

A Constrained Optimization Approach to Globally Consistent Mapping

Ranjith Unnikrishnan and Alonzo Kelly

Robotics Institute

Carnegie Mellon University

Pittsburgh, USA

email: ranjith@andrew.cmu.edu, alonzo@ri.cmu.edu

Abstract

Mobile robot localization from large-scale appearance mosaics has been showing increasing promise as a low-cost, high-performance and infrastructure free solution to vehicle-guidance in man-made environments. The generation of the globally consistent high-resolution mosaics crucial to this procedure suffers from the same problem of loop-closure in cyclic environments that is commonly encountered in all map-building procedures. This paper presents a batch solution to the problem of reliably generating globally consistent mosaics at low computational cost, that simultaneously exploits the topological constraints among the observations and minimizes the total residual in observed features. An extension to a general scalable framework that facilitates an incremental online mapping strategy is also presented, along with results using simulated data and from real indoor environments.

1 Introduction

Mosaic-based localization [5] is a technique employing real-time imagery to track motion over a previously stored high-resolution image of a known environment. Naturally-occurring features in a composite mosaic built using images from a downward-looking camera encode position information corresponding to the robot pose when the feature observation was made. Template matching between observed and predicted views provide absolute position fixes to dampen the otherwise unbounded growth of errors originating in a primary position estimation system such as odometry. Previous papers on the subject of mosaic-based position estimation have discussed the feasibility and implementation issues [5] of this approach, and algorithms to construct the mosaic itself [10].

1.1 Challenges

Map construction requires a solution to the dual problems of recovering the motion history of the robot, as well as reconstructing the robot's environment based on observations made over the duration of the robot motion. In our application, the observations are in the form of images captured sequentially in one-dimensional swaths called *segments* by passing a calibrated downward looking camera over the floor. Recorded poses correspond to the 2D location and heading of each image as determined from

dead-reckoning. For high-speed navigation to be possible, the mosaic constructed using these images must satisfy the conditions of *local smoothness*, whereby the error in temporally adjacent images is bounded to some acceptably small value, and *global consistency*, whereby the position reported at a particular location becomes neither time nor path dependent, and is uniquely represented in the constructed model. These requirements are common to all map-building procedures, independent of the type of sensor used for the observations.

The observations made by the sensor in mosaic-based localization differ from those of sensors like laser range-finders and those used in beacon-based navigation systems, in that the features observed have low persistence in both temporal and spatial domains. The problem of generating accurate mosaics also typically involves the manipulation of thousands of images, and hence demands a time-efficient and scalable solution to be tractable.

1.2 Prior Work

Considerable literature exists on the field of image mosaicing and its applications, although only more recently have real-time and globally consistent solutions been addressed [9]. Real-time video mosaicing of the ocean floor for navigation, exploration and wreckage visualization has also received attention. Work by Fleischer et al. [2] used iterative smoother-follower techniques to reduce errors accumulated over an image chain, but no mention was made of the tractability of its extension to networks.

Lu and Milios [6] enforced global consistency in their work on map-building with laser range-scans, solving an overconstrained system of measurements to compute the required least-squares perturbation to absolute scan poses. Gutmann and Konolige [4] extended the approach to an online algorithm for mapping cyclic environments. A local registration step linked scans in a K-neighborhood of the last scan in time depending primarily on K. Loops were closed incrementally when a patch correlation scheme returned an unambiguous high match score, using the scheme of [6]. The complexity of this step, however, was $O(n^3)$ in the number of scans. Although marginal cost reductions could be achieved using sparse matrix techniques, it was still far too prohibitive for our purposes.

Work by Newman [8] on the problem of simultaneous

localization and mapping (SLAM) recognized that observations of relative positional relationships between landmarks had uncertainties that were uncorrelated to uncertainties in robot position. An overcomplete, and possibly inconsistent set of such observations formed the state vector in a *relative map*, that could be updated in time linear in the number of observations. When a consistent map was explicitly required, *geometric projection filtering* was done to repeatedly project the relative map estimates onto a constraint space described by the topological relationships between the observations. The constraint equations were chosen heuristically amongst a large but finite set, such that they spanned a large subset of the observations.

Previous work by the authors [10] described how the spanning set of constraint equations required to enforce global consistency could be easily extracted, and presented an algorithm for the same. Each *loop equation* in this set corresponded to an element of the fundamental cycle basis of the topological graph describing the network of observations. On the assumption of good local estimates from odometry, an efficient scalable solution to the problem of map-building in large cyclic environments was presented for observations with a short sensory horizon.

This paper presents a formulation that systematically relaxes the assumption of good initial estimates, while preserving the low time complexity of the algorithm in [10]. Our method employs the approach of total residual minimization between observations, while incorporating knowledge of prior state uncertainty and enforcing topological constraints to maintain map consistency. A framework for the solution that facilitates incremental online implementation is also proposed.

2 Residual at an Image Overlap




Figure 1: A linear mosaic of 5 images

Consider a linear sequence of images, such as that shown in Figure 1. The images overlap such that each image contains a portion of the scene in common with both the image before and after it in the sequence. Let each image be assigned with an index i with magnitude increasing in order of image capture. Let each image i be associated with its pose p_i^{i-1} expressed relative to the image preceding it in its sequence. The first image in each sequence has its pose expressed relative to the world frame w .

2.1 Non-temporally adjacent overlaps

Let i and j be two images that are not adjacent in sequence, and possibly even belong to different image sequences, but have features common to them. Such image pairs con-

stitute potential loop-forming overlaps in a network of images. Let at least two point features be identified in the images i and j . The residual of the m -th feature observed in both images can be expressed in terms of the relative pose between the images, as

$$z_{ij(m)} = p_{(i,m)}^i - R_j^i p_{(j,m)}^j - t_j^i$$

where $z_{ij(m)} \in \mathbb{R}^2$ is the feature residual, $p_{(i,m)}^i \in \mathbb{R}^2$ is the position of the m -th feature in the i -th image (subscript) represented in the coordinate frame of image i (superscript), R_j^i is the 2-by-2 rotation matrix relating the coordinate frame of image j with respect to that of image i , and $t_j^i \in \mathbb{R}^2$ is the corresponding displacement between the two frames.

On linearizing, this can be expressed as

$$\Delta z_{ij(m)} = -H_{(j,m)}^i \Delta p_j^i$$

where $H_{(j,m)}^i \in \mathbb{R}^{2 \times 3}$ is the Jacobian relating the change in position of feature m in image j with change in relative pose of image j with respect to image i , and $\Delta p_j^i \in \mathbb{R}^3$ represents the linearized change in relative pose.

Collecting the terms for feature residuals in non-temporally adjacent overlapping images, where $j \neq i + 1$, gives a system of the form

$$\Delta z_{ij} = H_{ij} \Delta x_{ij}$$

which can be solved iteratively by the least-norm left-pseudoinverse solution as

$$\Delta x_{ij} = (H_{ij}^T H_{ij})^{-1} H_{ij}^T \Delta z_{ij}$$

Here the matrix to be inverted for each pair of overlapping images is a 3-by-3 matrix, and hence its computation is inexpensive. Each such solution from feature matching provides an initial estimate for the relative pose between non-temporally overlapping images.

Note that this formulation leaves open the scope of using robust estimation techniques, such as iteratively re-weighted least squares (IRLS), commonly encountered in computer vision literature for outlier rejection. There exists a large body of literature on the topic, and further discussion is avoided with the understanding that the techniques can be suitably incorporated in each term involving feature residual minimization using appropriate weighting coefficients.

2.2 Temporally adjacent overlaps

For overlaps between images $i-1$ and i that are adjacent in sequence, the initial estimates of relative pose p_i^{i-1} are provided by dead-reckoning. We can however concatenate similar equations for residuals occurring at each temporally adjacent overlap, to give a system of equations in the linearized observer form

$$\Delta \underline{z} = H \Delta \underline{x}$$

where the concatenated change vector of change in residual is

$$\Delta \underline{z} = [\Delta z_{01(0)} \ \Delta z_{01(1)} \ \dots \ \Delta z_{ij(m)} \ \dots \ \Delta z_{N-1N(M)}]^T$$

and the state vector of change in relative pose is given by

$$\Delta \underline{x} = [\Delta \rho_1^0 \ \Delta \rho_2^1 \ \dots \ \Delta \rho_j^i \ \dots \ \Delta \rho_N^{N-1}]^T$$

2.3 Error Modeling

In registered non-temporally adjacent image pairs, the errors in estimates of feature residuals arise primarily from image noise which leads to inaccuracies in specification of point correspondences. This error can be treated as zero mean noise, uncorrelated between residual measurements. We model this error in a loop-closing overlap between images i and j as a diagonal matrix with entries σ^2 , our estimate of variance in pixel noise. Then the covariance matrix describing non-systematic error in estimate of relative pose between images i and j at loop-closure is given by

$$P_{ij} = \sigma^2 (H_{ij}^T H_{ij})^{-1}$$

For temporally adjacent images, say $i-1$ and i , the covariance $P_{i-1,i}$ can be estimated using a suitable choice of odometry model. Hence the *a priori* covariance of the state vector consisting of the concatenated set of relative poses between temporally adjacent images and relative poses between loop-closing non-temporally adjacent overlaps is given by the block diagonal matrix P_0 as

$$P_0 = \text{diag}(P_{01}, P_{12}, \dots, P_{ij}, \dots, P_{N-1N})$$

The matrix is block diagonal and non-singular, as errors between relative poses are inherently decoupled, and the uncertainty in estimate of non-temporally adjacent relative poses is derived solely from feature matching and is independent of poses of other images participating in the loop.

3 Loop analysis




Figure 2: Two-cycled network of 5 image segments

Consider the two cycled network shown in figure 2, composed of 5 overlapping segments. The image pairs (a,h), (i,j), (b,c), (d,e), (k,l) and (f,g) form 6 non-temporally adjacent overlapping image pairs. Let us assume that indices a through h are in increasing order of magnitude. As per our convention, every image pose is referenced with




Figure 3: Topology graph corresponding to a two-cycled network.

respect to its immediate predecessor having a smaller image index. Transforms represented by upper-case T refer to unknowns to be determined, but whose initial estimates are available either from odometry, in the case of temporally adjacent images, or from feature matching described earlier, in the case of overlaps between non-temporally adjacent images. A subscript index refers to the frame being referenced, and a superscript index refers to the frame with respect to which the reference is made.

Hence the homogenous transform relating coordinate frame of image b to that of image a is given by

$$T_b^a = T_{a+1}^a \cdot T_{a+2}^{a+1} \cdot T_{a+3}^{a+2} \dots T_b^{b-1}$$

Figure 3 shows an alternate representation of the network, termed a *topology graph*, having the property of having the same number of cycles as the original network of observations. We define a topology graph as a combination of a vertex set v , and an edge set e . Each vertex $v \in v$ is composed of a set of sub-nodes Φ . Each sub-node $n \in \Phi$ belonging to a vertex v represents a unique “footprint” of features associated with v , i.e. the images corresponding to each subnode of a vertex constitute non-temporally adjacent overlaps with common features. Edges between sub-nodes correspond to observed pose relationships.

The homogenous transform expressing a base frame relative to itself, as observed along any cycle in the graph, say $(a-b-c-d-e-f-g-h-a)$, should equal the identity. Mathematically,

$$T_a^a = T_b^a \cdot T_c^b \cdot T_d^c \cdot T_e^d \cdot T_f^e \cdot T_g^f \cdot T_h^g \cdot T_a^h = I_{3 \times 3}$$

Each equation is a function of the relative poses between images participating in the associated cycle. The set of such equations constitutes the constraint equations that are to be satisfied for consistency. In our example, there are three such equations corresponding to the three possible cycles $(a-b-c-d-e-f-g-h-a)$, $(a-i-j-k-l-f-g-h-a)$, and $(i-b-c-d-l-k-j-i)$ in the graph. As detailed in [10], the linearized form of any two of these three equations form an indepen-

dent basis in the space of constraints. The set of cycles that are used form an independent and complete cycle cover, termed in graph theory as a *fundamental cycle basis*.

4 Formulation as a Constrained Estimation

Let each equation in the chosen set of L constraint equations be of the form

$$g_l(X) = b_l$$

where X is the concatenated set of relative poses between temporally adjacent images and non-temporally adjacent overlapping images.

The maximum likelihood solution to X is then a minimizer of the Lagrangian

$$J(X, \lambda) = \frac{1}{2} \sum_{i,j} \sum_m z_{ij(m)}^T z_{ij(m)} + \frac{1}{2} \sum_{i,j} (\rho_j^i - \hat{\rho}_j^i)^T P_{ij}^{-1} (\rho_j^i - \hat{\rho}_j^i) + \sum_{l=1}^L \lambda_l^T [g_l(X) - b]$$

or in matrix representation as

$$J(X, \lambda) = \frac{1}{2} z(X)^T z(X) + \frac{1}{2} (X - \hat{X})^T P_o^{-1} (X - \hat{X}) + \lambda^T [g(X) - b]$$

The first term in J penalizes the sum of square residuals of feature observations in all overlaps, and the second term penalizes inverse covariance weighted perturbations of relative poses from their estimates. These two terms together constitute a full non-linear optimization that facilitates linearization of the estimation problem at every step using all feature observations and the current best estimate of the entire history of the robot trajectory. The third term containing the vector of Lagrange multipliers λ enforce the loop equations as hard constraints that are to be met by the posterior estimate upto first order accuracy.

The true X is a combination of the estimate value \hat{X} and a perturbation ΔX . Hence the cost function can be rewritten as

$$J(\Delta X, \lambda) = \frac{1}{2} z(\hat{X} + \Delta X)^T z(\hat{X} + \Delta X) + \frac{1}{2} \Delta X^T P_o^{-1} \Delta X + \lambda^T [g(\hat{X} + \Delta X) - b]$$

Linearizing and differentiating to extract the Lagrange conditions at a point of constrained minima gives

$$\frac{\partial J}{\partial \Delta X} = H^T z(\hat{X}) + H^T H \Delta X + P_o^{-1} \Delta X + G^T \lambda = 0$$

and

$$\frac{\partial J}{\partial \lambda} = g(\hat{X}) + G \Delta X - b = 0$$

which together yield the system

$$\begin{bmatrix} H^T H + P_o^{-1} & G^T \\ G & 0 \end{bmatrix} \begin{bmatrix} \Delta X \\ \lambda \end{bmatrix} = \begin{bmatrix} -H^T z(\hat{X}) \\ b - g(\hat{X}) \end{bmatrix}$$

Note that this is similar to the basic form of the Variable State Dimension filter [7], and that recursive partitioning

can be done to take advantage of the sparseness of the system matrix to be inverted. The matrix $H^T H + P_o^{-1}$ is block diagonal with a block size of 3, and can hence be inverted in time linear in the number of images. The solution to the system can thus be obtained directly by solving for λ in

$$G(H^T H + P_o^{-1})^{-1} G^T \lambda = [g(\hat{X}) - b] - G(H^T H + P_o^{-1})^{-1} H^T z(\hat{X})$$

and substituting back to solve for ΔX in

$$(H^T H + P_o^{-1}) \Delta X = -H^T z(\hat{X}) - G^T \lambda$$

The above update is performed iteratively until convergence is achieved as per some suitable criterion. The Levenberg-Marquardt procedure can be used to modify this update rule by adding a value μ to the diagonal elements of the linearized system before inversion. This value of μ is increased to smoothly switch towards gradient descent, or decreased to switch towards Gauss-Newton update, depending on whether the value of the cost function J increases or decreases respectively.

Computational Complexity: The state vector ΔX consists of $n_i - 1$ relative pose terms between n_i temporally adjacent images, and n_k terms corresponding to k non-temporally adjacent overlaps. If we denote the total number of relative poses in the state vector by n , the total number of loops by l , and the average number of features considered in each overlap by p (typically chosen to be 6), it can be shown that λ can be computed in $O(np + nl^2 + l^3)$ and ΔX is subsequently computable in $O(np + nl)$ time. The computational complexity of each iteration in the algorithm is hence bounded by $O(np + nl^2 + l^3)$ where typically $l \ll n$.

5 Constrained estimation as a geometric projection operation

Back-substitution of the value of λ yields a solution to ΔX of the form

$$\Delta X = -[I - KG][H^T H + P_o^{-1}]^{-1} H^T z(\hat{X}) + K[b - g(\hat{X})]$$

where

$$K = [H^T H + P_o^{-1}]^{-1} G^T [G[H^T H + P_o^{-1}]^{-1} G^T]^{-1}$$

The unconstrained posterior estimate of state obtained on only minimizing feature residuals is given by the weighted left pseudoinverse solution

$$\Delta X_r = -[H^T H + P_o^{-1}]^{-1} H^T z(\hat{X})$$

with posterior covariance given by $P_r = [H^T H + P_o^{-1}]^{-1}$

This implies that the constrained estimate of required perturbation in relative pose is of the form

$$\Delta X = [I - KG] \Delta X_r + K[b - g(\hat{X})]$$

which is the constrained estimator form of the Extended Kalman filter, with $K = P_r G^T [G P_r G^T]^{-1}$ equivalent to the Kalman gain matrix. The constraints consisting of the set of fundamental loop equations are equivalent to artificial observations with zero measurement noise [8].

Hence the Lagrange formulation of the map-building

problem is a constrained estimator that (a) projects the feature residual vector $z(\hat{X})$ into the column space of H to yield a least-squares unconstrained left pseudoinverse estimate of perturbation ΔX_r , and (b) projects the unconstrained estimate with covariance P_r onto a linearized hyperplane approximation of the non-linear constraint surface $g(X) = b$ evaluated at the unconstrained estimate.

6 Incremental Online Mapping

The ability to decouple the residual minimization and consistency enforcement steps using a Kalman filter suggests the possibility of an incremental map-building strategy. At each step of incorporation of a new sensor reading, there are a finite number of possible cases, depending on the change enforced in the topology graph.

Let there be n elements in the state vector X with covariance given by matrix $P^{(-)}$. Let a new observation be made, and the state element corresponding to the relative pose between the current observation and the temporally preceding observation be x_{n+1}

Case A: No feature is observed that is present in a non-temporally adjacent observation. In this case there is no change in the topology graph. The current observation is then simply an extension of a sequence of observations whose associated relative poses are independent. The state vector X is simply augmented by the new relative pose x_{n+1} , and the diagonal of $P^{(-)}$ by its covariance from an odometry model.

Case B: A feature is observed that is present in a non-temporally adjacent observation with no new loop formed. Here an additional node is added to the topology graph. In addition to augmenting the state vector and covariance matrix as in case A, another relative pose between the non-temporally adjacent overlapping images is added to X and its covariance matrix from feature matching is added as a diagonal element in $P^{(-)}$.

Case C: A feature is observed that is present in a non-temporally adjacent observation and a loop is closed. The state vector and covariance matrix are augmented as in Case B, and a subsequent step of constraint enforcement is carried out. We identify the newly added cycle in the topology graph as the shortest cycle containing the newly added edge, with the weight of each graph edge set proportional to the number of images lying in the edge. The state vector and covariance matrix can then be partitioned as

$$X = \begin{bmatrix} x_m \\ x_l \\ x_a \\ x_b \end{bmatrix} \quad P^{(-)} = \begin{bmatrix} P_{mm} & P_{ml} & 0 & 0 \\ P_{lm} & P_{ll} & 0 & 0 \\ 0 & 0 & P_{aa} & 0 \\ 0 & 0 & 0 & P_{bb} \end{bmatrix}$$

where

x_m consists of the n_m relative pose elements that are part of a previously assimilated cycle in the topology graph, i.e. part of a cycle other than the one newly created,

x_l consists of the n_l relative pose elements that are part of

both, a previously assimilated cycle in the topology graph, as well as the newly created cycle,

x_a consists of the n_a relative pose elements that are only part of the newly created cycle and none other, and

x_b consists of the n_b relative pose elements that are not part of any cycle.

Then P_{aa} and P_{bb} are both block diagonal elements, where each sub-block is a 3-by-3 matrix.

The constrained posterior estimate of the state vector can then be expressed as a projection onto a linearized subspace of the constraint equation. Let the constraint corresponding to new cycle be $g(X) = b$. Since g is not a function of state elements in x_m and x_b , its Jacobian with respect to the state can be written as

$$G = \begin{bmatrix} 0 & G_l & G_a & 0 \end{bmatrix}$$

The Kalman gain matrix is then given by

$$K = P^{(-)} G^T [G P^{(-)} G^T]^{-1}$$

and the matrix to be inverted can be written as

$$S = G P^{(-)} G^T = [G_l P_{ll} G_l^T + G_a P_{aa} G_a^T]$$

and can be shown to be non-singular. With careful implementation, the updated estimate can be computed in $O(n_m n_l + n_l^2 + n_a)$ time, which is minimized by the choice of shortest cycle. The covariance update step is more expensive at $O((n_m + n_l + n_a)^2)$ cost, but the step need only be performed at points of loop closure in the course of the robot's motion.

7 Results

Figures 4-5 show results with a toy problem of 122 images, using a vehicle simulated at 1/10th scale and moving at 0.5m/s. The actual images used in this experiment



Figure 4: Mosaic constructed using only initial pose estimates from dead-reckoning. Note the discrepancy in alignment of the circled feature at the crossover point (lower left corner).




Figure 5: Mosaic after consistency enforcement

are extracted from a larger aerial photograph. Initial estimates of vehicle pose are obtained from a simulated differential heading based dead-reckoning system. The odometers on the left and right wheels are each calibrated to 95% accuracy in wheel velocity measurements. Figure 4 illustrates the indication of a systematic rightward drift by dead-reckoning measurements, leading to distortion and inconsistency in the generated map on traversing a loop. Figure 5 shows the map after consistency was enforced with our algorithm.

Figure 6 shows results with a 170m long vehicle guideway in our indoor laboratory, mapped using our test vehicle with a total of 1836 images. The mosaic built using only initial pose estimates from dead-reckoning had 4 intersections that were initially unresolved.

8 Conclusions and Future Work

What we have presented is a map-building algorithm that enforces global consistency by exploiting the topology of observations made with sensors having small spatial field of view, while minimizing observed feature residuals. The algorithm is $O(np + nl^2 + l^3)$ in complexity, where typically the number of loops l is much smaller than the number of images n , and hence scales well to large datasets in the environments common to our application. In our work, we have generated reliable cyclic mosaics ranging from 80m to 800m in total length, both in laboratory and factory environments. The mosaics have been used by our test vehicles at speeds exceeding normal operating conditions for up to 40hr stretches without error requiring human intervention.

Ongoing work includes application of the algorithm in domains with other sensor modalities, including building dense maps with laser-range data, and eliminating the need to survey positions of laser reflectors for the installation of bearing-only guidance systems in factories. Consis-




Figure 6: Map of indoor vehicle guideway. Regions of the final mosaic at the intersections are highlighted.

tent odometry error modeling and incorporation of online auto-calibration techniques are also being pursued.

9 References

- [1] H.F. Durrant-Whyte, "Integration of Disparate Sensor Observations", *Proc. IEEE Intl. Conf. Robotics and Automation*, 1986, p. 1464
- [2] S.D. Fleischer, H.H. Wang, S.M. Rock, M.J. Lee, "Video Mosaicing along arbitrary vehicle paths", *Proc. Symposium on Autonomous Underwater Vehicle Technology*, 1996, pp. 293-299
- [3] J. Guivant and E. Nebot, "Optimization of the Simultaneous Localization and Map Building Algorithm for Real Time Implementation", *IEEE Trans. on Robotics and Automation*, June 2001, 17(3), pp. 242-257
- [4] J.-S. Gutmann and K. Konolige, "Incremental Mapping of Large Cyclic Environments", *Proc. Computational Intelligence in Robots and Automation (CIRA)* 1999, pp. 318-325
- [5] A. Kelly, "Mobile Robot Localization from Large-Scale Appearance Mosaics", *Intl. Journal of Robotics Research*, 19(11), 2000
- [6] F. Lu and E. Milios, "Globally consistent range scan alignment for environment mapping", *Autonomous Robots*, 4:pp.333-349, 1997
- [7] P.F. McLauchlan, "The Variable State Dimension Filter applied to Surface-based Structure from Motion", *CVSSP Technical Report VSSP-TR-4/99*
- [8] P. Newman, "On the Structure and Solution of the Simultaneous Localization and Map Building Problem", *Ph.D. thesis submitted in March 1999, Australian Centre for Field Robotics, Univ. of Sydney*.
- [9] H.S. Sawhney, S. Hsu and R. Kumar, "Robust video mosaicing through topology inference and local-to-global alignment", *Proc. European Conf. on Computer Vision (ECCV)*, Freiburg, Germany, vol.2, pp.103-119, June 1998.
- [10] R. Unnikrishnan and A. Kelly, "Mosaicing Large Cyclic Environments for Visual Navigation in Autonomous Vehicles", *Proc. IEEE Conf. on Robotics and Automation (ICRA)*, May 2002.