. http://leibniz.cs.huji.ac.il/tr/acc/2003/HUJI-CSE-LTR-2003-71_paper.ps.gz
- [2] R. Gupta and R.I. Hartley. Linear pushbroom cameras. *IEEE TPAMI*, 19(9):963–975, September 1997.
- [3] R.I. Hartley and A. Zisserman. *Multiple View Geometry*. Cambridge University Press, 2000.
- [4] H. Ishiguro, M. Yamamoto, and S. Tsuji. Omni-directional stereo. *IEEE TPAMI*, 14(2):257–262, February 1992.
- [5] T. Pajdla. Stereo with oblique cameras. *IJCV*, 47(1-3):161–170, April 2002.
- [6] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, 1988.
- [7] S. Peleg, M. Ben-Ezra, and Y. Pritch. Omnistereo: Panoramic stereo imaging. *IEEE TPAMI*, 23(3):279–290, March 2001.
- [8] S. Peleg and J. Herman. Panoramic mosaics by manifold projection. In *Proc. of CVPR'97*, pages 338–343, June 1997.
- [9] S. Peleg, B. Rousso, A. Rav-Acha, and A. Zomet. Mosaicing on adaptive manifolds. *IEEE TPAMI*, 22(10):1144–1154, October 2000.
- [10] P. Rademacher and G. Bishop. Multiple-center-of-projection images. In *SIGGRAPH*, volume 32, pages 199–206, 1998.
- [11] S.M. Seitz. The space of all stereo images. In *Proc of ICCV'01*, pages I: 26–33, 2001.
- [12] H.Y. Shum and L.W. He. Rendering with concentric mosaics. In *SIGGRAPH*, pages 299–306, 1999.
- [13] H.Y. Shum and R. Szeliski. Stereo reconstruction from multiperspective panoramas. In *Proc. of ICCV'99*, 14–21, 1999.

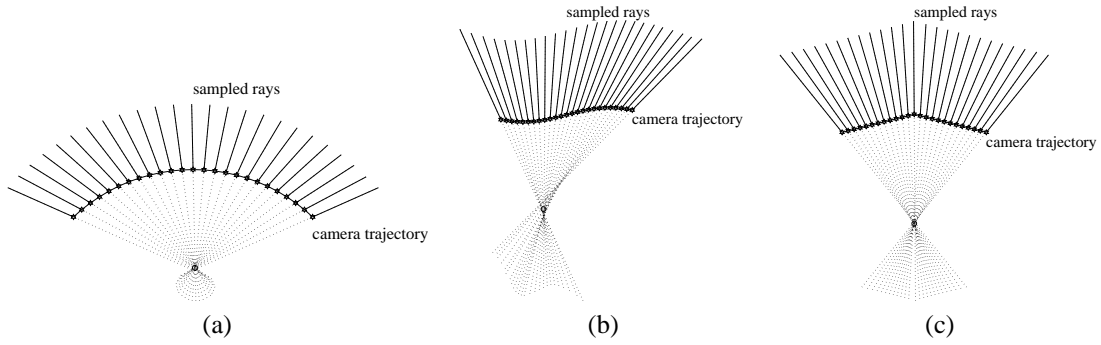


Figure 7: **Generating least-distorted mosaics with non-linear camera trajectories.** The illustrations show a top view of the camera trajectories and the planes of sampled rays, as computed by a numerical discrete optimization. In all cases, the optimal sampling is obtained when the sampled rays intersect in a line.




Figure 8: **A comparison between different strip sampling methods.** (a1) and (a2) are two rectified input images. Mosaic (b) was generated by the optimal sampling function, mosaic (c) by the constant sampling function (pushbroom mosaic) and mosaic (d) by a non-linear sampling function $\phi(t) = \sqrt{t}$.

- [14] R. Swaminathan, M.D. Grossberg, and S.K. Nayar. A perspective on distortions. In *Proc. of CVPR'03*, II: 594–601, 2003.
- [15] D.N. Wood, A. Finkelstein, J.F. Hughes, C.E. Thayer, and D.H. Salesin. Multiperspective panoramas for cel animation. *SIGGRAPH*, 31:243–250, 1997.
- [16] J.Y. Zheng and S. Tsuji. Panoramic representation for route recognition by a mobile robot. *IJCV*, 9(1):55–76, October 1992.
- [17] Z. Zhu, E.M. Riseman and A.R. Hanson. Parallel-perspective stereo mosaics. In *Proc. of ICCV'01*, pages II: 345–352, 2001.
- [18] A. Zomet, D. Feldman, S. Peleg, and D. Weinshall. Mosaicing new views: The crossed-slits projection. *IEEE TPAMI*, 25(6):741–754, June 2003.