Incremental estimation without specifying a-priori covariance matrices for the novel parameters

Christian Beder
Computer Science Department
Kiel University, Germany
beder@mip.informatik.uni-kiel.de

Richard Steffen
Department of Photogrammetry
Bonn University, Germany
rsteffen@uni-bonn.de

Abstract

We will present a novel incremental algorithm for the task of online least-squares estimation. Our approach aims at combining the accuracy of least-squares estimation and the fast computation of recursive estimation techniques like the Kalman filter.

Analyzing the structure of least-squares estimation we devise a novel incremental algorithm, which is able to introduce new unknown parameters and observations into an estimation simultaneously and is equivalent to the optimal overall estimation in case of linear models. It constitutes a direct generalization of the well-known Kalman filter allowing to augment the state vector inside the update step. In contrast to classical recursive estimation techniques no artificial initial covariance for the new unknown parameters is required here. We will show, how this new algorithm allows more flexible parameter estimation schemes especially in the case of scene and motion reconstruction from image sequences.

Since optimality is not guaranteed in the non-linear case we will also compare our incremental estimation scheme to the optimal bundle adjustment on a real image sequence. It will be shown that competitive results are achievable using the proposed technique.

1. Introduction

Least-squares parameter estimation is a well-known technique in computer vision, which has been widely accepted as post-processing step for all structure-from-motion algorithms aiming at highly accurate results. An extensive overview on the current state of the art of bundle adjustment is given in [21].

Many previous works show the versatile applicability of least-squares estimation methods for instance in the areas of 3D reconstruction from image sequences (e.g. [13]), camera calibration (e.g. [4]), vehicle navigation (e.g. [17]) or image mosaicing (e.g. [10]).

Because bundle adjustment is very expensive in computation time, especially in case of large image sequences, a lot of faster solutions to solve the normal equation system were proposed over the last decades. Those solutions use fast algorithms for solving the equation systems by factorization (cf. [20],[12],[3]) or use the special design of the structure-from-motion equations by dividing the parameters into a structure and a motion part (cf. [21],[5]). However, while this techniques are vital for solving large problems, they are not the scope of this paper, as our approach is not intended to be specific for a certain estimation problem.

Other solutions are based on choosing optimal keyframes (cf. [1]) and compute a local adjustment for a subset of parameters. In contrast to our approach, those do not consider all observations and loose sight of some correlations to old parameters (cf. [18], [16] [15], [23]).

Another way of exploiting the special design of the structure-from-motion problem for image sequences is the so-called Variable State-Dimension Filter (cf. [9], [11]), which is closely related to our approach, as it also includes novel parameters incrementally. The first key difference is, that we only need to invert a small sub-matrix of the normal equation matrix corresponding to the newly introduced parameters and update the existing parameters recursively. In contrast to [11] we are able to deal with arbitrary correlations within the parameter vector. The second difference is, that we only retain the parameter vector and its covariance matrix between the steps, so that a Kalman filter like prediction step between the updates is straightforward. For instance removing obsolete parameters from the estimation is a simple matter of canceling out rows and columns of the respective matrices instead of having to compute the Schur complement.

It has also been proposed in the literature to efficiently solve the structure-from-motion task by formulating it as a recursive estimation problem and using Kalman filter based approaches to compute a solution. This approaches (e.g.

978-1-4244-2340-8/08/$25.00 ©2008 IEEE

[8], [6],[19] and [2]) are very fast in computation time compared to bundle adjustment but their accuracy is lower due to a built up of linearization errors for previous states. Our goal in this work is to combine the accuracy of the bundle adjustment with the fast computation times of those recursive estimation algorithms. We will compare our approach to [2] in section 4.

Another recursive Kalman filter based approach for the incremental structure-from-motion problem has been presented in [14], where the motion parameters are estimated using a particle filter and the structure parameters are estimated using a Kalman filter enabling it to handle large maps efficiently.

Our approach is based on least-squares adjustment, but extends it in several ways:

- It is possible to include new parameters into the estimation incrementally without having to specify an artificial a priori covariance matrix for them. This is in contrast to the classical recursive estimation techniques.

- It is possible to estimate the newly introduced parameters and their covariances as well as update the old parameters and their covariances separately, taking into account all mutual correlations. The whole normal equation matrix has not to be inverted again in order to compute this updates.

- No history of observations has to be maintained. Instead the parameter vector and its covariance matrix is built up incrementally using only newly acquired observations. The parameter vector and its covariance is the only information required from previous steps, which in contrast to [9] allows to easily implement a prediction step between the updates like for instance eliminate parameters from the estimation by simply canceling out rows and columns.

The paper is organized as follows: In section 2 the generic least-squares adjustment is analyzed and our novel incremental technique will be derived. We will then briefly show in section 3, how the structure-from-motion problem for image sequences can be easily integrated into the generic framework presented in section 2. Results on a real image sequence will be show in section 4, where we will compare the proposed method with the optimal gold standard method of overall bundle adjustment and the Kalman filter based method of [2]. Finally we will conclude and give an outlook on some possible future work.

2. Incremental least-squares estimation

We will first show, how classical least-squares adjustment works. Given a set of observations $\boldsymbol{l}_1$ together with their covariance matrix C_{11} that depend on a set of unknown parameters $\boldsymbol{p}_1$ according to the known linear model function

$$\tilde{\boldsymbol{l}}_1 = A_{11}\tilde{\boldsymbol{p}}_1 \tag{1}$$

the best linear unbiased estimate of the parameters is obtained as (cf. [7])

$$\hat{\boldsymbol{p}}_1^{(-)} = C^{(-)} A_{11}^T C_{11}^{-1} \boldsymbol{l}_1 \tag{2}$$

with its covariance being the inverse of the normal equation matrix

$$C^{(-)} = (A_{11}^T C_{11}^{-1} A_{11})^{-1} \tag{3}$$

If the model function is not linear, its Taylor expansion has to be used and the estimation process must be iterated.

Now we want to add new uncorrelated observations $\boldsymbol{l}_2$ having the covariance C_{22} and new additional unknown parameters $\boldsymbol{p}_2$ in a later stage. Hence, the previous model equation (1) must be augmented and reads then as

$$\begin{pmatrix} \tilde{\boldsymbol{l}}_1 \\ \tilde{\boldsymbol{l}}_2 \end{pmatrix} = \begin{pmatrix} A_{11} & 0 \\ A_{21} & A_{22} \end{pmatrix} \begin{pmatrix} \tilde{\boldsymbol{p}}_1 \\ \tilde{\boldsymbol{p}}_2 \end{pmatrix} \tag{4}$$

Observe that the new observations $\boldsymbol{l}_2$ may depend on the newly introduced parameters $\boldsymbol{p}_2$ via A_{22} as well as on the old parameters $\boldsymbol{p}_1$ via A_{21}.

Using again equations (2) and (3) for this augmented model equation the best linear unbiased estimate of the old as well as the newly introduced parameters is now given by

$$\begin{pmatrix} \hat{\boldsymbol{p}}_1^{(+)} \\ \hat{\boldsymbol{p}}_2 \end{pmatrix} = C^{(+)} \begin{pmatrix} A_{11}^T C_{11}^{-1} \boldsymbol{l}_1 + A_{21}^T C_{22}^{-1} \boldsymbol{l}_2 \\ A_{22}^T C_{22}^{-1} \boldsymbol{l}_2 \end{pmatrix} \tag{5}$$

The covariance of this augmented parameter vector is the inverse of a symmetric matrix with the following block structure

$$C^{(+)} = \begin{pmatrix} P & Q \\ Q^T & R \end{pmatrix}^{-1} \tag{6}$$

using the substitutions

$$P = A_{11}^T C_{11}^{-1} A_{11} + A_{21}^T C_{22}^{-1} A_{21} \tag{7}$$
$$Q = A_{21}^T C_{22}^{-1} A_{22} \tag{8}$$
$$R = A_{22}^T C_{22}^{-1} A_{22} \tag{9}$$

Observe that the new parameter vector $\boldsymbol{p}_2$ has been introduced without having to specify an artificial initial covariance matrix for it.

Now we will further analyze this expression. First note that P is the sum of an invertible matrix (cf. equation (3)) and a dyadic product, so that its inverse may be computed as (cf. [7] and equation (3))

$$P^{-1} = F C^{(-)} \tag{10}$$

with the substitution

$$F = I_d - C^{(-)} A_{21}^T (C_{22} + A_{21} C^{(-)} A_{21}^T)^{-1} A_{21} \quad (11)$$

It is therefore possible to invert the blockmatrix as follows (cf. [7])

$$C^{(+)} = \begin{pmatrix} K & L \\ L^T & M \end{pmatrix} = \begin{pmatrix} P & Q \\ Q^T & R \end{pmatrix}^{-1} \quad (12)$$

with the substitutions

$$\begin{aligned} K &= P^{-1} + P^{-1} Q (R - Q^T P^{-1} Q)^{-1} Q^T P^{-1} && (13) \\ &= (I_d + F C^{(-)} Q M Q^T) F C^{(-)} && (14) \\ L^T &= -(R - Q^T P^{-1} Q)^{-1} Q^T P^{-1} && (15) \\ &= -M Q^T F C^{(-)} && (16) \\ M &= (R - Q^T P^{-1} Q)^{-1} && (17) \\ &= (R - Q^T F C^{(-)} Q)^{-1} && (18) \end{aligned}$$

Observe, that a matrix inversion is only needed for the computation of the covariance matrix of the novel parameters M, which is usually small. Obviously, M has to have full rank, which implies that the observations $\boldsymbol{l}_2$ are sufficient to estimate the novel parameters $\boldsymbol{p}_2$.

Putting everything together and evaluating equation (5) the old parameters update according to

$$\begin{aligned} \hat{\boldsymbol{p}}_1^{(+)} &= (I_d + F C^{(-)} Q M Q^T) F (\hat{\boldsymbol{p}}_1^{(-)} + C^{(-)} A_{21}^T C_{22}^{-1} \boldsymbol{l}_2) \\ &\quad - F C^{(-)} Q M A_{22}^T C_{22}^{-1} \boldsymbol{l}_2 && (19) \end{aligned}$$

having the new covariance matrix K given in equation (14). The new parameters are given by

$$\begin{aligned} \hat{\boldsymbol{p}}_2 &= -M Q^T F \hat{\boldsymbol{p}}_1^{(-)} - M Q^T F C^{(-)} A_{21}^T C_{22}^{-1} \boldsymbol{l}_2 \\ &\quad + M A_{22}^T C_{22}^{-1} \boldsymbol{l}_2 && (20) \end{aligned}$$

having the covariance M given in equation (18). The mutual covariance between the updated old parameters and the new parameters is given by L in equation (16).

Observe that none of the update equations (19), (20), (14), (18) and (16) contains any reference to past observations (i.e. C_{11} and $\boldsymbol{l}_1$) or constraints (i.e. A_{11}). Instead all information is encoded in the parameter vector $\hat{\boldsymbol{p}}_1^{(-)}$ and its covariance matrix $C^{(-)}$, so that the past observations do not need to be retained. After the estimation of the new augmented parameter vector and its covariance matrix is completed it may be modified or truncated for further processing like in the prediction step of the Kalman filter enabling the introduction of a motion model.

Also note, that if A_{22} and $\boldsymbol{p}_2$ have size zero, then everything boils down to the update step of the classical Kalman-filter (cf. [22]). Hence, the presented method is a direct generalization of the Kalman filter update equations.

3. Structure from motion

Although the incremental estimation technique presented in the previous section is applicable to all estimation problems with online acquirement of new observations and parameters, it is especially well suited for the joint structure and motion recovery from image sequences, where new camera positions and new scene points appear at each new frame.

If the j^{th} scene point is seen by the i^{th} camera, its homogeneous coordinates are given by

$$\mathbf{x}^{(ij)} = \mathsf{K}^{(i)} R^{(i)} (I_3 | - \boldsymbol{Z}^{(i)}) \mathbf{X}^{(j)} \quad (21)$$

We will assume the calibration matrices $\mathsf{K}^{(i)}$ of the cameras to be known. Furthermore we will assume the rotation matrices $R^{(i)}$ to depend on the quaternion $(1, \boldsymbol{q}^{(i)})$ and the scene points $\mathbf{X}^{(j)} = (\boldsymbol{X}^{(j)}, 1)$ to be not at infinity and normalized. The coordinates of the image points are measured in an Euclidean frame, i.e. $\boldsymbol{x}^{(ij)} = \mathbf{x}_{1:2}^{(ij)} / \mathbf{x}_3^{(ij)}$. Hence, equation (21) is phrased in an Euclidean frame as a function $\boldsymbol{f} : I\!R^3 \times I\!R^3 \times I\!R^3 \mapsto I\!R^2$ like this

$$\boldsymbol{x}^{(ij)} = \boldsymbol{f}(\boldsymbol{q}^{(i)}, \boldsymbol{Z}^{(i)}, \boldsymbol{X}^{(j)}) \quad (22)$$

A Taylor expansion of this at appropriate initial values $\left(\bar{\boldsymbol{q}}^{(i)}, \bar{\boldsymbol{Z}}^{(i)}, \bar{\boldsymbol{X}}^{(j)}\right)$ yields

$$\underbrace{\boldsymbol{x}^{(ij)} - \boldsymbol{f}(\bar{\boldsymbol{q}}^{(i)}, \bar{\boldsymbol{Z}}^{(i)}, \bar{\boldsymbol{X}}^{(j)})}_{\boldsymbol{l}^{(ij)}} \quad (23)$$

$$\approx \underbrace{\begin{pmatrix} \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{q}^{(i)}} & \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{Z}^{(i)}} & \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{X}^{(j)}} \end{pmatrix}}_{A^{(ij)}} \underbrace{\begin{pmatrix} \boldsymbol{q}^{(i)} - \bar{\boldsymbol{q}}^{(i)} \\ \boldsymbol{Z}^{(i)} - \bar{\boldsymbol{Z}}^{(i)} \\ \boldsymbol{X}^{(j)} - \bar{\boldsymbol{X}}^{(j)} \end{pmatrix}}_{\boldsymbol{p}^{(ij)}}$$

which is the linear form needed to construct the components $\boldsymbol{l}_1$, $\boldsymbol{p}_1$ and $\boldsymbol{A}_{11}$ of the model equation (1) and estimated the improved parameters as shown above. This process can be iterated using the improved parameters as new initial values and is commonly referred to as bundle adjustment. We will not focus on how to obtain good initial values here.

Now at each new frame an additional camera position and rotation vector has to be introduced. Furthermore, novel points become visible in the new frame. Those new unknowns constitute the entries of the parameter vector $\boldsymbol{p}_2$ in the augmented model equation (4).

Some of the old points remain visible in the new frame. Furthermore, the scene points introduced in the new frame should have been visible in previous frames to be included stably into the estimation (cf. [1]). Their image coordinates appear in the observation vector $\boldsymbol{l}_2$. For each such entry in the observation vectors rows of A_{21} and A_{22} have to be filled in. The entries going into A_{21} are those concerning the functional dependence with previous frames or previous scene points. The entries going into A_{22} are those concerning the functional dependence with the new frame and the newly introduced scene points.

Having filled in the vectors l_2 and the matrices A_{21} and A_{22} one can estimate the parameter updates as depicted in the previous section and obtains improved estimates for the parameters that can be used as new initial values for the Taylor expansions. This process is then iterated until convergence.

4. Results

The incremental technique presented in the previous sections is only equivalent to the overall adjustment in the linear case. This is because the incremental method is unable to re-linearize the functional dependence for past frames and retains only the parameter vector and its covariance matrix, which is a problem also known from the extended Kalman filter.

In order to assess the performance of the presented technique for the non-linear structure-from-motion problem, we used the well-known rotating dinosaur sequence depicted in figure 1, where ground-truth camera calibration and orientation data were available. We extracted point features and tracked them across the sequence. For reference we computed an overall bundle adjustment, which is the optimal solution in terms of reprojection error of the tracked features. We also compared our approach to the iterative extended Kalman filter based solution proposed in [2].

To initialize the extended Kalman filter and the incremental method, we computed a bundle adjustment only for the first five frames of the sequence. Both methods initialized new object points at the centroid of the point cloud and were iterated until convergence. While the extended Kalman filter method is able to predict its novel camera pose, the novel camera poses for the incremental method were initialized using a simple linear extrapolation from the previous two frames to generate the initial values. Furthermore, we used large isotropic initial covariance matrices for novel parameters in the extended Kalman filter, which need not to be specified for our incremental method.

We then added the remaining frames one by one using an overall bundle adjustment, the iterated extended Kalman filter approach of [2] and our incremental technique proposed in the previous sections. The resulting camera positions and scene points for the overall bundle adjustment are shown in figure 2. The camera positions and scene points for the iterated extended Kalman filter approach are shown in figure 3 and finally the results from our approach are depicted in figure 4. As expected, the best result is achieved by the overall bundle adjustment, while the quality of the incremental adjustment does not achieve this quality due to its inability

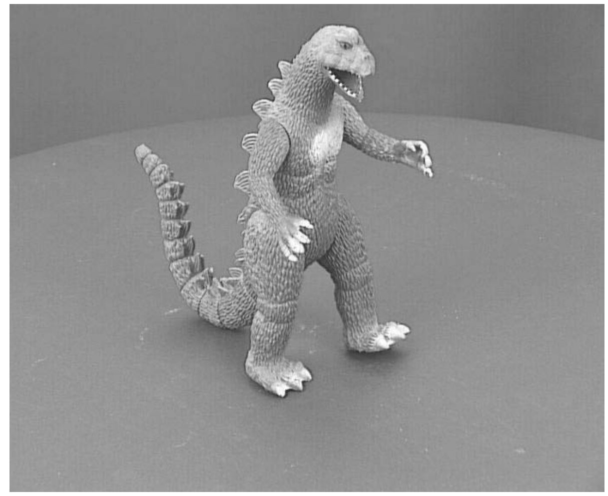

Figure 1. A single frame of the well-known rotating dinosaur sequence. The sequence consists of 36 images rotated in 10° steps around the dinosaur. Ground-truth for the camera calibration, position and rotation is available and will be used to quantify the performance of the presented methods.

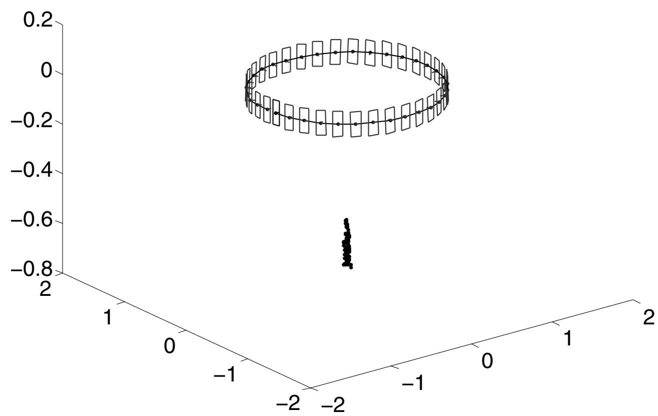

Figure 2. Bundle adjustment solution for the camera poses and scene points.

to re-linearize at previous camera positions. This is also a problem for the extended Kalman filter, which performed even a bit worse in this scenario.

To quantify the quality of the results, we compared the computed camera positions and orientations with the available ground-truth. The angular errors of the camera orientations as well as the position errors of the camera projection centers are depicted in figure 5 for each frame. As expected the overall bundle adjustment performed best. The average errors of the sequential method are a little worse with an angular error of up to approximately 3° and a position error of up to approximately $0.1m$. Comparing the proposed incremental method with the iterated extended Kalman filter

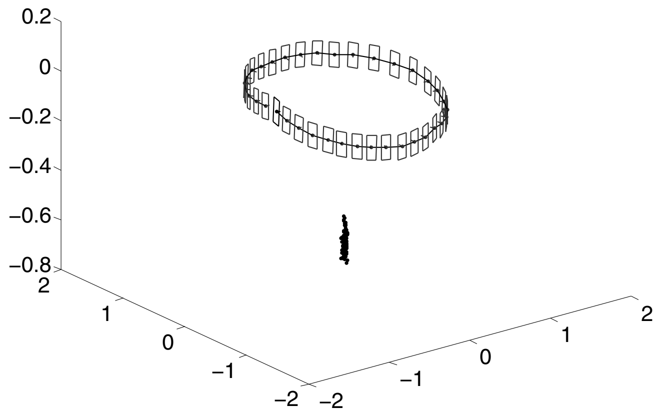

Figure 3. Iterated extended Kalman filter based solution for the camera poses and scene points.

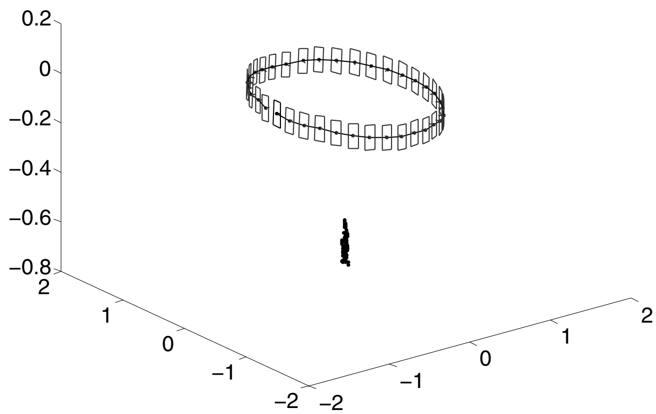

Figure 4. Our incremental method solution for the camera poses and scene points.

ter approach of [2] we can see that the angular error goes up to approximately 10° and the position error goes up to approximately $0.3m$ along the sequence.

5. Conclusion

We have presented an algorithm that combines the accuracy of bundle adjustment and the fast computation of recursive estimation techniques. The presented technique is a direct generalization of the well-known Kalman filter allowing to introduce novel observations as well as novel parameters inside the update step. Thereby no artificial a-priori covariance matrix has to be specified for such novel parameters.

Our approach is closely related to [9] and [11], but we only need to invert a very small sub-matrix of the normal

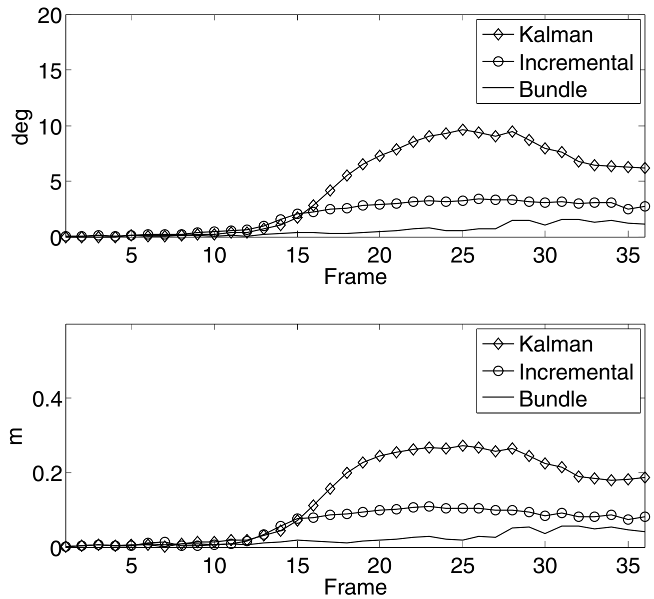

Figure 5. Comparison of bundle adjustment, incremental estimation and iterated extended Kalman filter based estimation on the dino sequence. *Top:* Angular distance of the camera pose to the ground truth plotted against frame number. *Bottom:* Distance of the camera center to the ground truth plotted against frame number.

equation matrix corresponding to the newly introduced parameters and update the old parameters in a recursive manner taking all correlations into account.

In contrast to classical recursive adjustment schemes for image sequences (e.g. [2]), the proposed method does not require any artificial a priori covariance matrix for newly introduced parameters and the results do not depend on a motion model. However, a motion model can be easily included. Specifically this means, that only uncertainties of observations have to be supplied, which is conceptually much clearer than having to introduce a-priori uncertainties of the parameters to be estimated. This enables more transparent estimation schemes, that do not rely on such prior information on the uncertainty of the parameters.

Because the parameter vector and its covariance is the only information needed to be maintained after the estimation, we can easily process it between the update steps using error propagation like it is common practice with Kalman filter methods. Specifically it is a simple matter of canceling out rows and columns from the covariance matrix and the parameter vector in order to remove obsolete parameters from the estimation. In contrast to [9] no Schur complement is needed for this elimination step.

We evaluated the proposed incremental technique on a real image sequence and compared the performance in terms of achieved accuracy to the gold standard method of overall bundle adjustment and the Kalman filter based method of [2]. In case of linear problems the incremental estimation and the overall least-squares adjustment are equivalent and the results suggested a competitive performance for the non-linear problem of structure-from-motion.

Our future work will focus on analyzing the internal structure of the design matrices in order to be able to exploit their sparsity resulting from the structure-from-motion problem for image sequences. Thereby we expect to speed up the computation times again and become competitive with current optimized state of the art bundle adjustment implementations.

References

[1] C. Beder and R. Steffen. Determining an initial image pair for fixing the scale of a 3d reconstruction from an image sequence. In K. Franke, K.-R. Müller, B. Nickolay, and R. Schäfer, editors, *Pattern Recognition*, number 4174 in LNCS, pages 657–666. Springer, 2006.

[2] A. J. Davison. Real-time simulataneous localisation and mapping with a single camera. In *Proceeding of the 9th International Conference on Computer Vision*, pages 674–679, 2003.

[3] C. Engels, H. Stewenius, and D. Nister. Bundle adjustment rules. In *Photogrammetric Computer Vision (PCV)*, September 2006.

[4] C. S. Fraser, M. R. Shortis, and G. Ganci. Multisensor system self-calibration. In S. F. El-Hakim, editor, *Proceedings of SPIE Vol. 2598, p. 2-18, Videometrics IV, Sabry F. El-Hakim; Ed.*, pages 2–18, Sept. 1995.

[5] A. Grün. An optimum algorithm for on-line triangulation. In *Symposium of the Commission III of the ISPRS*, 1982.

[6] S. Julier and J. Uhlmann. A new extension of the kalman filter to nonlinear systems. *Int. Symp. Aerospace/Defense Sensing, Simul. and Controls*, 3, 1997.

[7] K.-R. Koch. *Parameter estimation and hypothesis testing in linear models*. Springer, 1988.

[8] L. Matthies, R. Szelinski, and T. Kanade. Kalman filter-based algorithms for estimating depth from image sequences. *International Journal of Computer Vision*, 3(3):209–238, 1989.

[9] P. McLauchlan. A batch/recursive algorithm for 3D scene reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 738–743, 2000.

[10] P. F. McLauchlan and A. Jaenicke. Image mosaicing using sequential bundle adjustment. *Image Vision Computing*, 20:751 – 759, 2002.

[11] P. F. McLauchlan and D. W. Murray. A unifying framework for structure and motion recovery from image sequences. In *Proceedings of the International Conference on Computer Vision*, pages 314–320, 1995.

[12] C. Meyer. *Matrix Analysis and Applied Linear Algebra with Solutions Manual*. Society for Industrial and Applied Mathematics, SIAM, Philadelphia, 2000.

[13] R. Mohr, B. Boufama, and P. Brand. Accurate projective reconstruction. In *Proceedings of the Second Joint European - US Workshop on Applications of Invariance in Computer Vision*, pages 257–276, London, UK, 1994. Springer-Verlag.

[14] M. Montemerlo, S. Thrun, D. Koller, and B. Wegbreit. Fastslam: a factored solution to the simultaneous localization and mapping problem. In *Eighteenth national conference on Artificial intelligence*, pages 593–598, Menlo Park, CA, USA, 2002. American Association for Artificial Intelligence.

[15] E. Mouragnon, F. Dekeyser, P. Sayd, M. Lhuillier, and M. Dhome. Real time localization and 3d reconstruction. In *CVPR '06: Proceedings of the 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 363–370, Washington, DC, USA, 2006. IEEE Computer Society.

[16] E. Mouragnon, M. Lhuillier, M. Dhome, F. Dekeyser, and P. Sayd. Monocular vision based slam for mobile robots. In *ICPR '06: Proceedings of the 18th International Conference on Pattern Recognition*, pages 1027–1031, Washington, DC, USA, 2006. IEEE Computer Society.

[17] N. Seunderhauf, K. Konolige, S. Lacroix, and P. Protzel. Visual odometry using sparse bundle adjustment on an autonomous outdoor vehicle. In P. Levi and M. Schanz, editors, *Autonome Mobile Systeme 2005*, pages 157–163. Springer, 2005.

[18] D. Steedly and I. A. Essa. Propagation of innovative information in non-linear least-squares structure from motion. In *ICCV*, pages 223–229, 2001.

[19] R. Steffen and C. Beder. Recursive estimation with implicit constraints. In F. Hamprecht, C. Schnörr, and B. Jähne, editors, *Proceedings of the DAGM 2007*, number 4713 in LNCS, pages 194–203. Springer, 2007.

[20] P. Sturm and B. Triggs. A factorization based algorithm for multi-image projective structure and motion. In *Proceedings of the European Conference on Computer Vision*, pages 709–720. Springer, 1996.

[21] 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, 2000.

[22] G. Welch and G. Bishop. An introduction to the kalman filter. Technical report, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA, 1995.

[23] Z. Zhang and Y. Shan. Incremental motion estimation through local bundle adjustment. Technical Report MSR-TR-01-54, Microsoft Research, 2001.