A Robust Approach to Simultaneous Pose Estimation and Map-Updating in Dynamic Environments

Gregor Pavlin
Department of General Mechanics
Graz University of Technology, A-8010 Graz
email: pavlin@mech.tu-graz.ac.at

Reinhard Braunstingl
Department of General Mechanics
Graz University of Technology, A-8010 Graz
email: braun@mech.tu-graz.ac.at

Abstract

We present a robust and efficient approach to simultaneous pose tracking and updating of the world models in changing environments. The estimation of the robot's pose and the corresponding uncertainties is based on very simple models which are used in the context of extra information extracted from the observation of the localization process. A special attention is payed to the consistency of the estimator. The presented pose estimator is very robust and reactive at the same time, since the lack of modelling details is compensated through on-line learning of the pose errors. In addition, geometric consistency of the portions of the world model is estimated with the help of the Bayes' rule. The estimated consistency serves for selective updating of the world model and the robot's pose. The presented approach allows reliable navigation also in environments which do not contain any static objects. The only requirement is that a sufficient amount of objects remains static over a certain period of time.

1 Introduction

For reliable navigation and execution of complex tasks over longer periods of time, a mobile robot must be able to estimate its pose[1] with respect to its environment. The simplest way of pose estimation is integration of odometric data which, however, is associated with unbounded errors, resulting from uneven floors, limited resolution of encoders, wheel slippage etc. Clearly, over longer runs pure odometric pose estimation is not reliable. Therefore, a mobile robot must be able to localize, i.e. estimate its pose also with respect to an internal world model (map) by using the information obtained with its exteroceptive sensing system. However, map-based pose estimation in typical indoor environments is quite challenging. Namely, such environments are populated by humans and certain objects can be moved easily (e.g. doors, desks, boxes, etc.) which inevitably results in partially inconsistent maps; geometric inconsistencies in maps as well as other unpredictable aspects of the interaction between the robot and its environment (e.g. sensor noise and failures) can result in great pose errors.

1. A pose in a 2D reference frame is completely defined by a vector $\mathbf{p}=(x, y, \theta)^T$.

Often approaches to pose estimation are based on the Kalman-Filter. These approaches work well as long as the used information can be described by simple statistics well enough. The lack of relevant information is compensated by using models of various processes. However, often such model-based approaches require assumptions about the physical world which are often violated and, therefore, the process is not optimal or it can even diverge. An interesting solution to this problem was proposed in [2], where observation of the pose corrections is used for updating of the covariance matrices. However, in contrast to our approach, this approach seems to be vulnerable to significant geometric inconsistencies of the world models, since inconsistent information can influence the estimated covariance matrices. Markow localization, on the other hand, is very robust since it supports multi modal probability representation. Especially, it is suitable also for the global localization. However, appropriate discretization represents a substantial problem. Some approaches to Markow localization are based on very coarse tessellation, which results in low localization accuracy such as [6]. On the other hand, the approach presented in [1] is based on dense grids which, however, requires substantial computing power.

In contrast to common approaches to pose tracking, the presented approach is based on very simple models and the robustness and reactivity are achieved through the use of appropriate information. In addition, the proposed approach allows robust updating of world models, which is indispensable for the operation in dynamic environments.

The key to robust and efficient pose tracking and map updating, however, are different uncertainty models used for the matching and filtering of inconsistent information, respectively. Matching process is based on simple bounded pose uncertainty models while filtering of information is based on adaptive estimation of localization errors. Namely, the presented *Adaptive Pose Estimator* (APE) supports on-line learning of the localization uncertainties by observing and averaging pose updates over a certain time-window (fig. 1).

Moreover, simple heuristic rules are used for detection of inconsistencies of the estimation process itself, which can result from significant odometric errors (thresholds etc.) and changing environments.

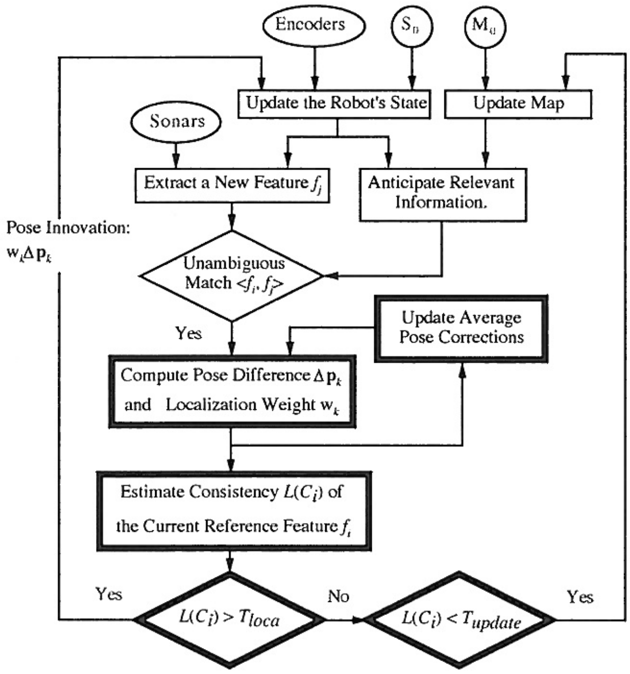

Fig. 1: Simultaneous localization and map-updating process. The localization is based on the *Adaptive Pose Estimator.* The Decision about correcting either the robot's or the map's state is based on the results of the *Consistency Filter.* $L(C_i)$ denotes the likelihood that the reference feature f_i is consistent. S_0 and M_0 denote the initial robot's state and the initial map, respectively.

In addition, the presented APE allows efficient implementation of *Consistency Filters* [4], which can be used for reliable estimation of the map's consistency. Namely, due to a very selective determination, the *localization weights* can be used as evidence for cumulative estimation of consistency of the geometric information stored in maps. The estimated consistency allows reliable decision making between updating the robot's state and updating of portions of the world model (fig. 1).

2 Consistency of the Localization Process

In general, localization requires focusing on relevant information and filtering of unreliable information. Therefore, estimation of the greatest possible position/orientation errors is indispensable for the determination of the space in which the relevant information can be located. There are different sources of possible pose errors and we need appropriate models of the propagation of such errors through the navigation process. In general, possible pose errors can be taken into account through the pose uncertainties, i.e. sets of possible poses of the robot and other objects (see fig. 2) with respect to the world model. In the presented work, pose uncertainties of a modelled object in a 2D world model are represented by a rectangle π_V and an interval π_o of possible orientations (see fig. 2). Such uncertainties are used primarily for determination of different kinds of validation gates and they are based on simple error models.

In general, a localization process is consistent if the true pose of the robot is within the space defined by the estimated pose uncertainty intervals. If this is not the case the localization process usually fails. Namely, if the matching and weighting validation gates are smaller than the actual discrepancy between the estimated and the true pose, either the matching process does not yield matches or computed innovations are ignored. Of course, the estimator's consistency can be improved by using modelled uncertainties which are much larger than the actual errors. However, by using such "loose" models inconsistent information can influence the localization process to a great extent; inconsistent information originates from mobile objects or objects which cannot be perceived reliably. Therefore, the presented *Adaptive Pose Estimator* (APE) simultaneously maintains different models of uncertainties for the matching and filtering processes, respectively.

3 Matching

The matching is based on validation gates which represent the space in which relevant information can be found. A currently extracted feature f_j can be matched with the mapped feature f_i if both features are of the same type and simple conditions are satisfied. Namely, the vector ${}^{ij}\mathbf{x}$, connecting the centers of f_i and f_j, as well as the orientation discrepancy ${}^{ij}\theta$ between f_i and f_j must be within the validation gate, defined relative to the center of the feature f_j

$$\left|{}^{j}\mathbf{n}\,{}^{ij}\mathbf{x}\right| < {}^{ij}_{M}U^{n};\quad \left|{}^{j}\mathbf{t}\,{}^{ij}\mathbf{x}\right| < {}^{ij}_{M}U^{t};\quad \left|{}^{ij}\theta\right| < {}^{ij}_{M}U^{\theta}. \qquad (1)$$

In (1) the vector ${}^{j}\mathbf{n}$ is parallel and the unit vector ${}^{j}\mathbf{t}_i$ is normal to the currently obtained feature f_j, respectively. ${}^{ij}_{M}U^{n}$, ${}^{ij}_{M}U^{t}$ and ${}^{ij}_{M}U^{\theta}$, on the other hand, define the relative pose uncertainties between the features f_i and f_j. Note, that the matching validation gates are based on loose uncertainties; i.e. assumed worst case errors are greater than the actual errors. Possible error models are presented in [5].

4 Pose Estimation

As soon as the robot obtains an unambiguous match the corresponding pose update relative to the reference frame is computed through alignment of the mapped and observed features f_i and f_j, respectively. For each localization step the robot uses only a single pair of linear features. The orientation discrepancy $\Delta\theta_k$ is identical to the orientation discrepancy between the matched features while the translation discrepancy $\Delta\mathbf{x}_k$ is identical to the normal distance between the matched features, which would be obtained if the robot's orientation would be changed by $\Delta\theta_k$ (fig. 2).

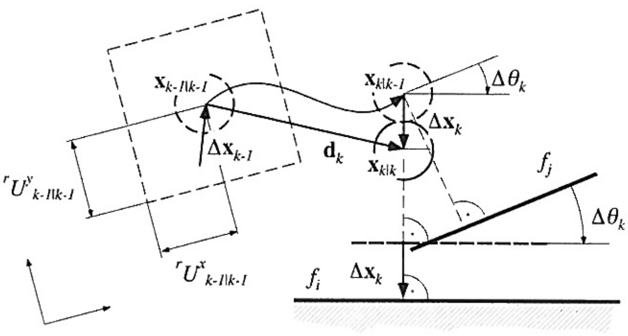

Fig. 2: Localization step based on linear features. f_i denotes the reference (mapped) feature while f_j represents the currently extracted feature. Dashed rectangle on the left represents position uncertainties which correspond to the position $\mathbf{x}_{k-1|k-1}$ resulting from the k-1[th] localization step.

The pose $\mathbf{p}_{k|k} = (x_{k|k}, y_{k|k}, \theta_{k|k})^T$ is updated according to

$$\mathbf{p}_{k|k} = \mathbf{p}_{k|k-1} + \Delta\mathbf{p}_k w_k, \quad (2)$$

where w_k denotes the localization weight corresponding to the k-th localization step (see section 5), $\mathbf{p}_{k|k-1}$ and $\mathbf{p}_{k|k}$ denote the pose vectors prior to and after the k-th pose correction while $\Delta\mathbf{p}_k = (\Delta x_k, \Delta y_k, \Delta\theta_k)^T$ denotes the estimated pose discrepancy between the matched features f_i and f_j. In addition, with each correction the robot's pose uncertainties ${}^rU^x_{k|k-1}$, ${}^rU^y_{k|k-1}$ and ${}^rU^\theta_{k|k-1}$ are updated according to

$${}^rU^\xi_{k|k} = (1-w_k)\,{}^rU^\xi_{k|k-1} + w_k({}^iU^\xi + \min({}^{ir}U^\xi_{k|k}, {}^{ir}U^\xi_{k|k-1})), \quad (3)$$

where $\xi \in \{x, y, \theta\}$, ${}^{ir}U^\xi_{k|k-1}$ and ${}^{ir}U^\xi_{k|k}$ define the relative pose uncertainties between the robot and the reference f_i prior to and after the k-th localization step, respectively; with such pose uncertainties we take into account the worst case sensing errors. Moreover, ${}^iU^\xi$ define pose uncertainties associated with the reference feature[1] f_i.

Note, that estimation of bounded uncertainties is crucial for determination of the *matching validation gates*.

1. As a new feature f_i is extracted and accumulated into the map, it is associated with the current robot's pose uncertainties: ${}^iU^\xi = {}^rU_k^\xi$.

5 Adaptive Filtering of Information

The localization process, however, is vulnerable to inconsistent information. Therefore, with the APE we introduce a simple and very efficient adaptive approach to filtering of inconsistent information. Namely, the robot tracks the average translation/orientation corrections $\bar{c}_k^t$ and $\bar{c}_k^\theta$ which are based on a sequence of position/orientation discrepancies $\Delta\mathbf{x}_k$ and $\Delta\theta_k$

$$\begin{aligned} \bar{c}^t_{k+1} &= |\Delta\mathbf{x}_k|\lambda w_k + (1-\lambda w_k)\bar{c}^t_k \\ \bar{c}^\theta_{k+1} &= |\Delta\theta_k|\lambda w_k + (1-\lambda w_k)\bar{c}^\theta_k \end{aligned} \quad (4)$$

where λ is a learning factor and w_k denotes the localization weight (see below). $\bar{c}_k^t$ and $\bar{c}_k^\theta$ can be viewed as a kind of learned localization uncertainty intervals which normally converge toward certain intervals. Therefore, $\bar{c}_k^t$ and $\bar{c}_k^\theta$ are used for the determination of the influence of a certain information on the localization process. Namely, the current pose discrepancies between the matched features and the corresponding average corrections are compared with the help of the mappings $w_k^t = \mu(|\Delta\mathbf{x}_k|, \bar{c}_k^t)$ and $w_k^\theta = \mu(|\Delta\theta_k|, \bar{c}_k^\theta)$; the function μ: $\Re^2 \rightarrow [0,1]$ yields rather low weights if the current discrepancies $|\Delta\mathbf{x}_k|$ and $|\Delta\theta_k|$ are significantly greater than the average corrections $\bar{c}_k^t$ and $\bar{c}_k^\theta$. In the presented work we use a very simple mapping

$$\mu(c,\bar{c}) = \begin{cases} 1, if\ k_c\bar{c} > c \vee t_{res} > c \\ 1 - \dfrac{c - k_c\bar{c}}{(k_p - k_c)\bar{c}}, if\ k_p\bar{c} > c > k_c\bar{c} \\ 0, if\ k_p\bar{c} < c \end{cases}, \quad (5)$$

where $\bar{c}$ denotes the average position/orientation correction and k_c and k_p are appropriate scaling factors. Finally, the localization weight w_k is defined as

$$w_k = \min(w_k^t, w_k^\theta). \quad (6)$$

By using the *localization weights* defined by (6) a significant portion of the inconsistent information is eliminated from the localization process.

6 Maintenance of the Estimator's Consistency

However, often the weighing validation gates $\bar{c}_k^t$ and $\bar{c}_k^\theta$, determined with (4), converge to relatively small intervals. Therefore, the estimator could easily become inconsistent if the odometric or localization errors would suddenly increase significantly. For example, significant odometric errors could be introduced through passing a threshold, turning on a carpet etc. In such cases the resulting discrepancies $|\Delta \mathbf{x}_k|$ and $|\Delta\theta_k|$ are usually much greater than the localization bounds $\bar{c}_k^{\,t}$ and $\bar{c}_k^{\,\theta}$ and, therefore, the corresponding localization weight w_k is small or zero. Consequently, the required corrections would be simply ignored and the robot could get lost.

Therefore, the APE attempts to detect such inconsistencies of the pose estimator and temporarily increase the localization bounds $\bar{c}_k^{\,t}$ and $\bar{c}_k^{\,\theta}$. The detection of inconsistencies is based on the observation of subsequent pose discrepancies $\Delta\mathbf{p}_k$ and $\Delta\mathbf{p}_{k+1}$, which were based on different reference features and both were out of the localization bounds; i.e. either $|\Delta\mathbf{x}_k| > \bar{c}_k^{\,t} \wedge |\Delta\mathbf{x}_{k+1}| > \bar{c}_{k+1}^{\,t}$ or $|\Delta\theta_k| > \bar{c}_k^{\,\theta} \wedge |\Delta\theta_{k+1}| > \bar{c}_{k+1}^{\,\theta}$.

If a significant odometric or localization error takes place and the map is consistent, the subsequent pose discrepancies satisfy relations

$$\frac{\Delta\mathbf{x}_k \Delta\mathbf{x}_{k+1}}{|\Delta\mathbf{x}_k \Delta\mathbf{x}_{k+1}|} \ge \cos(\alpha) \ge 0 \quad \text{and} \quad k_{sim} \le \frac{\Delta\theta_k}{\Delta\theta_{k+1}} \le \frac{1}{k_{sim}} \tag{7}$$

Relations (7) are satisfied if the subsequent pose discrepancy vectors have similar directions and subsequent orientation corrections are similar to a certain degree; the factor $0<k_{sim}<1$ defines the required similarity between the subsequent orientation discrepancies.

With the help of (7) the APE can detect the filter's inconsistency and adapt to different odometric errors automatically. After a temporary increase of the uncertainty intervals the localization validation gates normally get smaller as the robot navigates relative to a consistent world model. Clearly, if a significant portion of the perceived objects in a certain neighborhood would be moved in the same way, the rules (7) could erroneously indicate an odometric error. However, in this case, the error is usually compensated as soon as the robot enters an area corresponding to a consistent portion of the world model.

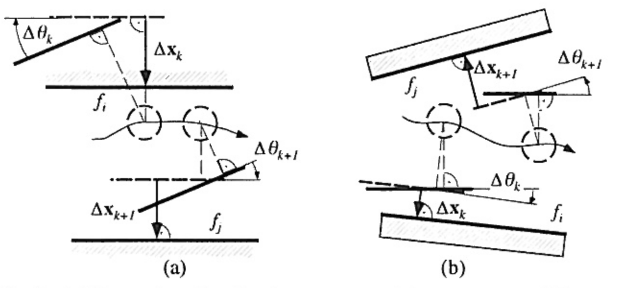

Fig. 3: a) Odometric or localization errors result in a sequence of discrepancies, which are out of bounds $\bar{c}_k^{\,t}$ and $\bar{c}_k^{\,\theta}$ and satisfy relations (8), if the map is consistent. (b) Conditions (8) are not satisfied since the map is inconsistent.

7 Maintenance of the Map's Consistency

With the *Geometric Consistency Filters* [4] we can estimate geometric consistency of the portions of the world model. The *Geometric Consistency Filter* basically accumulates evidence and determines the likelihood $L(C_j|\{w\}_k)$ that the reference feature f_j is consistent. Since the presented localization process is very selective, the *localization weights* (see (6)) represent reliable evidence which is fused with the help of the Bayes' rule

$$L(C_j|\{w\}_{k+1}) = \frac{w_i L(C_j|\{w\}_k)}{w_i L(C_j|\{w\}_k) + (1-w_i) L(\neg C_j|\{w\}_k)}, \tag{8}$$

where $L(\neg C_j|\{w\}_k) = 1 - L(C_j|\{w\}_k)$. Each w_i can be interpreted as a likelihood, that current discrepancy vector $\Delta\mathbf{p}_i$ could be observed, given that the reference feature f_j is consistent. It should be noted that the Bayes' rule is used rather in a qualitative manner as a practical tool for information fusion. Namely, assumptions about the independence of subsequent pieces of evidence are partially violated; pose discrepancy $\Delta\mathbf{p}_i$ influences $\Delta\mathbf{p}_{i+1}$ to a certain extent. While these dependencies cannot be eliminated completely, they can be reduced significantly by weighting the pose corrections (see (2)). In most cases dependencies can be reduced such, that the updating process converges to correct values; consistent features are usually associated with $L(C_j|\{w\}_k)$'s which converge to *1*, while $L(C_j|\{w\}_k)$'s of inconsistent features converge to *0*.

If the likelihood $L(C_j|\{w\}_k)$ is smaller than the threshold T_{update}, the robot's pose is not updated while the reference feature f_r is removed and the currently extracted feature is accumulated into the world model (see fig. 1).

8 Experimental Results

The proposed approach to the localization and map updating has been verified in typical office environments. The robot initially generated a simple world model (fig. 4) by using a ring of 24 Polaroid sonars and the wheel encoders [3]. With the help of the generated world model the robot was able to estimate its pose without artificial beacons.

Moreover, the implementation was simplified by assuming that the learned uncertainties $\bar{c}^{\,t}$ are isotropic; i.e. they are assumed to be uniform in all directions. This is in general not true and temporary inconsistencies of the APE can be introduced through this assumption. However, the APE normally adapts very quickly to such situations and, therefore, the mentioned simplification does not influence the localization significantly.

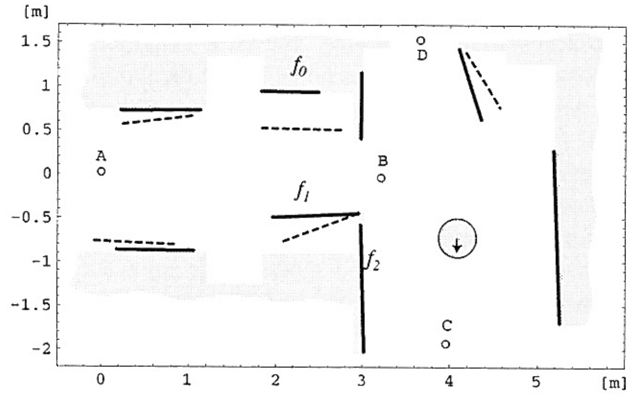

Fig. 4: A simple world model consisting of linear features was generated autonomously by the robot. Solid and dashed lines correspond to the initial and final states of the workspace, respectively. Grey regions represent the initial layout of the workspace. The objects corresponding to the features f_0 and f_1 were moved simultaneously.

While the pose estimation was based on the presented APE, also "raw" discrepancies $|\Delta \mathbf{x}_k|$ and $\Delta\theta_k$ between the matched features were recorded. Initially, the used world model was geometrically consistent and, as the robot moved in the known environment, both the pose corrections resulting from the APE and discrepancies $|\Delta \mathbf{x}_k|$ and $\Delta\theta_k$ were relatively small. After a while, however, geometric inconsistencies were introduced gradually by moving two doors and desks (dashed lines in fig. 4). These objects were moved in such a way, that they were still within the matching validation gates during the following navigation.

Diagrams (a) and (b) in fig. 5 show "raw" discrepancies $|\Delta \mathbf{x}_k|$ and $\Delta\theta_k$ between the matched features. After inconsistencies in the map were introduced, however, a significant portion of the discrepancies $|\Delta \mathbf{x}_k|$ and $\Delta\theta_k$ was erroneous. Diagrams (c) and (d) in fig. 5, on the other hand, show corrections computed with the APE during the same experiment.

Up to the localization step 400 the used world model was geometrically consistent. The peaks between the steps 200 and 400 in fig. 5 (a-b) correspond to corrections following significant abrupt odometric errors, which were introduced through motion over thresholds and thick carpets. After the step 400, however, geometric inconsistencies were introduced by gradually moving some objects in the workspace (fig. 4), which is evident from the huge pose discrepancies $|\Delta \mathbf{x}_k|$ and $\Delta\theta_k$.

However, It is evident that the APE could cope with inconsistencies in the world model; erroneous "raw" pose discrepancies between the matched features could be filtered out and the accuracy of the APE was not significantly influenced. Note also peaks in diagrams (c) and (d) in fig. 5, which correspond to corrections after abrupt and large odometric errors.

Moreover, with the help of the *Consistency Filter*, the world model was updated on-line, which resulted in reduced geometric inconsistencies and relatively small discrepancies between the matched features; e.g. between the steps 630 and 850 in fig. 5 (a-b).

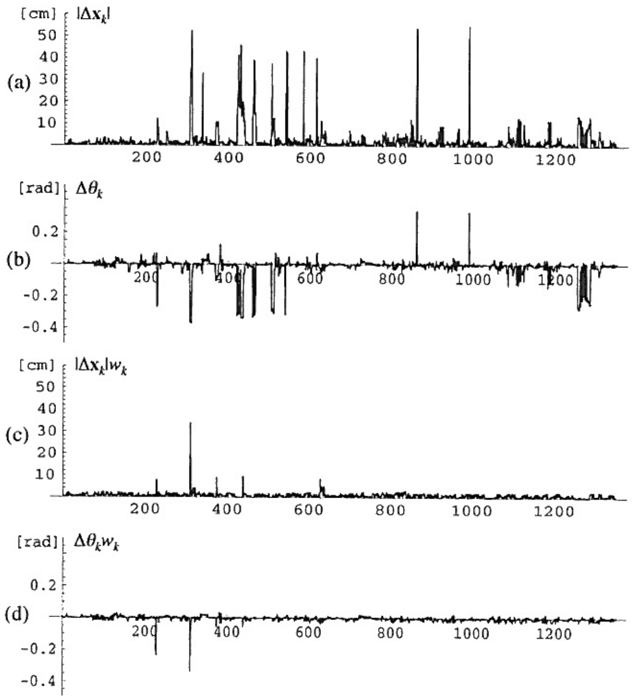

Fig. 5: **(a-b)** Raw discrepancies between the matched features. **(c-d)** Pose corrections based on the APE were not significantly influenced by inconsistencies introduced after the step 400. Innovations based on inconsistent information were set to 0.

Selective filtering of relevant evidence was based on the learned average pose corrections, which were in average relatively "tight" (see fig. 6). Therefore, the unmodelled odometric and localization errors occasionally resulted in temporarily inconsistent estimator. However, APE could detect such situations with the help of (7) and adapt quickly; the peaks in fig. 6 correspond to adapted average pose corrections which resulted from detection of the APE's inconsistency.

Moreover, the deviation between the robot's true and the estimated pose was controlled at different control points in the workspace. The robot was centered at each control point by the operator and the discrepancy was recorded. It turns out that the deviations between the true and the estimated poses are bounded and they are relatively small. In average, the pose discrepancy remained below 3 cm and 4°

which is evident also from the small pose corrections (fig. 5-c and fig. 5-d).

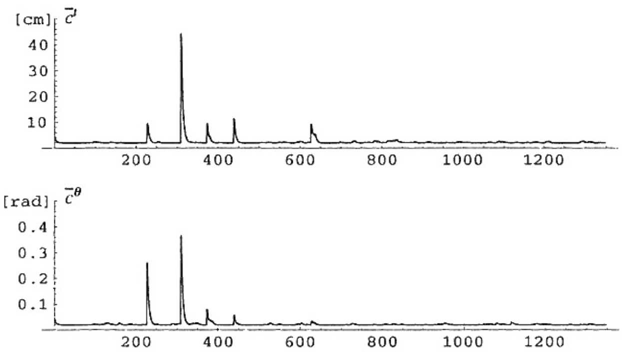

Fig. 6: Magnitudes of average translation and orientation corrections. Between the steps 200 and 400 the average corrections were increased due to detection of estimator's inconsistency resulting from odometric errors.

The consistency of each feature in the world model (fig. 7) was estimated on-line with the help of the *Consistency Filter* (8); each localization step yielded a piece of evidence. If the likelihood of consistency $L(C_i|\{r\}_k) < 0.1$, the corresponding reference feature f_i was removed and the currently extracted feature was accumulated instead. In this way the world model could be adapted to new configurations of the workspace on-line. Note, that all features observed between the control points *A* and *B* were gradually moved during the experiment (fig. 4). However, the accuracy of the pose estimation was not influenced significantly.

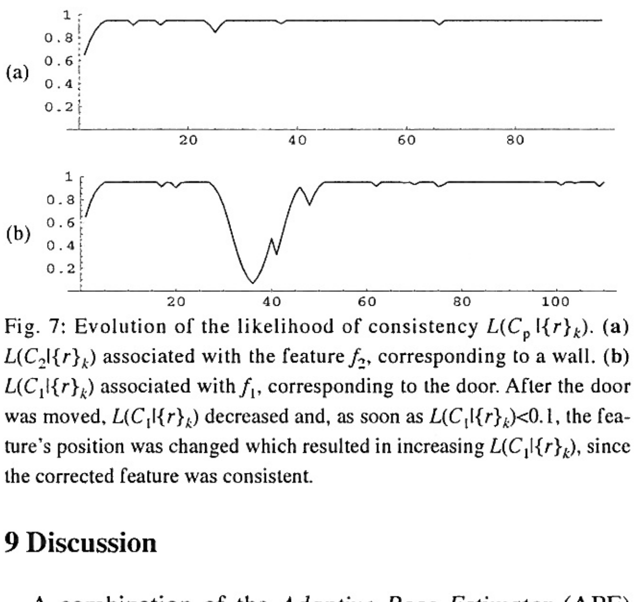

Fig. 7: Evolution of the likelihood of consistency $L(C_p|\{r\}_k)$. (a) $L(C_2|\{r\}_k)$ associated with the feature f_2, corresponding to a wall. (b) $L(C_1|\{r\}_k)$ associated with f_1, corresponding to the door. After the door was moved, $L(C_1|\{r\}_k)$ decreased and, as soon as $L(C_1|\{r\}_k)<0.1$, the feature's position was changed which resulted in increasing $L(C_1|\{r\}_k)$, since the corrected feature was consistent.

9 Discussion

A combination of the *Adaptive Pose Estimator* (APE) and the *Consistency Filter* allows robust localization in partially unpredictable indoor environments and on-line updating of internal world models. In general, the APE can cope with different kinds of unpredictable aspects of the environments such as mobile objects, sensing and odometric errors etc. This is achieved in a very simple way. Namely, the presented approach uses simple models in conjunction with the information obtained through observation of the localization process. In contrast to common approaches to pose tracking, the APE requires only few coarse assumptions about properties of the world.

Also, the *localization weights* determined by the *Adaptive Pose Estimator* provide useful evidence about the consistency of the portions of the world model. By fusing such evidence with the *Consistency Filter*, we are able to isolate information which is inconsistent and even improve the consistency of the world models. The knowledge about consistency of information can further enhance robustness of the APE, because possible erroneous interpretations of the constellations can be significantly reduced.

Finally, it should be noted that the APE is relevant especially for the robots using simple sensors, which require clustering of data along relatively long paths in order to extract reliable geometric features. Consequently, odometric errors can introduce significant geometric inconsistencies within the sets of subsequently obtained features which, in turn, can result in devastating errors if the pose is estimated by solving a set of equations. In addition, it can be shown, that the pose errors remain bounded if the robot's pose is corrected with one linear reference feature at a time, given that not all mapped features are parallel.

Acknowledgments

This work is fully supported by FWF, grant P12861-MAT.

References

[1] D. Fox, W. Burgard, S. Thrun, A.B. Cremers: Position Estimation for Mobile Robots in Dynamic Environments, In *J. of Artificial Intelligence Research*, 1999.

[2] L. Jetto, S. Longhi, G. Venturini: Development and Experimental Validation of an Adaptive Extended Kalman Filter for the Localization of Mobile Robots, In *IEEE Transactions on Robotics and Automation*, VOL. 15, No. 2, April 1999.

[3] G. Pavlin, R. Braunstingl: Context-Based Feature Extraction With Wide-Angle Sonars, In proc. of *International Conf. on Intelligent Robots and Systems*, Takamatsu, Japan, 2000.

[4] G. Pavlin, R. Braunstingl: Estimating Consistency of Geometric World Models Through Observation of the Localization Process, In proc. of *International Conference on Robotics and Automation, ICRA'2001*, Seoul, Korea, 2001.

[5] G. Pavlin, R. Braunstingl: Robust Pose Estimation for Mobile Robots in Unpredictable Environments, *IFAC Workshop on Mobile Robot Technology*, Jejudo, Korea, 2001.

[6] R. Simmons, S. Koenig: Probabilistic Navigation in Partially Observable Environments, In proceedings of *AAAI International Joint Conference on Artificial Intelligence*, 1995.