In Proceedings of the 1995 Oceans Conference, MTS/IEEE, October 1995

IMPROVED REAL-TIME VIDEO MOSAICKING OF THE OCEAN FLOOR

Stephen D. Fleischer * Richard L. Marks † Stephen M. Rock ‡
Stanford University Aerospace Robotics Laboratory
Durand Building 250
Stanford, California 94305

Michael J. Lee §
Monterey Bay Aquarium Research Institute
160 Central Avenue
Pacific Grove, California 93950

ABSTRACT

This paper presents a new method for the real-time creation of video mosaics of the ocean floor from a semi-autonomous underwater vehicle. Our approach uses iterative smoother-follower techniques to reduce the propagation of image alignment errors within the mosaic. This improvement relaxes previous constraints on mosaic size and shape, thereby enabling real-time, autonomous video mosaicking along unconstrained vehicle paths. Simulation results which demonstrate the advantages of our approach and preliminary experimental work are presented.

INTRODUCTION

A visual mosaic of the ocean floor is created by aligning several camera images taken at regular intervals to form a composite image of the scene. Recent innovations in visual sensing and control will soon give scientists the ability to produce these mosaics of the underwater environment in real-time, using cameras mounted on autonomous underwater vehicles [7].

While this technology will provide scientists with an excellent tool for exploring the ocean depths, there are limitations on what can be accomplished with current systems. The most significant problem that remains unsolved in real-time mosaic creation is the accumulation of alignment errors in the chain of overlapping camera images (see Figure 1). Since the error in aligning two consecutive images is small but finite, the error in positioning the final image relative to the initial image will increase without bound as the length of the image chain increases.

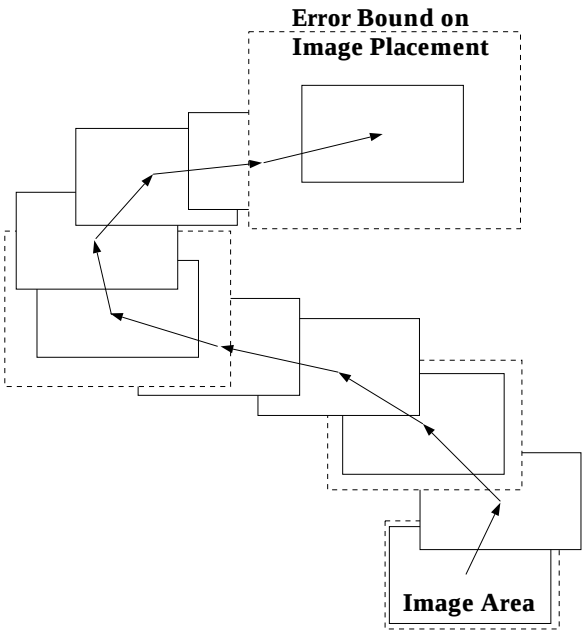

Figure 1: **Error Propagation in Image Chain**

For an image chain of length n, the error variance in the global position of the final image is proportional to n. Similarly, if there is a bias present in the local measurements between consecutive images, the bias error in the final image position will be proportional to n.

This unbounded error propagation has a fundamental impact on the mosaicking process; it constrains the

*Doctoral Candidate, Department of Aeronautics and Astronautics, fleisch@sun-valley.Stanford.EDU
†Doctoral Candidate, Department of Aeronautics and Astronautics, rlmarks@sun-valley.Stanford.EDU
‡Associate Professor, Department of Aeronautics and Astronautics, rock@sun-valley.Stanford.EDU
§Senior Research Engineer, Monterey Bay Aquarium Research Institute, lemi@hp850.Mbari.ORG

types of mosaics which can be created. For any vehicle path which loops back upon itself, the resulting mosaic will show poor image alignment at the crossover point. Thus, unconstrained mosaics are severely limited by path length and shape, making it difficult to map large areas of the ocean floor in this fashion. Although it is possible to improve mosaic quality by aligning non-consecutive images in the chain, this requires that the vehicle be constrained to follow a specific coverage pattern. As a result, constrained mosaicking utilizes all of the vehicle's resources; the scientist cannot control the vehicle motion while a mosaic is autonomously created.

Measurement errors also prohibit the use of vision hardware for global position sensing. In the ideal case, given the global position of the initial image, local measurements between successive images in the chain may be summed to determine the global vehicle position at all times. However, in the presence of uncertainties in the local measurements, the global position error is unbounded as the length of the image chain increases. Thus, vehicle navigation using vision as the only sensor becomes an impossible task.

This paper will present a method to optimally reduce the propagation of errors along loops in the image chain. This new method will enable the real-time, autonomous creation of a video mosaic along an unconstrained vehicle path. Furthermore, it will lay the groundwork for navigation from video by extracting global position data from the vision sensor.

BACKGROUND

In recent literature, several methods have been proposed to map the ocean floor using autonomous underwater vehicles and to navigate within unknown underwater environments. The most striking difference among these approaches is the choice of sensors. Several organizations have successfully integrated long or short/ultra-short baseline acoustic positioning networks with differential GPS to perform such tasks as autonomous station-keeping and biological sample collection[5, 13]. While this technique shows excellent promise for AUV navigation, it does not provide the scientist with any visual or topographical map of the ocean floor. Recent efforts in the area of sonar research have shown only limited success in real-time mapping of the ocean floor. Side-scan sonar has been used to create qualitative elevation maps consisting of interesting features in the terrain[6], but according to the authors, the nature of side-scan sonar precludes the creation of highly accurate range maps. FM subbottom penetrating sonar provides an accurate representation of the stratified layers below the sea-bed[14], which may be useful if a surface terrain model is not required by the scientist. In both cases, navigation from sonar has been proposed, but this may prove to be a difficult task without accurate surface range data.

Navigation from video has been accomplished on a land-based mobile robot, using a panoramic representation of the video data along a desired route through an unkown environment[15]. A visual map is produced as the robot is guided along the prescribed route, and the robot then autonomously navigates back along the identical path. This work is analogous to recent underwater research in constrained video mosaicking, in which a rectangular video mosaic of the ocean floor is created by correlating images in adjacent columns[7]. By directing an autonomous underwater vehicle to follow a fixed coverage pattern, automatic real-time creation of constrained mosaics in a test-tank environment has been demonstrated. This research effort has also achieved unprecedented experimental results in producing single-line unconstrained mosaics of the ocean floor, using *Ventana*, a remotely operated vehicle built by the Monterey Bay Aquarium Research Institute (MBARI).

APPROACH

To enable the task of unconstrained mosaicking and to move closer to navigation from video, we have developed an approach based on the theory of optimal estimation and smoother-follower techniques. Before deriving this approach in detail, it is worthwhile to review our decision to utilize vision as our primary sensor for these tasks.

Sensor Selection

While cameras are obviously essential for the creation of visual maps of the ocean floor, the choice of which sensor to use for alignment and consolidation of the images into a single mosaic is a difficult one. We have decided to use our vision sensor to accomplish image consolidation for several reasons.

Intuitively, it is possible to achieve greater accuracy by directly measuring what we wish to control. For our case, measuring the displacement between two images through visual correlation is the best way to properly align them within a mosaic. Another advantage of vision is the measurement accuracy which can be achieved through standard image correlation techniques. Using specialized hardware[11, 12], results of sub-pixel resolution in image displacement measurements have been demonstrated[9]. This corresponds to approximately 1 cm accuracy in vehicle position determination for a typical range of 1-2 meters above the ocean floor. Finally, since all of the required equipment can be carried on-board the vehicle (cameras, comput-

ers, vision processors), the vehicle is capable of exploring unknown, unstructured environments. This is a serious limitation of acoustic positioning systems, which require accurate placement of a network of transceivers around the area of interest.

In order to achieve navigation from video, the vision sensor must provide global position data to the vehicle navigation system. Although image correlation is essentially a local sensing technique, we have been able to infer the vehicle's global position and orientation from the local vision measurements. As a result, our vision system is able to compute the same information typically provided by acoustic transceiver networks. This computation will be detailed in the following sections.

System Geometry

Our method for unconstrained mosaic creation does not require a dynamic model for the underwater vehicle to estimate its current position and orientation. Instead, our estimates of image and vehicle position depend only on the sensor data and the system geometry. We will derive dynamic equations of state for the mosaicking process, based on the system geometry.

Our system consists of a stereo camera pair mounted on a semi-autonomous underwater vehicle, looking downward at the ocean floor. For simulation purposes, we have assumed the cameras are attached a fixed distance away from the vehicle center of mass, and they are oriented parallel to the vertical axis of the vehicle. While the actual vehicle provides pan-and-tilt actuators which enable the cameras to achieve different orientations than the vehicle, our assumption still holds if the vehicle frame is re-defined to coincide with the camera orientation, centered at the vehicle center of mass.

As the vehicle travels along its commanded path, our vision system takes image snapshots of the ocean floor at regular intervals along the path. The geometry for two consecutive snapshots is depicted in Figure 2. The vehicle frame (V) is defined by the position and orientation of the vehicle when the previous image snapshot was taken. The vehicle position and orientation for the current snapshot are represented by the V' frame. Similarly, the M and M' frames represent the image area corresponding to the cameras in the V and V' frames, repectively. The world frame (W) is defined so that the XY plane is aligned with the ocean floor, and the origin coincides with the initial vehicle position. Since the vision sensor is capable of producing a detailed range map of each image, the ocean floor need not be restricted to a planar surface. However, we will make the assumption of planarity for two reasons: the calculation of simulated range data is straightforward under this restriction; and non-planar surfaces may significantly degrade the accuracy of visual correlation results, depending on the severity of the surface curvature[8].

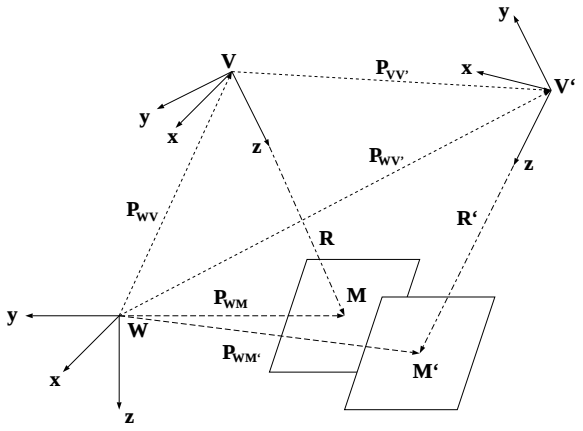

Figure 2: **System Geometry**

By correlating successive image snapshots, our vision sensor provides position and orientation measurements of frame M' relative to frame M. Therefore, we must transform these results to the world frame, in order to determine the global position and orientation of each image snapshot. Using the standard rules of geometric transformations:

$$\begin{aligned} {}^{W}P_{WM'} &= {}^{W}P_{WM} + {}^{W}P_{MM'} \\ &= {}^{W}P_{WM} + {}^{W}_{M}T\ {}^{M}P_{MM'} \end{aligned} \qquad (1)$$

where ${}^{A}P_{BC}$ depicts a vector from point B to C, represented in frame A, and ${}^{E}_{D}T$ represents a transformation matrix from frame D to frame E.

Since our vehicle is under active control in all degrees of freedom, we can assume the control system is attempting to regulate the vehicle orientation to coincide with the world frame. Therefore, we can make the small angle assumption for the euler angles relating frames W, V, and V'. Since frames M and M' have the same orientations as frames V and V', respectively (since the cameras are fixed relative to the vehicle), every angle in the system can be assumed to be small. Under this assumption, the multiplication of transformation matrices is approximately equal to the addition of the euler angles for these transformations:

$${}^{W}_{M'}T = {}^{W}_{M}T\ {}^{M}_{M'}T \Rightarrow {}^{W}_{M'}Q = {}^{W}_{M}Q + {}^{M}_{M'}Q \qquad (2)$$

where ${}^{E}_{D}Q$ is a vector of Euler angles representing the transformation from frame D to frame E.

When a new snapshot is taken and added to the image chain, the V' and M' frames become the new V and M frames, respectively, and the V' and M' frames will be redefined by the next image snapshot. Since

this process is repeated for each image in the mosaic, we can define frame V/M to be the vehicle/image state at step k, and frame V'/M' to be the corresponding state at step $k+1$. If we define the image state, $x(k)$, and the image local displacements, $u(k)$, as follows:

$$x(k) = \begin{pmatrix} {}^{W}P_{WM} \\ {}^{W}_{M}Q \end{pmatrix} \qquad u(k) = \begin{pmatrix} {}^{M}P_{MM'} \\ {}^{M}_{M'}Q \end{pmatrix} \quad (3)$$

we may combine equations (1), (2), and (3) to arrive at the following discrete dynamic equations of motion for the image state:

$$\begin{aligned} x(k+1) &= x(k) + \begin{pmatrix} {}^{W}_{M}T(k) & 0 \\ 0 & I \end{pmatrix} u(k) \\ \text{and} \qquad & x(0) = 0,\ {}^{W}_{M}T(0) = I \end{aligned} \quad (4)$$

For the purpose of global vehicle positioning, the vehicle state can be calculated directly from the image state:

$$\begin{aligned} {}^{W}P_{WV}(k) &= {}^{W}P_{WM}(k) - {}^{W}_{M}T(k)\ {}^{M}R(k) \\ {}^{W}_{V}T(k) &= {}^{W}_{M}T(k) \end{aligned} \quad (5)$$

The only question which remains is when the snapshots should be taken. To produce the highest quality mosaic, subsequent images should be recorded at fixed visual intervals[7]. In other words, a new snapshot will be taken when the current image has moved a specified distance in the XY plane of the M frame. As a result, the discrete, independent variable (k) in the above equations of motion can be defined as the displacement magnitude between snapshots, rather than the usual definition of time between snapsots. Although this displacement is constant between images in the ideal case, its variance will affect the accuracy of the equations of motion. This source of error will be discussed in more detail during the presentation of simulation results.

Optimal Estimation Theory

In attempting to reduce the error propagation along an image chain, it would seem intuitive to utilize some version of Kalman filtering to smooth the errors as additional sensor measurements are recorded. This real-time approach is applicable to dynamic systems of the form [1]:

$$\begin{aligned} x(k+1) &= A(k)x(k) + B(k)\mu(k) \\ z(k) &= C(k)x(k) + D(k)\nu(k) \end{aligned} \quad (6)$$

where $\mu(k)$ and $\nu(k)$ are process and sensor noise variables, respectively. However, if our state equations (4) are rewritten in similar form:

$$\begin{aligned} x(k+1) &= x(k) + \begin{pmatrix} {}^{W}_{M}T(k) & 0 \\ 0 & I \end{pmatrix} u(k) \\ z(k) &= u(k) - \nu(k) \end{aligned} \quad (7)$$

where $z(k)$ is the sensor measurement, it becomes evident that the state, $x(k)$, is completely unobservable by the sensor, $z(k)$. As a result, it is impossible for Kalman filtering to reduce the errors in image alignment within a mosaic[4, 1].

Thus, it would seem that there is no way to bound the errors on unconstrained mosaicking. However, we can gain additional information whenever the image chain loops back upon itself. By correlating the nth image with the jth image as well as the $(n-1)th$ image (see Figure 3), we gain another measurement of the nth image global state. Furthermore, this new measurement is more accurate, since the jth image occurs earlier in the image chain and thus its global state measurement has a lower variance.

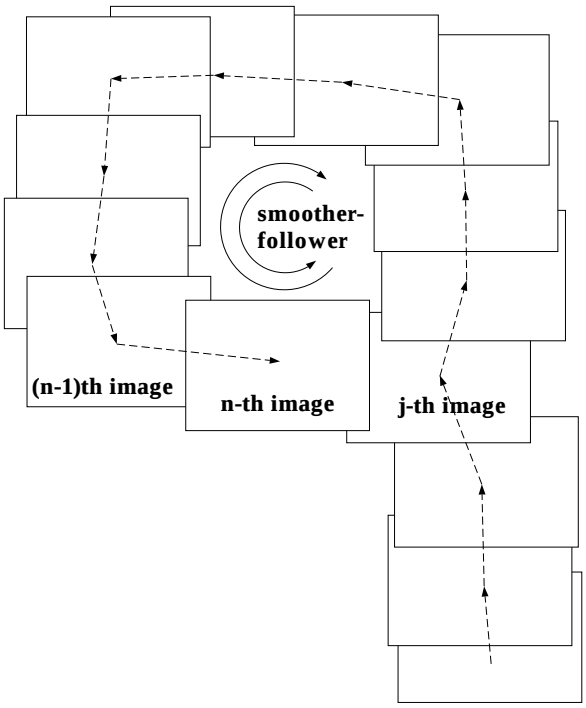

Figure 3: **Geometry for a Closed Loop Image Chain**

By utilizing an iterative smoother-follower, the additional information gained at the crossover point can be propagated along the entire loop to minimize the errors along that section of the image chain.

If we isolate the measurements along the loop between image j and image n, the equations of motion (7) are valid along this path. These equations can be combined as follows:

$$
\begin{aligned}
x(k+1) &= x(k) + \begin{pmatrix} {}^{W}_{M}T(k) & 0 \\ 0 & I \end{pmatrix} z(k) \\
&+ \begin{pmatrix} {}^{W}_{M}T(k) & 0 \\ 0 & I \end{pmatrix} \nu(k) \\
\text{for } k &= 0, \ldots, (n-j) \qquad (8)
\end{aligned}
$$

with known initial and final states equal to the jth and nth image states, respectively. To propagate the new information throughout the loop, we can use an optimal smoother-follower[1]. However, this state equation (8) is not in the exact form of a standard smoother-follower, so we must modify the performance index to be optimized. In order to minimize the magnitude of the estimated image local displacements, $u(k)$ (which will minimize the sensor noise, $\nu(k)$), while maintaining the accuracy of the endpoint positions, we have chosen the following modified performance index:

$$
\begin{aligned}
J &= \frac{1}{2}[x(0) - x_o]^T S_o [x(0) - x_o] \\
&+ \frac{1}{2}[x(n) - x_f]^T S_f [x(n) - x_f] \\
&+ \frac{1}{2} \sum_{k=0}^{n-j-1} u^T(k) R_d(k) u(k) \qquad (9)
\end{aligned}
$$

where S_o and S_f are the information (defined as the inverse of the error covariance) matrices for the initial (image k) and final (image n) states, respectively, and $R_d(k)$ is the information matrix for each sensor measurement. The numerical values of these matrices are based on estimated error distributions of the local image displacement measurements. The next section provides the estimated values used in our simulation work. This performance index differs from the standard index in that no weighting matrix has been assigned to the output vector, $z(k)$, since it has been absorbed into the state equation. In addition, the constraint equation (8) contains a known input vector, $z(k)$, which is not present in the standard state equation (6).

Although the derivation of this type of smoother-follower algorithm is beyond the scope of this paper, the results of simulations running this algorithm are quite interesting, and these will be presented in the following section. Before proceeding, it is worthwhile to examine the equation of state (8) more closely. Since the coefficients of $z(k)$ and $\nu(k)$ are actually functions of the state $x(k)$, they will change after running the smoother-follower algorithm. As a result, it is necessary to iteratively calculate the coefficients and smooth the data until the performance index (9) converges to a minimum value. This technique has been implemented for our simulation work.

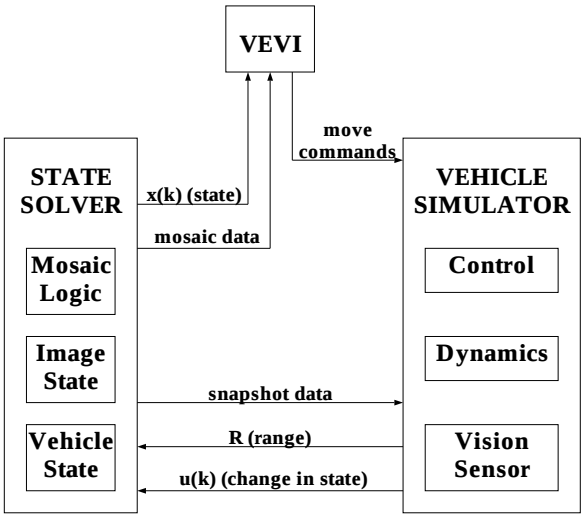

Figure 4: **System Diagram for Simulation**

SIMULATION WORK

Figure 4 represents the three main subsystems in our simulation. The vehicle subsystem includes a component for simulating the output of the vision sensor, namely, the difference in state between consecutive image frames. We have combined our sensor signal with a model of the errors associated with our vision correlation hardware. According to recent experimental data[9], the error distributions for the x, y, z, and yaw (ψ) state components are approximately uniform, with the following variances:

$$
E[\begin{pmatrix} x \\ y \\ z \\ \theta \end{pmatrix} (\; x \quad y \quad z \quad \theta \;)] =
$$

$$
\begin{pmatrix} 0.0032R & 0 & 0 & 0 \\ 0 & 0.00201R & 0 & 0 \\ 0 & 0 & 0.0055R & 0 \\ 0 & 0 & 0 & 0.0085 \end{pmatrix} \qquad (10)
$$

where R is the range from the camera to the image area. For our simulation, we have assumed these four degrees of freedom are measured by the vision sensor, and the pitch (θ) and roll (ϕ) are measured by additional sensors on the vehicle. Therefore, only the states in the above equation (10) need to be processed by our smoother-follower algorithm. Since the smoother-follower assumes Gaussian distributions for all random variables, we have modelled the above uniform distributions as Gaussian with identical means and variances. In addition to the sensor component, the vehicle subsystem contains a "perfect" control/dynamics component, which simply moves the vehicle to the

commanded endpoint, based on the noisy sensor data.

The VEVI (Virtual Environment Vehicle Interface) is a 3-D graphical user interface created by NASA Ames and modified to monitor and control the actions of our underwater vehicle[2, 3]. Position and orientation updates of both the vehicle and the images in the mosaic are received by the VEVI and displayed in the 3-D model of the underwater environment. Position commands are sent from the interface to the vehicle to control its path while an unconstrained mosaic is autonomously created.

The state solver subsystem computes the current vehicle and image state based on the incoming sensor data. This subsystem is also responsible for controlling mosaic creation. The mosaic logic component determines when the next image snapshot should be acquired, and sends a command to the vision hardware on board the vehicle. All mosaic data is stored in this component, and the smoother-follower algorithm is started whenever a crossover point is detected in the image chain.

There are two sources of error which our unconstrained mosaicking method has not taken into account. The first is the error on the range measurement from camera to image area. If we use stereo correlation to measure the absolute range, instead of measuring the change in range since the last image snapshot, the error will remain approximately constant. As a result, this error is negligible to first order, and we have ignored it in our computations.

The second source of error is the uncertainty in the independent variable, k, which is the (approximately) constant displacement magnitude between consecutive images. This uncertainty can be modelled quite simply by incorporating it into the error on the image states. However, instead of attempting to estimate the relative magnitude of this uncertainty, we have chosen to defer this issue and measure its magnitude during the experimental phase of our investigations.

To demonstrate the merit of our new approach to unconstrained mosaicking, the following figures illustrate a typical vehicle path and its corresponding mosaic. The vehicle follows a simple rectangular path in the XY plane while maintaining a constant heading. Figure 5 depicts the actual image position. The remaining two lines show the estimated position based on the sensor data, before and after the data has been filtered by the smoother-follower. Clearly, the endpoints of the loop have the smallest error, since these points are known to the greatest degree of accuracy, while the error around the loop has been minimized. This can be seen directly in Figure 6, which shows the standard deviation of computed image position before and after the smoother-follower processing. While the original variance increases without bound, our method clearly bounds the variance around closed loops in the image chain. In order to provide a visual comparison of the effect of this algorithm, we have taken an image (Figure 7) and created two mosaics, based on the estimated and filtered data in Figure 5. The results are shown in Figures 8 and 9.

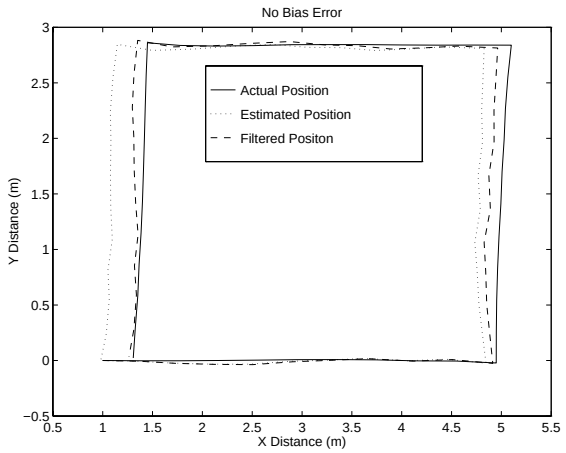

Figure 5: **Image Position within a Rectangular Mosaic**

When compared to the actual image position, the filtered data is more accurate than the estimated position based purely on noisy sensor data, particularly at the endpoints of the loop.

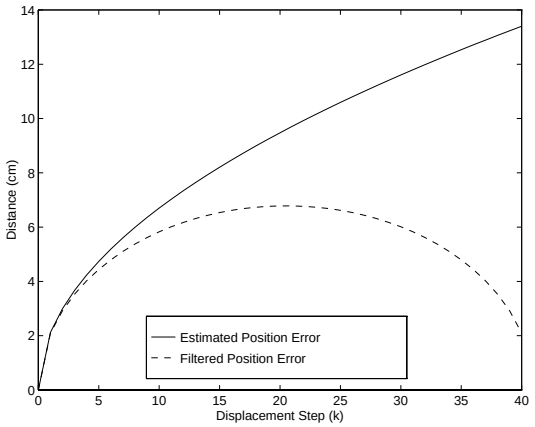

Figure 6: **Standard Deviation of State in X Direction**

Around any closed loop in the mosaic, our smoother-follower algorithm minimizes the variance, subject to the constraints of the equations of motion.

Another significant advantage of our method is that it removes unmodelled bias present in the local image displacement measurements. For instance, if we include a bias in the model of our vision sensor error,

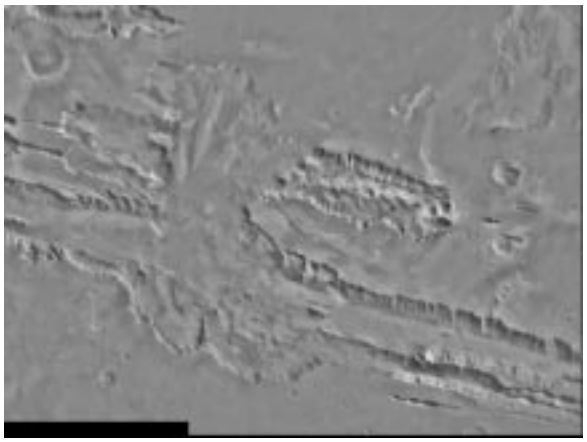

Figure 7: **Original Image for Mosaic Simulations**

Using this image as a simulated ocean floor, pseudo-mosaics can be created with our simulation data by mimicking a a vehicle traversal along a rectangular path over this image[10].

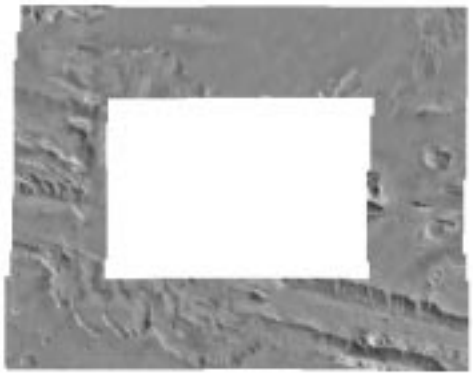

Figure 8: **Mosaic based on Estimated Image Position**

The discrepancy in image alignment, due to errors in the image displacement measurements, can be clearly seen at the crossover point (lower left corner).

Figure 10 shows an even more dramatic error growth along the image chain. As long as the estimated final position is sufficiently close to the actual crossover point for image correlation, the smoother-follower will successfully minimize the errors in image alignment. As before, Figures 11 and 12 provide a visual comparison of two mosaics created from the estimated and filtered data.

The potential benefits of this approach to unconstrained mosaicking are fully realized when the vehicle makes multiple passes over the same area of the ocean floor. During each pass, our algorithm further refines the mosaic data, thereby providing a continual

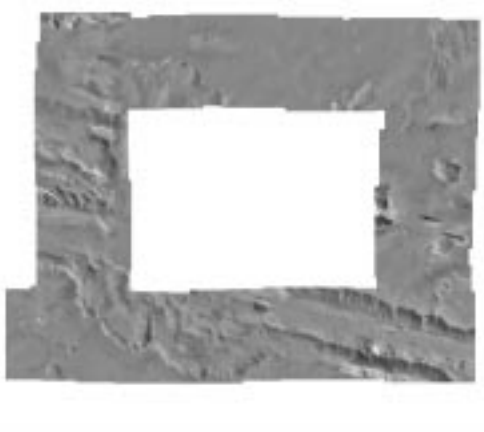

Figure 9: **Mosaic based on Filtered Image Position**

Compared to the previous mosaic, the image alignment is more accurate, particularly at the endpoints of the image chain (lower left corner).

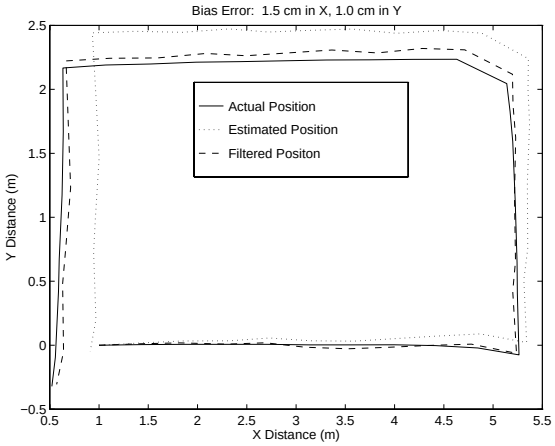

Figure 10: **Image Position within a Rectangular Mosaic (Bias)**

When compared to the actual image position, the filtered data is more accurate than the estimated position based purely on noisy sensor data, particularly at the endpoints of the loop.

improvement in the visual quality of the mosaic and the accuracy of the global position information.

EXPERIMENTS IN PROGRESS

The *OTTER* robot (Oceanographic Technology Testbed for Engineering Research) is being used for experimental testing of our approach to unconstrained video mosaicking (Figure 13). *OTTER* is roughly 2 meters long, 1 meter wide, and has a dry mass of 145 kg. It

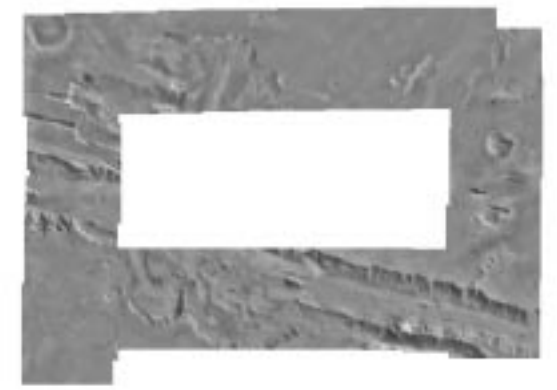

Figure 11: **Mosaic based on Estimated Image Position (Bias)**

The discrepancy in image alignment, due to errors in the image displacement measurements, can be clearly seen at the crossover point (lower left corner).

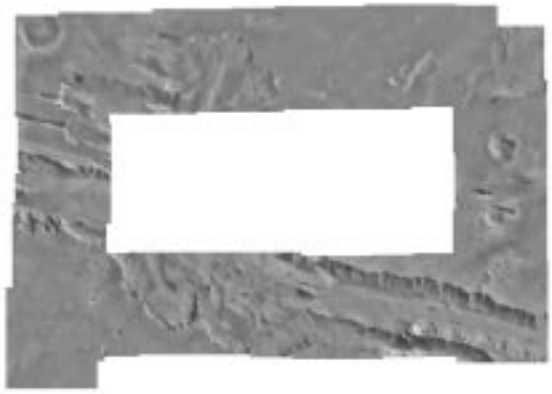

Figure 12: **Mosaic based on Filtered Image Position (Bias)**

Compared to the previous mosaic, the image alignment is more accurate, particularly at the endpoints of the image chain (lower left corner).

is made up of three pressure housings surrounded by 8 ducted thrusters and covered by a fiberglass shell. One housing holds two independent VME card cages with 68040 single board computers for control and sensor processing. The other two hold NiCad batteries which provide approximately 750 W-hrs of power. Currently, a tether is used to trickle charge the batteries and provide ethernet communications. The sensor suite includes custom real-time vision processing boards, pitch and roll gravity sensors, a small intertial measurement unit with 3 accelerometers and 3 rate gyros, a flux-gate compass, and a pressure depth sensor. Two black/white CCD video cameras are mounted as a stereo pair on a custom pan/tilt unit. Main propulsion is provided by two 2 hp brushless DC variable reluctance motors. Six 1/2 hp VR motors are used for lateral and vertical motions as well as attitude control. The real-time vision hardware that performs the digital filtering and correlation is produced by Teleos Research [11, 12].

Figure 13: *OTTER* **Underwater Vehicle**

Currently, we have been able to make single-column unconstrained and multiple-column constrained mosaics. Figure 14 shows an example of a constrained mosaic created in the test tank at MBARI. To create this mosaic, the vehicle followed a specific coverage pattern, so that each image could be correlated with an image from the previous column. This enabled uniform image alignment throughout the mosaic.

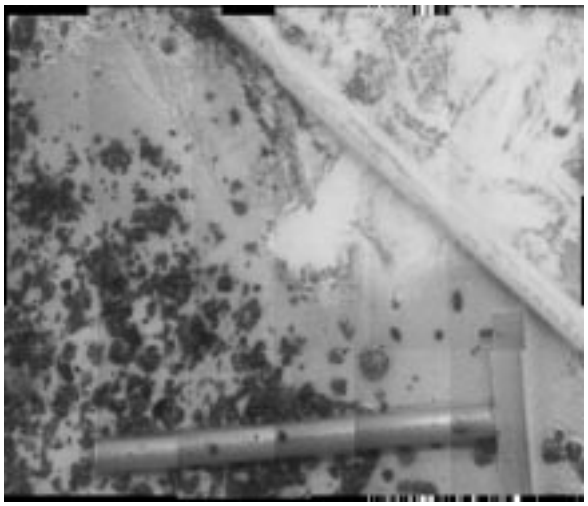

Figure 14: **Multiple-Column, Constrained Mosaic**

Using the same vision processing hardware, Figure 15 depicts a single-column unconstrained mosaic created using *Ventana*, MBARI's remotely operated vehicle. This mosaic was created autonomously while the pilot attempted to fly the vehicle in a straight line over the ocean floor.

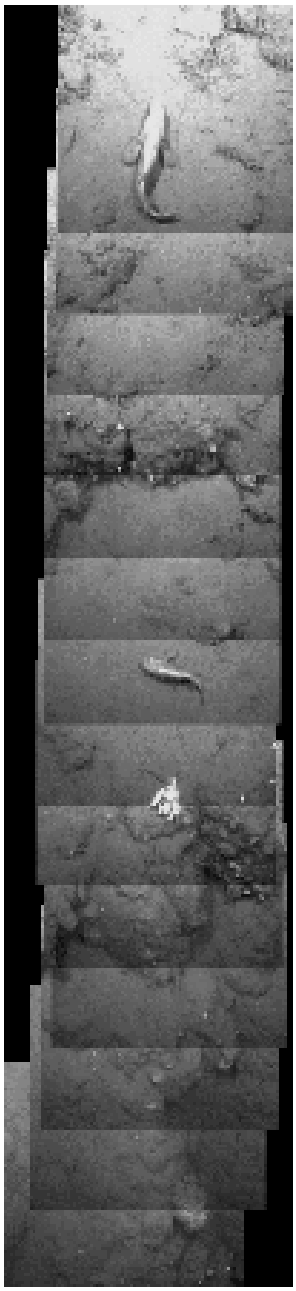

Figure 15: **Single-Column, Unconstrained Mosaic**

Our current task is to experimentally verify our simulation work by creating an unconstrained, closed-loop mosaic and minimizing the errors around that loop. The next logical step is to extend this work to include unconstrained mosaics with multiple crossover points.

CONCLUSIONS

We have presented a new technique for the real-time creation and on-line improvement of unconstrained video mosaics. To achieve this, we have extended the theory of smoother-follower estimation to account for unobservable states in the output and non-zero control input terms in the state equation. Our simulations have demonstrated the feasibility of this method, which will provide scientists with a new capability for undersea exploration, namely, the ability to remotely explore the marine environment while the underwater vehicle autonomously maps the ocean floor.

By reducing the global position error of each image, we have taken another step toward autonomous vehicle navigation from video. After creating the video mosaic of the ocean floor, we can correlate the current camera image with the stored mosaic to determine our precise vehicle position and orientation. Once we have accomplished this goal, the scientist will simply point-and-click within the virtual environment interface to direct the vehicle to an interesting area within the underwater environment.

REFERENCES

[1] Bryson, A. Dynamic Optimization with Uncertainty. Paper in progress, March 1993.

[2] Fleischer, S. D., Rock, S. M., and Lee, M. J. Underwater Vehicle Control from a Virtual Environment Interface. In *Proceedings of the Symposium on Interactive 3D Graphics* (Monterey, CA, April 1995), SIGGRAPH.

[3] Fong, T. A Computational Architecture for Semi-autonomous Robotic Vehicles. In *Proceedings of AIAA Computing in Aerospace 9 Conference* (San Diego, CA, October 1993), AIAA.

[4] Franklin, G., Powell, J., and Workman, M. *Digital Control of Dynamic Systems*. Addison-Wesley Pub. Co., Reading, MA, 1990.

[5] Illman, J., Milburn, H., and Macdonald, R. Integrated system for navigation and positioning of an rov for scientific exploration. In *Proc. IEEE Oceans* (1993), pp. 499–503.

[6] Langer, D., and Hebert, M. Building Qualitative Elevation Maps From Side Scan Sonar Data For Autonomous Underwater Navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation* (Sacramento, CA, April 1991).

[7] Marks, R., Rock, S., and Lee, M. Real-time video mosaicking of the ocean floor. *IEEE Journal of Oceanic Engineering 20*, 3 (July 1995), 229–241.

[8] Marks, R. L., Lee, M. J., and Rock, S. M. Visual sensing for control of an underwater robotic vehicle. In *Proceedings of IARP Second Workshop on Mobile Robots for Subsea Environments* (Monterey, May 1994), IARP.

[9] Marks, R. L., Wang, H. H., Lee, M. J., and Rock, S. M. Automatic visual station keeping of an underwater robot. In *Proceedings of IEEE Oceans 94 Osates* (Brest, France, September 1994), IEEE.

[10] NASA Ames Research Center. *Photo of the Valles Marineris on Mars.* Courtesy of the Explorer Mission Public FTP Image Archives.

[11] Nishihara, H. Prism: a practical realtime imaging stereo matcher. In *SPIE* (1983), pp. 134–142.

[12] Nishihara, H. Practical realtime imaging stereo matcher. *Optical Engineering 23*, 5 (1984), 536–545. Also in *Readings in Computer Vision: issues, problems, principles, and paradigms.*

[13] Parthiot, F., and Denis, J. A better way to navigate on deep sea floors. In *Proc. IEEE Oceans* (1993), pp. 494–498.

[14] Steer, B., J. Klosk and, P. G., LeBlanc, L., and Schock, S. Towards Sonar Based Perception and Modelling for Unmanned Untethered Underwater Vehicles. In *Proceedings of the IEEE International Conference on Robotics and Automation* (Atlanta, GA, May 1993).

[15] Zheng, J., and Tsuji, S. Panoramic representation for route recognition by a mobile robot. *Int. Journal of Computer Vision 9*, 1 (1992), 55–76.