

Efficient Multi-hypotheses Unscented Kalman Filtering for Robust Localization

Gregor Jochmann, Sören Kerner, Stefan Tasse, and Oliver Urbann

Robotics Research Institute
Section Information Technology
TU Dortmund University
44221 Dortmund, Germany

Abstract. This paper describes an approach to Gaussian mixture filtering which combines the accuracy of the Kalman filter and the robustness of particle filters without sacrificing computational efficiency. Critical approximations of common Gaussian mixture algorithms are analyzed and similarities are pointed out to particle filtering with an extremely low number of particles. Known techniques from both fields are applied in a new combination resulting in a multi-hypotheses Kalman filter which is superior to common Kalman filters in its ability of fast relocalization in kidnapped robot scenarios and its representation of multi-modal belief distributions, and which outperforms particle filters in localization accuracy and computational efficiency.

1 Introduction

Localization is a central aspect for autonomous robots playing soccer as well as in most other application fields. The higher level decision making relies on an accurate knowledge of the robot's position, e.g. positioning a defending player to block its own goal or to support a team mate, or kicking the ball into the right direction even when goal posts are temporarily occluded.

Most localization algorithms which have been applied in RoboCup contexts are Bayesian algorithms such as particle or Kalman filters [11]. In general, particle filters are favored when belief distributions are expected to be multi-modal and sensor information is uncertain and ambiguous, while otherwise Kalman filters are expected to produce more accurate and smooth results. Gutmann and Fox express this common consensus in [4]: “Markov localization is more robust than Kalman filtering while the latter can be more accurate than the former”. While hybrid solutions, special variants and problem-specific adaptations always have the potential to outperform pure implementations of the general methods, this seems to be the general trend and particle filters are often the method of choice for handling multi-modal belief distributions [11]. However, multiple model Kalman filters have been well established in other fields of research [1] and are well suited for such tasks.

The Kalman filter as an estimator for linear Gaussian systems has been adapted with variants such as the Extended or Unscented Kalman filter. Both

maintain the convenient Gaussian representation throughout the filter update steps by linearizing the process and measurement functions. However, such Gaussian approximations perform unsatisfactory or even diverge completely when the non-linearity in the system becomes too severe. This is obvious especially for multi-modal systems.

Weighted sums of Gaussians offer an approximation for those non-linear non-Gaussian systems [1]. The resulting filters are referred to as Gaussian sum or Gaussian mixture filters. The high number of Gaussians necessary to appropriately approximate any given belief distribution leads to an increase in computational complexity so that pruning of the belief representation becomes an important issue in practical real-time implementations on mobile platforms. Due to this, multiple-model Kalman filter implementations lose some of their generality, computational efficiency and theoretical elegance.

The same is true for particle filter implementations which aim at high performance and applicability on limited platforms. Those usually operate with an extremely low number of particles to be efficient enough to operate for example on the Aibo or the Nao. This paper's main contribution is to point out and apply techniques originally introduced in particle filtering contexts to multiple-model Kalman filtering.

This has been implemented for a RoboCup Standard Platform League scenario, i.e. for humanoid robots with highly uncertain odometry in a dynamic soccer scenario where most landmarks are ambiguous field features, occlusion frequently occurs and false positives are likely to be generated from observations of the audience. Situations with various degrees of similarity to robot kidnapping happen due to frequent struggles and shoving among the robots and interventions by the referees. This work is related to [10] in terms of the application scenario and the general idea of using a multiple-model Kalman filter to address the correspondence problem for ambiguous landmarks and false positive observations, but the proposed solution differs in the overall approach as well as various implementation details.

2 Gaussian Mixture Filtering

An overview of Gaussian Mixture filtering and its applicability to state estimation in domains of multi-modal probability distributions is given in this chapter. The general Bayes filter convention under the Markov assumption [11] is used in the following discussion in order not to commit to any specific form of non-linearity approximation to the Kalman filter.

The initial probability distribution for the n -dimensional state x_0 is expressed by the prior belief $bel(x_0)$. Each discrete time step the estimation is updated using the following equations.

$$\overline{bel}(x_t) = \int p(x_t|u_t, x_{t-1})bel(x_{t-1})dx_{t-1} \quad (1)$$

$$bel(x_t) = \eta p(z_t|x_t)\overline{bel}(x_t) \quad (2)$$

Equation 1 describes the prediction step or *process update*, in which the past belief is updated using the known control input u_t and the process model, which is expressed by the conditional probability $p(x_t|u_t, x_{t-1})$. This predicted posterior belief $\overline{bel}(x_t)$ is corrected in equation 2 using the measurement z_t and the sensor model $p(z_t|x_t)$. This is commonly referred to as *measurement or sensor update*.

The standard Kalman algorithm provides an optimal estimation in case of linear Gaussian systems, which however is rarely given for real problems. Non-linearity in the process and sensor model is handled by linearization in order to apply the familiar Kalman filter equations, which then do not yield the optimal solution but only an approximation. This can be done either using Jacobians in the Extended Kalman filter or the unscented transform in the Unscented Kalman filter.

If the model shows only certain non-linear characteristics around the current estimation, this linearization is enough to allow an appropriate estimation. Representing the belief state with Gaussian mixtures as in equation 3 can also improve the estimation in those cases, but the real benefit becomes obvious in cases where the models show multi-modal characteristics.

$$bel(x_t) = \sum_{i=1}^N \alpha_i \frac{1}{(2\pi)^{n/2} |P_i|^{1/2}} e^{(-1/2(x_t - \mu_i)^T P_i^{-1} (x_t - \mu_i))} \quad (3)$$

Here μ_i and P_i are the means and covariances of the individual Gaussians and α_i are the weights which sum up to 1 over all N models. Note that N might change during operation of the filter, and the explicit dependence on time in the indexes of those parameters is dropped for simplicity.

Gaussian mixtures introduce several changes into the filtering process. Note that in most applications only a subset of these aspects is actually implemented [10,9,3,12,5], i.e. the one or two aspects most crucial to the estimation process, and a strict limit to the number of separate Gaussians is enforced to maintain acceptable processing time.

The initial belief is obviously easier represented by Gaussian mixtures than by a single Gaussian. This can be done either by a regular distribution parameterized to fit an a-priori belief as proposed in [1], or by generating the initial belief from the first sensor information as done in [3] and [5]. Re-localization from kidnapped robot scenarios is closely related, since allowing for the possibility of robot kidnapping means assigning a small probability to the case that the robot is repositioned without any knowledge, which is similar to the initial global position finding. Gaussian mixtures might also be applied for a better representation of highly uncertain odometry as for walking robots, but due to the drawbacks of the increase of terms in the mixture this is rarely implemented.

2.1 Sensor Update and Correspondence Problem

The most focused on filtering aspect for applying Gaussian mixtures is the sensor update. While Gaussian mixtures can be beneficial in modeling a spread-out

belief for uncertain measurements, the major reason for using them is the update with ambiguous landmarks. In single-Gaussian Kalman filters this case is handled by choosing the correspondence with maximum likelihood and proceeding by linearizing a sensor model for a unique landmark update. This results in a Kalman filter operating “more as a maximum likelihood estimator than as a minimum variance estimator and the mean follows (hopefully) one of the peaks of the density function” [1].

Gaussian mixtures allow to construct a sensor model with one Gaussian term for each possible correspondence. In [9] for example, Gaussian mixture sensor models are used to avoid an exclusive correspondence choice for the likelihood calculation in Monte Carlo localization. In the Kalman approach the sensor update using this model results in applying all possible correspondences to all hypotheses maintained by the current belief prediction $\overline{bel}(x_t)$ as done in [10]. The terms in $bel(x_t)$ therefore increase by the factor M_s which is the number of terms in the sensor model.

The weights α_i are recursively updated according to [1] by multiplication with the probability of the measured m -dimensional innovation $\eta = (z_t - \hat{z}_i)$ by

$$\alpha_i = \nu \alpha'_i \left(\frac{1}{\sqrt{(2\pi)^m |P_\eta|}} e^{-\frac{1}{2} \eta^{-1} P_\eta^{-1} \eta} \right) \quad (4)$$

where \hat{z}_i is the expected observation for the fixed correspondence according to the i 'th model, P_η is the sum of the measurement and the prediction covariance, and ν is a normalization factor. To improve the lack of robustness to outliers, Quinlan and Middleton add a term ϵ_0 expressing a static probability of the observation being an outlier, i.e. a false positive in the measurement process such as echoes in sonar data or incorrectly classified objects in a vision system [10]. A further discussion of the implications of this will be given in section 3.

2.2 Pruning the Belief Representation

The multiplicative increase per time step in the number of terms in the Gaussian mixture introduced by the sensor update and potentially also the process update is countered by pruning the resulting belief representation, for which several different methods have been proposed. In general, a suitable Gaussian mixture representation with a specified maximum number of terms needs to be found which approximates a given probability distribution. This is similar to finding the initial parametrization, only in this case the target distribution is already given as a Gaussian mixture, only with a higher number of terms.

In most practical applications, sophisticated re-parametrization strategies as described in [12] and [6] achieve excellent results, but are often too time-consuming. Instead of applying iterative optimization or regression procedures, simple heuristics are used to reduce the number of terms in the Gaussian mixture. This is done by combining multiple terms into one and by neglecting terms with small enough weighting factors whenever possible [1,10].

3 Problems Resulting from Efficiency-Related Trade-offs and Model Limitations

Many of the problems classical Kalman filters face in realistic application scenarios can be addressed by the extension to use sums of Gaussians. Ensuring efficient computation however introduces compromises which together with some of the common assumptions described in the previous section potentially nullify certain aspects which make the Gaussian mixture representation desirable in the first place. Two potential problems related to pruning and to the handling of false positives shall be analyzed in the following as a motivation for the alternative approach described in section 4 which differs in contrast to common implementations like [1,10].

3.1 Influence of Pruning on Quality of Estimation

The exponential growth of Gaussian terms necessitates aggressive pruning as described in section 2. This pruning however does not only lead to a slight loss of accuracy, but may potentially undo some of the central benefits of using Gaussian mixtures.

One significant problem shall be illustrated using the following simplified example. Consider a robot which is well localized with a current belief state consisting of a single Gaussian, and whose position is altered at a time t_0 , e.g. by a collision with another agent changing the robot's orientation significantly without it perceiving it. In the following time steps the robot repeatedly makes ambiguous observations which correspond to one of two possible landmarks in its map, as visualized in figure 1(a). In this scenario, let each blue link correspond to the correspondence choice c_1 and each red link to choice c_2 . c_1 is the real observed landmark, but c_2 is the choice which fits the (wrong) initial belief state $bel(x_{t_0})$ better. Figure 1(b) visualizes the model splitting resulting from each multi-modal sensor update without any pruning.

In this situation, updates based on the correspondence choice c_2 result in small innovations, while those based on choice c_1 initially produce large innovations, which only decrease when frequent updates using c_1 shift the corresponding mean closer to the true position. In this scenario, continuous multi-modal updates obviously do not resolve the localization ambiguity, but lead to a belief state $bel(x_{t_0+n})$ which includes a term $\alpha_i N_i$ which is the result of a series of "right" choices and which, after the corresponding mean is close to the true position, also has a high weight. In a situation like this, few additional complementary observations of different landmarks might resolve the ambiguity and leave the true position as the most probable estimate in the belief. Thus the Gaussian mixture filtering allows for more than just mere maximum likelihood estimation if extensive pruning does not interfere with this characteristic.

Assuming equal a-priori probabilities for both data associations and assuming simple uni-modal process updates, the weight factors α_i of the different terms in the belief of the following time steps $t_0 + \delta$ is changed exclusively by the sensor update in equation 4. Since initially the correspondence choices c_1 do not fit the

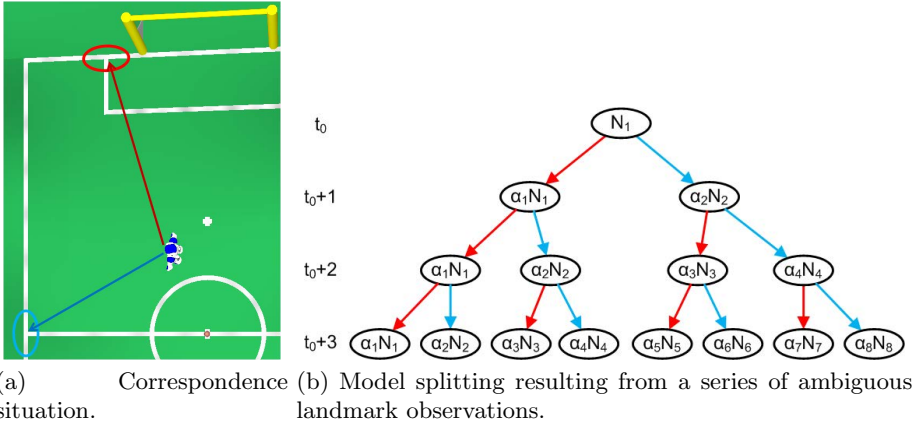


Fig. 1. Example situation with ambiguous observation correspondence

belief as the alternative choices do, the weight α_8 in figure 1(b) is much smaller than α_1 at time step $t_0 + 3$. This makes paths in a tree of correspondence choices, which do not instantly lead to maximum likelihood estimates, ideal candidates for pruning techniques as described in [10]. So this kind of aggressive pruning does not only decrease the estimation quality, but actually removes one of the most significant advantages of Gaussian mixture filtering, i.e. the possibility to maintain different hypotheses of which some may temporarily be unlikely, but still allow the observation of the influence of new measurements on regions of the state space away from the maximum likelihood estimate. Note that this is also essential for re-localization in kidnapped-robot scenarios.

3.2 Integrations of Explicit False Positive Handling

False positives are incorrect measurements that violate the assumption of Gaussian distributed errors insofar as they are not only inaccurate measurements from a known feature, but originate from some other unmodeled source in the environment unrelated to any known feature and its position. In scenarios such as the RoboCup Standard Platform League those false positive observations are quite common since no barrier exists between the field and its surroundings. This frequently causes false perceptions from the robot's image processing when the clothings of people standing close to or directly on the edge of the field show the same outline and color of expected features such as goal posts.

The common compensation is to enlarge the tail-end of an otherwise Gaussian distribution as done in [10] where a probability is assumed of ϵ_0 that an observation is an outlier, and equation 4 is adapted as follows:

$$\alpha_i = \nu \alpha'_i \left((1 - \epsilon_0) \frac{1}{\sqrt{(2\pi)^m |P_\eta|}} e^{-\frac{1}{2} \eta^{-1} P_\eta^{-1} \eta} + \epsilon_0 \right). \quad (5)$$

This change prevents the weighting of Gaussians to drop too much by an update with a single outlier. While this approach seems intuitive and is often applied to the weighting functions in particle filters, it does not affect the actual Kalman update with the outlier observation itself. Instead the implicit assumption is that outliers appear randomly and therefore do not systematically influence the system estimation. As can be seen from the SPL scenario example, this is not necessarily true and might lead to seriously biased estimation errors.

The means for an alternative method for handling such false positives is already provided by the sensor update step described in section 2.1. Instead of incorporating the possibility for false positive observations into each sensor update for each correspondence, a more natural alternative is to handle false positives as just another correspondence alternative. The underlying assumption is that each landmark observation, not just the inherently ambiguous ones, either corresponds to one of the known locations of such landmarks on the map or to another source not included in the map. In SLAM contexts these additional observation origins might be mapped and used for further localization purposes, resulting in a multiple-hypothesis SLAM approach similar to [2]. In localization tasks however such observations corresponding to unmapped landmarks can simply be discarded, i.e. no update is performed at all to the term generated by this correspondence choice.

This approach provides the possibility of a position tracking unbiased by false positives, thus more robust especially in situations where false positive observations do not occur randomly.

4 An Alternative Approach to Multiple-hypotheses Kalman Filtering

As argued in section 3.1, operating Gaussian mixture filters with a strictly low limit on the number of terms and the consequential aggressive pruning deprives such filters of much of their multiple-hypothesis tracking potential. This situation shows parallels to certain particle filter implementations. Both algorithms have originally been designed under the assumption of enough hypotheses to appropriately cover the state space. For particle filters with extremely low numbers of particles several techniques have been proposed to compensate for the low state space coverage [7,8]. The established policy in this case consists of two measures. The first one is to limit the influence of single inconclusive measurements (for details see for example [8]). The second is to accept the impossibility to track all important hypotheses in the exponentially growing number of paths, while at the same time providing the means to recover from the neglect to model those which rise in importance again in the future. Particle filters add new particles which are uncorrelated to the current belief during the resampling step, either randomly distributed over the state space, or more efficiently by drawing few particles directly from the sensor model, called *sensor resetting* [7].

In the context of Gaussian mixture Kalman filters this mainly means a modification of the sensor update step. First of all, only the maximum likelihood




Fig. 2. Tracking only maximum-likelihood correspondence choices for ambiguous landmarks, relying on sensor resetting to generate terms close to those corresponding to the neglected paths leading to conclusive estimates

update is applied, explicitly taking into account also false positive measurement possibilities, in which case the current mean stays unchanged. Consequently it is obvious that the weight update needs to be adjusted. Following the temporal smoothing idea for particle filters [8] the weight update for hypotheses can be adjusted so as not to be influenced that much by different degrees of outlier measurements. Instead it is possible to weight different models based exclusively on how many observations can be conclusively explained by them. But more importantly an additional new sensor update in parallel to the old one will perform the same function as the sensor resetting part of the particle filter’s resampling step, i.e. introducing new Gaussian terms based not on the previous state estimate, but on the recent sensor measurements’ sensor model only. Note that this is not equivalent to model splitting since the new terms are not correlated to the old ones.

The resulting filter frequently injects new low-weighted models into regions with high probability based on the last observations, which might be expected to rise in weight in case these hypotheses also fit future observations. This however relieves the necessity to track multiple paths of correspondence choices for updates of existing models as described in section 2.1 and illustrated in figure 1(b). In [10] many multi-modal sensor updates basically result in uni-modal ones due to the aggressive pruning. While this already corresponds to implicit maximum-likelihood updates, it is now possible to explicitly do exactly this. As a consequence both the process update and the old sensor update can be applied with the uni-modal maximum-likelihood choice, explicitly neglecting alternative paths in the decision tree of correspondence choices, but relying on the sensor resetting functionality to pick up those paths which lead to conclusive estimates as illustrated in figure 2 relating to the example in section 3.1.

This application of the sensor resetting concept also solves the problem that common Kalman implementations have concerning the kidnapped robot problem. Both a sudden change of robot orientation with consequential wrong data association as described in the example in 3.1 and a real teleportation event will be handled accordingly.

5 Evaluation

The multiple-hypotheses Kalman filter approach described in this paper has been implemented for the humanoid robot Nao which is produced by Aldebaran Robotics and used in the RoboCup *Standard Platform League* (SPL). This robot contains a x86 AMD Geode 500 MHz CPU and 256 MB SDRAM. The SPL environment consists of a soccer field of specified dimensions and colored goals, which can be used for localization together with the field markings, i.e. lines, corners, and center circle. Most measurements are ambiguous: Observing a single goal post leaves at least two choices, while a field line crossing can be associated with 6 true positions on the field in case of a T-crossing or with 8 positions in case of an L-crossing. Even more correspondences are possible when allowing incompletely/uncertainly classified crossings, e.g. in case of occlusion or for the observation of two perpendicular lines whose crossing is outside of the image. One important feature of the SPL environment is that no barrier exists between the soccer field and its surroundings, which frequently includes colorfully clothed audience. So a localization algorithm needs to be robust not only to noisy measurements due to staggering robots but also to frequent false positives. This is the setup in which the proposed multiple-hypotheses Kalman filter is evaluated.

The software running on the robot includes the manufacturer's middleware NaoQi as well as the robot control software consisting of a motion thread running at 100 Hz and a cognition thread running at 15 Hz, both with all regular SPL soccer modules activated. To evaluate the localization quality a camera system is mounted above the field, detects markers attached to the robot, and provides ground truth data to an additional module in the robot's software.

The presented approach is evaluated in experiments on real robots in a typical SPL scenario and compared against a particle filter solution utilizing sensor resetting, temporal smoothing, and lazy resampling, and which has been in use in RoboCup competitions up to the development of the presented multiple-hypotheses Kalman filter. For all experiments both localization algorithms run in parallel on the robot, thus working with the exact same input from image processing and ensuring the comparability of the results.

A first experiment evaluates the re-localization ability after several teleportation events. In figure 3 each red mark on the time axis indicates that the robot has been picked up to be placed at a random new position on the field. The illustrated performance represents typical behavior for both filters, which expectedly show similar behavior since both get their relocalization ability from the same principles. Between seconds 90 and 110 the multiple-hypotheses Kalman filter generates a hypothesis close to the correct position, but does not rate it high enough in those 20 seconds to output it as the likeliest location. Similarly the particle filter jumps between different clusters after second 170. This shows that the proposed filter is able to perform at the same level as a particle filter implementation tuned especially to handle such situations. Note that this is classically regarded as a particle filter's specialty and similar performance has not been reported for related Kalman variants.




Fig. 3. Ability for relocalization from kidnapped-robot situations




Fig. 4. Comparison of localization performance in an SPL scenario

Additional experiments have been set up to quantitatively evaluate the localization quality. In all experiments the robot is placed on the field without prior knowledge of its position and the movement is started after a fixed time. Figure 4 shows a ground truth paths on the field and the localization result of the multiple-hypotheses unscented Kalman filter and the particle filter. The Kalman filter's characteristic of smooth and accurate position tracking is clearly visible and the filter's output in form of its strongest hypothesis is superior to the particle filter's most probable cluster.

Figure 5 shows one run where many persons walked around the field thus occluding field features and provoking false positives at the same time, e.g. blue jeans sometimes are falsely recognized as blue goal posts. While both localization approaches show diminished results, the multiply-hypotheses UKF clearly produces estimations closer to the real robot path.

A comparison of the different algorithms' runtime is given in figure 6. It has to be noted that the presented measurement is not unbiased, since larger runtimes tend to be influenced more by random threading issues and thus the larger module update time of the particle filter might include more motion thread update cycles. Nevertheless, the multiple-hypotheses Kalman filter is clearly much more efficient. While this does not allow to compare this implementation of a Gaussian mixture Kalman filter to the one proposed by Quinlan and Middleton [10], it can be argued to be more efficient since a similar comparison to a state of




Fig. 5. Comparison of localization performance in a worst case situation with much occlusion and many false positive observations




Fig. 6. Runtime comparison between the multiple-hypotheses Kalman filter and the particle filter

the art particle filter in [10] showed only slightly better runtime and averaged around a third of the image processing time, which would still be in the range of multiple milliseconds. The average runtime of the approach presented here is 0.4 ms, which obviously eliminates the localization problem as a computational bottleneck. Most of the time is spend on image processing, the remaining time difference goes to infrastructure and other tasks such as ball tracking or behavior decisions. The periodic tendencies in the measurements in figure 6 are a result of the robot's head motion which also searches for the ball in front of the robot's feet where only little localization information can be extracted.

6 Conclusion

This paper presents an approach for Gaussian mixture filtering which utilizes techniques from particle filtering to incorporate valuable aspects of both filter strategies. The resulting multiple-hypotheses unscented Kalman filter is superior to common Kalman filters in its ability of fast relocalization in kidnapped robot scenarios and its representation of multi-modal belief distributions, and it outperforms particle filters in localization accuracy and computational efficiency. A direct comparison to the approach proposed in [10] or similar classical Gaussian mixture filters has yet to be done.

Future work will focus on a collaborative localization and tracking strategy for a team of robots. First results have been presented in [2], but those were particle filter based and have not been efficient enough to run at high frame rates on the Nao. Building a similar system with a multiple-hypotheses Kalman filter as a basis promises real-time performance, but involves additional difficulties concerning the stochastic soundness of the overall filtering scheme.

References

1. Alspach, D., Sorenson, H.: Nonlinear Bayesian Estimation using Gaussian Sum Approximations. *IEEE Transactions on Automatic Control* 17(4), 439–448 (1972)
2. Czarnetzki, S., Rohde, C.: Handling Heterogeneous Information Sources for Multi-Robot Sensor Fusion. In: *Proceedings of the 2010 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems, MFI 2010, Salt Lake City, Utah*, pp. 133–138 (September 2010)
3. Duckett, T., Nehmzow, U.: Mobile Robot Self-Localisation using Occupancy Histograms and a Mixture of Gaussian Location Hypotheses. *Robotics and Autonomous Systems* 34(2-3), 117–129 (2001)
4. Gutmann, J.S., Fox, D.: An Experimental Comparison of Localization Methods continued. In: *2002 IEEE/RSJ International Conference on Intelligent Robots and Systems*, vol. 1, pp. 454–459 (2002)
5. Kaune, R.: Gaussian Mixture (GM) Passive Localization using Time Difference of Arrival (TDOA). In: Fischer, S., Maehle, E., Reischuk, R. (eds.) *GI Jahrestagung, LNI*, vol. 154, pp. 2375–2381. GI (2009)
6. Krauthausen, P., Hanebeck, U.D.: Regularized Non-Parametric Multivariate Density and Conditional Density Estimation. In: *Proceedings of the 2010 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems, MFI 2010, Salt Lake City, Utah*, pp. 180–186 (September 2010)
7. Lenser, S., Veloso, M.: Sensor Resetting Localization for Poorly Modelled Mobile Robots. In: *Proceedings of IEEE International Conference on Robotics and Automation, ICRA 2000*, vol. 2, pp. 1225–1232 (2000)
8. Nisticò, W., Hebbel, M.: Temporal Smoothing Particle Filter for Vision Based Autonomous Mobile Robot Localization. In: *Proceedings of the 5th International Conference on Informatics in Control, Automation and Robotics, ICINCO*, vol. RA-1, pp. 93–100. INSTICC Press (2008)
9. Pfaff, P., Plogemann, C., Burgard, W.: Gaussian Mixture Models for Probabilistic Localization. In: *IEEE International Conference on Robotics and Automation, ICRA 2008*, pp. 467–472 (May 2008)
10. Quinlan, M.J., Middleton, R.H.: Multiple Model Kalman Filters: A Localization Technique for RoboCup Soccer. In: Baltes, J., Lagoudakis, M.G., Naruse, T., Ghidary, S.S. (eds.) *RoboCup 2009. LNCS (LNAI)*, vol. 5949, pp. 276–287. Springer, Heidelberg (2010)
11. Thrun, S., Burgard, W., Fox, D.: *Probabilistic Robotics (Intelligent Robotics and Autonomous Agents)*. The MIT Press (2005)
12. Upercroft, B., Kumar, S., Ridley, M., Ong, L.L., Durrant-whyte, H.: Fast Re-parameterisation of Gaussian Mixture Models for Robotics Applications. In: Barnes, N., Austin, D. (eds.) *Proceedings of the Australasian Conference on Robotics and Automation, ACRA. Australian Robotics and Automation Association* (2004)