

Hybrid, High-Precision Localisation for the Mail Distributing Mobile Robot System MOPS

Kai O. Arras[†], Sjur J. Vestli[‡]

[†] Autonomous Systems Lab, Swiss Federal Institute of Technology (EPFL), CH-1015 Lausanne

[‡] Institute of Robotics, Swiss Federal Institute of Technology (ETHZ), CH-8092 Zurich

Abstract

This paper describes the new localisation algorithms under implementation for the mail distributing mobile robot, MOPS, of the Institute of Robotics, Swiss Federal Institute of Technology Zurich. Using geometric primitives as features, we employ consistent probabilistic feature extraction, clustering, matching and estimation of the vehicle position and orientation. The extracted features and their first-order covariance estimates are used, together with a world model, by an extended Kalman filter so as to get an optimal estimate of MOPS' current pose vector and the associated uncertainty. The line extraction consists of an initial segmentation, based on a feature-independent compactness measure in the model space, and a subsequent probabilistic clustering step. This yields a highly accurate and efficient localisation.

1. Introduction

The Mobile Post System (MOPS), see figure 1, is a wheeled mobile robot employed to distribute mail in a new building at the Swiss Federal Institute of Technology in Zürich. An important aspect of this application is that MOPS operates in a standard office environment. A general overview of the system and the task can be found in [1].

A critical factor influencing MOPS' ability to solve the mail distribution tasks is his knowledge, at all times, of the instantaneous position and orientation within the building. Should such information not be available to the rest of MOPS' control system it would not be possible to distribute mail. In this paper we will describe recent major improvements to the localisation sub-system of MOPS' control software.

In section "Robot System" we describe the complete robot system being used for the mail distribution application. In particular we will describe the new sensor system. To enable MOPS to manipulate the mail, a 3-dof mechanism is mounted on top of the mobile platform. The whole system is modular; the manipulator can be replaced with another module to fulfil other service tasks in an office building than mail distribution.




Figure 1: The Mobile Post System MOPS. The upper part of the system constitutes a 3-DOF manipulator / storage mechanism for two mail boxes. The lower part is a MECOS mobile robot platform equipped with two laser scanners, 12 ultrasound sensors, 2 triangulation sensors, a tactile bumper, infrared communication system, wireless modem, speech output system and a VME-based multi-processing robot controller for 6 servo axis (2 used for propulsion, 3 for manipulation, 1 spare).

Since we are using a method based on feature matching (a feature is typically a wall, a corner, a cylinder or a distinctive point somewhere in the environment) we will enter

into the details of the feature extraction performed on the robots extroperceptive sensor data in the chapter "Probabilistic Feature Extraction". The extracted features are later on used to localize the robot with the assistance of an extended Kalman filter and an on-board object-based map of MOPS' environment.

2. Robot System

For the task described above a non-holonomic mobile robot platform with two diametrically opposed drive wheels and a free-wheeling castor (the same kinematics as the robot Newt, as described in [2]) is being used, see figure 1. The two drive wheels are located at one end (the front end) of the vehicle, and the support castor at the other end.

The basic platform contains a multi-processing controller for servo motors (DC, EC or AC) based around the VME bus, currently the robot is equipped with one 200 MHz PowerPC 604 processing board (previously a dual processor 68040 system was used). The drive wheels constitutes two servo axis, and the remaining 4 axis are consequently available for a manipulating mechanism. Such a mechanism has been integrated on top of the platform. The programming is done in the object-oriented real-time development environment XOberon for general, high-performance mechatronic systems [3].

The basic robot platform is equipped with two scanning laser range-finders (front, rear), twelve ultrasonic range finders, and a bumper system with six segments. A radio modem enables wireless communication and various indicator lights and speech output module informs the environment about the robots immediate action. An infra-red communication system is used to communicate with the building infra-structure (elevator, pigeon hole blinds, etc.).

Recently the laser range finders has been upgraded to the latest model, the LMS200 from the company SICK [4]. This laser range finder is based on a time-of-flight principle and therefore delivers range data of exceptional linearity. The scanning mechanism of the sensor allows a polar range image to be generated between 0 and 180 degrees with a maximum angular resolution of 0.25°. Recording one complete scan takes only a few tens of milliseconds (mirror velocity 483 or 118 radians/s depending on angular resolution), and due to its high bandwidth serial communication interface all the data can be transferred to MOPS' navigation computer in real-time. The radial resolution of the sensor system is 10 mm. The maximum range of the sensor is configurable, in this application a maximum range of 8 m is used.

In addition to providing range data the sensor can also detect retro-reflective beacons. The detection of reflective beacons is done simultaneously with the range measurement and the information on the presence/absence of beacons is transmitted as intensity information with the range data. There is no difference in range noise between meas-

urements lying on such beacons and range measurement on normal surfaces. The beacons can be standard reflective foils such as 3M Scotchlite Diamond Grade. Reliable detection of reflective beacons can be achieved on the range interval of 1 to 3 metres.

3. Probabilistic Feature Extraction

In order to achieve high estimation accuracy, consistent propagation of the first and second moments from single range reading to all system stages involved in the position estimation is essential. Extraction of geometrical primitives from noisy 1D range data has to furnish the first two moments of the model parameter estimates. A Kalman filter for extraction provides this [10], [14], but includes the problem of initialising the filter each time a new feature segment is detected. By working on polar co-ordinates, and not with an intermediate Cartesian representation [15], it is directly possible to determine the first-order covariance estimate.

Using the range and angular information (ρ and θ) as well as the intensity information from the robots laser range finders we extract the Cartesian co-ordinates of the beacons and the parameters (r and α) of the line model in eq. (1).

$$\rho \cos(\theta - \alpha) - r = 0 \quad (1)$$

3.1 Estimating a line and its covariance matrix using polar co-ordinates

Feature extraction can be subdivided into two questions: *a) which* of the range measurements in the range image are generated by the sought structure in the environment (segmentation), and *b) how* do these points contribute to the estimation of the model parameters and their covariance matrix (model fit).

The answer to question b) is, in our case, the estimation of the line parameters α , r in the general least squares sense using polar co-ordinates minimising the perpendicular distance from the points to the line. It can be shown that the solution is as given in eq. (2) and eq. (3).

$$\tan 2\alpha = \frac{2 \sum_{i < j} w_i w_j \rho_i \rho_j \sin(\theta_i + \theta_j) + \sum (w_i - \sum w_j) w_i \rho_i^2 \sin 2\theta_i}{2 \sum_{i < j} w_i w_j \rho_i \rho_j \cos(\theta_i + \theta_j) + \sum (w_i - \sum w_j) w_i \rho_i^2 \cos 2\theta_i} \quad (2)$$

$$r = \frac{\sum w_i \rho_i \cos(\theta_i - \alpha)}{\sum w_i} \quad (3)$$

Where ρ_i is the range value of a single measurement, θ_i its angle and w_i the associated weight. Instead of implementing eq. (2) and (3) directly, we choose to use its identical weighted Cartesian form which is clearly more efficient. See also [5].

However, if we want to determine the covariance matrix

we need to know eq. (2) and (3). By writing the model parameters as random variables represented by their first-order Taylor series expansion about the mean point and assuming independence of radial and angular noise, the first-order covariance estimate is given in eq. (4).

$$C_l = m_{\rho\theta} C_{\rho\theta} m_{\rho\theta}^T \quad (4)$$

Where $m_{\rho\theta}$ is the Jacobian matrix containing all partial derivatives of the model parameters with respect to all measurements r and q involved in the estimation.

The weights can be chosen as the inverse of the variances giving most weight to data points with low variability. By measuring the signal strength of the reflected beam knowing its relation to the variance of radial noise, one can determine a weight for each range reading. However not all commercially available laser scanners provide this information. In this case the relationship of range variance to distance can be experimentally identified. It is obvious that all other sources which influence range variance (surface properties or beam splitting) are no longer taken into account.

Question *a*) is a segmentation problem. In this work we employ a probabilistic feature-independent fidelity criterion in the model space and a subsequent clustering step.

The model is fitted into n_f neighbouring points and the covariance matrix is computed. This is done for all points of the scan obtaining the same number of spatially diffused Gaussian distributions in the model space. When adjacent groups of range readings lie on the same landmark, their associated points constitute a cluster of these distributions in the model space corresponding to that landmark. Feature extraction is now the task of finding these clusters. In a general case a clustering problem of this size where no a priori knowledge is available would lead to impracticably high computation times under real time conditions. Here the fact can be exploited that points on the same landmark are *almost always* consecutive points. Due to this underlying regularity in the acquisition process, a compactness measure in the model space, c_i , is defined according to eq. (5) and applied to n_m adjacent points.

$$c_i = \sum_j (x_j - x_w)^T (C_j + C_w)^{-1} (x_j - x_w) \quad (5)$$

Where $j = i - (n_m - 1)/2, \dots, i + (n_m - 1)/2$ and x_w is the weighted mean and C_w its covariance matrix, see eq. (7) and (8) of section 3.3.

Low distance indicates that the points involved have high model fidelity. If c_i is plotted against the measurement index for all i , regions of low value can be expected at the corresponding index places of the sought clusters. A threshold c_m , found by deciding on the appropriate distribution of c , is applied cutting off the regions of low distance. A contributing segment is now defined to be the set of measurement points those representations in the model space satisfy $c_i < c_m$.

3.2 Clustering

After the completion of this step there will be, in general, more than one segment corresponding to one single landmark. We have an oversegmented range image. Associating these segments with each other is done in a clustering step. Otherwise valuable information would be lost, for example in a matching procedure of a Kalman filter-based localization task, where only the observed segment which is closest to the prediction will be considered, whereas the others belonging to the same landmark would be ignored. On the other hand, two landmarks which differ only slightly in one or more of their model parameters, should be identified as being distinct within the limits which are given by sensor noise.

The matching problem, i.e. finding correspondence of objects represented by randomly distributed vectors in some model space, also called *data* or *measurement association*, arises, e.g., in the target tracking domain [6]. The goal is to produce an assignment from a measurement to a predicted target in the presence of clutter causing false measurements. A validation region around the prediction in the model space is used to discard or confirm a measurement. When dealing with multiple targets or clutter its likely that a measurement falls in two validation regions or several measurements lie in a single validation region requiring more complex association concepts. In mobile robotics this problem arises not only if the Kalman filter is employed for localization [7] but generally when dealing with uncertain spatial information [8], [9], e.g., when matching features in visual images [10], [11]. However, opposed to this problem, it is always a pairwise correspondence which is sought, and not the unsupervised assignment of multiple estimates to an unknown reference.

Clustering techniques provide this treatment of data, and having gained a number of segments like previously explained or any other scheme providing propagation of the first two moments, the number of points in the model space has been significantly reduced. Starting from n , where n can be large, just the number of segments, n_s , remain. Hence, from the viewpoint of clustering techniques, segmentation can be perceived as being a preprocessing step in order to reduce input information leading to small clusters at the same locations where the original clusters have been observed before. An agglomerative hierarchical clustering algorithm [12] which permits an efficient implementation is utilized and due to its simplicity a short outline is given:

Having computed the $n_s \times n_s$ symmetric distance matrix D , the procedure starts with each point as a separate cluster.

- Find the minimal distance d_{ij} of clusters Q_i and Q_j in D and while $d_{ij} < d_{\alpha_e}$ proceed to
- Merge Q_i and Q_j obtaining Q_{ij} The number of clusters is decreased by one.
- Update D by calculating the new distances from Q_{ij} to all other clusters. Only one column and row is changed.

Then go back to the first step

The problem of determining the distance between two clusters is not straightforward and there are several propositions for its solution. Often it remains a matter of specific experimentation to select the most appropriate method. However, with the covariance estimates, there is additional information on the cluster elements available which can be exploited for eliminating the usual ambiguities. One needs not to deal with clusters consisting in more than one element since in the merging step, two distributions are instantly fused by making a new model fit with the combined set of measurements in the image space. Hence, the distance between two clusters is always the distance of two normally distributed vectors in the model space and we can refer back to the matching problem where the Mahalanobis distance, see eq. (6), is widely employed for that purpose.

$$d_{ij}^2 = (x_i - x_j)^T (C_i + C_j)^{-1} (x_i - x_j) \quad (6)$$

If x_i and x_j belong to the same landmark, d_{ij}^2 has chi-square distribution and an appropriate threshold must be chosen. It is interesting to note that if the points in the model space were truly normally distributed and therefore completely described by the first two moments, merging could be efficiently done with eq. (7) and (8) without leaving the model space. Experiments showed, however, that iterated fusion leads to accumulating deviations from the (true) fit alternative, which for large structures where many segments have been merged, became non-negligible.

The final estimates of the parameters are directly available after exiting the clustering algorithm. The final segments are obtained by combining measurement points of adjacent segments if they belong to the same line.

3.3 Probabilistic Beacon Extraction

By using the intensity information of the LMS200, the segmentation problem for the beacons is already solved.

In general, several measurement points lie on a retro-reflective beacon and the issue of observational uncertainty is addressed in the same way as for measurement points not lying on a retro reflector: the radial noise is assumed to be given in relationship to the distance and the angular uncertainty is assumed to be zero. The overall variability in radial direction is thus the weighted mean as found through eq. (7) and eq. (8).

$$x_w = C_w \sum C_i^{-1} x_i \quad (7)$$

$$C_w^{-1} = \sum C_i^{-1} \quad (8)$$

Where x_i is a (scalar) range reading and C_i is the (scalar) covariance estimate of range reading i .

The angular beacon location uncertainty is estimated according to the ad-hoc assumption that the centre of the beacon is within the angular interval spanned by the reflector

range readings and the two adjacent points with a 99% probability. With $\Delta\theta$ as the angular resolution, n_b representing the number of measurement points on the beacon and assuming Gaussian distributions, the mean and the variance of the angle is given as in eq. (9) and (10).

$$\sigma_{b\theta} = \frac{1}{\alpha_{0.99}} \Delta\theta (n_b + 1) \quad (9)$$

$$\mu_{b\theta} = \frac{1}{n_b} \sum \theta_{bi} \quad (10)$$

Further, assuming independence of radial and angular errors we find the beacon location and covariance as by eq. (11), (12) and (13).

$$\mu_{bp} = \left(\sum \sigma_{bpi}^{-2} \right)^{-1} \sum \rho_{bi} / \sigma_{bpi}^2 \quad (11)$$

$$\mu_{b\theta} = \frac{1}{n_b} \sum \theta_{bi} \quad (12)$$

$$C_b = \begin{bmatrix} \left(\sum \sigma_{bpi}^{-2} \right)^{-1} & 0 \\ 0 & \left(\frac{1}{\alpha_{0.99}} \Delta\theta (n_b + 1) \right)^{-2} \end{bmatrix} \quad (13)$$

4. Hybrid Position Estimation

In this section we describe how to use the extracted lines, their covariance estimates, and beacons for the task of accurately calculating MOPS' Cartesian position in the global reference frame.

Consider a localization scheme which relies on an extended Kalman filter to estimate the vehicles pose vector $x(k) = [x(k) \ y(k) \ \psi(k)]$ and the corresponding covariance estimate $P(k)$. We can then use the result of [6] that the Kalman filter gain can be written as in eq. (14)

$$W(k+1) = P(k+1|k)H^T S^{-1}(k+1) \quad (14)$$

which allows us to compute the updated position estimate as in eq. (15)

$$\hat{x}(k+1|k+1) = \hat{x}(k+1|k) + W(k+1)v(k+1) \quad (15)$$

With a variance of the position estimate as in eq. (16).

$$P(k+1|k+1) = P(k+1) - W(k+1)S(k+1)W^T(k+1) \quad (16)$$

The above notation corresponds to what was used in [7]. The innovations $v_i(k+1)$ is contained in the innovations vector $v(k+1)$ and are the differences between the observed model parameter vectors, $m_i(k+1)$, and their predictions, $\hat{m}_i(k+1)$, both in the sensor coordinate system, eq. (17).

$$v_i(k+1) = m_i(k+1) - \hat{m}_i(k+1) \quad (17)$$

The innovation covariance matrices, given in eq. (18),

quantify the error propagation from the position and orientation uncertainty to the target prediction $\hat{m}_i(k+1)$ and add to that the feature specific observational uncertainty C_{mi} of each observed target.

$$S_i(k+1) = H_i P(k+1|k) H_i^T + C_{mi} \quad (18)$$

In eq. (18) the matrix H_i denotes the linearised measurement model ${}^S\hat{m}_l(\cdot)$, for lines, and ${}^S\hat{m}_b(\cdot)$ for beacons, see also eq. (19) (for lines) and eq. (20) (for beacons).

$${}^S\hat{m}_l(k+1) = h_l^W m_l, \hat{x}(k+1|k) \quad (19)$$

$${}^S\hat{m}_b(k+1) = h_b^W m_b, \hat{x}(k+1|k) \quad (20)$$

The operation performed in eq. (19) and (20) is equivalent to a coordinate system transformation world-to-sensor for the model parameters ${}^W m$.

The innovations are obtained by matching observations to predicted targets. The matching criterion is given as elliptical validation regions in eq. (21).

$$v_i(k+1) S_i^{-1}(k+1) v_i(k+1)^T \leq \chi_{\alpha, 2}^2 \quad (21)$$

Similar to previous work on this topic, [7], we only accept a match when a single observation falls into a single validation region. Experiments with our mobile robot platforms performing real-time localization while moving (see also next section) showed that such simple and conservative matching strategy is sufficient even in complex real-world environments under the conditions that the localisation cycle is fast and that the feature extraction is precise. It is important to note that this matching is done in the unmodified model space, i.e. only with r and α , opposed to the more unreliable matching of single segments. This map organisation and matching strategy is motivated by the subsequent clustering which is part of the extraction process.

5. Implementation and Results

In this section we will present extraction results based on scans from the Acuity AccuRange3000LV. The sensor was configured to have a medium angular resolution of nearly 1° on 360° and is of relative low quality (figure 2 and figure 3). Currently the algorithms are being implemented on the MOPS system with the LMS200 sensors.

Each measurement receives a standard deviation based on a linear relationship to the distance, indicated by 4σ -intervals. Experiments with two other commercially available laser scanners (Leuze RS3 and SICK PLS100) showed that it is not crucial for the performance of the extraction to have elaborated sensor models for the range variances [5].

The localization was performed by a robot system with a 68020 CPU running at 20MHz equipped with a Leuze RS3. Its implementation is reduced such that only the first statistical moments are propagated. For all distance measures their Euclidian equivalents have been taken, the thresholds




Figure 2: 360°-scan of the Acuity AccuRange3000 LV. A scan with medium angular resolution while radial noise is relatively large. The simulated standard deviations are indicated by 4σ -intervals.




Figure 3: (a) Extracted lines after the agglomerative hierarchical clustering. Corners which are found according to the scheme in [5] are depicted by crosses and their 99% error ellipse. (b) Detail from (a). Corner probabilities are diffused in direction of the walls which is to be expected with perpendicular corners.

were informally determined. The overall performance especially with respect to robustness against outliers and discrimination property in the clustering step is clearly inferior compared to the presented probabilistic version.

Due to this, a preprocessing which eliminates spurious points and thins out the raw range data became necessary.

Firstly the extraction accuracy for lines was identified. The robot was put into a corner about two meters away from a metal cupboard and a large cardboard box, where 300 extractions were made. The hit rate was 100%, the standard deviation in r was found to be 5 mm, standard deviation in α was 0.2° . Computation time including preprocessing was not greater than 240 ms. Then the extraction was used for a map-based localization task with an extended Kalman filter, *solely relying on the line extraction*. Position update while moving is accomplished through a good interplay with the underlying position controller for non-holonomic kinematic configurations [13]. The robot can be manually pushed away from its given trajectory to which it responds with smoothly returning to its path within a few position update cycles. The localization accuracy which was immediately obtained without further optimization of any vehicle parameters was below one centimetre when the robot positioned itself at a desired location. As a further estimate for its precision, the noise of the position estimates has been identified when the robot was stationary. Standard deviation in x and y was found to be 6 mm, standard deviation in θ was 0.05° . The average localization cycle time during a journey through lab and corridors was 445 ms.

Recent experiments on the current MOPS hardware leads to expect that this over all processing time will be lower by a factor of 50 (i.e. around 10 ms).

6. Conclusions and Future Work

This paper has described the new localisation algorithms being implemented on the MOPS for its task of distributing the mail.

We have shown that the system has a highly accurate localisation system, which in turn enables the robot to "dock onto" the pigeon holes and to load/unload boxes of mail. Furthermore the position update process is tolerant towards unmodelled obstacles and it is also fast enough for utilisation in the real world.

Currently the system relies upon an a-priori map. Investigations are being undertaken with the goal of eliminating this costly, and unnecessary step.

The use of a vision system for the recognition of "important" scenes, e.g. for reading the room numbers, is also being examined.

7. References

- [1] Vestli S. J., Tschichold-Gürman N., "MOPS, a System for Mail Distribution in Office Type Buildings", Proc. of the 1996 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS'96), Osaka, Japan, November 4-8, 1996
- [2] Muir P. F., Neuman C. P., "Kinematic Modelling of Wheeled Mobile Robots", CMU, the Robotics Institute, Technical Report, CMU-RI-TR-86-12, 1986.
- [3] Diez D., Vestli S. J., "D'nia an object oriented real time system", Real-Time Magazine, Real-Time Consult Publ., rue de la Justice 23, 1070 Brussels, Belgium, August, 1995.
- [4] N.N., "LMS/LMI, Definition der Telegramme zwischen Benutzerschnittstelle und LMS- oder LMI-System über RS 422/RS 232", Internal Document, August 1997, SICK AG, Sebastian-Kneipp-Str. 1, 79183 Waldkirch, Germany.
- [5] Arras K.O., Siegwart R.Y., "Feature Extraction and Scene Interpretation for Map-Based Navigation and Map Building", Proc. of SPIE, Mobile Robotics XII, Vol. 3210, 1997.
- [6] Bar-Shalom Y., Fortmann T.E., "Tracking and Data Association", Mathematics in Science and Engineering, Vol. 179, Academic Press Inc., 1988.
- [7] Leonard J. J., Durrant-Whyte H. F., "Directed Sonar Sensing for Mobile Robot Navigation", Kluwer Academic Publishers, 1992.
- [8] Smith R., Self M., Cheeseman P., "Estimating Uncertain Spatial Relationships in Robotics", in Autonomous Robot Vehicles, Cox I.J., Wilfong G.T. (eds.), Springer, 1990.
- [9] Hebert P., Betge-Brezetz S., Chatila R., "Probabilistic Map Learning: Necessity and Difficulties, in Reasoning with Uncertainty in Robotics, Dorst L., van Lambalgen M., Voorbraak F. (eds.), Lecture Notes in Artificial Intelligence, Vol. 1093, Springer, 1995.
- [10] Ayache N., Faugeras O.D., "Maintaining Representation of the Environment of a Mobile Robot", in Autonomous Robot Vehicles, Cox I.J., Wilfong G.T. (eds.), Springer, 1990
- [11] Hummerl R., "Uncertainty Reasoning in Object Recognition by Image Processing", in Reasoning with Uncertainty in Robotics, Dorst L., van Lambalgen M., Voorbraak F. (eds.), Lecture Notes in Artificial Intelligence, Vol. 1093, Springer, 1995.
- [12] Kaufman L., Rousseeuw P.J., Finding Groups in Data: An Introduction to Cluster Analysis, Wiley Series in Probability and Mathematical Statistics, 1990.
- [13] Kaiser O., Pfiffner R., Vestli S., Astolfi A., "Positionsregelung für nicht-holonome mobile Roboter", 11. Fachgespräche Autonome Mobile Systeme, Germany, 1995.
- [14] Taylor R.M., Probert P.J., "Range Finding and Feature Extraction by Segmentation of Images for Mobile Robot Navigation", Proc. of the IEEE Int. Conf. on Robotics and Automation, Minneapolis, 1996.
- [15] Vandorpe J., Van Brussel H., Xu H., "Exact Dynamic Map Building for a Mobile Robot using Geometrical Primitives Produced by a 2D Range Finder", Proc. of the IEEE Int. Conf. on Robotics and Automation, Minneapolis, 1996.