

Robust Mosaicing Using Zernike Moments

Fady Badra
Virtual Reality Laboratory
McGill University
Montreal, Que., Canada H3A 2A7
e-mail: badraf@ee.mcgill.ca

Ala Qumsieh, Gregory Dudek
Center for Intelligent Machines
McGill University
Montreal, Que., Canada H3A 2A7
e-mail: {qumsieh,dudek}@cim.mcgill.ca

Abstract

This paper describes a new approach to image registration. Our intended application is the creation of image mosaics. Traditional approaches based are based on cross-correlation and only cope with rotation and scaling using a computationally costly search process. By using Zernike moments to compute the registration parameters, we are able to compensate not only for translation between images, but also for arbitrary amounts of rotation and for large amounts of scaling with a computationally efficient methodology. An attention-like operator based on edge density is used to compute a single candidate common feature between pairs of images being registered. Experimental results illustrate the performance of the method.

1 Introduction

In this paper, we present a new algorithm for automatic image mosaicing. Our approach uses Zernike moments to compute the relative transformations among the images. Translational registration, rotational registration and zooming can all be recovered at the same time.

Image mosaicing has a great deal of applications in medicine, digital terrain mapping and in autonomous navigation of robots. In each of these examples, smaller overlapping images are used to construct an image with a far larger field of view than could be obtained with a single photograph. Potential and existing applications include tasks such as segmentation, object recognition, shape reconstruction, motion tracking, stereo mapping and character recognition. In medical image analysis, relevant research areas include diagnostic medical imaging, such as tumor detection and disease localization and biomedical research including classification of microscopic images of blood cells, cervical smears, and

chromosomes. Registration methods can also be applied to civilian and military applications, agriculture, geology, oceanography, oil and mineral exploration, and target location and identification.

Zernike moments belong to a set of orthogonal polynomials, which allow independent moment invariance to be constructed easily to an arbitrarily high order. Such moments are generally superior to other types of moments in terms of noise sensitivity, information redundancy and image representation ability. Thus, Zernike moments offer a neat and complete solution to recover rotational and scaling parameters without the need for extensive correlation and search algorithms.

The next section presents a detailed review of the related work. In section III, we discuss some previous work and background on the problem of mosaicing. Our approach is discussed in section IV, while section V demonstrates some of the results. As a conclusion, we investigate the robustness and significance of our method along with a discussion of potential future applications.

2 Background

The main problem is to find the overlapping region between two neighboring images and to paste them together. The problem can be summarized as follows. Given a sequence of spatially overlapping images, find the transformation between adjacent images, align them, and use the results to construct a larger compound view (See Figure 1). The process of finding the relative transformation between any two images is called *image registration*, while the process of generating panoramic 3D views is called *image mosaicing* or *stitching*. Using *homogeneous coordinates*[9], the general transformation can be linearly decomposed into 3 independent components, which are translation, rotation and scaling.

The general matrix, M_{gen} , governing these transformations has the following form

$$M_{gen} = \begin{bmatrix} s_x r_{11} & s_x r_{12} & t_x \\ s_y r_{21} & s_y r_{22} & t_y \\ 0 & 0 & 1 \end{bmatrix} \quad (1)$$

where s_i denotes the scaling in the i -th direction, r is the rotational component, and t_i denotes translation in the i -th direction. The problem can thus be reduced to recovering the coefficients of each of these three components independently.




Figure 1: General mosaicing problem

There is no direct method available for automatically preparing image mosaics. The simplest approach used in practice to prepare a mosaic is to assemble the photographs manually. The key problem in automating the process lies in developing a good algorithm to reduce the seam between images of neighboring regions. In essence, this is an instance of the correspondence problem. The problem of spatial assembling of these images is solved manually by selecting a few points of interest in the overlapping region, setting the correspondence manually and then obtaining the necessary transformation.

Efficient methods have been developed to automatically solve for the translational shift. Consider the case when there is a large overlap between the images. Szeliski[6] has proposed a method for image registration by directly minimizing the discrepancy in intensities between pairs of images. An optimal solution is found using the Levenberg-Marquardt non-linear minimization algorithm. This method has the advantage of not requiring any feature points and being statistically optimal in the vicinity of the true solution, but it needs *a priori* approximation of the relative transformation between the images. Moreover, the Levenberg-Marquardt algorithm is computationally expensive, and does not necessarily converge to the desired minimum.

On the other hand, if the overlap between the images is small, a good method was devised where hierarchical matching is used to avoid local minima[10]. Kuglin and Hines[3] proposed another

method called *phase correlation*, which gave good results when the camera motion is very large. This method relies on the fact that a translation in the spatial domain corresponds to a phase shift in the frequency domain. Although the results were quite accurate, phase correlation requires a memory capacity that grows with the log of the image area, which makes it slow when handling larger images.

In the case of rotational shift, few methods proved to be efficient. Among the best methods we have the work of Dani and Chaudhuri[2]. Their general method consists of three stages. In the first stage, features in the image are computed. In the second stage, feature points in the reference image, often referred to as control points, are corresponded with feature points in the data image. In the last stage, a spatial mapping is determined using these matched feature points. Reassembling of one image onto the other is performed by applying the spatial mapping and interpolation. This method works well for up to 15 degrees of rotation using angles between edge points.

De Castro and Morandi[1] used the fact that rotation is invariant with the Fourier transform; rotating an image rotates the Fourier transform of that image by the same angle, ϕ . If we know the angle, then we can rotate the cross-power spectrum and determine the translation by phase correlation. However, since ϕ is not known, the phase of the cross-power spectrum is computed as a function of the rotation angle estimate. By first determining the angle ϕ which makes the inverse Fourier transform of the phase of the cross-power spectrum the closest approximation to an impulse, the translation is determined simply as the location of the pulse. Applying the method to a zero-padded image showed that it is very costly because of the difficulty in testing for each angle ϕ .

Zoghiami and Faugeras[4] used corners to build a 2D mosaics from a set of images. In addition to requiring an overlap of at least 50%, the method could successfully recover rotations only around the optical axis of the camera.

Peleg and Herman[5] used manifold projections to generate a panoramic view. This method is well suited to images with small degrees of rotations and does not address the zooming problem. While Rousso, Peleg and Finci[11] devised a method to create image mosaics using generalized strips. Basically, strips of the images, that are perpendicular to the optical flow pattern, are extracted. These strips are then warped to ensure the continuity of the image, and finally pasted successively. Results using this technique are quite satisfactory, but re-

quire the precomputing of optical flow fields, which are usually noisy and not very reliable. In addition, flow computation assumes dense temporal sampling. Small rotations are accounted for, but large rotations can not be handled.

Thus, existing mosaicing methods have strong limitations on imaging conditions and distortions are very common. No method was given that could simultaneously solve for lateral translation, rotation about any point, and zooming.

In attempting to propose a complete solution, we make the following assumptions:

- There is no warping of images (no elastic deformation).
- The subimages do not exhibit a high level of repetitive similarity, for example, a sinusoidal or a checker board variation in intensity.
- Intensity in the image does not change arbitrarily, but there may be a change in overall contrast due to changes in illumination or camera parameters.
- There is at least 10% overlap between adjacent images.

In the case where there is no overlap of the scenes in two such images, the registration cannot be accomplished without *a priori* knowledge of the pictured scene. Heuristic approaches such as analytic continuation of curvilinear elements in spatially adjacent photographs may be used under such circumstances, but this is beyond the scope of this paper.

This article presents a new method to compute the registration function between two images requiring only very little overlap (10–20%) and arbitrary rotation. Our technique will also take care of magnification and zooming. It is very fast, completely automatic without any human interaction and can handle any number of images.

3 Theory

3.1 Fourier Transforms

Kuglin and Hines[3] proposed a method to align two images which are shifted relative to one another. This method is called phase-correlation, and relies on the translation property of the Fourier Transform; basically, a lateral shift in spatial domain translates into a phase shift in the Fourier domain. The method states that given two images f_1 and f_2 which differ only by a shift (dx, dy) , their respective Fourier Transforms can be found

$$F_1 = |F_1|e^{-i\phi_1} \quad (2)$$

$$F_2 = |F_2|e^{-i\phi_2} \quad (3)$$

Multiplying F_1 by F_2^* , and dividing by the magnitudes, results in a function that depends only on the difference of the phases of the two input images

$$\begin{aligned} \Delta &= \frac{|F_1||F_2|e^{-i(\phi_1-\phi_2)}}{|F_1||F_2|} \\ &= e^{-i(\phi_1-\phi_2)} \end{aligned} \quad (4)$$

Now, taking the Inverse Discrete Fourier Transform (IDFT) of Δ results in a delta function whose peak corresponds to the relative spatial shift between the two images.

Although this method does not work very well in the case of different scaling, and large rotations, it still gives good results for small amounts of rotation. In the ideal case where the relative transformation between the images is a pure translation, phase-correlation proved to be extremely accurate to within a single pixel.

3.2 Zernike Moments

Zernike Polynomials constitute a complete set of complex polynomials that are orthogonal over the interior of the unit circle, $x^2 + y^2 = 1$. Since these polynomials are complex, and denoting them by $V_{nm}(x, y)$, they can be expressed by a radial and an angular components

$$V_{nm}(x, y) = V_{nm}(\rho, \theta) = R_{nm}(\rho)e^{im\theta} \quad (5)$$

where

n is a non-negative integer,

m is an integer obeying the constraints $n - |m|$ is even and $|m| \leq n$,

ρ is the magnitude of the vector from the origin to the point $(x, y) = x^2 + y^2$,

θ is the angle between the vector and the x-axis, in a counterclockwise direction.

The Radial component $R_{nm}(\rho)$ is defined as

$$R_{nm}(\rho) = \sum_{s=0}^{\frac{n-|m|}{2}} \frac{(-1)^s (n-s)!}{s! \left(\frac{n+|m|}{2} - s\right)! \left(\frac{n-|m|}{2} - s\right)!} \rho^{n-2s} \quad (6)$$

It should be noted that the radial polynomials are symmetric with respect to m , i.e. $R_{nm}(\rho) = R_{n,-m}(\rho)$. The orthogonality of the Zernike Polynomials is expressed by the following relationship

$$\int \int_{x^2+y^2 \leq 1} V_{nm}^*(x, y) V_{pq}(x, y) dx dy = \frac{\pi}{n+1} \delta_{np} \delta_{mq} \quad (7)$$

where

$$\delta_{ab} = \begin{cases} 1 & a = b \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

The projection of the image function $f(x, y)$ onto this orthogonal set of polynomials define the Zernike moments. More specifically, the Zernike moment of order n , with repetition m , has the form

$$A_{nm} = \frac{n+1}{\pi} \int \int_{x^2+y^2 \leq 1} f(x, y) V_{nm}^*(\rho, \theta) dx dy \quad (9)$$

In the case of a digital image, which is composed of a discrete number of pixels, the double integral simply reduces to a double sum

$$A_{nm} = \frac{n+1}{\pi} \sum_x \sum_y f[x, y] V_{nm}^*(\rho, \theta) \quad , \quad x^2 + y^2 \leq 1 \quad (10)$$

Note that the equation of the Zernike moment holds only within the unit disk. In order to correctly compute the Zernike moments of an image, the image has to be mapped first onto the unit circle. Only pixels lying within the unit circle are used.

Once the Zernike moments are computed, the orthogonality of the Zernike basis allows the use of these moments to reconstruct the image. In order to perfectly recover the input image, an infinite number of moments has to be computed. Since this is not possible, only a discrete replica $\hat{f}(x, y)$ can be recovered by computing the moments up to a given moment n_{max} . $\hat{f}(x, y)$ is computed as follows

$$\hat{f}(x, y) = \sum_{n=0}^{n_{max}} \sum_m A_{nm} V_{nm}(\rho, \theta) \quad (11)$$

Now consider an image $f(r, \theta)$, and a rotated replica $f_r(r, \theta)$ rotated by an angle α . Those two images are related as such

$$f_r(r, \theta) = f(r, \theta - \alpha) \quad (12)$$

By the virtue of their definition, Zernike moments of both of these images are coupled by a simple relationship. In polar coordinates, where $x = \rho \cos \theta$ and $y = \rho \sin \theta$, we can express the Zernike moments of the original image as

$$\begin{aligned} A_{nm} &= \frac{n+1}{\pi} \int_0^{2\pi} \int_0^1 f(\rho, \theta) V_{nm}^*(\rho, \theta) \rho d\rho d\theta \quad (13) \\ &= \frac{n+1}{\pi} \int_0^{2\pi} \int_0^1 f(\rho, \theta) R_{nm}(\rho) e^{-im\theta} \rho d\rho d\theta \end{aligned}$$

The equations for the rotated image are

$$A_{nm}^r = \frac{n+1}{\pi} \int_0^{2\pi} \int_0^1 f(\rho, \theta - \alpha) R_{nm}(\rho) e^{-im\theta} \rho d\rho d\theta \quad (14)$$

Letting $\beta = \theta - \alpha$, we can relate both sets of Zernike moments as follows

$$\begin{aligned} A_{nm}^r &= \frac{n+1}{\pi} \int_0^{2\pi} \int_0^1 f(\rho, \beta) R_{nm}(\rho) e^{-im(\beta+\alpha)} \rho d\rho d\beta \\ &= \left[\frac{n+1}{2\pi} \int_0^{2\pi} \int_0^1 f(\rho, \beta) R_{nm}(\rho) \cdot e^{-im\beta} \rho d\rho d\beta \right] \\ &\quad \times e^{-im\alpha} \\ A_{nm}^r &= A_{nm} e^{-im\alpha} \end{aligned}$$

Equation 15 shows that, under simple rotation, Zernike moments merely acquire a phase shift. Therefore, if the complex values of A_{nm} and A_{nm}^r could be computed, then the relative angle of rotation α between the two images can be easily recovered. Equation 15 also has an additional advantage. By taking the ratio of the magnitudes of A_{nm} and A_{nm}^r , the relative scaling, s , of the two images can be approximated to within a reasonable degree of accuracy

$$s = \frac{|A_{nm}|}{|A_{nm}^r|} \quad (16)$$

4 Algorithm Details

Our approach to image registration is based on selecting a preliminary correspondence, and then refining it using the Zernike moment computation. If the initial correspondence proves incorrect, an alternative correspondence is attempted. The process is summarized in the following pseudo-code, followed by a detailed explanation.

1. Select potential correspondences using interest operator, define points of points p_1 and p'_1 .
2. Calculate initial estimate $t_0 = (p_{1x} - p'_{1x}, p_{1y} - p'_{1y})$
3. Compute Zernike moments about p_1 and p'_1 .
4. Estimate R_n and S_n .
5. Align Images.
6. Compute t_{n+1} using FFT.
7. if $|t_{n+1} - t_n| < \delta \rightarrow$ END, else go back and recompute moments (goto step 3)

As noted above, the iterative solution for both translation and rotation is based on an initial estimate $t_0 = (p_1 - p'_1)$ to define a common center of rotation; that is, a pair of points that are projections of the same point in the scene. The alignment of the projections of such points provides a preliminary estimate of translation. Prior work on feature-based

correspondence has used a set of corner points under this type of assumption to solve to entire mosaic problem [4]. Corner points, however, are based on specific assumptions regarding the scale and structure of distinctive image content. We have evaluated the use of an alternative "interest operator" to define a small number of points which are candidates for preliminary correspondence.

We propose the use of an attention operator inspired by observations of human visual saccades [14]. Visual saccades are drawn to regions of unusual image content, and in particular regions with anomalous edge density. Similarly, a robust attention operator can be defined computationally which defines interest points as those regions of the image whose local edge-element density is maximal over the image [15]. While a detailed consideration of this operator is outside the scope of this paper, it appears to be robust and effective over a wide range of scene types.

A disk of a predefined radius is mapped around each of the points, and the Zernike moments within each disk are computed. Initial estimates, R_0 and S_0 , of the relative rotation and scaling of the two images are then calculated. This allows one of the images to be brought into correspondence with the other by applying the suitable scaling and rotation. Given this correction for scale and rotation, we can refine the translation estimate t_0 , to produce a more precise estimate t_1 . This is computed from phase correlation of the underlying images, a technique with several desirable robustness properties [13]. If the refined translation estimate differs substantially from the previous one, the scaling and rotation is re-evaluated. This process repeats until the two images align perfectly, or until the alignment error is below a threshold δ .




Figure 2: Blending function used to eliminate intensity differences between images.

In order to increase the speed and performance of the algorithm, we need discrete Fast Fourier Trans-

forms (FFT)[8]. The DFT needs $O(N^2)$ steps to be computed, while the FFT can be computed in only $O(N \log_2 N)$ steps. For the case of a 640×480 image, $N = 307200$. The ratio of the steps needed to calculate FFT to those needed to calculate DFT is 1 : 55982.58, which shows the immense difference between the two approaches.

Phase correlation is used to find corresponding peaks in the region of overlap, thus, there is no need to apply an FFT to the rest of the images. By assuming that the region of overlap covers at most 50% of the image, only half of the image needs to be processed. This further cuts down on the length of processing time, and increases the resolution of the FFT, thus accounting for smaller overlaps between images.

The program iterates until a suitable transformation is recovered. The termination criterion for the iteration process is satisfied when the difference between successive estimates of the lateral shift is below a certain threshold. Accurate results can be achieved if the threshold is lowered and the program is left to iterate further. On the other hand, resolution could be traded off for speed by simply increasing the threshold, and performing fewer iterations.

The program chooses the first image as the anchor image, and all other images are stitched relative to it. If the relative transformations are found with respect to this anchor image, major distortions may be caused by accumulated rotation, translation and scaling. This would reduce the resolution of the final image as we go away from the anchor image. In order to overcome the problem, all transformation coefficients are recovered first, without any stitching. Then, all the images are aligned relative to the anchor image, and stitched together.

Due to various lighting effects, successive images can have substantially different average intensities. These variations might not be apparent to the naked eye when the images are observed separately, but when stitched together, the resulting image can have apparent discontinuities at the locations corresponding to the images boundaries. In order to compensate for such discontinuities, a blending function was used to merge images together, and smoothly blend in the intensity variations. The blending function itself has the form of a ramp and is depicted in figure 2. The intensity value at each pixel in the region of overlap is a weighted sum of the values at the corresponding points in the original images. The closer the point to an image, the more it is affected by the value at the corresponding point in that image.

5 Results and Discussion

This section discusses the validity and robustness of our algorithm. It also presents some of the panoramic mosaics that were automatically generated using our approach. As can be seen later on, our algorithm can easily handle images with little overlap, large rotations about any axis, and scale differences.

5.1 Robustness

In order to verify the robustness of our algorithm, a set of tests was devised. A set of images with different scales and rotations were generated. These values were chosen so as to test the limitations of using Zernike moments for rotation and scale recovery. Figures 3 and 4 show the results.




Figure 3: Plot of the actual angles of rotation versus the calculated ones. All angles are in degrees.




Figure 4: Plot of the actual scale factors versus the calculated ones.

In figure 3, close examination shows that Zernike moments gave almost accurate results for *any* angle of rotation. Only shown are values corresponding to positive angles. Since negative angles can be thought of as reversing the order of the images, Zernike moments can recover any angle in the range of $-\pi$ to π .

In figure 4, we can see that the recovered values were very accurate for a scale factor of $\pm 70\%$. Beyond that limit, Zernike moments gave erroneous results that varied with the contents of the input images.

5.2 Pure Translation

The simplest case of image mosaicing is when the transformation between relative images is a pure translation. In this case, we only need to compute the relative (x, y) shifts between neighbouring images. Figure 5 shows such a case.

As can be seen from these images, the algorithm correctly detected that there is no relative rotation or scaling in the images. The final constructed images are flawless and extremely accurate.

5.3 Translation and Rotation

Various attempts by others to solve the rotation problem proved successful for only "small rotations" [11], while others assumed a rotation only about the optical axis of the camera[4]. Here, we show that our algorithm can successfully recover *any* angle of rotation about *any* axis. The trick is that once a corresponding pair of points has been identified in the two images, this point is taken as the center of rotation. Figures 6 and 7 show some results.

The relative rotation between the images in figures 6 and 7 range between 35 to 50 degrees. In figure 6, the relative rotation between the last two images is exactly 42 degrees. The algorithm computed this angle and found it to be 41.89 degrees. While in figure 7, the relative rotation between the final two images is exactly 39 degrees. The algorithm computed this rotation to be 39.03 degrees. As apparent from the final images, our algorithm is extremely accurate for any angle of rotation, and the results are almost perfect.

5.4 Zooming

Other known mosaicing methods assumed no difference in scale among images. Among the most elaborate solutions, [11] provides only "preliminary" results. In this section, we present results for the zooming problem using our method (See equation 16). We show that our algorithm can efficiently and accurately recover zoom scales between images of up to 150%. Moreover, the algorithm can detect zoom scales even if there is a relative rotation between the images. In fact, both the rotation angle and the zoom scale are computed *at the same time*.




Figure 5: A mosaic of an old building involving translation only.




Figure 6: Shown above are the three original images. The translation and phase difference between them is apparent. The resulting mosaic is shown on the bottom.




Figure 7: Another mosaic of a famous painting showing both translation and rotation recovery.

During dilation, the resolution of the image increases, while the field of view becomes smaller. On the other hand, during zooming out, the resolution decreases, while the field of view becomes larger. This creates a difference in resolution between images of different scales. In order to minimize the difference in resolution between images, all images are scaled down to the size of the smallest one, before being stitched. Linear and fractal interpolations can also be used to reduce this resolution discrepancy.

In figure 8, The two images are of the same size. There is a difference factor of 2 between the relative scales. Only parts of the images bounded by the boxes were used to generate the final mosaic on the left. The reason for cropping the images is that, due to the scale difference, one of the images is totally contained in the second one. As can be seen, the algorithm correctly detected the scale difference.

6 Conclusions & Future Work

In this paper, we have introduced a complete solution to the image mosaicing problem. The presented algorithm could quickly and satisfactorily solve for the relative translation, rotation and scaling between the input images. The quality of the final images is extremely good, and discrepancies at image boundaries are nonexistent. The algorithm is extremely fast, and could generate mosaics in real time. It could be used with a robot, where images taken by the robot are quickly stitched together as they are obtained.

The algorithms utilized Zernike moments to estimate the rotation and scaling differences among the images. It has been shown in studies comparing several types of moments (regular, complex, Legendre, Zernike, pseudo-Zernike) that Zernike moments are superior to other types of moments in terms of image representation capability, information redundancy and noise sensitivity [12].

Future application of image mosaics could be in dynamic virtual reality (sometimes called telepresence), which composites video from multiple sources in real time to create the illusion of being in a dynamic (and perhaps reactive) 3D environment. An example application might be to view a 3D version of a concert or a sporting event with control over the camera shots. Another application could be virtual classroom. Building such dynamic 3D models at frame rates is beyond the processing power of today's high performance workstations, but it could be achieved using a collection of these machines or a special purpose stereo hardware. Other important issues are those relating to the development of

a navigation system, this include design of a suitable scene representation, storage of images and so on.

References

- [1] E. D. Castro and C. Morandi, *Registration of translated and rotated images using finite Fourier transform*, IEEE Trans on Pattern Analysis and Machine Intelligence, PAMI-9, 700-702, 1997.
- [2] P. Dani and S. Chadhuri, *Automated assembling of images: Image montage preparation*, Pattern Recog., 28(3):431-445, 1995.
- [3] C. D. Kuglin and D.C. Hines, *The phase correlation image alignment method*, In Conference on Cybernetics and Society, pp. 163-165, New York, 1975.
- [4] I. Zoghلامي, O.Faugeras and R. Deriche, *Using geometric corners to build a 2D mosaic from a set of images*, Computer Vision and Pattern Recog. pp 420-425, Puerto Rico, June 1997.
- [5] S. Peleg and J. Herman, *Panoramic Mosaics by Manifold Projection*, Computer Vision and Pattern Recog. pp 338-343, Puerto Rico, June 1997.
- [6] R. Szeliski, *Image mosaicing for tele-reality applications*, Technical Report CRL 94/2, DEC-CRL, May 1994.
- [7] L. G. Brown, *A survey of image registration techniques*, ACM computing Surveys, Vol. 24, No. 4, pp. 325-376, December 1992.
- [8] W.H. Press et al., *Numerical Recipes in C: The Art of Scientific Computing*, 2nd edition, Cambridge Univ. Press, Cambridge, England, 1992.
- [9] J.D.Foley et al., *Computer Graphics: Principles and Practice*, 2nd edition, Addison-Wesley, Reading, Mass., 1990.
- [10] A. Witkin, D.Terzopoulos, and M.Kass, *Signal Matching through Scale Space*, International Journal of Computer Vision, Vol. 1, pp. 133-144, 1987.
- [11] B. Rousso, S. Peleg, and I. Finci, *Mosaicing with Generalized Strips*, DARPA Image Understanding Workshop, May 1997.
- [12] C. H. Teh and R. T. Chin, *On Image Analysis by Methods of Moments*, IEEE Trans. on Pattern Analysis and Machine Intelligence, PAMI-10, No. 4, pp. 496-513, July 1988.
- [13] D. J. Fleet and A. D. Jepson, *Computation of Component Velocity from Local Phase Information*, International Journal of Computer Vision, IJCV-5, No. 1, pp.77-104, August 1990.
- [14] D. Noton and L. Stark, *Eye Movements and Visual Perception*, Scientific American, Vol. 224, No. 6, pp. 33-43, June 1971.
- [15] R. Sim and G. Dudek *Navigation by the Stars: A Method for Vision-based Robot Localization*, CIM Technical Report, CIM TR-7634, September, 1997.



Figure 8: A panoramic image showing scaling. The two images on the left are the original images. Only the subimages shown inside the box were used to generate the final image on the right.