Clustered Stochastic Optimization for Object Recognition and Pose Estimation

Juergen Gall Bodo Rosenhahn
Hans-Peter Seidel

MPI–I–2007–4-001 March 2007

Authors' Addresses

Juergen Gall
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken
Germany

Bodo Rosenhahn
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken
Germany

Hans-Peter Seidel
Max-Planck-Institut für Informatik
Stuhlsatzenhausweg 85
66123 Saarbrücken
Germany

Acknowledgements

Our research is funded by the Max-Planck Center for Visual Computing and Communication.

Abstract

We present an approach for estimating the 3D position and in case of articulated objects also the joint configuration from segmented 2D images. The pose estimation without initial information is a challenging optimization problem in a high dimensional space and is essential for texture acquisition and initialization of model-based tracking algorithms. Our method is able to recognize the correct object in the case of multiple objects and estimates its pose with a high accuracy. The key component is a particle-based global optimization method that converges to the global minimum similar to simulated annealing. After detecting potential bounded subsets of the search space, the particles are divided into clusters and migrate to the most attractive cluster as the time increases. The performance of our approach is verified by means of real scenes and a quantative error analysis for image distortions. Our experiments include rigid bodies and full human bodies.

Keywords

Pose Estimation, 3D Model Registration, Stochastic Optimization, Object Recognition, Interacting Simulated Annealing

Contents

1 Introduction

Finding the 3D position and rotation of a rigid object in a set of images from calibrated cameras without any initial information is a difficult optimization problem in a 6-dimensional space. The task becomes even more challenging for articulated objects where the dimensionality of the search space is much higher, e.g., a coarse model of an human skeleton has already 24 degrees of freedom (DoF) yielding a 30-dimensional space. Although the initial pose is essential for many state-of-the-art model-based tracking algorithm, e.g. [3, 4, 21], relatively little attention was paid to the initialization of rigid and articulated models. A manual initialization is usually required, which is time demanding and assumes some expertise on the model and on the world coordinate system.

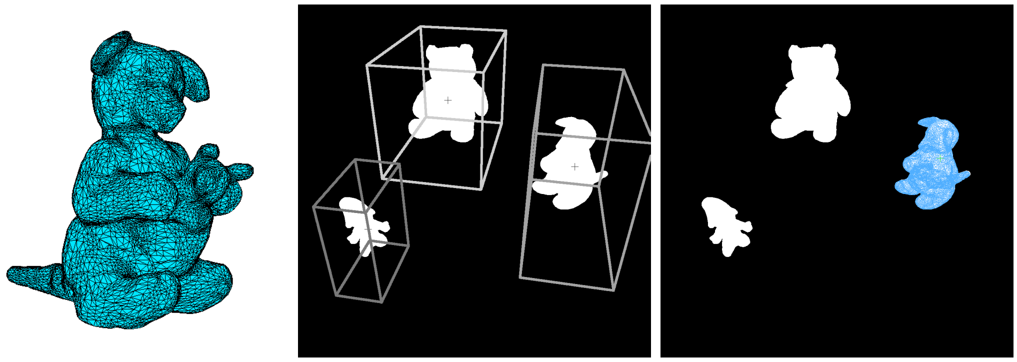

Figure 1.1: **From left to right:** *a)* 3D model of object. *b)* Potential bounded subsets of the search space. *c)* Projection of the mesh. The pose is correctly estimated.

Depending on the image features, there are several techniques for pose estimation in the literature. Edge-based approaches, e.g. [16, 17, 1], align curves or lines of the model to detected edges. They work best for homogeneous objects, however, textured objects and cluttered background typically involve many edges that are not related to the model. Texture-based approaches [12, 15] use correspondences between the textured model and an image for pose estimation. Separate from the fact that they require textured surfaces for self-initialization, the texture needs to be registered to the model beforehand, i.e., a manual initialization is done for the texture acquisition during preprocessing. Furthermore, experiments were only done for rigid models and low dimensional spaces.

Our approach for solving the initialization problem estimates the pose of rigid and articulated objects by minimizing an energy function based only on the silhouette information. Although we are not restricted to silhouettes, the object region has the advantage that it is an appearance independent feature that can be easily extracted from a single frame, e.g. by background subtraction. Since an initial guess is not available, local optimization algorithm like iterative closest point (ICP) [2, 23] are not suitable for this task. For finding the exact pose, we use a novel particle-based global optimization, called *interacting simulated annealing* [10], that converges to the global optimum similar to simulated annealing [14]. In order to deal with multiple objects, we extend the work in [10] by clustering the particles with respect to previously detected bounded subsets of the search space.

After a brief introduction to interacting simulated annealing in Section 2, we give details of our method in Section 3. In Section 4, some extensions for human bodies are explained. The experimental results are discussed in Section 5 followed by a brief conclusion.

2 Interacting Simulated Annealing

Interacting particle systems are well-known as particle filter [8] and approximate a distribution of interest η_t by

$$\eta_t^n := \sum_{i=1}^{n} \pi^{(i)} \delta_{X_t^{(i)}}, \tag{2.1}$$

where $X_t^{(i)}$ are n random variables, termed particles, weighted with $\pi^{(i)}$ and δ is the Dirac measure. In the case of *interacting simulated annealing (ISA)*, the distribution is proportional to a Boltzmann-Gibbs measure

$$g_t(dx) = \exp\left(-\beta_t\, V(x)\right)\, \lambda(dx), \tag{2.2}$$

where $V \geq 0$ is the energy function to minimize, β_t is an annealing parameter that increases with t, and λ is the Lebesgue measure.

Each iteration of Algorithm 1 consists of two steps, namely *Selection* and *Mutation*. During a selection step, the particles are weighted according to the Boltzmann-Gibbs measure (2.2) and then resampled with respect to their weights, where particles with a great weight generate more offspring than particles with lower weight. In the second step, the particles evolve according to a Markov kernel K_t. There are several ways to specify the parameter ϵ_t for the resampling, see for instance [13, 10]. In our experiments, we set $\epsilon_t = 1/\sum_{k=1}^{n} \pi^{(k)}$. If the particle is not accepted, a particle is selected, e.g., by multinomial sampling. An overview of various resampling schemes can be found in [7].

It has been shown that interacting particle systems approximate a Feynman-Kac model [18], i.e., η_t^n converges to a Feynman-Kac distribution η_t as the number of particles n goes to infinity. Assumed that the annealing scheme β_t increases slow enough and the search space is bounded, the flow of Feynman-Kac distributions η_t, which is approximated by interacting simulated annealing, becomes concentrated in the region of global minima of V as t tends

Algorithm 1 Interacting Simulated Annealing Algorithm

1. Initialization
   - Sample $x_0^{(i)}$ from η_0 for all i
2. Selection
   - Set $\pi^{(i)} \leftarrow \exp(-\beta_t V(x_t^{(i)}))$ for all i
   - For i from 1 to n:

     Sample κ from $U[0,1]$

     If $\kappa \leq \epsilon_t \pi^{(i)}$ then

     $\star$ Set $\check{x}_t^{(i)} \leftarrow x_t^{(i)}$

     Else

     $\star$ Set $\check{x}_t^{(i)} \leftarrow x_t^{(j)}$ with probability $\frac{\pi^{(j)}}{\sum_{k=1}^{n} \pi^{(k)}}$
3. Mutation
   - Sample $x_{t+1}^{(i)}$ from $K_t(\check{x}_t^{(i)}, \cdot)$ for all i and go to step 2

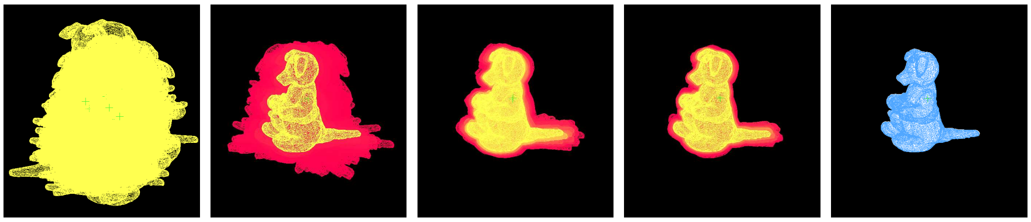

Figure 2.1: Particles at $t = 0, 5, 10, 15$ and 19 for ISA. Particles with a higher weight are brighter, particles with a lower weight are darker. The particles converge to the pose with the lowest energy as t increases. **Most left:** Equally weighted particles after initialization. **Most right:** Estimate after 20 iterations.

to infinity. For more details on the mathematical properties of interacting simulated annealing, we refer to [18].

In [10], the authors evaluated several annealing schemes. In our experiments we used a polynomial scheme, i.e.

$$\beta_t = (t+1)^b \quad \text{for some } b \in (0,1), \tag{2.3}$$

with $b = 0.7$. As Markov kernels K_t, we propose a modified dynamic variance scheme in Section 3.3.

3 Clustered Optimization

3.1 Initial Subsets

Having a binary image for each camera view, where pixels that belong to the foreground are set to 1 else to 0, the pixels are first clustered with repsect to the 8-neighbor connectivity. In order to make the system more robust to noise, clusters covering only a very small area are discarded. In the next step, the 4 corners of the bounding box of each cluster are determined and the projection ray for each corner is calculated. The projection rays are represented as Plücker lines [22], i.e., the 3D line is determined by a normalized vector d and a moment m such that $x \times n = m$ for all x on the line. Provided that two projection rays from different views are not parallel, the midpoint p of the shortest line segment between the two rays l_1 and l_2 is unique and can be easily calculated. If the minimum distance between l_1 and l_2 is below a threshold, p is regarded as a corner of a convex polyhedron. After 8 corners of the polyhedron are detected for two clusters from two different views, the bounding cube is calculated as shown in Figure 1.1 b). In the case of more than two available camera views, each pair of images – starting with the views containing the most clusters – is checked until a polyhedron is found. The corners are similarly refined by calculating the midpoint of the shortest line segment between a ray from another view and a corner of a polyhedron. The resulting bounding cubes provide the initial bounded subsets of the search space.

3.2 Particles

Since we know the 3D model, the pose is determined by a vector in $\mathbb{R}^{6+\sharp joints}$, i.e., each particle is a $6+\sharp joints$-dimensional random vector. The rigid body motion M is represented by the axis-angle representation given by the 6D vector $(\theta\omega, t)$ with $\omega = (\omega_1, \omega_2, \omega_3)$ and $\|\omega\|_2 = 1$. The transformation from

$(\theta\omega, t)$ to M is given by

$$M = \begin{pmatrix} \exp(\theta\hat{\omega}) & t \\ 0 & 1 \end{pmatrix}, \qquad \hat{\omega} = \begin{pmatrix} 0 & -\omega_3 & \omega_2 \\ \omega_3 & 0 & -\omega_1 \\ -\omega_2 & \omega_1 & 0 \end{pmatrix}, \tag{3.1}$$

where $\exp(\theta\hat{\omega})$ can be efficiently computed by the Rodriguez formula [19].

The inverse mapping of the exponential $\log(R)$ can be also computed [19]. When R is the identity matrix, θ is set to 0. For the other case, θ and the rotation axis ω are given by

$$\theta = \cos^{-1}\left(\frac{trace(R) - 1}{2}\right), \qquad \omega = \frac{1}{2\sin(\theta)} \begin{pmatrix} r_{32} - r_{23} \\ r_{13} - r_{31} \\ r_{21} - r_{12} \end{pmatrix}. \tag{3.2}$$

In order to estimate the pose, the mean of a set of rotations $r^{(i)}$ weighted by $\pi^{(i)}$ is required. This can be done by finding a geodesic on the Riemannian manifold determined by the set of 3D rotations. When the geodesic starting from the mean rotation in the manifold is mapped by the logarithm onto the tangent space at the mean, it is a straight line starting at the origin, see [20]. The tangent space is called exponential chart. Hence, the weighted mean rotation $\bar{r}$ satisfies

$$\sum_i \pi^{(i)} \left(\bar{r}^{-1} \star r^{(i)}\right) = 0, \tag{3.3}$$

where $r^{(j)} \star r^{(i)} := \log\left(\exp(r^{(j)}) \cdot \exp(r^{(i)})\right)$ and $r^{-1} := \log\left(\exp(r)^T\right)$. The weighted mean can thus be estimated by

$$\hat{r}_{t+1} = \hat{r}_t \star \left(\frac{\sum_i \pi^{(i)} \left(\hat{r}_t^{-1} \star r^{(i)}\right)}{\sum_i \pi^{(i)}}\right). \tag{3.4}$$

3.3 Mutation

Due to multiple objects as shown in Figure 1.1, each particle belongs to a certain cluster C given by the bounding cubes and denoted by $x^{(i,C)}$. At the beginning, a small number of particles is generated with different orientations located in the center of the cube for each cluster. The complete set of particles is initialized by randomly assigning each particle the values of one of the generated particles. Afterwards, each particle is independently diffused by a normal distribution with mean $x^{(i,C)}$. The covariance matrix is identical for all clusters except for the translation where the variance is given by the edge length divided by 6 such that over 99.5% of the particle centers are inside the cube.

The dynamic variance scheme for the mutation step is implemented by cluster dependent Gaussian kernels $K_t^{(C)}$ with covariance matrices

$$\Sigma_t^{(C)} := \frac{d}{|C|-1} \sum_{\substack{i=1\\ i\in C}}^{n} (x_t^{(i,C)} - \mu_t)_\rho \, (x_t^{(i,C)} - \mu_t)_\rho^T, \quad \mu_t := \frac{1}{|C|} \sum_{\substack{i=1\\ i\in C}}^{n} x_t^{(i,C)}, \tag{3.5}$$

where $d > 0$ and $((x)_\rho)_k = \max(x_k, \rho)$ for a $\rho > 0$.

3.4 Selection

Since each particle defines the pose of the model, the fitness of a particle x can be evaluated by the difference between the original image and the template image that is the projected surface of the model. For this purpose, we apply a signed Euclidean distance transformation [9] on the silhouette image I_v and on the template $T_v(x)$ for each view v. The energy function is defined by $V(x) := \frac{\alpha}{r} \sum_{v=1}^{r} V_{d,v}(x)$ with

$$V_{d,v}(x) := \frac{1}{2|T_v^+(x)|} \sum_{p \in T_v^+(x)} |T_v(x,p) - I_v(p)| + \frac{1}{2|I_v^+|} \sum_{p \in I_v^+} |T_v(x,p) - I_v(p)|, \tag{3.6}$$

where $T_v^+(x) := \{p \in T_v(x) \ : \ T_v(x,p) > 0\}$, r is the number of views, and $\alpha > 0$ is a normalization constant.

The resampling step is cluster independent, i.e., the particles migrate to the most attractive cluster where the particles have more weight and give more offspring. At the end, there are no particles left where the silhouettes do not fit the model, see Figure 1.1 c).

4 Human Bodies

While for rigid bodies the correlation between the parameters is neglected due to computational efficiency, correlation between connected joints in the human skeleton are incorporated. That is, the correlation of the joints that belong to the same skeleton branch, e.g. the left leg, are calculated in the dynamic variance scheme (3.5) while correlations with joints to other branches are set to zero.[1]

In order to focus the search on poses with higher probabilities, prior knowledge is incorporated into the energy function as soft constraint. The probability of a pose p_{pose} is estimated by a Parzen-Rosenblatt estimator with Gaussian kernels [11, 5] over a set of subsamples from different motions from the CMU motion database [6]. Since the dependency between the joints of the upper body and the joints of the lower body is low, the sample size can be reduced by splitting p_{pose} up into two independent probabilities p^u_{pose} and p^l_{pose}, respectively. Hence, the energy function is extended by

$$V(x) := \frac{\alpha}{r} \sum_{v=1}^{r} V_{d,v}(x) - \frac{\eta}{2} \ln\left(p^u_{pose}(x) p^l_{pose}(x)\right), \tag{4.1}$$

where $\eta > 0$ regulates the influence of the prior. Moreover, the mean and the variance of the joints in the training data is used to initialize the particles. To get rid of a biased error from the prior, the final pose is refined by ICP [2, 23] that is initialized by the estimate of ISA.

[1] Samples from a multivariate normal distribution $\mathcal{N}(\mu, \Sigma)$ can be drawn via a Cholesky decomposition $\Sigma = AA^T$: $x = \mu + Az$ where z is drawn from $\mathcal{N}(0, I)$.

5 Results

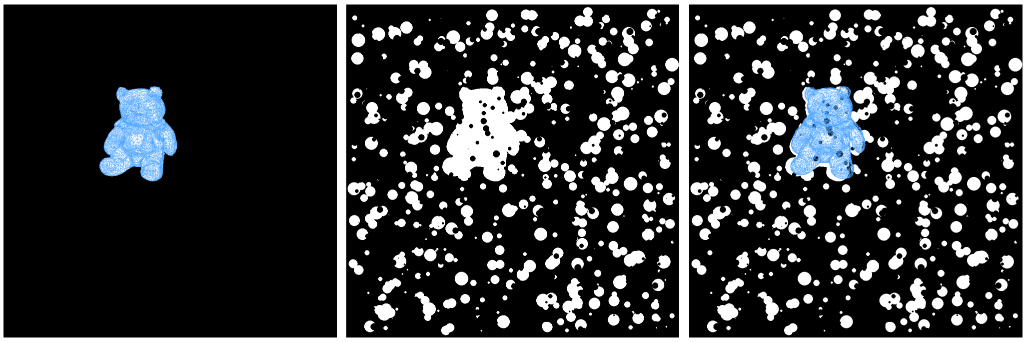

Figure 5.1: **From left to right:** *a)* Estimated pose without noise. The error is less than $1mm$ (median). *b)* Silhouettes are randomly distorted by 500 white and 500 black circles. *c)* Median estimate with error less than $4cm$.

For the error analysis, the silhouettes of the bear were generated by projecting the model for 3 different views and the Euclidean distance between the estimated position and the exact position was measured to get the error of our method without extrinsic factors like camera calibration and segmentation. Each simulation was repeated 25 times and the average errors for different numbers of particles and iterations are plotted in Figure 5.3. The estimates for 200 particles and 30 iterations are very accurate with a median error less than $1mm$, see Figure 5.1 a). The influence of distorted silhouettes is simulated by randomly drawing first a fixed number of white circles and then black circles. Holes, dilatation and erosion are typically for background subtraction and change the outcome of the Euclidean distance transform. The diagrams in Figure 5.3 show that our method performs also well for distorted silhouettes. In the case of 500 white and 500 black circles, the error of the median estimate shown in Figure 5.1 is still less than $4cm$. Results for a real scene with background subtraction are shown in Figure 5.2. The perfor-

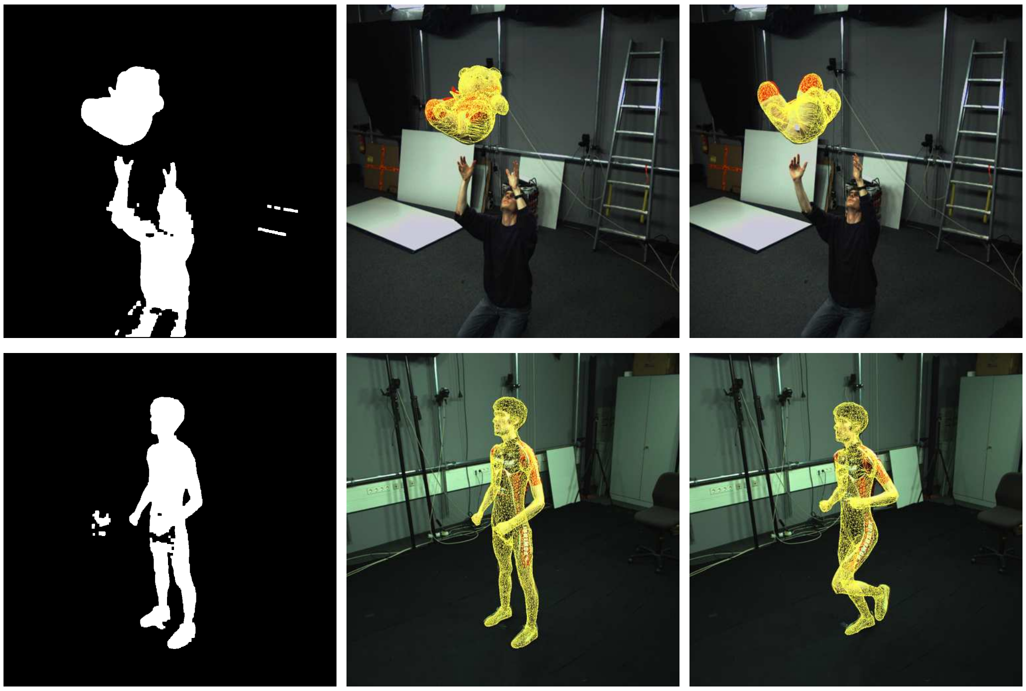

Figure 5.2: Estimates for a real scene. 3 views were segmented for the bear and 4 views for the human (Only one is shown). **Most left:** Silhouettes from background subtraction.

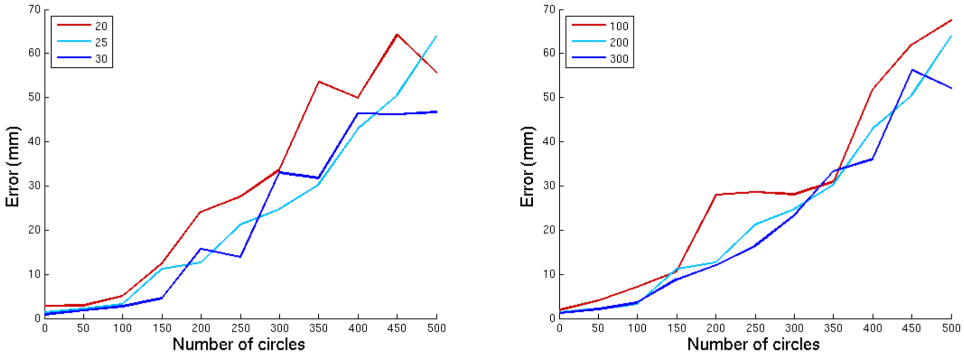

Figure 5.3: Average error of the estimates for different numbers of iterations and 200 particles *(left)* and for different numbers of particles and 25 iterations *(right)*. 200 particles and 25 iterations are sufficient for rigid bodies.

Figure 5.4: Estimates for 12 poses from a motion sequence from the CMU database (The estimated poses of the human model are projected onto the silhouette images). Each row shows one of the three views.

mance for a human body with 30 DoF was tested by generating silhouettes for 12 poses from a sequence of the CMU database that was not used for the prior. The estimates are given in Figure 5.4. The average error of the joints for 400 particles and 40 iterations was 1.05°.

6 Discussion

We proposed an accurate and robust approach, which relies on a global optimization method with clustered particles, for estimating the 3D pose of rigid and articulated objects with up to 30 DoF. It does not require any initial information about position or orientation of the object and solves the initial problem as it occurs for tracking and texture acquisition. Our experiments demonstrate that the correct pose is estimated when multiple objects appear. It could also be extended to the case when the object is not visible by rejecting estimates with an high energy. In general, our method can be easily modified for certain applications, e.g., by including prior as we did for humans. Other possibilities are multi-cue integration and exploitation of an hierarchical structure, however, these features are object specific and not suitable for a general solution.

Bibliography

[1] A. Ansar and K. Daniilidis. Linear pose estimation from points or lines. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25(5):578–589, 2003.

[2] P. Besl and N. McKay. A method for registration of 3-d shapes. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.

[3] M. Bray, P. Kohli, and P. Torr. Posecut: Simultaneous segmentation and 3d pose estimation of humans using dynamic graph-cuts. In *European Conf. on Computer Vision*, volume 3952 of *LNCS*, pages 642–655. Springer, 2006.

[4] T. Brox, B. Rosenhahn, D. Cremers, and H.-P. Seidel. High accuracy optical flow serves 3-d pose tracking: Exploiting contour and flow based constraints. In *European Conf. on Computer Vision*, volume 3952 of *LNCS*, pages 98–111. Springer, 2006.

[5] T. Brox, B. Rosenhahn, U. Kersting, and D. Cremers. Nonparametric density estimation for human pose tracking. In K. Franke, R. Mueller, B. Nickolay, and R. Schaefer, editors, *Pattern Recognition (DAGM)*, volume 4174 of *LNCS*, pages 546–555. Springer, 2006.

[6] CMU. Graphics lab motion capture database. http://mocap.cs.cmu.edu/.

[7] R. Douc, O. Cappe, and E. Moulines. Comparison of resampling schemes for particle filtering. In *Int. Symposium on Image and Signal Processing and Analysis*, pages 64–69, 2005.

[8] A. Doucet, N. de Freitas, and N. Gordon, editors. *Sequential Monte Carlo Methods in Practice*. Springer, New York, 2001.

[9] P. Felzenszwalb and D. Huttenlocher. Distance transforms of sampled functions. Technical Report TR2004-1963, Cornell Computing and Information Science, 2004.

[10] J. Gall, J. Potthoff, C. Schnoerr, B. Rosenhahn, and H.-P. Seidel. Interacting and annealing particle systems – mathematics and recipes. *Journal of Mathematical Imaging and Vision*, 2007. To appear.

[11] J. Gall, B. Rosenhahn, T. Brox, and H.-P. Seidel. Learning for multi-view 3d tracking in the context of particle filters. In *Int. Symposium on Visual Computing (ISVC)*, volume 4292 of *LNCS*, pages 59–69. Springer, 2006.

[12] J. Gall, B. Rosenhahn, and H.-P. Seidel. Robust pose estimation with 3d textured models. In *IEEE Pacific-Rim Symposium on Image and Video Technology*, volume 4319 of *LNCS*, pages 84–95. Springer, 2006.

[13] J. Gall, B. Rosenhahn, and H.-P. Seidel. *Human Motion - Understanding, Modeling, Capture and Animation*, chapter An Introduction to Interacting Simulated Annealing. Springer, 2007. To appear.

[14] S. Kirkpatrick, C. G. Jr., and M. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983.

[15] V. Lepetit, J. Pilet, and P. Fua. Point matching as a classification problem for fast and robust object pose estimation. In *IEEE Conf. on Computer Vision and Pattern Recognition*, volume 2, pages 244–250, 2004.

[16] D. Lowe. Three-dimensional object recognition from single two-dimensional images. *Artificial Intelligence*, 31(3):355–395, 1987.

[17] D. Lowe. Fitting parameterized three-dimensional models to images. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(5):441–450, 1991.

[18] P. D. Moral. *Feynman-Kac Formulae. Genealogical and Interacting Particle Systems with Applications.* Springer, New York, 2004.

[19] R. Murray, Z. Li, and S. Sastry. *A Mathematical Introduction to Robotic Manipulation.* CRC Press, Boca Raton, FL, 1994.

[20] X. Pennec and N. Ayache. Uniform distribution, distance and expectation problems for geometric features processing. *Journal of Mathematical Imaging and Vision*, 9(1):49–67, 1998.

[21] B. Rosenhahn, T. Brox, and J. Weickert. Three-dimensional shape knowledge for joint image segmentation and pose tracking. *Int. Journal of Computer Vision*, 73(3):243–262, 2007.

[22] J. Stolfi. *Oriented Projective Geometry: A Framework for Geometric Computation*. Academic Press, Boston, 1991.

[23] Z. Zhang. Iterative point matching for registration of free-form curves and surfaces. *Int. J. of Computer Vision*, 13(2):119–152, 1994.

Below you find a list of the most recent technical reports of the Max-Planck-Institut für Informatik. They are available by anonymous ftp from `ftp.mpi-sb.mpg.de` under the directory `pub/papers/reports`. Most of the reports are also accessible via WWW using the URL `http://www.mpi-sb.mpg.de`. If you have any questions concerning ftp or WWW access, please contact `reports@mpi-sb.mpg.de`. Paper copies (which are not necessarily free of charge) can be ordered either by regular mail or by e-mail at the address below.

Max-Planck-Institut für Informatik
Library
attn. Anja Becker
Stuhlsatzenhausweg 85
66123 Saarbrücken
GERMANY
e-mail: `library@mpi-sb.mpg.de`

MPI-I-2007-5-001	G. Ifrim, G. Kasneci, M. Ramanath, F.M. Suchanek, G. Weikum	NAGA: Searching and Ranking Knowledge
MPI-I-2007-2-001	A. Podelski, S. Wagner	A Method and a Tool for Automatic Veriication of Region Stability for Hybrid Systems
MPI-I-2006-5-006	G. Kasnec, F.M. Suchanek, G. Weikum	Yago - A Core of Semantic Knowledge
MPI-I-2006-5-005	R. Angelova, S. Siersdorfer	A Neighborhood-Based Approach for Clustering of Linked Document Collections
MPI-I-2006-5-004	F. Suchanek, G. Ifrim, G. Weikum	Combining Linguistic and Statistical Analysis to Extract Relations from Web Documents
MPI-I-2006-5-003	V. Scholz, M. Magnor	Garment Texture Editing in Monocular Video Sequences based on Color-Coded Printing Patterns
MPI-I-2006-5-002	H. Bast, D. Majumdar, R. Schenkel, M. Theobald, G. Weikum	IO-Top-k: Index-access Optimized Top-k Query Processing
MPI-I-2006-5-001	M. Bender, S. Michel, G. Weikum, P. Triantafilou	Overlap-Aware Global df Estimation in Distributed Information Retrieval Systems
MPI-I-2006-4-010	A. Belyaev, T. Langer, H. Seidel	Mean Value Coordinates for Arbitrary Spherical Polygons and Polyhedra in $\mathbb{R}^3$
MPI-I-2006-4-009	J. Gall, J. Potthoff, B. Rosenhahn, C. Schnoerr, H. Seidel	Interacting and Annealing Particle Filters: Mathematics and a Recipe for Applications
MPI-I-2006-4-008	I. Albrecht, M. Kipp, M. Neff, H. Seidel	Gesture Modeling and Animation by Imitation
MPI-I-2006-4-007	O. Schall, A. Belyaev, H. Seidel	Feature-preserving Non-local Denoising of Static and Time-varying Range Data
MPI-I-2006-4-006	C. Theobalt, N. Ahmed, H. Lensch, M. Magnor, H. Seidel	Enhanced Dynamic Reflectometry for Relightable Free-Viewpoint Video
MPI-I-2006-4-005	A. Belyaev, H. Seidel, S. Yoshizawa	Skeleton-driven Laplacian Mesh Deformations
MPI-I-2006-4-004	V. Havran, R. Herzog, H. Seidel	On Fast Construction of Spatial Hierarchies for Ray Tracing
MPI-I-2006-4-003	E. de Aguiar, R. Zayer, C. Theobalt, M. Magnor, H. Seidel	A Framework for Natural Animation of Digitized Models
MPI-I-2006-4-002	G. Ziegler, A. Tevs, C. Theobalt, H. Seidel	GPU Point List Generation through Histogram Pyramids
MPI-I-2006-4-001	A. Efremov, R. Mantiuk, K. Myszkowski, H. Seidel	Design and Evaluation of Backward Compatible High Dynamic Range Video Compression
MPI-I-2006-2-001	T. Wies, V. Kuncak, K. Zee, A. Podelski, M. Rinard	On Verifying Complex Properties using Symbolic Shape Analysis
MPI-I-2006-1-007	H. Bast, I. Weber, C.W. Mortensen	Output-Sensitive Autocompletion Search
MPI-I-2006-1-006	M. Kerber	Division-Free Computation of Subresultants Using Bezout Matrices
MPI-I-2006-1-005	A. Eigenwillig, L. Kettner, N. Wolpert	Snap Rounding of Bzier Curves
MPI-I-2006-1-004	S. Funke, S. Laue, R. Naujoks, L. Zvi	Power Assignment Problems in Wireless Communication

MPI-I-2005-5-002	S. Siersdorfer, G. Weikum	Automated Retraining Methods for Document Classification and their Parameter Tuning
MPI-I-2005-4-006	C. Fuchs, M. Goesele, T. Chen, H. Seidel	An Emperical Model for Heterogeneous Translucent Objects
MPI-I-2005-4-005	G. Krawczyk, M. Goesele, H. Seidel	Photometric Calibration of High Dynamic Range Cameras
MPI-I-2005-4-005	M. Goesele	?
MPI-I-2005-4-004	C. Theobalt, N. Ahmed, E. De Aguiar, G. Ziegler, H. Lensch, M.A. Magnor, H. Seidel	Joint Motion and Reflectance Capture for Creating Relightable 3D Videos
MPI-I-2005-4-003	T. Langer, A.G. Belyaev, H. Seidel	Analysis and Design of Discrete Normals and Curvatures
MPI-I-2005-4-002	O. Schall, A. Belyaev, H. Seidel	Sparse Meshing of Uncertain and Noisy Surface Scattered Data
MPI-I-2005-4-001	M. Fuchs, V. Blanz, H. Lensch, H. Seidel	Reflectance from Images: A Model-Based Approach for Human Faces
MPI-I-2005-2-004	Y. Kazakov	A Framework of Refutational Theorem Proving for Saturation-Based Decision Procedures
MPI-I-2005-2-003	H.d. Nivelle	Using Resolution as a Decision Procedure
MPI-I-2005-2-002	P. Maier, W. Charatonik, L. Georgieva	Bounded Model Checking of Pointer Programs
MPI-I-2005-2-001	J. Hoffmann, C. Gomes, B. Selman	Bottleneck Behavior in CNF Formulas
MPI-I-2005-1-008	C. Gotsman, K. Kaligosi, K. Mehlhorn, D. Michail, E. Pyrga	Cycle Bases of Graphs and Sampled Manifolds
MPI-I-2005-1-007	I. Katriel, M. Kutz	A Faster Algorithm for Computing a Longest Common Increasing Subsequence
MPI-I-2005-1-003	S. Baswana, K. Telikepalli	Improved Algorithms for All-Pairs Approximate Shortest Paths in Weighted Graphs
MPI-I-2005-1-002	I. Katriel, M. Kutz, M. Skutella	Reachability Substitutes for Planar Digraphs
MPI-I-2005-1-001	D. Michail	Rank-Maximal through Maximum Weight Matchings
MPI-I-2004-NWG3-001	M. Magnor	Axisymmetric Reconstruction and 3D Visualization of Bipolar Planetary Nebulae
MPI-I-2004-NWG1-001	B. Blanchet	Automatic Proof of Strong Secrecy for Security Protocols
MPI-I-2004-5-001	S. Siersdorfer, S. Sizov, G. Weikum	Goal-oriented Methods and Meta Methods for Document Classification and their Parameter Tuning
MPI-I-2004-4-006	K. Dmitriev, V. Havran, H. Seidel	Faster Ray Tracing with SIMD Shaft Culling
MPI-I-2004-4-005	I.P. Ivrissimtzis, W.-. Jeong, S. Lee, Y.a. Lee, H.-. Seidel	Neural Meshes: Surface Reconstruction with a Learning Algorithm
MPI-I-2004-4-004	R. Zayer, C. Rssl, H. Seidel	r-Adaptive Parameterization of Surfaces
MPI-I-2004-4-003	Y. Ohtake, A. Belyaev, H. Seidel	3D Scattered Data Interpolation and Approximation with Multilevel Compactly Supported RBFs
MPI-I-2004-4-002	Y. Ohtake, A. Belyaev, H. Seidel	Quadric-Based Mesh Reconstruction from Scattered Data
MPI-I-2004-4-001	J. Haber, C. Schmitt, M. Koster, H. Seidel	Modeling Hair using a Wisp Hair Model
MPI-I-2004-2-007	S. Wagner	Summaries for While Programs with Recursion
MPI-I-2004-2-002	P. Maier	Intuitionistic LTL and a New Characterization of Safety and Liveness
MPI-I-2004-2-001	H. de Nivelle, Y. Kazakov	Resolution Decision Procedures for the Guarded Fragment with Transitive Guards
MPI-I-2004-1-006	L.S. Chandran, N. Sivadasan	On the Hadwiger's Conjecture for Graph Products
MPI-I-2004-1-005	S. Schmitt, L. Fousse	A comparison of polynomial evaluation schemes
MPI-I-2004-1-004	N. Sivadasan, P. Sanders, M. Skutella	Online Scheduling with Bounded Migration