



Robust factorization

Aanæs, Henrik; Fisker, Rune; Åström, Kalle; Carstensen, Jens Michael

Published in:

I E E E Transactions on Pattern Analysis and Machine Intelligence

Link to article, DOI:

[10.1109/TPAMI.2002.1033213](https://doi.org/10.1109/TPAMI.2002.1033213)

Publication date:

2002

Document Version

Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):

Aanæs, H., Fisker, R., Åström, K., & Carstensen, J. M. (2002). Robust factorization. *I E E E Transactions on Pattern Analysis and Machine Intelligence*, 24(9), 1215-1225. <https://doi.org/10.1109/TPAMI.2002.1033213>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Robust Factorization

Henrik Aanæs, Rune Fisker,
Kalle Åström, *Member, IEEE Computer Society*, and Jens Michael Carstensen

Abstract—Factorization algorithms for recovering structure and motion from an image stream have many advantages, but they usually require a set of well-tracked features. Such a set is in general not available in practical applications. There is thus a need for making factorization algorithms deal effectively with errors in the tracked features. We propose a new and computationally efficient algorithm for applying an arbitrary error function in the factorization scheme. This algorithm enables the use of robust statistical techniques and arbitrary noise models for the individual features. These techniques and models enable the factorization scheme to deal effectively with mismatched features, missing features, and noise on the individual features. The proposed approach further includes a new method for Euclidean reconstruction that significantly improves convergence of the factorization algorithms. The proposed algorithm has been implemented as a modification of the Christy-Horand factorization scheme, which yields a perspective reconstruction. Based on this implementation, a considerable increase in error tolerance is demonstrated on real and synthetic data. The proposed scheme can, however, be applied to most other factorization algorithms.

Index Terms—Robust statistics, feature tracking, perspective reconstruction, Euclidean reconstruction, structure from motion.

1 INTRODUCTION

THE reconstruction of structure and motion of a rigid object from an image stream is one of the most studied problems within computer vision. A popular way of addressing this problem is to extract and track features through the image sequence and then limit the problem to estimating the structure and motion of these tracked features. A family of effective and popular algorithms for solving this estimation problem are the so-called factorization algorithms, see, e.g., [6], [7], [16], [17], [18], [20], [22], [28].

These factorization algorithms work by linearizing the camera observation model and give good results rapidly and without an initial guess for the solution. Hence, the factorization algorithms are good candidates for solving the structure and motion problem, either as a full solution or as an initialization to other algorithms such as bundle adjustment, see, e.g., [25], [30].

The factorization algorithms assume that the correspondence or feature tracking problem has been solved. The correspondence problem is, however, one of the most difficult fundamental problems within computer vision. No perfect and truly general solution has yet been presented. For most practical purposes, one must deal with erroneously tracked features as input to the factorization algorithm. This fact poses a considerable challenge to factorization algorithms, since they implicitly assume independent identically distributed Gaussian noise on the 2D features (the 2-norm is used as error function on the 2D features). This noise assumption based on the 2-norm is known to perform rather

poorly in the presence of outliers induced by such erroneous data. These errors typically arise from mismatching 2D features or from a 2D feature being absent due to occlusion. It is common for badly tracked features to disturb the estimation of structure and motion considerably.

Previous attempts have been made at addressing this problem. Irani and Anandan [16] assume that the noise is separable in a 3D feature point contribution and a frame contribution. In other words, if a 3D feature point has a relatively high uncertainty in one frame, it is assumed that it has a similar high uncertainty in all other frames. However, large differences in the variance of the individual 2D feature points is critical to the implementation of robust statistical techniques that can deal with feature point noise, missing features, and feature mismatch in single frames. Morris and Kanade [20] propose a bilinear minimization method as an improvement on top of a standard factorization. The bilinear minimization incorporates directional uncertainty models in the solution. However, the method does not implement robust statistical techniques. Tomasi and Kanade [28] and Jacobs [17] address the problem of missing data points by the use of heuristics. Attempts at solving similar linear problems, in the presence of missing and erroneous data, have also been made, e.g., [24].

Here, we propose a combined approach that deal effectively with missing features and is robust towards errors in the matching of the 2D features in a factorization framework. This is achieved by allowing for an arbitrary noise model on the 2D features—i.e., we are not restricted to a Gaussian model. Hereby, the proposed approach is capable of dealing effectively with mismatched features or outliers by the use of robust statistics. Arbitrary noise models also deal with missing 2D features by emulating them as being located arbitrarily in the image with very high noise variance.

The proposed approach is implemented as an improvement to the factorization algorithm of Christy and Horand [6]. The Christy-Horand algorithm has the advantage that it assumes a perspective camera model as opposed to a linearized version, e.g., [28]. The Christy-Horand algorithm

• H. Aanæs and J.M. Carstensen are with Informatics and Mathematical Modelling, DTU, Building 321, DK-2800 Lyngby, Denmark. E-mail: {haa, jmc}@imm.dtu.dk.

• R. Fisker is with 3Shape ApS, Bredgade 58, 2-3. sal, 1260 Copenhagen K, Denmark. E-mail: rune@3shape.com.

• K. Åström is with the Centre for Mathematical Sciences, Lund University, PO Box 118, SE-221 00 Lund, Sweden. E-mail: kalle@maths.lth.se.

Manuscript received 30 Oct. 2000; revised 22 Aug. 2001; accepted 12 Feb. 2002.

Recommended for acceptance by Z. Zhang.

For information on obtaining reprints of this article, please send e-mail to: tpami@computer.org, and reference IEEECS Log Number 113076.

with the proposed approach incorporated deals efficiently with real and simulated data containing large feature errors.

The presentation is organized by giving an overview of the factorization algorithms in Section 2 followed by a discussion of how to deal with erroneous data in Section 3. In Section 4, a new numerical algorithm for estimating the optimal subspace of a matrix with weighted entries is proposed. In Sections 5 and 6, schemes improving the robustness of the factorization approach are presented, followed by experimental results in Section 7.

2 FACTORIZATION OVERVIEW

As a courtesy to the reader and to introduce notation, a short overview of the factorization algorithm is presented. For a more thorough introduction, the reader is referred to [5]. All the factorization methods cited utilize some linearization of the pinhole camera with known intrinsic parameters:

$$s_{ij} \begin{bmatrix} x_{ij} \\ y_{ij} \\ 1 \end{bmatrix} = \begin{bmatrix} \mathbf{a}_i^T & t_i^x \\ \mathbf{b}_i^T & t_i^y \\ \mathbf{c}_i^T & t_i^z \end{bmatrix} \begin{bmatrix} \mathbf{P}_j \\ 1 \end{bmatrix}, \quad (1)$$

where the 3D feature, \mathbf{P}_j , is projected in frame i as (x_{ij}, y_{ij}) , $\mathbf{t}_i = [t_i^x, t_i^y, t_i^z]^T$ is the appropriate translation vector and \mathbf{a}_i^T , \mathbf{b}_i^T , and \mathbf{c}_i^T are the three rows vectors of the rotation matrix. The used/approximated observation model can thus be written as:

$$\begin{bmatrix} x_{ij} \\ y_{ij} \end{bmatrix} = \mathbf{M}_i \mathbf{P}_j, \quad (2)$$

where \mathbf{M}_i is the 2×3 “linearized motion” matrix associated with frame i .

When n features have been tracked in k frames, $i \in [1 \dots k]$ and $j \in [1 \dots n]$, the observations from (2) can be combined to:

$$\mathbf{S} = \mathbf{M} \mathbf{P}, \quad (3)$$

where \mathbf{M} is a $2k \times 3$ matrix composed of the \mathbf{M}_i and \mathbf{P} is a $3 \times n$ matrix composed of the \mathbf{P}_j . Thus, the elements of \mathbf{S} are given by:

$$\mathbf{S} = \begin{bmatrix} x_{11} & \cdots & x_{1n} \\ \vdots & \ddots & \vdots \\ x_{k1} & \cdots & x_{kn} \\ y_{11} & \cdots & y_{1n} \\ \vdots & \ddots & \vdots \\ y_{k1} & \cdots & y_{kn} \end{bmatrix}.$$

The solution to this linearized problem is then found as the \mathbf{M} and \mathbf{P} that minimize:

$$\mathbf{N} = \mathbf{S} - \mathbf{M} \mathbf{P}, \quad (4)$$

where \mathbf{N} is the residual between the model, $\mathbf{M} \mathbf{P}$, and the data, \mathbf{S} . The residuals, \mathbf{N} , are usually minimized using the Frobenius norm. This is equivalent to minimizing the squared Euclidean norm of the reprojection error, i.e., the error between the measured 2D features and the corresponding reprojected 3D feature. Thus, the objective function is:

$$\min_{\mathbf{M}, \mathbf{P}} \|\mathbf{S} - \mathbf{M} \mathbf{P}\|_F^2 = \min_{\mathbf{M}, \mathbf{P}} \sum_{j=1}^n \|\mathbf{S}_j - \mathbf{M} \mathbf{P}_j\|_2^2, \quad (5)$$




Fig. 1. Overview of the Christy-Horad factorization algorithm. A linearized—e.g., paraperspective-solution—is iterated into a perspective.

where \mathbf{S}_j and \mathbf{P}_j denote the j th column of \mathbf{S} and \mathbf{P} , respectively. In this case, the solution to \mathbf{M} and \mathbf{P} can be found through the singular value decomposition, SVD, of \mathbf{S} . It is noted that, for any invertible 3×3 matrix, \mathbf{A} :

$$\mathbf{M} \mathbf{P} = \mathbf{M} \mathbf{A} \mathbf{A}^{-1} \mathbf{P} = \tilde{\mathbf{M}} \tilde{\mathbf{P}}. \quad (6)$$

Hence, the solution is only defined up to an affine transformation. In [6], Euclidean reconstruction is achieved by estimation of an \mathbf{A} such that the rotation matrices, $[\mathbf{a}_i \ \mathbf{b}_i \ \mathbf{c}_i]^T$, are as orthonormal as possible. Further details and an improved approach are presented in Section 6.

2.1 Christy-Horad Factorization

The factorization algorithm that the proposed method extends on is that of Christy and Horad [6], see Fig. 1. The Christy-Horad algorithm is an extension to the work in [22], [28] partly by incorporating the work of [8]. The algorithm has the advantage of iteratively achieving a solution to the original nonlinearized problem, thus achieving perspective reconstruction. This is achieved by iteratively solving a linearized version of the problem and then modifying the data to approach the perspective camera. The update formula for the data, i.e., \mathbf{S} or x_{ij}, y_{ij} is [6]:

$$\begin{bmatrix} \tilde{x}_{ij} \\ \tilde{y}_{ij} \end{bmatrix} = \left(\begin{bmatrix} x_{ij} \\ y_{ij} \end{bmatrix} - \begin{bmatrix} x_{o_i} \\ y_{o_i} \end{bmatrix} \right) (1 - \epsilon_{ij}), \quad (7)$$

where $(\tilde{x}_{ij}, \tilde{y}_{ij})$ is the updated data, (x_{o_i}, y_{o_i}) is the object origin projected onto frame i , and ϵ_{ij} is the scaled depth defined as:

$$\epsilon_{ij} = \frac{\mathbf{c}_i^T \cdot \mathbf{P}_j}{t_i^z}. \quad (8)$$

For further details, the reader is referred to [5], [6].

3 DEALING WITH IMPERFECT DATA

3.1 Types of Errors

A framework for dealing with imperfect data should be geared toward the types of errors expected. Three types of errors have been identified, the first two originating from [31].

- **Bad Feature Locations.** The locations of a 2D feature is distorted. If this distortion is anisotropic and/or the covariance structure varies, the Frobenius norm is suboptimal. This is the problem addressed in [16], [20].
- **False Matches.** There has been a mismatch of 2D features, i.e., two 2D feature originating from different 3D features have been matched as originating from the same 3D feature.

- **Missing Feature.** The projection of a 3D feature can not be found or matched, e.g., due to occlusion or simply by “breakdown” of the feature matching algorithm. This is addressed in [17], [28].

Our proposed method deals with all these types of errors by attaching an uncertainty to the individual 2D features. It is implemented by associating a weighting structure to \mathbf{S} . This mainly effects the solution of (5), but, in the case of Christy-Horand, it also effects the object frame origin (see Section 5). When weights are introduced, (5) becomes:

$$\min_{\mathbf{M}, \mathbf{P}} \sum_{j=1}^n \|\mathbf{V}_j(\mathbf{S}_j - \mathbf{M}\mathbf{P}_j)\|_2^2, \quad (9)$$

where \mathbf{V}_j is a $2k \times 2k$ weighting matrix representing the weights of the j th column of \mathbf{S} . The minimization of (9) is the subject of Section 4. The $\mathbf{V}_j^T \mathbf{V}_j$ is seen to be the inverse covariance structure of \mathbf{S}_j and (9) is equivalent to minimizing the Mahalanobis distance. The variance, Σ_j , of the noise on \mathbf{S}_j is incorporated by:

$$\Sigma_j^{-1} = \mathbf{V}_j^T \mathbf{V}_j. \quad (10)$$

An approach with this uncertainty formulation is seen to deal with the three types of identified errors:

- **Bad Feature Locations.** Assuming Gaussian noise, the weights can be constructed to incorporate the uncertainty structure of the 2D features as shown above. This approach can even deal with arbitrary Gaussian noise. With the presented formulation of the weights-one \mathbf{V}_j matrix per column-covariance between the x and y coordinates can be expressed. An extension to general covariance between all features would require a 3D tensor formulation of the weights.
- **False Matches.** A mismatched 2D feature can be down weighted, even approaching zero, leaving it out of the optimization to any desired extent.
- **Missing Features.** This is equivalent to predicting that the missing 2D feature is located somewhere in the image, but that the uncertainty of the prediction is very high.

The information of missing features should be apparent from the feature matching algorithm and should be directly expressed in Σ_j . If prior knowledge about the distribution of bad feature location is present, this can also be expressed in Σ_j . However, this not a requirement. *False Matches* are almost by nature unknown a priori so robust statistical techniques are employed to deal with these errors. Note that, in the absence of prior knowledge, the Σ_j is initialized as the identity matrix.

3.2 Adaptive Weighting

As described in [14], [29], and [31], among others, false matches give rise to outliers in the data. Here, outliers are understood as 2D features, where the residual, i.e., the distance between the original 2D feature and it's reprojected counterpart, is large. If the percentage of false matches is relatively low (considerably lower than 50 percent), false matches can be dealt with effectively by diminishing the effect of outliers. In practice, this is done via a robust error function. Popular robust functions are the truncated quadratic and Huber M-estimator (sometimes called Huber Norm), see Fig. 2. For a detailed discussion, the reader is referred to [2]. The use of robust error functions efficiently deals with the problem of false matches.




Fig. 2. Three popular error functions with $k = 3$ for the truncated quadratic and the Huber M-estimator.

These error functions are implemented via Iteratively Reweighted Least Squares (IRLS). IRLS works by iteratively fitting the model to the data by minimizing the weighted least squares of the residuals. These weights are then altered such that the residuals are reweighted according to the desired error function and not the 2-norm. In this approach, the weights, w_{ij} , from the IRLS are collected in:

$$\mathbf{W}_j = \begin{bmatrix} w_{1j} & & & & \\ & \ddots & & & \\ & & w_{nk} & & \\ & & & w_{1j} & \\ & & & & \ddots \\ & & 0 & & & w_{nk} \end{bmatrix} \quad (11)$$

and incorporated by letting:

$$\mathbf{V}_j^T \mathbf{V}_j = \mathbf{W}_j^T \mathbf{W}_j.$$

In case of a Gaussian prior on the 2D features with covariance Σ_j , the \mathbf{V}_j are given by:

$$\mathbf{V}_j^T \mathbf{V}_j = \mathbf{W}_j^T \Sigma_j^{-1} \mathbf{W}_j.$$

As an example, the reweighting formula for the truncated quadratic is:

$$w_{ij} = \begin{cases} 1 & \|N_{ij}\|_{\Sigma_{ij}} < k \\ \sqrt{\frac{k^2}{N_{ij}^2}} & \|N_{ij}\|_{\Sigma_{ij}} > k, \end{cases} \quad (12)$$

where N_{ij} is the residual on datum ij , w_{ij} is the corresponding weight, and k is a user-defined constant relating to the image noise. Here, $\|\cdot\|_{\Sigma_{ij}}$ denotes the Mahalanobis distance induced by Σ_j . If no prior is available, the 2-norm is used. The parameter k is an indication of the general image noise in the image. For a detailed discussion of how to choose k , refer to [13]. It is noted that the experimental results (see Fig. 20) show that the proposed method is rather robust toward the choice of k .

It is also noted that almost arbitrary error functions can be implemented via the IRLS approach. *This allows an arbitrary*




Fig. 3. Overview of the proposed reweighting scheme. This is employed to deal effectively with false matches.

noise model on the bad feature locations by enabling the implementation of the induced error function. The scheme for combining IRLS with the weighted factorization is illustrated in Fig. 3.

4 SEPARATION WITH WEIGHTS

The main computational problem in the proposed approach is to determine a solution to the weighted least squares problem (9). The solution to (9) is \mathbf{M} and \mathbf{P} given \mathbf{S} and \mathbf{V}_j . Note that an SVD cannot be applied as for (5). To solve (9), a method similar to the idea in the Christy-Horad factorization algorithm [6] is proposed. This method is generally known as surrogate modeling, see, e.g., [3]. Surrogate modeling works by applying a computationally “simpler” model to iteratively approximate the original “hard” problem.

The best known example of surrogate modeling is probably the Newton optimization method. Here, a second order polynomial is approximated to the objective function in each iteration and a temporary optimum is achieved. This temporary optimum is then used to make a new second order approximation and, thus, a new temporary optimum. This is continued until convergence is achieved.

Here, (5) is used to iteratively approximate (9), getting a temporary optimum, which in turn can be used to make a new approximation. The approximation is performed by modifying the original data, \mathbf{S} , such that the solution to (5) with the modified data, $\tilde{\mathbf{S}}$, is the same as (9) with the original data. By letting $\tilde{\mathbf{S}}$ denote modified data, the goal is to obtain:

$$\begin{aligned}
 \min_{\mathbf{M}, \mathbf{P}} \sum_{j=1}^n \|\mathbf{V}_j(\mathbf{S})_j - \mathbf{M}\mathbf{P}_j\|_2^2 &= \\
 \min_{\mathbf{M}, \mathbf{P}} \sum_{j=1}^n \mathbf{N}_j^T \mathbf{V}_j^T \mathbf{V}_j \mathbf{N}_j &\stackrel{def}{=} \\
 \min_{\mathbf{M}, \mathbf{P}} \sum_{j=1}^n \tilde{\mathbf{N}}_j^T \tilde{\mathbf{N}}_j &= \\
 \min_{\mathbf{M}, \mathbf{P}} \sum_{j=1}^n \|\tilde{\mathbf{S}}_j - \mathbf{M}\mathbf{P}_j\|_2^2, &
 \end{aligned} \tag{13}$$




Fig. 4. A geometric illustration of how the data is modified in Steps 3 and 4 of the proposed algorithm for separation with weights.

where $\mathbf{N} = [\mathbf{N}_1 \mathbf{N}_n]$ denotes the residuals:

$$\mathbf{N}_j = \mathbf{S}_j - \mathbf{M}\mathbf{P}_j.$$

Hereby, the subspace, \mathbf{M} , can be found via SVD and \mathbf{P} via the normal equations once \mathbf{M} is known. Let q denote the iteration number, then the algorithm goes as follows:

1. **Initialize.** $\tilde{\mathbf{S}}^0 = \mathbf{S}$, $q = 1$.
2. **Estimate Model.** Get \mathbf{M}^q by the singular vectors corresponding to the three largest singular values of $\tilde{\mathbf{S}}^{q-1}$, via SVD. Get \mathbf{P}^q from:

$$\forall j: \quad \mathbf{P}_j^q = [\mathbf{M}^{qT} \mathbf{V}_j^T \mathbf{V}_j \mathbf{M}^q]^{-1} \mathbf{M}^{qT} \mathbf{V}_j^T \mathbf{V}_j \cdot \mathbf{S}_j.$$

3. **Calculate Residuals.** $\mathbf{N}^q = \mathbf{S} - \mathbf{M}^q \cdot \mathbf{P}^q$.
4. **Modify Data.**

$$\begin{aligned}
 \forall j: \tilde{\mathbf{N}}_j^q &= \mathbf{V}_j \mathbf{N}_j^q \\
 \tilde{\mathbf{S}}^q &= \mathbf{M}^q \mathbf{P}^q + \tilde{\mathbf{N}}^q.
 \end{aligned}$$

5. **If Not Stop.** $q = q + 1$, go to 2. The stop criteria is:

$$\|\mathbf{N}^q - \mathbf{N}^{q-1}\|_\infty < \text{tolerance}.$$

As illustrated in Fig. 4 the data, \mathbf{S} , is modified such that the Frobenius norm of the modified residuals, $\tilde{\mathbf{N}}_j^q$, are equal to the norm of the original residuals, \mathbf{N}_j^q , in the norm induced by the weights, \mathbf{V}_j . The last part of Step 2 ensures that the residual, \mathbf{N}_j , is orthogonal to \mathbf{M} in the induced norm since $\mathbf{M}^q \mathbf{P}_j^q$ is the projection of \mathbf{S}_j onto \mathbf{M}^q in the induced norm.

4.1 Separation Performance with Given Weights

A quasi-Newton method could also be used to solve (9). One of the most effective quasi-Newton methods is the Broyden [4], Fletcher [10], Goldfarb [11], and Shanno [23] method (BFGS). However, this is not recommended since the proposed method is faster and more reliable partly because with a “standard” optimization method, the problem is very likely to become ill-conditioned due to the potentially large differences in weights.

To illustrate this, the proposed method and the BFGS were tested against each other, see Table 1. The \mathbf{S} matrix was formed by (3) with noise added from a compound

TABLE 1
Computational Time Comparison of the Proposed Algorithm with MatLab's BFGS (fminu()),—Denotes that the Optimization Did Not Converge Due to Ill-Conditioning

\mathbf{S} $k \times n$	Noise Level	Proposed Method	BFGS	Flop Ratio
20x40	0.02	1.20e+07	2.32e+08	19.33
20x40	0.10	1.58e+07	5.81e+08	36.73
20x40	0.50	5.50e+07	4.22e+08	7.67
40x40	0.02	7.20e+07	1.99e+09	27.58
40x40	0.10	1.15e+08	3.64e+09	31.73
40x40	0.50	3.59e+08	—	—
80x40	0.02	5.17e+08	1.78e+10	34.41
80x40	0.10	8.00e+08	7.08e+10	88.52
80x40	0.50	2.30e+09	8.74e+10	37.93

The computational time is measured in Flops.

Gaussian distribution. The compound distribution consisted of two Gaussian distributions, one with a standard deviation 10 times larger than the other. The fraction of the larger varying Gaussian is the *Noise Level*. It is seen that the proposed method performs better than BFGS and that the BFGS approach did not converge for $\mathbf{S} = 40 \times 40$ and *Noise Level* = 0.5, due to ill-conditioning.

A formal proof of convergence seems to be infeasible. However, this is a common problem for most numerical optimization schemes. The convergence has been followed closely during the tests of the algorithm. These empirical results show that the convergence has been very good.

A thorough investigation of the numerical properties of the algorithm is presented in [21]. Here, the proposed method for separation is also compared to other alternatives besides BFGS with favorable results. Among these are the types of algorithms where \mathbf{M} is estimated given \mathbf{P} and vice versa in an iterative manner, like the one presented in [24].

5 THE OBJECT FRAME ORIGIN

In order to achieve robust factorization, the calculation scheme of the object frame origin, (x_{o_i}, y_{o_i}) , applied in the Christy-Horad algorithm needed to be improved. The object frame origin is the point around which the camera model is linearized and errors here will propagate to the rest of the estimation. Making the estimation origin robust is essential to dealing efficiently with erroneous data. In the Christy-Horad algorithm, the 2D features corresponding to an arbitrary 3D feature are chosen. This scheme is error prone, with erroneously tracked features, e.g., if the chosen feature corresponds to a false feature match.

A natural alternative, would be to iteratively calculate the center of mass from the estimated structure and reproject it onto the estimated camera positions. This however turns out to be *unstable*.




Fig. 5. Overview of the proposed approach. This approach also deals with estimating the object frame origin robustly.

We still choose to use the center of mass, but calculated from the 2D features and the reprojected 3D features in a similar manner as (7). Combining 2D features and the reprojected 3D features allows us to diminish the influence of the 2D features with high uncertainty, while stabilizing the estimation with the reliable 2D features. The formula for the object frame origin of frame i is given by:

$$\begin{bmatrix} x_{o_i} \\ y_{o_i} \end{bmatrix} = \frac{\sum_{j=1}^n \begin{bmatrix} \tilde{x}_{ij} \\ \tilde{y}_{ij} \end{bmatrix} (1 + \epsilon_{ij})}{n + \sum_{j=1}^n \epsilon_{ij}}, \quad (14)$$

where ϵ_{ij} is defined in (8) and $(\tilde{x}_{ij}, \tilde{y}_{ij})$ is a weighted mean between the 2D feature and the reprojected 3D feature defined by.

$$\begin{bmatrix} \tilde{x}_{ij} \\ \tilde{y}_{ij} \end{bmatrix} = \gamma_{ij} \begin{bmatrix} x_{ij} \\ y_{ij} \end{bmatrix} + (1 - \gamma_{ij}) \begin{bmatrix} \hat{x}_{ij} \\ \hat{y}_{ij} \end{bmatrix}. \quad (15)$$

Here, γ_{ij} denotes the weight. This weighted mean ensures robustness toward outliers since outliers are partly replaced by the estimate of the model. If $\mathbf{V}_j^T \mathbf{V}_j$ is *diagonal*, then the weight γ_{ij} is given by:

$$\gamma_{ij} = \frac{1}{\pi} \arctan(\alpha(v_{ij} - \frac{1}{k'})) + \frac{1}{2}, \quad (16)$$

where v_{ij} is the i th diagonal element of \mathbf{V}_j . The “cutoff” value k' serves the same function as k in the truncated quadratic and the Huber M-estimator. It is seen that (16) is a smooth version of the step function around $\frac{1}{k'}$, where α controls how smooth this approximation is.

If $\mathbf{V}_j^T \mathbf{V}_j$ is *not diagonal*, then its principal components are used. So, instead of v_{ij} , the eigenvalues of $\mathbf{V}_j^T \mathbf{V}_j$ and a linear combination of (x_{ij}, y_{ij}) , are used in (14) and (15). This linear combination corresponds to the eigenvectors of $\mathbf{V}_j^T \mathbf{V}_j$. With this modification to the approach, the final flow chart is seen in Fig. 5.



Fig. 6. A sample frame from the Hotel sequence with 197 tracked features some of these are mismatched.

6 EUCLIDEAN RECONSTRUCTION

The objective of Euclidean reconstruction is to estimate \mathbf{A} in (6) such that the estimated \mathbf{a}_i , \mathbf{b}_i , and \mathbf{c}_i of (1) are as orthonormal as possible. In the paraperspective case [22], which is the linearization used by Christy and Horaud [6], the M_i s composing \mathbf{M} are given by:

$$\mathbf{M}_i = \frac{1}{t_i^z} \begin{bmatrix} \mathbf{a}_i^T - x_{oi}\mathbf{c}_i \\ \mathbf{b}_i^T - y_{oi}\mathbf{c}_i \end{bmatrix} = \begin{bmatrix} \mathbf{I}_i^T \\ \mathbf{J}_i^T \end{bmatrix},$$

where (x_{oi}, y_{oi}) is the object origin projected in frame i .

Since the paraperspective approximation is obtained by linearizing $\frac{1}{t_i} \mathbf{c}_i^T \cdot \mathbf{P}_j$, the orthonormal constraints are restricted to \mathbf{a}_i and \mathbf{b}_i . With $\mathbf{Q} = \mathbf{A}\mathbf{A}^T$ these constraints can be formulated as [6], [22]:

$$\begin{aligned} \forall i \quad & \mathbf{a}_i^T \mathbf{Q} \mathbf{a}_i = \mathbf{b}_i^T \mathbf{Q} \mathbf{b}_i \Rightarrow \\ \forall i \quad & \frac{\mathbf{I}_i^T \mathbf{Q} \mathbf{I}_i}{1 + x_{oi}^2} - \frac{\mathbf{J}_i^T \mathbf{Q} \mathbf{J}_i}{1 + y_{oi}^2} = 0 \\ \forall i \quad & \mathbf{a}_i^T \mathbf{Q} \mathbf{b}_i = 0 \Rightarrow \\ \forall i \quad & \mathbf{I}_i^T \mathbf{Q} \mathbf{J}_i - \frac{x_{oi}y_{oi}(\mathbf{I}_i^T \mathbf{Q} \mathbf{I}_i)}{2(1 + x_{oi}^2)} - \frac{x_{oi}y_{oi}(\mathbf{J}_i^T \mathbf{Q} \mathbf{J}_i)}{2(1 + y_{oi}^2)} = 0. \end{aligned}$$

With noise, this cannot be achieved for all i and a least squares solution is sought. In order to avoid the trivial null-solution, the constraint $\mathbf{a}_1^T \mathbf{Q} \mathbf{a}_1 = \mathbf{b}_1^T \mathbf{Q} \mathbf{b}_1 = 1$ is added [6], [22] and the problem is linear in the elements of \mathbf{Q} .

Unfortunately it is impossible to reconstruct \mathbf{A} if \mathbf{Q} has negative eigenvalues. This problem indicates that an unmodeled distortion has overwhelmed the third singular value of \mathbf{S} [22]. This is a fundamental problem when the factorization method is used on erroneous data.

To solve this problem, we propose to parameterize \mathbf{Q} as:

$$\mathbf{Q}(\mathbf{e}, \lambda) = \mathbf{R}(\mathbf{e}) \begin{bmatrix} \lambda_1^2 & 0 & 0 \\ 0 & \lambda_2^2 & 0 \\ 0 & 0 & \lambda_3^2 \end{bmatrix} \mathbf{R}(\mathbf{e})^T, \quad (17)$$

where $\mathbf{R}(\mathbf{e})$ is a rotation matrix with the three Euler angles, denoted by \mathbf{e} . The term $\mathbf{a}_1^T \mathbf{Q} \mathbf{a}_1 = \mathbf{b}_1^T \mathbf{Q} \mathbf{b}_1 = 1$ is replaced by




Fig. 7. A section of the Hotel sequence illustrating where feature 64 is mismatched. The correct position of feature 64 is at the end of the residual vector (bottom left of the image).




Fig. 8. A section of the Hotel sequence where feature 64 is not erroneous. Notice how little effect the false match error has with the proposed method.

$\det(\mathbf{A}) = 1$ such that the overall scale of \mathbf{A} is much more robust and less sensitive to noise in a particular frame.

Hence, the estimation of \mathbf{Q} is a nonlinear optimization problem in six variables, with a guaranteed symmetric positive definite \mathbf{Q} . Our experience shows that this approach to the problem is well-behaved with a quasi-Newton optimization method.

7 EXPERIMENTAL RESULTS

In order to illustrate the proposed algorithms ability to deal with the errors identified, it has been applied to a set of real images. To provide a more systematic test, it was then applied to a set of simulated data, where arbitrary errors could be induced.

7.1 Real Data

The proposed approach was run on three sequences, demonstrating different properties. The first sequence was the hotel sequence [19] with accompanying features, see Fig. 6. Some of these features were mismatched. To illustrate the handling of mismatched features, notice the different position of feature 64 in Fig. 7 and Fig. 8. The lines in the images denote the residual between the tracked 2D features and the back-projected 3D estimates. The correct position of a feature is determined by the location with the greatest support.

The proposed approach with the truncated quadratic error function was applied to the hotel sequence ($k = 3$, $k' = 2$, $\alpha = 20$). The result is shown in Fig. 7 and Fig. 8. Two things should be noted here. First, the mismatched feature 64 does not effect the estimate of the other features. Second, the back-projections of the 3D estimate of feature 64 are located correctly in all images (including seven where the 2D feature itself is mismatched).

TABLE 2
Comparison of the 95 Percent Smallest Residuals

Error Function:	$\frac{1}{n} \sum_i Res _i$
Truncated Quadratic (robust)	1.94 <i>pixels</i> ²
2-Norm (non-robust)	3.40 <i>pixels</i> ²

It is noted that, without subpixel feature location, this number is highly unlikely to be lower than 1.




Fig. 9. Same section and frame as Fig. 8., but with the Christy-Horaud method. Notice the increased effect of the errors.




Fig. 10. Weight evolution for feature 64 plotted as a function of frame and iteration

Comparing with the Christy-Horaud approach, Fig. 9, it is seen that the mismatch error in Fig. 7 effects the overall estimation of feature 64's 3D position. To give a more quantitative evaluation, the residuals of the nonerroneous features were summed up—it is assumed that there were no more than 5 percent errors. It is seen from Table 2 that, due to the capability of implementing a robust error function, the fit to the nonerroneous data is improved significantly.

To illustrate the reweighting process, the evolution of the weights off the mismatched feature 64 is depicted in Fig. 10. After the first iteration, all the 2D features are down weighted since only an erroneous 3D estimate exists due to the mismatch features and the uniform weighting. In the following iterations, the correctly located features obtain increasing weights, whereas the weights of the mismatch feature decrease toward zero. The plot directly shows how the proposed method can be used to detect mismatched features.



Fig. 11. A sample image from the kitchen sequence.




Fig. 12. Close up on the kitchen sequence. Notice the location of feature 5.

To demonstrate the proposed approach's ability to deal with missing features, it was run on the kitchen sequence [19], see Fig. 11. Here, some of the accompanying features were missing, e.g., feature 5 in Fig. 13. Note that the back-projection of feature 5 in Fig. 13 is located correctly, indicating that the 3D estimate is correct. From the depicted residuals, which are hardly visible, it is seen that the missing features does not disrupt the structure and motion estimation. Thus giving the desired result in dealing with missing features.

The proposed modification to the Christy-Horaud algorithm should not considerably decrease its ability to deal with perspective reconstruction. To validate this, both approaches were applied to a sequence with considerable depth, see Fig. 14. The 2D features were carefully hand-tracked to ensure that the algorithms were directly comparable. The results were evaluated by comparing the two estimated structures via the standard shape distance measure from statistical shape analysis, the Procrustes distance [9], [12]. The Procrustes distance is obtained by normalizing the two structures and then applying the similarity transform [14] such that the mean squared error is minimized. The remaining mean squared error is then the Procrustes distance. The resulting Procrustes distance was 0.03. This implies that the Christy-Horaud




Fig. 13. Closeup on the kitchen sequence. A circle around a feature denotes it is the estimated location. The reason being that the respective feature is missing in this frame. Not that the estimated position of feature 5 correspond to its position in figure 12.



Fig. 14. A sample frame from an image sequence of Thorvaldsens Museum in Copenhagen with 20 hand tracked features through eight frames.




Fig. 15. The setup for the simulated data. A box with 100 features is "photographed" from eight views. These views are marked on the trajectory curve.

method with the proposed enhancement maintains its ability to deal with perspective data.

7.2 Simulated Data

To perform a more systematic test of the proposed algorithms ability to deal with errors in the tracked data, a simulated data set was created, see Fig. 15. Several kinds of errors were introduced into this data set, hereby testing the approach with respect to the three identified types of errors. The Huber M-estimator and truncated quadratic error functions were




Fig. 16. Errors for the simulated data corrupted with a compound distribution of two Gaussians. The abscissa denotes the likelihood, $[0...1]$, that samples are obtained from the higher varying Gaussian.




Fig. 17. Errors for corrupting the simulated data by removing the 2D features at random. The abscissa indicates how many percent of the data has been removed. No difference is seen between the two error functions.

applied with the parameters settings $k = 0.02^1$ and $k' = 2$ and we set $\alpha = 20$. The results were evaluated by the Procrustes distance [9], [12]. In the extreme cases where an estimation process did not converge, this is denoted by a missing measurement.

The first experiment consists of corrupting the 2D features of the simulated data by a compound distribution of two Gaussians, one with a standard deviation of 0.005 and the other 10 times larger. Two different error functions were applied: the Huber M-estimator and the truncated quadratic. For comparison, the Christy-Horad algorithm was also applied. From Fig. 16, it is seen that the choice of error functions has a considerable effect on the result and that the proposed approach is capable of implementing them.

1. It should be noted that k depends on the image noise and, hence, also on the image size. In this simulated data, the image size was unnaturally small (0.25×0.25) and, as such, $k = 0.02$ should not be seen as a guide line.




Fig. 18. Errors for the simulated data with swapped features emulating mismatched features. The data set upon which features were swapped was the original uncorrupted. The abscissa denoted the percentage of altered features.

In the second experiment, an increasing number of the 2D features were removed from the data set. These missing features were emulated as being located in the center of the frame in question with a weight equal to 10^{-6} of the weight of the normal 2D features. From the results, it is seen that the proposed approach is highly robust toward this type of error since up to 40 percent missing features corrupts the estimated structure by less than 10^{-2} , see Fig. 17. The approach of Christy and Horaud could not work on this data and was not included.

To test the tolerance to mismatched features, we emulated these by swapping 2D features of the simulated data. The results in Fig. 18 illustrate that the proposed approach also deals efficiently with mismatched features and has been shown to be robust toward the identified types of errors.

To challenge the proposed approach, all the experiments on the simulated data are performed with up to 35-40 percent errors. It should be noted that the algorithm works very well

TABLE 3
Results with Noise Exclusively on the Feature Christy-Horaud Uses as Object Frame Origin

Factorization	Percent	Mean Procrustes
Approach	Convergence	Distance
Christy-Horaud	60%	0.0153
Proposed Approach	100%	0.0055

This amounts to 1 percent of the data being erroneous.

with up to 10-20 percent errors. This is the amount of errors that the proposed approach is expected to work on.

To evaluate the effect of the proposed approach for estimating the object frame origin, noise was added exclusively to the feature that Christy-Horaud uses as object frame origin. The evaluation was performed on the simulated data and the noise was Gaussian with a standard deviation of 0.05. Five experiments were conducted comparing the Christy-Horaud approach to the proposed approach. The results are shown in Table 3. It is seen that significant improvement is obtained with respect to error as well as convergence. For validation purposes the same experiment was made, but, this time, noise was added to a different feature than the one used as the object origin by Christy-Horaud. In this case the two approaches gave similar results. Note also that the original approach for estimation of the object frame origin could be problematic to use with missing features since the 2D feature chosen as the object frame origin is likely to be missing in at least one frame.

In order to demonstrate the benefits of the proposed method for Euclidean reconstruction, the experiments with mismatched features were repeated *without* the proposed approach. Instead, the original method for Euclidean reconstruction proposed in [6], [22] was applied. The number of runs that did not converge has been summed up with and

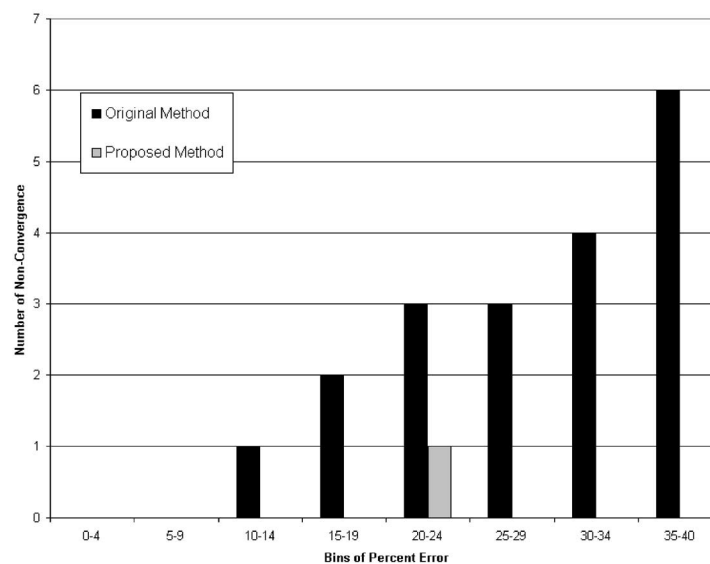


Fig. 19. The number of non-converging runs on the experiment of Fig. 18, using the Huber M-estimator as error function. The experiment was made with and without the proposed method for Euclidean reconstruction. The runs are pooled in bins of five to give a better overview.




Fig. 20. The experiment of Fig. 18 repeated with different parameter settings for k in the truncated quadratic error function.

without the proposed method. The change is dramatic, as seen in Fig. 19.

To illustrate, that the proposed method is not overly sensitive to the choice of error function and the involved parameters, these were varied. This was done with the truncated quadratic error function in the experiments with mismatched features, see Fig. 18. The k parameter was altered with ± 50 percent. From the results in Fig. 20, it is seen that the proposed approach is not overly sensitive to the choice of parameters in the error function.

8 CONCLUSION AND DISCUSSION

An approach for applying arbitrary error functions in the factorization algorithms for structure and motion has been presented. This ability has been exploited to implement robust statistical techniques, whereby errors in the 2D features have been dealt with effectively. The algorithm has been implemented as an extension to the Christy-Horaud factorization algorithm [6], whereby perspective reconstruction is achieved. The core of the approach is, however, so general, that it can also be applied to most other factorization schemes, like [1], [15], [26], [27], [22], [28]. This approach has been applied to both simulated and real data, where it is demonstrated that the algorithm deals well with all the expected types of errors, thus making the proposed approach a robust factorization method.

To further investigate the possibilities of the proposed approach, we aim at implementing it with other factorization algorithms, e.g., [15], hence making a robust factorization algorithm with uncalibrated cameras. The proposed method also allows for more elaborate noise structures incorporating covariance within and between images. This covariance structure has been implemented with the algorithm of Section 4 and sparsely tested, but should be thoroughly tested on real image data in order to evaluate its robustness. This is the goal of ongoing research.

ACKNOWLEDGMENTS

The authors would like to thank the Modeling by Videotaping Group at the Robotics Institute, Carnegie Mellon University, for supplying the hotel and kitchen data sequences and the accompanying tracked features [19]. They would also like to thank Radu Horaud and Claus Gramkow for valuable comments to this work and, lastly, Hans Bruun Nielsen for comments on the numerical aspects.

REFERENCES

- [1] R. Berthilsson and K. Åström, "Extension of Affine Shape," *J. Math. Imaging and Vision*, vol. 11, no. 2, pp. 119-136, 1999.
- [2] M.J. Black and A. Rangarajan, "On the Unification of Line Processes, Outlier Rejection, and Robust Statistics with Applications in Early Vision," *Int'l J. Computer Vision*, vol. 19, no. 1, pp. 57-91, 1996.
- [3] A.J. Booker, J.E. Dennis, P.D. Frank, D.B. Serafini, V. Torczon, and M.W. Trosset, "A Rigorous Framework for Optimization of Expensive Functions by Surrogates," *Structural Optimization*, vol. 17, no. 1, pp. 1-13, 1999.
- [4] C.G. Broyden, "The Convergence of a Class of Double-Rank Minimization Algorithms. ii. The New Algorithm," *J. Inst. of Math. and Its Applications*, vol. 6, no. 3, pp. 222-231, 1970.
- [5] S. Christy and R. Horaud, "Euclidian Shape and Motion from Multiple Perspective Views by Affine Iterations," Technical Report 2421, INRIA, Dec. 1994.
- [6] S. Christy and R. Horaud, "Euclidean Shape and Motion from Multiple Perspective Views by Affine Iteration," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 18, no. 11, pp. 1098-1104, Nov. 1996.
- [7] J. Costeira and T. Kanade, "A Multibody Factorization Method for Independently Moving Objects," *Int'l J. Computer Vision*, vol. 29, no. 31, pp. 159-179, 1998.
- [8] D. Dementhon and L. Davis, "Model-Based Object Pose in 25 Lines of Code," *Int'l J. Computer Vision*, vol. 15, no. 1-2, pp. 123-141, 1995.
- [9] I.L. Dryden and K.V. Mardia, *Statistical Shape Analysis*. Wiley, 1998.
- [10] R. Fletcher, "A New Approach to Variable Metric Algorithms," *Computer J.*, vol. 13, no. 3, pp. 317-322, 1970.
- [11] D. Goldfarb, "A Family of Variable Metric Methods Derived by Variational Means," *Mathematics of Computation*, vol. 24, pp. 23-26, 1970.
- [12] J.C. Gower, "Generalized Procrustes Analysis," *Psychometrika*, vol. 40, pp. 33-50, 1975.
- [13] F.R. Hampel, P.J. Rousseeuw, E.M. Ronchetti, and W.A. Stahel, *Robust Statistics*. John Wiley & Sons, 1986.
- [14] R. Hartley and A. Zisserman, *Multiple View Geometry*. Cambridge Univ. Press, 2000.
- [15] A. Heyden, R. Berthilsson, and G. Sparr, "An Iterative Factorization Method for Projective Structure and Motion from Image Sequences," *Image and Vision Computing*, vol. 17, no. 13, pp. 981-991, 1999.
- [16] M. Irani and P. Anandan, "Factorization with Uncertainty," *Proc. European Conf. Computer Vision '2000*, pp. 539-553, 2000.
- [17] D.W. Jacobs, "Linear Fitting with Missing Data for Structure-from-Motion," *Computer Vision and Image Understanding*, vol. 82, no. 1, pp. 57-81, 2001.
- [18] T. Kanade and D. Morris, "Factorization Methods for Structure from Motion," *Philosophical Transactions of the Royal Society of London*, vol. A, no. 356, pp. 1153-1173, 1998.
- [19] The Modeling by Videotaping Group at the Robotics Inst., Carnegie Mellon Univ., <http://www.ius.cs.cmu.edu/IUS/mbvc0/www/modeling.html>. 2000.
- [20] D. Morris and T. Kanade, "A Unified Factorization Algorithm for Points, Line Segments and Planes with Uncertainty Models," *Proc. Int'l Conf. Computer Vision '98*, pp. 696-702, Jan. 1998.
- [21] H.B. Nielsen and H. Aanæs, "Separation of Structure and Motion by Data Modification," technical report, Informatics and Mathematical Modeling, Technical University of Denmark, to appear.
- [22] C. Poelman and T. Kanade, "A Paraperspective Factorization Method for Shape and Motion Recovery," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 19, no. 3, pp. 206-218, Mar. 1997.

- [23] D.F. Shanno, "Conditioning of Quasi-Newton Methods for Function Minimization," *Mathematics of Computation*, vol. 24, pp. 647-656, 1970.
- [24] H.-Y. Shum, K. Ikeuchi, and R. Reddy, "Principal Component Analysis with Missing Data and Its Application to Polyhedral Object Modeling," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 17, no. 9, pp. 854-867, 1995.
- [25] C.C. Slama Editor, *Manual of Photogrammetry*, fourth edition, Am. Soc. of Photogrammetry, 1984.
- [26] G. Sparr, "Simultaneous Reconstruction of Scene Structure and Camera Locations from Uncalibrated Image Sequences," *Proc. Int'l Conf. Pattern Recognition '96*, pp. 328-333, 1996.
- [27] P. Sturm and B. Triggs, "A Factorization Based Algorithm for Multi-Image Projective Structure and Motion," *Proc. European Conf. Computer Vision '96*, vol. 2, pp. 709-720, 1996.
- [28] C. Tomasi and T. Kanade, "Shape and Motion from Image Streams under Orthography: A Factorization Method," *Int'l J. Computer Vision*, vol. 9, no. 2, pp. 137-154, Nov. 1992.
- [29] P.H.S. Torr and D.W. Murray, "The Development and Comparison of Robust Methods for Estimating the Fundamental Matrix," *Int'l J. Computer Vision*, vol. 24, no. 3, pp. 271-300, 1997.
- [30] B. Triggs, P.F. McLauchlan, R.I. Hartley, and A.W. Fitzgibbon, "Special Sessions—Bundle Adjustment: A Modern Synthesis," *Lecture Notes in Computer Science*, vol. 1883, pp. 298-372, 2000.
- [31] Z. Zhang, R. Deriche, O. Faugeras, and Q.T. Luong, "A Robust Technique for Matching Two Uncalibrated Images through the Recovery of the Unknown Epipolar Geometry," *J. Artificial Intelligence*, vol. 78, no. 1, pp. 87-119, 1995.



Henrik Aanaes received the MS degree in engineering from the Technical University of Denmark in 2000 and is, at present, pursuing the PhD degree in computer vision at the same institution. His current research interests include the geometry and algebra of multiple views, incorporating statistical methods—especially 3D model priors—into structure and motion algorithms and surface estimation from multiple views, including the reflectance and texture properties.



Rune Fisker holds the PhD degree in computer vision and image processing from the Technical University of Denmark which was obtained in 2000. He graduated from the Technical University of Denmark in 1996. After finishing his PhD, he joined 3Shape Inc., creating the next generation of 3D scanners and 3D modeling software. He has published more than 10 papers and filed two patents. His current research interests include 3D reconstruction, 3D modeling, shape analysis, deformable template models, image segmentation and applied statistics.



Kalle Åström received the BSc degree in mathematics in 1990, MSc degree in Engineering Physics in 1991, and PhD degree in mathematics in 1996 from Lund University, Sweden. His thesis was awarded Best Nordic PhD Thesis in pattern recognition and image analysis 1995-1996 at the Scandinavian Conference in Image Analysis, 1997. He has been a postdoctoral research fellow and subsequently, research associate in the Department of Mathematics, Lund University. He is the president of the Swedish Society for Automated Image Analysis. His teaching includes undergraduate and graduate courses in mathematics, image processing, and computer vision. His current research interests include stochastic analysis of low-level vision, computer vision for autonomous guided vehicles, geometry and algebra of multiple views of points, curves and surfaces. He is a member of the IEEE Computer Society.



Jens Michael Carstensen received the MS degree in engineering from the Technical University of Denmark in 1988 and the PhD degree in statistical image analysis from the Technical University of Denmark in 1992. He joined the Technical University of Denmark faculty in 1992 as an assistant research professor and is currently an associate professor of statistical image analysis. His current research interests include statistical methods for image analysis, Markov random fields, texture analysis, shape and appearance models, industrial and biotechnological applications.

► **For more information on this or any other computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.**