

# Localizing a Network of Non-Overlapping Cameras

Ali Rahimi      Trevor Darrell

{ali,trevor}@mit.edu

Computer Science and Artificial Intelligence Laboratory  
Massachusetts Institute of Technology  
Cambridge, MA

## Abstract

We describe a method for recovering the extrinsic calibration parameters of non-overlapping cameras in a multi-camera system. Non-overlapping camera networks can cover a much wider area than overlapping configurations, but calibrating them is more challenging. To compensate for the lack of overlap, our approach assumes that the calibration target follows a smooth trajectory over time, as described by a stochastic dynamics model. Our algorithm then tracks the target when it comes in the field of view of each camera, and searches for calibration parameters and target trajectories that are consistent with the observed tracks. We demonstrate the idea with a network of indoor wireless cameras.

## 1 Introduction

Networks of cameras have been shown to be an effective means for tracking people in surveillance settings. In these applications, the field of view of the cameras are usually made to overlap at least slightly, making it easier to track people as they move from the field of view of one camera to another, and in turn to recover the pose the cameras. Estimating these poses allows tracking to be performed in a globally consistent coordinate system. To instrument a wider area using fewer cameras, various authors have considered camera networks where the field of views of the cameras do not overlap [8, 12, 7]. Under these settings, cameras are placed in regions where interesting activity is expected to occur, and human motion is interpolated in uninstrumented regions.

We describe a calibration procedure for recovering the relative location and orientation of non-overlapping cameras in order to provide a globally consistent coordinate system for tracking. Our procedure requires a point calibration source that moves along the ground plane in smooth paths between the field of view of the

cameras. This path is not required to be straight, or to have constant velocity: we only assume that its evolution is consistent with a known stochastic dynamical model. Our main contribution is to show that a model of dynamics for the motion of the calibration target can compensate for the lack of overlap between the cameras.

When there is overlap between the fields of view of cameras, the appearance of the calibration target in a region of overlap provides information about the relative pose of the cameras. With enough such appearances, the relative orientation of the cameras can be recovered by taking advantage of stereopsis. Without information about the dynamics of the target, observing the calibration target where the cameras do not overlap only provides information about its pose relative to the camera that observed it, but not between the pose of two cameras.

For each camera, our algorithm recovers the two translation parameters and the rotation parameter in the ground plane. We rely on existing single-camera calibration techniques to recover the remaining extrinsic and intrinsic parameters before our algorithm is applied. The calibration task can be separated into these two tasks because single-camera calibration procedures can recover the pose of the camera relative to a local patch on the ground plane, up to a rotation and translation in the ground plane, in addition to recovering the intrinsic camera parameters. After such a procedure, our algorithm can be applied to globally align the coordinate system of the patches corresponding to each camera.

To perform this global alignment step, one could require that the target moves at constant velocity and in a straight line when it leaves the field of view of a camera. An intuitive approach would then be to have each camera track the target when the target passes through the camera's field of view, and estimate the target's velocity as it leaves the field of view. The travel time between the fields of view of two cameras, along with the

target’s estimated velocity can be used to estimate the distance between the two fields of views. When enough such distances are obtained between pairs of cameras, they can be combined to recover the relative location and orientation of all the cameras in the network [7].

Our algorithm relaxes the assumption that the target moves in straight lines and at constant velocity. This allows the calibration target to be more casually carried by a person moving along curved paths. Our algorithm searches for camera calibration parameters and a smooth target trajectory that are consistent with the tracks observed by each camera while the target was in its field of view. This search is framed as finding the maximum *a posteriori* (MAP) camera calibration parameters and target trajectory, with a prior that prefers smooth target trajectories.

We carry out the MAP estimation of calibration parameters and trajectories using the Newton-Raphson nonlinear optimization algorithm. An alternative approach might be to iteratively estimate camera calibration parameters, then interpolate the trajectory until convergence. This coordinate ascent optimization strategy is attractive because each of these steps can be carried out very quickly and is easy to implement. However, in Section 7 we show that these iterations will converge arbitrarily slowly as observation noise is *reduced*. Hence the convenient coordinate ascent strategy does not work for this problem. This has implications for coordinate ascent strategies in other settings as well.

Concurrently recovering calibration parameters and trajectories is reminiscent of Structure from Motion (SFM) and Simultaneous Localization and Mapping (SLAM) [3]. In SFM and SLAM, a moving camera measures the image-plane coordinate of scene features and recovers the trajectory of the camera and the 3D location of the scene features. Our problem is converse: A set of stationary cameras observe the image-plane coordinate of a moving feature (the target), and the goal is to recover the pose of the cameras and the trajectory of the feature. A coordinate ascent strategy that iterates between computing structure assuming motion, and vice version has been proposed in the SFM literature. Following the analogy, our non-convergence proof demonstrates that coordinate ascent cannot solve the Structure from Motion problem when only one feature is observed in each frame.

We have implemented our algorithm on a network of wireless PDAs equipped with cameras (see Figure 1). A real-time person tracker runs on each PDA which report trajectories to a central server that implements our algorithm.




Figure 1: We use Compaq IPAQs as wireless camera nodes in our network. The IPAQs are mounted on the ceiling, with their camera image plane parallel to the ground plane.

## 2 Related Work

Many methods explicitly designed for calibrating networks of cameras rely on overlapping fields of view [14, 15, 9, 6]. The method of Javed et al. [7] is one exception, where an idea similar to velocity extrapolation is used to find the projection of the field of view lines of one camera onto the field of view of cameras. Knowledge of these projections is tantamount to recovering calibration parameters, but this method requires people to walk in a straight line and at constant velocity outside camera fields of view. Our work is most closely related to [5], which shows how to calibrate a network of non-overlapping cameras using distant objects (stars) to recover orientation, and nearby objects (airplanes) to recover relative position. Whereas [5] relies on strict constraints on the motion of the calibration targets (the motion of stars is parabolic, nearby targets must move in a line with constant velocity), our targets may move freely. Furthermore, we require only one type of target, and this target may be the same type of target the network will ultimately track.

Various researchers have addressed the problem of maintaining consistent identity between multiple targets as they exit one field of view and enter another [8, 12] in non-overlapping multi-camera systems. These techniques focus on recovering the topology of the network, and do not attempt to recover metric pose of the cameras.

Various SFM and SLAM techniques, particularly those based on the Kalman Filter [10, 4, 2, 1, 13], have a dynamics model for the motion of the camera. The major difference between our problem and SLAM or SFM is that due to the lack of overlap, at most one camera can observe the target at any time. This corresponds to an SFM setup where the moving camera observes only one scene feature in each frame. Because this is rarely the case in SFM, the main advantage of a motion model is additional robustness against noisy measurements. In our situation, however, a motion model compensates for the lack of simultaneous measurement, and without it, calibration would not be possible. We cast our problem in an optimization framework similar to bundle adjustment in SFM [16].

## 3 Single-Camera Calibration

Before running our algorithm, each camera is calibrated so that it can map the image coordinate of the target to a local coordinate system laid on the ground-plane. The homography required to perform this mapping can be estimated for each camera individually [17]. For example, by laying a calibration object of known

size on the ground plane, the focal length, height, pitch, and yaw of the camera can be estimated. This local homography can be used to rectify the image location of the target in each camera so that the camera appears to be virtually fronto-parallel to the ground plane. This leaves a translation and roll in the ground plane to be estimated by our algorithm.

In the remainder of this paper, we presume that the homography for each camera has been recovered up to a rotation and a translation in the ground plane, and focus on aligning the local ground-plane coordinate system of each camera to the global ground plane coordinate system. In our experiments, we used overhead cameras, so that only the height and the focal length of the cameras needed to be estimated during the single-camera calibration phase.

## 4 Global Alignment

To globally calibrate the network, we jointly find the maximum *a posteriori* (MAP) estimate of the remaining calibration parameters and the trajectory of a point calibration target that moves smoothly in the network. Each camera tracks the point calibration target as it passes through its field of view. MAP estimation amounts to searching for calibration parameters and a trajectory that are consistent with the observations made by each camera and a dynamics model for the motion of the target. Since these can only be recovered up to a global rotation and translation in the ground plane, we fix the parameters of one of the cameras.

We assume that the state evolves according to linear Gaussian Markov dynamics. Define  $x_t$  to be the state of the target at time  $t$ . This state contains information about the location, velocity, or any other dynamic state of the target, and evolves with:

$$x_{t+1} = \mathbf{A}x_t + \nu_t, \quad (1)$$

where  $\nu_t$  is a zero-mean Gaussian random variable with covariance  $\Sigma_\nu$ .

In our experiments, we let  $x_t = [u_t; \dot{u}_t; v_t; \dot{v}_t]$ , where  $(u_t, v_t)$  is the ground-plane location of the target and  $(\dot{u}_t, \dot{v}_t)$  is its ground-plane velocity. We use the model parameters

$$\mathbf{A} = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad (2)$$

so that each  $x_{t+1}$  adds the velocities in  $x_t$  to the positions in  $x_t$ , and nudges the old velocities by Gaussian noise. We use a diagonal  $\Sigma_\nu$  whose position components have much smaller value than its velocity components (by a few orders of magnitude), so that the

velocities follow Brownian dynamics, and the resulting poses are nudged only by a small amount of Gaussian noise. Equation (1) defines a prior  $p(x)$  over a state trajectory  $x = \{x_t\}_{t=1}^T$ .

To extract the location components in  $x_t$ , we can multiply  $x_t$  by  $\mathbf{C}$ :

$$\mathbf{C} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}.$$

After taking into account the homography discussed in the previous section, camera  $i$  reports the location of the target in its own ground-plane coordinate system whenever the target is in its field of view. Let  $\theta^i$  and  $p^i$  denote the unknown rotation and translation of camera  $i$  with respect to the global ground-plane coordinate system. Let the index set of observations  $\mathcal{Z} = \{(t, i)\}$  be a collection of time index and camera pairs, where  $(t, i) \in \mathcal{Z}$  iff camera  $i$  sees the target at time  $t$ . Let  $\mu^i = [p^i; \theta^i]$  be the parameters of camera  $i$ , and define  $\mu = [\mu^1; \dots; \mu^N]$  to be the collection of all the camera parameters. Letting  $\mathbf{R}^i(\theta)$  denote the rotation matrix corresponding to a rotation of  $\theta^i$ , the location reported by a camera when the target is in its field of view is:

$$y_t^i = \pi(x_t; \mu^i) + \omega_t = \mathbf{R}^i(\mathbf{C}x_t - p^i) + \omega_t, \quad (3)$$

where  $\pi$  is the mapping between global and local camera coordinate systems. We assume  $\omega_t$ , the observation noise corrupting each measurement, is iid zero-mean and Gaussian with covariance  $\sigma_y^2 \mathbf{I}$ . Equation (3) defines a likelihood  $p(y|x, \mu)$  over trajectories and camera calibration parameters.

Any configuration of  $x$  and  $\mu$  will produce the same tracking measurements as the same configuration arbitrarily rotated and translated [16]. This ambiguity can be fixed by setting the parameters of one of the cameras to some arbitrary value. We use a prior that favors configurations where the parameters of the first camera are 0.

Using the likelihood defined by Equation (3) and the prior over  $x$  defined by Equation (1) and the gauge-fixing prior for the first camera, the most *a posteriori* probable trajectory and calibration parameters can be obtained by performing a nonlinear least squares optimization over  $\mu$  and  $x$ :

$$(x^*, \mu^*) = \arg \min_{x, \mu} \sum_{(t, i) \in \mathcal{Z}} \frac{1}{\sigma_y^2} \|y_t^i - \pi^i(x_t; \mu)\|^2 + \sum_{t=1}^T \|x_t - \mathbf{A}x_{t-1}\|_{\Sigma_\nu}^2 + \|\mu^1\|^2. \quad (4)$$

The first term in the cost function favors trajectories and parameters that would have generated the observed trajectory snippets. The second term favors

trajectories that are consistent with the given dynamics by favoring trajectories where  $x_t$  can be predicted from  $x_{t-1}$  using the transition matrix of the dynamics model, and penalizing the prediction error by with a mahalanobis distance defined by the covariance of the driving noise. The third term fixes the gauge by setting the parameters of the first camera to zero.

To find the optimal  $\mu$  and  $x$ , we use Newton-Raphson with a first order approximation to the Hessian. This involves it iteratively linearizing the non-linearity inside the first term of Equation (4), transforming it to a linear least squares problem. This least squares problem involves a minimization over  $3 * N + 4 * T$  variables with  $|\mathcal{Z}|$  rows, where  $N$  is the number of cameras and  $T$  is the number of time steps in the trajectory to be estimated. It is also sparse, with  $O(|\mathcal{Z}|)$  nonzero elements, and so can be solved efficiently using standard sparse least squares solvers. The appendix derives the quantities needed in each Newton-Raphson step.

## 5 Synthetic Results

Our goal with this synthetic experiment was to see if a disparity between the dynamics model and actual target behavior had adverse effects on trajectory estimation and calibration. We simulated a point target traveling in a square environment with elastic walls. The target's trajectory was generated by Brownian motion that deflected off of walls. The resulting trajectory was smoothed to yield the setup shown in Figure 2(a). The synthetic cameras were front-parallel to the ground plane, and measured the location of the target without noise within their coordinate system (i.e. up to a rotation and translation of the true location of the target).

To define the dynamics model, we used the  $\mathbf{A}$  matrix of equation (2), and set  $\Sigma_\nu = \begin{bmatrix} 10^{-4} & 1 & 10^{-4} & 1 \end{bmatrix}$ . Notice that this dynamics model is different from the model used for generating the synthetic trajectory, as it does not model the bouncing behavior near walls, the final smoothing step, or the underlying Brownian motion.

The Newton-Raphson iterations were initialized with all trajectory points and camera parameters at the origin, pointing to the right. Figures 2(b-c) show an intermediate iteration and the converged estimate.

The estimated locations are wrong by an average of 0.03 size units, or 1.4% of the size of the environment. These experiments show that when there is no observation noise, the sensor parameters and trajectories are recovered very accurately, even if the target's true dynamics don't match those used in estimation.

## 6 Real Data

We implemented our algorithm on a network of wireless PDAs equipped with cameras. The PDAs were mounted on the ceiling in an indoor lab environment. The camera image planes are approximately parallel to the ground-plane so that computing the local ground-plane location of the person does not require any calibration beyond finding the focal length of the cameras. A real-time person tracker runs on each PDA and reports to a base station the time-stamped location of a person with respect to the camera’s ground plane coordinate system every 250 ms. The person tracker uses background subtraction to extract the target and clusters the foreground pixels to compute the person’s location. The individual trackers do not need to filter or smooth the data, as the dynamics model automatically regularizes trajectories during the optimization procedure.

In our first experiment, we installed 4 cameras in an open area in our building. The fields of view of the cameras were about 1.5 meters on each side, and the cameras were 3-4 meters apart. One person walked between the cameras at varying velocities and served as the calibration target. Figure 3(a) illustrates the setup. The ground truth camera parameters were found by measuring the location and orientation of every camera by hand. To render the trajectory of the target in this figure, we set  $\mu$  to these ground truth camera parameters and optimized (4) for the trajectory only.

We used the same parameters for the dynamics model as in the synthetic case. The initial iterate for the optimizer was also the same as in the synthetic case. We set the observation noise  $\sigma_y^2$  to be very small, a factor of  $10^5$  smaller than the driving noise of the velocity. Figures 3(b-c) show the recovered trajectory and sensor locations. On average, the sensor were misplaced by 28 cm from the locations measured by hand. Cameras `ipaq3` and `ipaq10` are off by 50 cm. The rotation of `ipaq3` is off by  $8^\circ$ . That of `ipaq8` by  $9^\circ$  degrees. `ipaq10`’s rotation is off by  $0.7^\circ$ .

In the second experiment, we instrumented the hallway to the left of the open area with 2 additional cameras. The setup and results appear in Figure 4(a). The only way to reach `ipaq5` and `ipaq6` was to take a sharp right below `ipaq7`, outside its FOV. Because no camera ever witnessed this sharp right, and because straight trajectories are slightly favored by our dynamics model over turns, the estimated configuration did not recover the turn. Adding `ipaq9` (Figure 4(b)) provided enough information to recover the turn. This is because `ipaq5`, `ipaq7`, and `ipaq9` form a triangle, and the angle 5-7-9 becomes constrained.

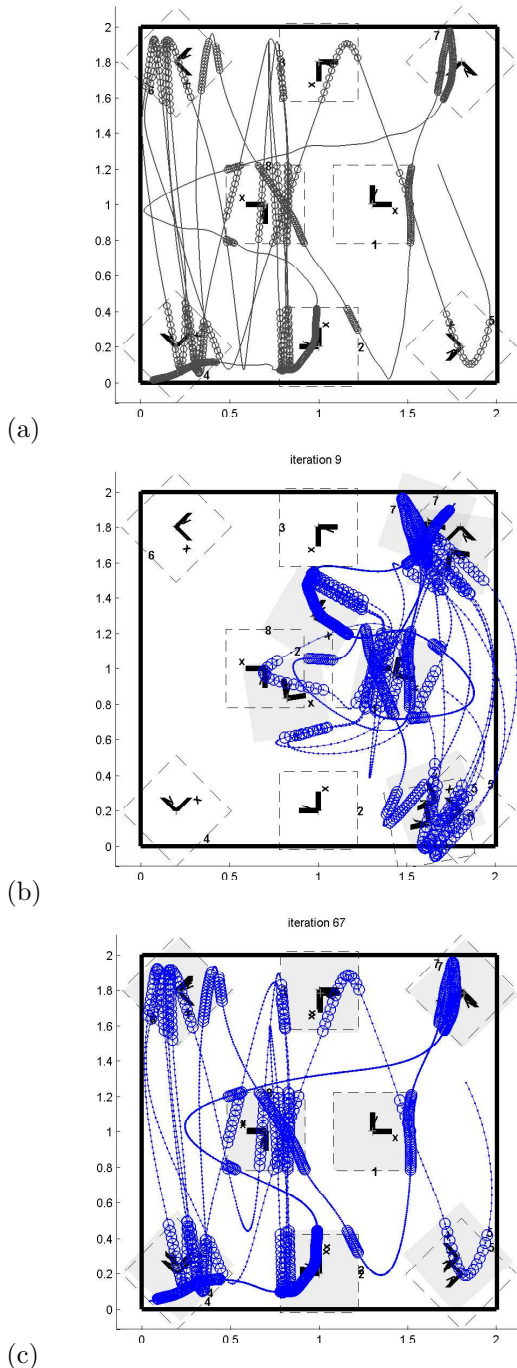


Figure 2: (a) 2,000 steps of a synthetic trajectory. True camera fields of view are depicted as dashed squares. Circles along the trajectories indicate time steps in which a synthetic camera observes the target’s location relative to its coordinate system. (b) After 9 iterations of the optimization procedure. The gauge is fixed by fixing sensor 1 at its true location and orientation. The blue (dark) path denotes the recovered trajectory. Gray squares are the recovered sensor fields of view. (c) Convergence after about 65 iterations. The sensor locations are estimated correctly.

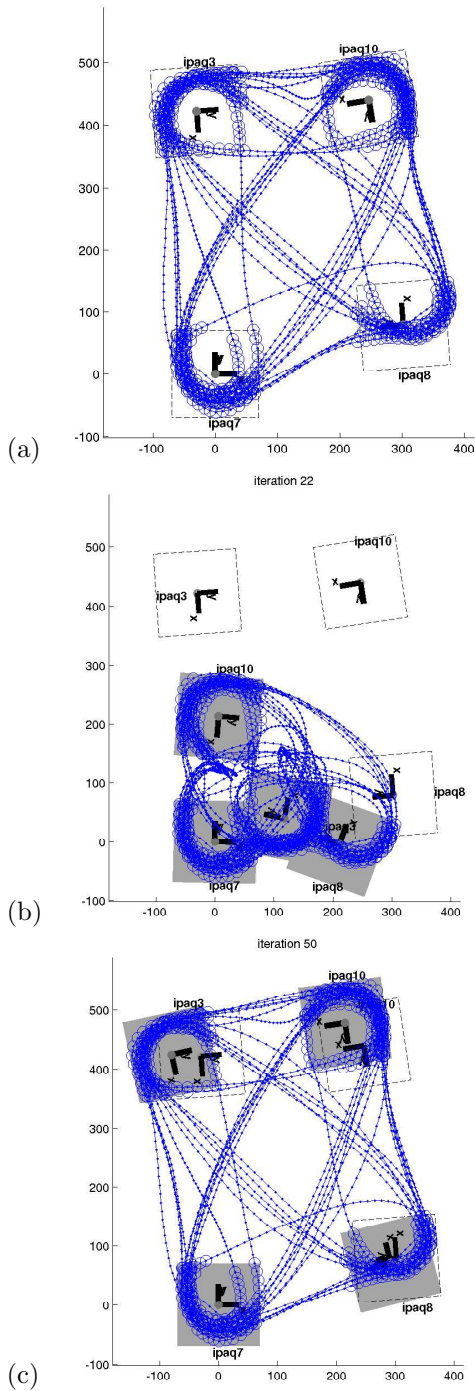


Figure 3: (a) The trajectory for the first experiment. Axes are labeled in centimeters. The coordinate system is fixed on ipaq7. (b) After 22 iterations. (c) Convergence after about 50 iterations.

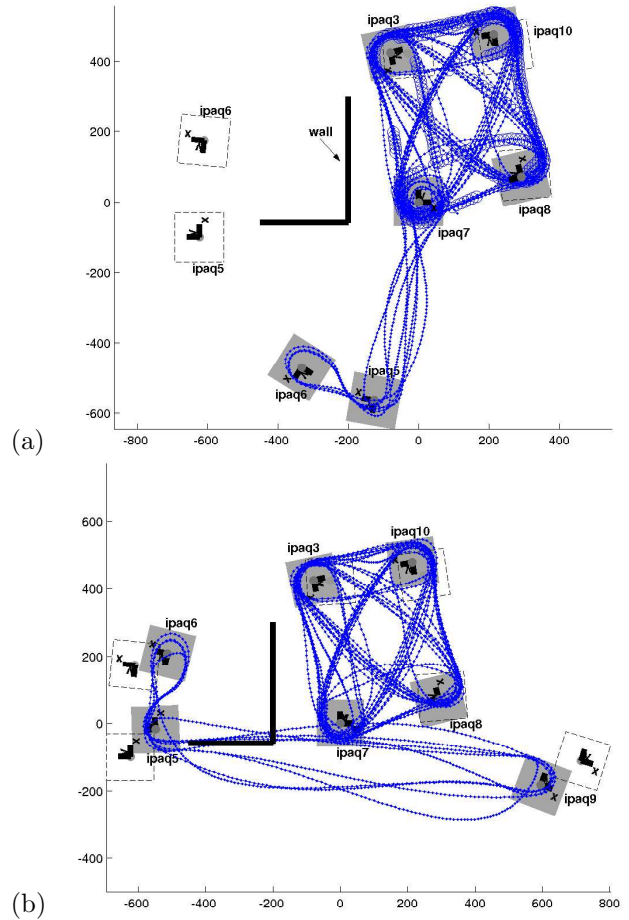


Figure 4: (a) A Wall forces people to take a sharp turn outside the FOV of ipaq7. This turn is never witnessed by any camera, so the algorithm does not deduce its existence. (b) Although ipaq9 doesn't observe the turn directly, its presence provides enough information to determine that ipaq7 and 6 cannot be below ipaq7.




Figure 5: Recovering curved trajectories.

In contrast to the velocity extrapolation idea discussed earlier, our dynamics model allows curved trajectories. Figure 5 shows that our method works even when people take a sharp turn when entering the FOV of ipaq3. Using velocity extrapolation, the distance between ipaq7 and ipaq3 was found to be 518 cm, whereas in reality, it was only 423 cm. Our method recovered this distance more accurately to be 415 cm.

## 7 Coordinate Ascent Cannot Recover Trajectories and Sensor Locations

It was tempting for us to attempt to recover the trajectory and calibration parameters by iteratively solving for one and then the other. In this section, we show that when there is no overlap between the fields of views of the cameras, a coordinate descent strategy will converge arbitrarily slowly as the observation noise is reduced. A corollary of the result is that sufficient overlap between the camera’s field of view will cause coordinate descent to converge.

*We demonstrate this convergence property only in the case when the camera angles are known to be zero. An early version of this proof assumed this and was cited by the CVPR version of this paper. When we attempted to extend its result to the case that handled angles, we found that the proof strategy is not viable.*

The coordinate ascent strategy has been suggested in the structure-from-motion literature, where starting from an initial guess about the camera trajectory, a scene structure can be easily obtained, and given this recovered scene structure, the camera trajectory can be easily estimated. In keeping with our analogy between Structure from Motion and the network calibration problem, our non-convergence result implies that

if the SFM camera only sees one feature at each time step, the structure of the scene cannot be recovered by coordinate descent.

Each iteration of coordinate descent first minimizing  $f$  over  $X$ , assuming  $\mu$  is fixed, and then over  $\mu$  by fixing  $X$  to the result of the minimization over  $X$ :

$$\mu^{(t+1)} = \mathcal{M}_\mu(x^{(t)}) = \arg \min_\mu f(\mu, x^{(t)}) \quad (5)$$

$$x^{(t+1)} = \mathcal{M}_X(\mu^{(t+1)}) = \arg \min_X f(\mu^{(t+1)}, x^{(t)}). \quad (6)$$

This is a very attractive optimization strategy because both optimizations are easy to carry out. Estimating the trajectory is a straightforward interpolation task that takes time linear in the length of the trajectory. Given the trajectory, estimating sensor locations is a simple alignment problem where the parameters of a camera can be estimated independently of the parameters of other camera.

Because each iteration in (5)-(6) can only decrease the value of  $f$ , we might expect that these iterations would converge to a minimum of  $f$ . We show that when the measurement noise is small compared to the uncertainty of the dynamics model ( $\sigma_{y|x}^2 \ll \text{tr} \Sigma_w$ ), each iteration of (5)-(6) produces an infinitesimal change in  $X$ , so that the algorithm will not converge in a practical amount of time if  $\sigma_{y|x}^2 \ll \Sigma_w$ .

Providing guarantees on the derivatives of  $\mathcal{M}_\mu$  and  $\mathcal{M}_X$  everywhere will result in guarantees on the maps  $\mathcal{M}_\mu$  and  $\mathcal{M}_X$  everywhere. Assuming the maps  $\mathcal{M}_\mu$  and  $\mathcal{M}_X$  are differentiable, we can linearize them the optimal point  $(x^*, \mu^*)$  using the mean value theorem:

$$\mu^{(t+1)} = \mathcal{M}_\mu(x^{(t)}) = \mathcal{M}_\mu(x^*) + \frac{\partial \mathcal{M}_\mu}{\partial x} \Big|_{x_0} (x^{(t)} - x^*) \quad (7)$$

$$x^{(t+1)} = \mathcal{M}_x(\mu^{(t+1)}) = \mathcal{M}_x(\mu^*) + \frac{\partial \mathcal{M}_x}{\partial \mu} \Big|_{\mu_0} (\mu^{(t+1)} - \mu^*) \quad (8)$$

The mean value theorem guarantees that the linearization is exact for an appropriate choice of  $x_0$  and  $\mu_0$ . In what follows, we use the shorthand  $\dot{\mathcal{M}}_x = \frac{\partial \mathcal{M}_x}{\partial \mu}(\mu_0)$  and  $\dot{\mathcal{M}}_\mu = \frac{\partial \mathcal{M}_\mu}{\partial x}(x_0)$  for simplicity.

Using the facts that the optima are also fixed points of the iteration ( $\mathcal{M}_\mu(x^*) = \mu^*$ , and  $\mathcal{M}_x(\mu^*) = x^*$ ), we can write the map  $\mu^{(t)} \rightarrow \mu^{(t+1)}$  in linear form by substituting (7) into (8):

$$\mu^{(t+1)} = \mathcal{M}_\mu(\mathcal{M}_x(\mu^{(t)})) = \mu^* + \dot{\mathcal{M}}_\mu \dot{\mathcal{M}}_x (\mu^{(t)} - \mu^*). \quad (9)$$

Subtracting  $\mu^{(t)}$  from both sides and taking norms gives:

$$\|\mu^{(t+1)} - \mu^{(t)}\| \leq \|I - \dot{\mathcal{M}}_\mu \dot{\mathcal{M}}_x\| \|\mu^{(t)} - \mu^*\|. \quad (10)$$

This expression provides a bound on the change in the trajectory iterate as a function of the distance of the current estimate to the optimal trajectory  $\mu^*$ .

However, we will demonstrate that as  $\sigma_{y|x}^2 \ll \Sigma_w$ ,  $\|I - \dot{\mathcal{M}}_\mu \dot{\mathcal{M}}_x\| \rightarrow 0$ , so that the iteration will not perturb  $\mu^{(t)}$  by a significant amount, and the coordinate ascent is infinitely slow.

When the observation noise is very small, a small change by  $dp^i$  of the location of sensor  $i$  will cause a shift by  $dp^i$  in the trajectory points observed by the sensor. Similarly, a shift by  $dx$  in the trajectory will cause a shift in each sensor by the average shift incurred in the trajectory segments it observed.

Consider a two-camera example with six trajectory steps. Suppose camera 1 observed trajectory steps 1,2,5,6 and camera 2 observes trajectory steps 3 and 4. The change in trajectory samples due to a shift  $dp$  in camera positions is:

$$dx = \dot{\mathcal{M}}_x(\mu)d\mu \approx \begin{bmatrix} I & 0 \\ I & 0 \\ 0 & I \\ 0 & I \\ I & 0 \\ I & 0 \end{bmatrix} \begin{bmatrix} dp^1 \\ dp^2 \end{bmatrix}. \quad (11)$$

and the change in camera position due to a shift  $dx$  in trajectory is

$$dp = \dot{\mathcal{M}}_\mu(x)dx \approx \begin{bmatrix} I/4 & I/4 & 0 & 0 & I/4 & I/4 \\ 0 & 0 & I/2 & I/2 & 0 & 0 \end{bmatrix} dx \quad (12)$$

So that  $\dot{\mathcal{M}}_\mu \dot{\mathcal{M}}_x \approx I$

## 8 Conclusion

We have shown how to calibrate a network of non-overlapping cameras. Our solution is framed as a joint MAP estimation camera pose parameters and the trajectory of a calibration target. The prior on the trajectory is a linear Gaussian Markov chain, which allows for both nonlinear paths and speed changes when the target is outside the field of view of the cameras. Information about the target's dynamics allows the system to reason about the behavior of the target when it is not visible, compensating for the lack of overlap between the cameras. We demonstrated our system with a network of battery-operated cameras that are easy to deploy.

## References

[1] N. Ayache and O. Faugeras. Maintaining representations of the environment of a mobile robot. *IEEE Tran. Robot. Automat.*, 5(6):804–819, 1989.

[2] A. Azarbayejani and A. Pentland. Recursive estimation of motion, structure, and focal length. *Pattern Analysis and Machine Intelligence (PAMI)*, June 1995.

[3] J. A. Castellanos, J. M. M. Montiel, J. Neira, and J. D. Tardis. The SPmap: A probabilistic framework for simultaneous localization and map building. *IEEE Transactions on Robotics and Automation*, 15:948–953, October 1999.

[4] A. Chiuso, P. Favaro, H. Jin, and S. Soatto. Structure from motion causally integrated over time. *Pattern Analysis and Machine Intelligence (PAMI)*, 24(4):523–535, April 2002.

[5] R. B. Fisher. Self-organization of randomly placed sensors. In *European Conference on Computer Vision (ECCV)*, volume 4, pages 146–160, 2002.

[6] B. K. P. Horn. Relative orientation. *International Journal of Computer Vision (IJCV)*, 4:59–78, January 1989.

[7] O. Javed, Z. Rasheed, O. Alatas, and M. Shah. *KNIGHT<sup>TM</sup>*: A real time surveillance system for multiple overlapping and non-overlapping cameras. In *International Conference on Multimedia and Expo*, July 2003.

[8] V. Kettner and R. Zabih. Counting people from multiple cameras. In *ICMCS*, volume 2, pages 267–271, 1999.

[9] S. Khan and M. Shah. Consistent labeling of tracked objects in multiple cameras with overlapping fields of view. *Pattern Analysis and Machine Intelligence (PAMI)*, 25(10), October 2003.

[10] P. F. McLauchlan. A batch/recursive algorithm for 3D scene reconstruction. *Computer Vision and Pattern Recognition (CVPR)*, 2:738–743, 2000.

[11] T.P. Minka. Old and new matrix algebra useful for statistics. Technical report, Media Lab, <http://www.media.mit.edu/~tpminka/papers/matrix.html>, 2001.

[12] H. Pasula, S. J. Russell, M. Ostland, and Y. Ritov. Tracking many objects with many sensors. In *International Joint Conferences on Artificial Intelligence (IJCAI)*, pages 1160–1171, 1999.

[13] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In *Uncertainty in Artificial Intelligence*, 1988.

- [14] C. Stauffer and K. Tieu. Automated multi-camera planar tracking correspondence modeling. In *Computer Vision and Pattern Recognition (CVPR)*, volume 1, June 2003.
- [15] G. P. Stein, R. Romano, and L. Lee. Monitoring activities from multiple video streams: Establishing a common coordinate frame. Technical Report AIM-1655, MIT AI Lab, 1999.
- [16] B. Triggs, P. McLauchlan, R. Hartley, and A. Fitzgibbon. Bundle adjustment – a modern synthesis. In W. Triggs, A. Zisserman, and R. Szeliski, editors, *Vision Algorithms: Theory and Practice*, LNCS, pages 298–375. Springer Verlag, 2000.
- [17] R.Y. Tsai. Multiframe image point matching and 3-d surface reconstruction. *Pattern Analysis and Machine Intelligence (PAMI)*, 5(2):159–173, March 1983.

## A Optimization Procedure

This section derives the Newton-Raphson steps required to solve (4). We rewrite Equation (4) in nonlinear least squares form, linearize the nonlinearity, and describe one Newton-Raphson iteration.

The dynamics prior term  $\sum_{t=1}^T \|x_t - \mathbf{A}x_{t-1}\|_{\Sigma_\nu}^2$  can be written as a quadratic form in terms of  $x$ , the column vector of stacked states. Denote the Cholesky factor of the inverse covariance  $\Sigma_\nu^{-1}$  of the driving noise in the dynamics model by  $\Sigma_\nu^{-\frac{1}{2}}$ .<sup>1</sup> Define the matrix  $\mathbf{G}$  whose  $t$ th row is:

$$[\mathbf{G}]_t = \left[ 0 \dots, \Sigma_\nu^{-\frac{1}{2}} \mathbf{A}, -\Sigma_\nu^{-\frac{1}{2}}, 0 \dots \right] \quad (13)$$

where the zeros pad the matrix to align its non-zero components with  $x_t$  and  $x_{t+1}$ . Then the dynamics prior term can be rewritten as  $\sum_{t=1}^T \|x_t - \mathbf{A}x_{t-1}\|_{\Sigma_\nu}^2 = x^T \mathbf{G}^T \mathbf{G} x$ .

Equation (4) can now be rewritten as

$$(x^*, \mu^*) = \arg \min_{x, \mu} r(x, \mu)^T r(x, \mu),$$

with

$$r(\mu, x) = \begin{bmatrix} r_y \\ r_x \\ r_\mu \end{bmatrix} = \begin{bmatrix} \sigma_y^{-1} (\mathbf{R}^i (\mathbf{C}x_t - p^i) - y_t^i) \\ \vdots \\ \mathbf{G}x \\ (\mu^1 - \mu_0) \end{bmatrix}.$$

<sup>1</sup>The Cholesky factor of a positive semidefinite matrix  $\mathbf{A}$  is an upper triangular matrix  $\mathbf{C}$  such that  $\mathbf{A} = \mathbf{C}^T \mathbf{C}$ .

The column vector  $r(\mu, x)$  is partitioned into three sections, corresponding to the three terms in (4). Each measurement  $(t, i) \in \mathcal{Z}$  introduces two elements into  $r_y$ .

To optimize a nonlinear least squares cost function such as  $r^T r$ , Newton-Raphson requires the Jacobian  $\mathbf{J}$  of  $r$ . Newton-Raphson maps an iterate  $\chi^{(t)} = [\mu^{(t)}; x^{(t)}]$  to the next iterate by solving a linear least-squares problem:

$$\chi^{(t+1)} = \chi^{(t)} - \arg \min_{\delta} \|\mathbf{J}\delta - r\|^2. \quad (14)$$

The process of computing  $r$ ,  $\mathbf{J}$ , and applying equation (14) is repeated until  $\chi$  converges to a fixed point.

For completeness, we derive  $\mathbf{J}$  here. It is block structured and very sparse:

$$\mathbf{J} = \nabla r(\mu, x) = \begin{bmatrix} \mathbf{J}_\mu & \mathbf{J}_{\mu\mathbf{x}} \\ \mathbf{0} & \mathbf{J}_\mathbf{x} \\ \mathbf{I} & \mathbf{0} \end{bmatrix}. \quad (15)$$

The left block column of  $\mathbf{J}$  corresponds to differentiating  $r$  with respect to  $\mu$ , and its right block column corresponds to differentiating it with respect to  $x$ .

If the  $z$ th measurement in  $r_y$  came from camera  $i$ , the derivative of  $r_y$  with respect to parameters  $p^i$  and  $\theta^i$  of this camera introduces a block row into  $\mathbf{J}_\mu$  at location  $(z, i)$ .

$$[\mathbf{J}_\mu]_{z,i} = -\sigma_y^{-1} \left[ \mathbf{R}^i, ((\mathbf{C}x_t - p^i)^T \otimes \mathbf{I}) \frac{d \text{vec}(\mathbf{R})}{d\theta} \right], \quad (16)$$

where we have used the identity  $\text{vec}(\mathbf{X}\mathbf{Y}) = (\mathbf{Y}^T \otimes \mathbf{I}) \text{vec}(\mathbf{X})$  [11], where  $\otimes$  is the Kronecker product,  $\text{vec}(\cdot)$  stacks elements of a matrix into a column, and  $\mathbf{I}$  is the  $2 \times 2$  identity matrix.

Differentiating the  $z$ th element of  $r_y$  with respect to the observed trajectory element  $x_t$  introduces a block into  $\mathbf{J}_{\mu\mathbf{x}}$  at location  $(z, t)$ :

$$[\mathbf{J}_{\mu\mathbf{x}}]_{z,t} = \sigma_y^{-1} \mathbf{R}^i \mathbf{C}, \quad (17)$$

where  $i$  is the number of the camera that made the  $z$ th observation.

Differentiating  $r_x$  with respect to  $\mu$  yields 0. Differentiating it with respect to  $x$  yields  $\mathbf{J}_\mathbf{x} = \mathbf{G}$ .

$r_\mu$  is only a function of the parameters of the first sensor. Its derivative with respect to  $\mu^1$  is  $\mathbf{I}$ , where  $\mathbf{I}$  is the  $3 \times 3$  identity matrix.