

PEOPLE TRACKING USING LASER RANGE SCANNERS AND VISION

Andreas Kräußling, Bernd Brüggemann, Dirk Schulz

Department of Communication, Information Processing and Ergonomics (FKIE)

Research Establishment for Applied Sciences (FGAN), Neuenahrer Straße 20, 53343 Wachtberg, Germany

Armin B. Cremers

Institute of Computer Science III, University of Bonn, Römerstraße 164, 53117 Bonn, Germany

Keywords: People tracking, crossing targets, laser range scanners, vision, HSV colour space.

Abstract: Tracking multiple crossing people is a great challenge, since common algorithms tend to lose some of the persons or to interchange their identities when they get close to each other and split up again. In several consecutive papers it was possible to develop an algorithm using data from laser range scanners which is able to track an arbitrary number of crossing people without any loss of track. In this paper we address the problem of rediscovering the identities of the persons after a crossing. Therefore, a camera system is applied. An infrared camera detects the people in the observation area and then a charge-coupled device camera is used to extract the colour information about those people. For the representation of the colour information the HSV colour space is applied using a histogram. Before the crossing the system learns the mean and the standard deviation of the colour distribution of each person. After the crossing the system relocates the identities by comparing the actually measured colour distributions with the distributions learnt before the crossing. Thereby, a Gaussian distribution of the colour values is assumed. The most probably assignment of the identities is then found using Munkres' Hungarian algorithm. It is proven with data from real world experiments that our approach can reassign the identities of the tracked persons stable after a crossing.

1 INTRODUCTION AND RELATED WORK

Multi-robot systems and service robots need to cooperate with each other and with humans in their environments. For this reason, they have to know about the locations and actions of the objects they want to interact with. Target tracking deals with the state estimation of one or more objects. It is a well studied topic in the field of aerial surveillance using radar devices (Bar-Shalom and Fortmann, 1988) and also in the area of mobile robotics. Here, mainly laser scanners are used for the purpose of people tracking (Prassler et al., 1999; Schulz et al., 2001; Fod et al., 2002; Romera et al., 2004; Zhao and Shibasaki, 2005; Bellotto and Hu, 2007). Due to the high resolution of laser scanners, which mostly cover a 180 degree field of view with 180 or 360 measurements, one target is usually the source of multiple returns within one laser scan. This conflicts with the assumption of punctiform targets used in the field of radar tracking.

There, each target is the origin of exactly one measurement. In contrast to that, using laser scanners, one needs to be able to assign the obtained measurements to extended targets.

A second important characteristic of tracking in the field of mobile robotics is the occurrence of crossing or interacting targets, for example two or more persons getting close to each other, so that they can no longer be distinguished by common tracking algorithms (Fortmann et al., 1983; Kräußling et al., 2005; Kräußling et al., 2007). In this article we present an approach to deal with this particular problem. The key idea of our approach is to adopt an algorithm for tracking punctiform objects in clutter, known from the radar community, for the purpose of reliably tracking extended objects with laser scanners. Several different methods for tracking punctiform crossing targets in clutter, i.e. tracking in the presence of false alarm measurements close to a target, have been developed over the last decades:

1. the MHT (Multi Hypothesis Tracker) introduced

by Reid in 1979 (Reid, 1979).

2. The JPDAF (Joint Probabilistic Data Association Filter) introduced by Fortmann, Bar-Shalom and Scheffe in 1983 (Fortmann et al., 1983).

These techniques can easily be extended to tracking extended objects as well, but there are several reasons, why such approaches are brittle:

- In most cases there are several measurements from the same target.
- Interacting objects might be indistinguishable over longer periods of time.
- Some of the objects might be occluded for some time.
- The objects can carry out abrupt manoeuvres, especially when they are crossing their paths.

These difficulties are well known in the mobile robotics community:

- Tracking moving objects whose trajectories cross each other is a very general problem ... Problems of this type cannot be eliminated even by more sophisticated methods ... (Prassler et al., 1999).
- Tracks are lost when people walk too closely together ... (Schumitch et al., 2006).

Due to these reasons, we have developed methods for tracking interacting people in laser data (Kräußling et al., 2004b; Kräußling et al., 2005; Kräußling et al., 2007). These methods have in common, that they employ a variant of the well known Viterbi algorithm (Viterbi, 1967; Forney Jr., 1973) in combination with geometrical properties of the people tracking problem, in order to achieve a high degree of robustness against track loss.

However, although the tracks are very rarely lost by these algorithms, they tend to confuse the assignment of the tracks to the individual persons being tracked after a crossing of paths has occurred. This happens because the distance measurements of the laser scanners do not provide direct information about the persons' identities. For this reason, additional cues are required, if we want to reliably distinguish between persons. Possible cues are:

1. Different colours and surface-textures of the pairs of trousers people wear might result in different intensities of the reflected laser beams.
2. Ultrasound or infrared signals uniquely identifying individuals, which are transmitted by special active badges the people wear (Schulz et al., 2003).
3. Different colours of the clothes people wear. This information can be exploited for the identification of the people using a camera network.

4. Differences in physiognomy like size and built of persons. These differences can again be detected using cameras (Schulz, 2006).

In this article we propose a technique to combine Viterbi-based tracking with person identification based on colour information. A calibrated setup consisting of an infrared and a CCD camera is used to learn colour histograms of the persons, while they are well separated during tracking. This information is then employed to correctly reassign person IDs to tracks after interactions have occurred. The new assignments are determined using the Hungarian algorithm, which computes the maximum likelihood assignment, based on the likelihood of colour observations. Our experiments show that this approach allows to track several interacting humans without loss of track and without accidental confusion of the track assignments.

The remainder of this paper is organised as follows. In Section 2 the combined method for tracking multiple interacting persons is described. It consists of the tracking method based on the Viterbi algorithm and an identity assignment method based on colour information. Section 3 presents experiments illustrating the robustness of our approach against loss of track as well as against errors in track assignment. We conclude in Section 4.

2 THE METHOD

In order to reliably keep track of several interacting persons, we have to solve two problems: the trajectories of the persons have to be estimated without losing track of the persons and we have to make sure, that we can always assign the individual trajectories to the correct person. We have developed the so called Cluster Sorting algorithm (CSA) to solve the first problem (Kräußling, 2006b). The CSA uses data from laser range scanners to estimate the trajectories of objects over time. The second problem is then solved by additionally using colour histograms extracted from camera images, in order to compute the most likely assignment of the trajectories to the persons being tracked. In the following, we will first explain the CSA in detail. Afterwards we will describe how the reassignment of tracks based on colour information can be integrated into the approach.

The Cluster Sorting algorithm estimates multiple trajectories using a hidden Gauß-Markov chain, where the tracking process is carried out using Kalman filters. Because laser range scanners return range measurements to any object in the surrounding of the robot, the measurements originating

from persons have to be discriminated from measurements of static objects. The CSA computes validation gates (Bar-Shalom and Fortmann, 1988) for this purpose, i.e. only measurements, which are close to the currently estimated positions of persons are being considered; we call those measurements the selected measurements. The distinction between measurements of different persons is possible based on the distance between selected measurements, as long as persons do not get close to each other; otherwise, persons share selected measurements. For this reason, the CSA deals with the selected measurements in two different ways:

(1) as long as the measurements of persons are well separated, it computes for each person the unweighted mean of all the selected measurements of that person. These means are then used to update the Kalman filters for the individual trajectories of each person; we call this procedure the Kalman Filter Algorithm (KFA). It has been shown in (Kräußling et al., 2005; Kräußling, 2006a) to be very fast and to provide good information about the position of the targets, but it cannot reproduce multi-modal probability distributions. Thus it is not able to handle multiple interacting people.

(2) When persons have selected measurements in common, the CSA no longer computes one single track for each person, but it starts to compute individual tracks for each selected measurement of each person using a variant of the Viterbi algorithm. The algorithm calculates for every old selected measurement a separate position estimate and validation gate. The new selected measurements are the ones which lie in at least one of those gates; we call this algorithm Viterbi-based algorithm (VBA); it has been introduced in (Kräußling et al., 2004a). The algorithm allows to represent multi-modal probability distributions to some extent, which is a major advantage when dealing with multiple interacting targets. The VBA is much more robust against track loss, when compared to the KFA, because the VBA maintains several hypotheses about a persons position, one hypothesis for each gating measurement. Common tracking algorithms like the KFA, in contrast, make a hard decision which measurement they use. In difficult situations, they tend to assign the same measurement to several objects. Algorithms with a random component like the SJPDAF (Schulz et al., 2001) occasionally choose for each track the path of a different person, so that no persons gets lost. But this behaviour is not stable (Kräußling and Schulz, 2006).

It remains to describe, how the CSA actually decides when to switch between the KFA and the VBA. The Cluster Sorting algorithm uses two classes of ob-

jects:

- single targets.
- clusters, which represent at least two interacting persons, i.e. humans that are moving very close to each other.

Single targets are tracked with the KFA, since there is no need for representing multi-modal probability distributions. Clusters are tracked with the VBA, since multi-modal distributions have to be represented. This approach guarantees that none of the objects that are associated with the cluster is lost. This fact is important especially when the objects split and start to move separately again.

Three different events have to be regarded when tracking multiple interacting people:

1. The merging of two single persons. This means that two single targets get very close to each other. This is the case, if at least one measurement is located in the validation gates of both targets. Then the algorithm stops to track the two single targets with the KFA and starts tracking a cluster, which contains both targets, using the VBA. Therefore, it uses the measurements located in the validation gates of at least one target.
2. The merging of a single human and a cluster. This means that a single person and a cluster get very close to each other. This happens, if at least one measurement is located in the validation gates of the person and the cluster. In this case the algorithm stops to track the single human and the cluster separately. Instead it starts tracking a combined cluster. Therefore, it uses the measurements located in the validation gates of either the single target or the previously considered cluster or both.
3. The merging of two clusters. This means that two clusters get very close to each other. This is the case, if at least one measurement is located in the validation gates of both clusters. If this is true, the algorithm stops to track the two clusters and starts tracking a combined cluster. Therefore, it uses the measurements located in the validation gates of at least one of the previously considered clusters.

Note, that whenever a merging takes place, the algorithm remembers the humans which correspond to the newly combined cluster.

For each tracked cluster, we also have to decide, if it has split into single person tracks again. Whether clusters are split depends on three conditions:

1. The position estimates corresponding to the measurements in the validation gates are separated into subclusters. For this purpose, we select the first estimate, which then is associated with the

first subcluster. For all other estimates associated with the cluster, the Euclidean distance to the first estimate is calculated. If this distance is below a certain threshold, the estimate is associated with the first subcluster. In our experiments, we set the threshold to 150 cm , which corresponds to the maximum distance between the legs of a walking person. We then have to consider the estimates, for which the Euclidean distance to the first subcluster exceeds this manually chosen threshold. Using the same procedure we applied for building the first subcluster, we now construct subclusters until all estimates are associated with one of these smaller clusters. If the number of subclusters equals the number of humans which were merged into this cluster, the first condition for the dispersion of the cluster is fulfilled. Then, we proceed with step 2.

2. We now check the pairwise distance between the subclusters. If the distance is above a manually chosen bound, we regard these clusters as separated. We choose the value of that bound to be 300 cm . The second condition is fulfilled, if the number of pairs of separated subclusters equals $\frac{n(n-1)}{2}$. Thereby, n is the number of single persons associated with the cluster. Hence, we are checking if all subclusters are pairwise separated.
3. Above this, we can separate single subclusters from the cluster to indicate them in the graphics. This follows the same logic as in step 1. Note, the algorithm is not able to determine, how many targets are represented by a single subcluster.

If conditions 1 and 2 are met, the n subclusters are associated with the n single targets therefrom tracked by the KFA. When separating targets from clusters, we cannot guarantee if the target association is the same as before merging the targets into the cluster. Thus, a possible solution to this problem, which uses colour information will now be proposed.

To obtain the colour information of the persons, a charge-coupled device camera (CCD camera) is used. In order to recognise which parts of the picture of the CCD camera belong to the persons, we employ an infrared camera. The two cameras are mounted in parallel on the robot, with only a small displacement; this allows us to easily correlate infrared and CCD images. The setup is shown in Figure 1. Since the temperature of the persons is in a small, well defined range, it is easy to identify the regions in the images of the CCD camera which originate from humans. Next, the persons which are detected by the camera system have to be assigned to the persons which are tracked by the laser scanners and the tracking algo-

Table 1: The values of the hue in the HSV colour space.

Hue	red	yellow	green	cyan	blue	magenta
Degrees	0	60	120	180	240	300




Figure 1: The robot equipped with the laser scanners and the cameras.

rithm. For this purpose, we exploit the fact that the camera system enumerates the persons in clockwise direction. Therefore, we arrange the persons tracked by the tracking algorithm in clockwise direction, too.

The colour information from the CCD camera is represented using the HSV colour space (Gonzalez and Woods, 1992), where H stands for hue, S for saturation and V for value. In our experiments we only used the hue, because it is fairly independent from the illumination and a stable characteristic of the persons. The hue values range from 0 to 360 degrees. The mapping between the values of the hue and six basic colours, we use for identification, is shown in Table 1.

The colour characteristics are learnt during the experiment before the crossings. We assume that the average relative frequencies of the hue for person i follow a Gaussian distribution with mean μ_i and variance Σ_i . Let $\mu_{i,k}$ and $\Sigma_{i,k}$ be the learnt mean and the learnt variance at time step k and let $y_{i,k}$ be the measurement corresponding to person i at time step k ,

$$\mu_{i,k} = \frac{\sum_{j=1}^k y_{i,j}}{k}, \quad (1)$$

$$\Sigma_{i,k} = \frac{\sum_{j=1}^k (y_{i,j} - \mu_{i,k})(y_{i,j} - \mu_{i,k})^T}{k-1}. \quad (2)$$

This mean and this variance can be learnt on-line

without storing previous values $y_{i,l}$, $l < k$, according to

$$\Delta_{k+1} = \mu_{i,k+1} - \mu_{i,k} \quad (3)$$

$$\mu_{i,k+1} = \frac{k}{k+1}\mu_{i,k} + \frac{y_{i,k+1}}{k+1} \quad (4)$$

$$\begin{aligned} \Sigma_{i,k+1} = & \frac{(y_{i,k+1} - \mu_{i,k+1})(y_{i,k+1} - \mu_{i,k+1})^\top}{k} + \\ & + \frac{k-1}{k}\Sigma_{i,k} + \Delta_{k+1}\Delta_{k+1}^\top. \end{aligned} \quad (5)$$

As soon as the tracking algorithm detects a crossing, the algorithm stops to learn the colour characteristics and the actual values $\mu_{i,k}$ and $\Sigma_{i,k}$ are assigned to the persons as the characteristics μ_i and Σ_i . As soon as the tracking algorithm detects the end of the crossing, the algorithm reassigns the identities to the persons. If $y_{j,k}$ is the measurement of the person that is assigned to track j at time step k , then the probability $p_{i,j,k}$, that the track j at time step k belongs to the person, which has been assigned the identity i before the crossing, is

$$\begin{aligned} p_{i,j,k} = & \frac{1}{(\det(2\pi\Sigma_i))^{1/2}} \cdot \\ & \cdot \exp\left\{-\frac{1}{2}(y_{j,k} - \mu_i)^\top \Sigma_i^{-1} (y_{j,k} - \mu_i)\right\}. \end{aligned} \quad (6)$$

Because the colour measurements of different points in time are independent, the probability $p_{i,j,k_1:k_2}$, that the track j from time step k_1 to time step k_2 originates from person i is

$$p_{i,j,k_1:k_2} = \prod_{k=k_1}^{k_2} p_{i,j,k}. \quad (7)$$

The tracks j , which are calculated by the tracking algorithm, are usually interchanged during a crossing. Thus, let m be the total number of persons associated with the cluster being split and let σ be a permutation of the person IDs $1, \dots, m$. Then, the probability that the IDs have been interchanged during the crossing according to the permutation σ given the measurements from time step k_1 to time step k_2 is

$$\Pi_\sigma = \prod_{l=1}^m p_{l,\sigma(l),k_1:k_2}. \quad (8)$$

The best reassignment $\hat{\sigma}$ of the tracks j to the learnt persons i is the one, that maximises the probability Π_σ . To compute $\hat{\sigma}$ we interpret the negative log-likelihoods $\log(p_{i,j,k_1:k_2})$ as the marginal assignment costs of assigning track j to person i after a crossing. The negative log-likelihoods $\log(\Pi_\sigma)$ then constitute the assignment cost of a complete assignment (permutation) σ . The minimum cost assignment $\hat{\sigma}$ is then calculated using the well known Hungarian algorithm (Munkres, 1957).

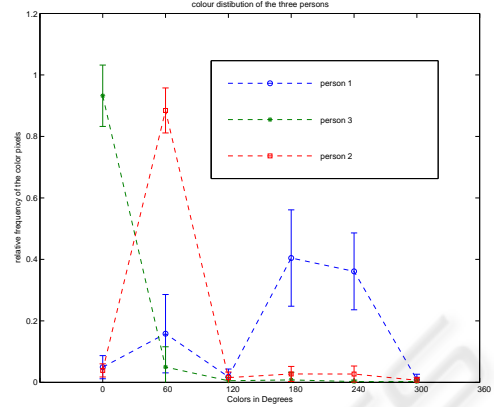


Figure 2: Colour distributions of the three subjects.

3 EXPERIMENTS

We conducted experiments with three persons in a real world scenario in our laboratory. The first person was wearing a blue cardigan and a blue pair of trousers. The second person was wearing a yellow cardigan and blue pair of trousers. The third person was wearing a red shirt and a red pair of trousers.

Figure 2 shows the corresponding colour distributions. The first subject has his maximum in the blue domain, the second has it in the yellow domain and the third in the red one. Thus, the measured colour distributions show a good coincidence with the real colours.

The experiments were accomplished with a B21 robot platform shown in Figure 1. On the top the camera system is mounted. The left camera is the infrared camera and the right camera is the CCD camera. There are two laser range scanners armed back to back at the robot, so that there is a 360 degree field of view.

The number of possible permutations of three objects is $3! = 6$. Thus, we conducted six experiments, for each permutation one experiment. We defined position 1 as the right upper part of the surveillance area, position 2 as the left upper part of the surveillance area and position 3 as the lower middle part. The persons are indexed in the order they appear in the surveillance area. In the six experiments the person indexed 1 occupied position 1 before the crossing, the person indexed 2 occupied position 2 and the person indexed 3 occupied position 3. After the crossing they occupied different positions corresponding to the six possible permutations. The experiment for the permutation $123 \mapsto 123$ is described in detail.

At first the three persons are occupying their start

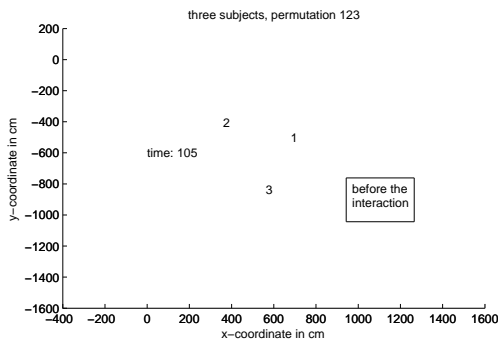


Figure 3: Three subjects before the crossing, permutation 123 \mapsto 123.

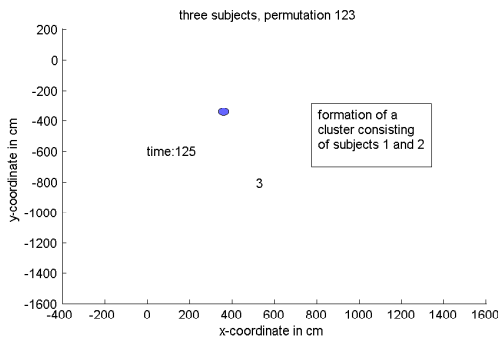


Figure 4: Formation of a cluster consisting of person 1 and 2, permutation 123 \mapsto 123.

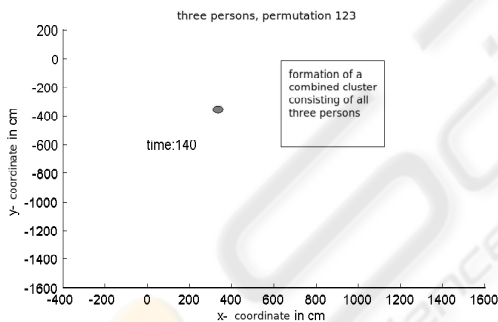


Figure 5: Formation of a cluster consisting of all three persons, permutation 123 \mapsto 123.

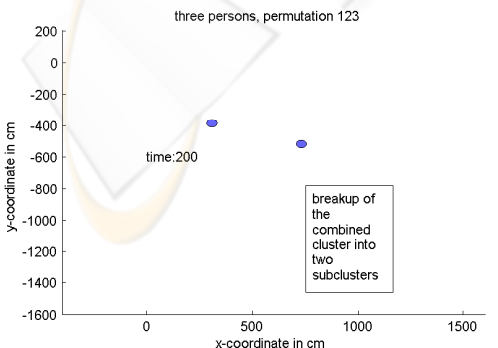


Figure 6: Disaggregation of the combined cluster into two subclusters, permutation 123 \mapsto 123.

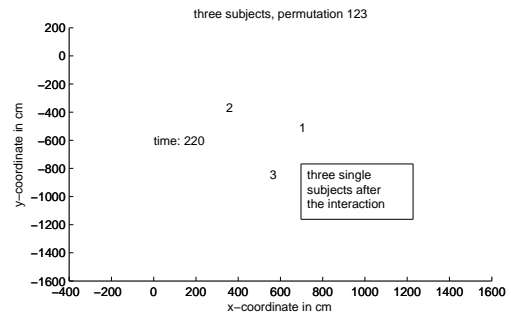


Figure 7: Three persons after the crossing, permutation 123 \mapsto 123.

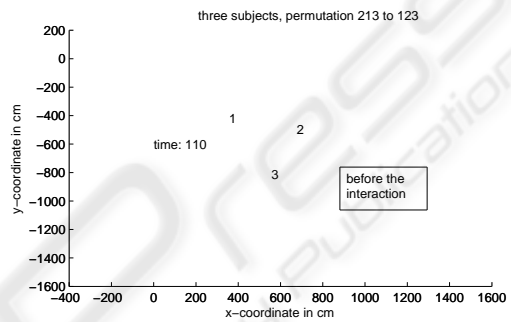


Figure 8: Three persons before the crossing, permutation 213 \mapsto 123.

positions and are indexed in the order of their appearance in the surveillance area; this is illustrated in Figure 3. Within this figure, the numbers corresponding to a person are drawn at the location computed by the tracking algorithm. In Figure 4 the persons 1 and 2 interact and merge into a cluster. The clusters are represented by ellipses within the figure. In the next step person 3 joins the other two and the algorithm merges them into a single cluster as shown in Figure 5. After some time, the group splits up into two subclusters (see Figure 6), and finally the three persons walk on their own again. This situation is illustrated in Figure 7. As can be seen, the algorithm tracks the three persons without loss of track and reassigns the identities correctly after the interaction.

Next, we investigated the question, whether the algorithm still works well, when the starting positions are interchanged. For this purpose we used the permutation 213 \mapsto 123, which means for instance, that the initial position of person number 2 is position number 1. Figures 8 and 9 show the starting and the end positions respectively with the assigned identities. Obviously the identities are in this case also reassigned correctly.

Finally, we examined the case, whether the algorithm can deal with several consecutive permutations. Therefore, we used the two consecutive permutations 123 \mapsto 213 and 213 \mapsto 231. Figure 10 shows the ini-

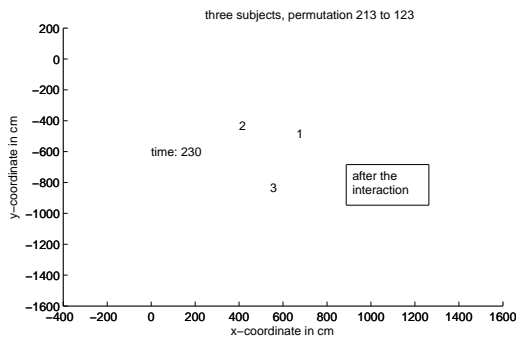


Figure 9: Three persons after the crossing, permutation 213 \mapsto 123.

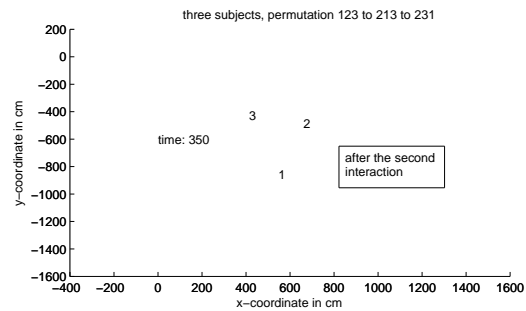


Figure 12: Three persons after the second crossing, permutation 123 \mapsto 213 \mapsto 231.

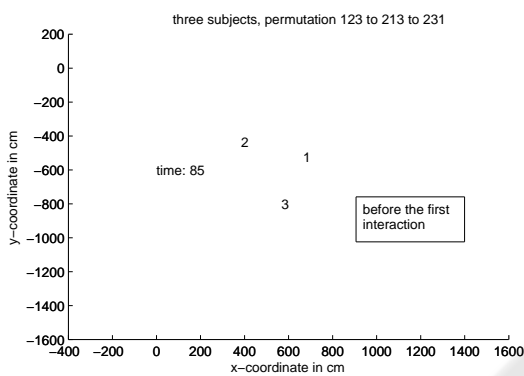


Figure 10: Three persons before the first crossing, permutation 123 \mapsto 213 \mapsto 231.

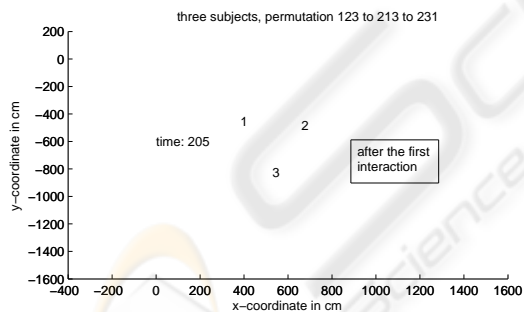


Figure 11: Three persons after the first crossing, permutation 123 \mapsto 213 \mapsto 231.

tial positions. Figure 11 shows the reassigned identities after the first crossing and Figure 12 shows the reassigned identities after the second crossing. It can easily be recognised that the identities are reassigned correctly after each crossing.

4 CONCLUSIONS

In this article we investigated the problem of tracking multiple interacting humans. There are two difficulties, which have to be addressed to solve this challenging problem:

1. The robot should not loose track of a person and
2. the robot should always assign the correct identity to the individual persons, especially after a group of interacting persons split.

We proposed a hybrid approach to solve these two problems, a laser-based tracking approach is applied to keep track of the trajectories of persons, and colour information about the persons' clothes is employed to disambiguate between the persons being tracked.

The proposed tracking algorithm reliably keeps track of several persons, even in very difficult situations, like the crossing of tracks and interactions of persons. Robustness against track loss is achieved by applying a switching approach, where a simple Kalman filter is used as long as tracks are well separated and a variant of the Viterbi algorithm, which tracks individual laser measurements independently, takes over as soon as persons get close to each other. A clustering technique is then employed to assign the measurement tracks to individual person tracks again, when the persons split up again.

However, during the Viterbi phase, the assignment of the persons identities to tracks is lost. For this reason, we employ the camera information to correctly reassign the persons IDs to the individual tracks after the crossing. The current implementation uses a combination of an infrared and a CCD camera for this purpose. The camera system provides the robot with colour information about the tracked persons. The most likely assignment of the identities is then found by using the Hungarian algorithm.

Our experiments show that this approach is able to reliably track multiple interacting persons without

interchanging the individual tracks even in challenging situations. But of course there are still possibilities for future research. The approach will run into problems, if the colour distributions of the peoples' clothes become too similar. This could be remedied by additionally taking size and shape information into account, like in (Schulz, 2006). In rare situations it is also still possible that the tracking algorithm loses track of an individual person, e.g. if a human moves away while it is in the shadow of the other persons during a crossing. This drawback could be overcome by coordinating a team of robots in order to keep full coverage of the scene.

REFERENCES

- Bar-Shalom, Y. and Fortmann, T. (1988). *Tracking and Data Association*. Academic Press.
- Bellotto, N. and Hu, H. (2007). People tracking with a mobile robot: A comparison of Kalman and particle filters. In *Proceedings of the 13th IASTED International Conference on Robotics and Applications*, pages 388–393. International Association of Science and Technology for Development.
- Fod, A., Howard, A., and Mataric, M. J. (2002). Laser-based people tracking. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3024–3029.
- Forney Jr., G. D. (1973). The Viterbi algorithm. *Proceedings of the IEEE*, 61(3):268–278.
- Fortmann, T. E., Bar-Shalom, Y., and Scheffe, M. (1983). Sonar tracking of multiple targets using joint probabilistic data association. *IEEE Journal of Oceanic Engineering*, 8(3).
- Gonzalez, R. C. and Woods, R. E. (1992). *Digital Image Processing*. Addison-Wesley.
- Kräußling, A. (2006a). Tracking extended moving objects with a mobile robot. In *Proceedings of the 3rd IEEE Conference on Intelligent Systems, also to be published in the Post Conference Volume (IEEE IS-06): Intelligent Techniques and Tools for Novel System Architectures, Series: Studies of Computational Intelligence*, Springer.
- Kräußling, A. (2006b). Tracking multiple objects using the Viterbi algorithm. In *Proceedings of the 3rd International Conference on Informatics in Control, Automation and Robotics (ICINCO)*, also to be published in the *Springer book of best papers of ICINCO 2006*, pages 18–25.
- Kräußling, A., Schneider, F. E., and Wildermuth, D. (2004a). Tracking expanded objects using the Viterbi algorithm. In *Proceedings of the 2nd IEEE Conference on Intelligent Systems*.
- Kräußling, A., Schneider, F. E., and Wildermuth, D. (2004b). Tracking of extended crossing objects using the Viterbi algorithm. In *Proceedings of the 1st International Conference on Informatics in Control, Automation and Robotics (ICINCO)*, pages 142–149.
- Kräußling, A., Schneider, F. E., and Wildermuth, D. (2005). A switching algorithm for tracking extended targets. In *Proceedings of the 2nd International Conference on Informatics in Control, Automation and Robotics (ICINCO)*, pages 126–133.
- Kräußling, A., Schneider, F. E., Wildermuth, D., and Sehestedt, S. (2007). *A switching algorithm for tracking extended targets*, volume II, pages 117–128. Springer, informatics in control, automation and robotics edition.
- Kräußling, A. and Schulz, D. (2006). Tracking extended targets — a switching algorithm versus the SJPDAF. In *Proceedings of the 9th IEEE International Conference on Information Fusion, CD-ROM*.
- Munkres, J. (1957). Algorithms for assignment and transportation problems. *Journal of the Society for Industrial and Applied Mathematics*, 5(1).
- Prassler, E., Scholz, J., and Elfes, E. (1999). *Tracking people in a railway station during rush-hour*, volume 1542, pages 162–179. Springer Lecture Notes, computer vision systems edition.
- Reid, D. B. (1979). An algorithm for tracking multiple targets. *IEEE Trans. Automatic Control*, 24:843–854.
- Romera, M. M., Vazquez, M. A. S., and Garcia, J. C. G. (2004). Tracking multiple and dynamic objects with an extended particle filter and an adapted k-means clustering algorithm. In *Proceedings of the 5th IFAC/EURON Symposium on Intelligent Autonomous Vehicles, CD-ROM*.
- Schulz, D. (2006). A probabilistic exemplar approach to combine laser and vision for person tracking. In *Proceedings of Robotics: Science and Systems II*.
- Schulz, D., Burgard, W., Fox, D., and Cremers, A. B. (2001). Tracking multiple moving objects with a mobile robot. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Schulz, D., Fox, D., and Hightower, J. (2003). People tracking with anonymous and ID-sensors using Rao-Blackwellised particle filters. In *Proceedings of the 18th International Joint Conference on Artificial Intelligence (IJCAI)*.
- Schumitch, B., Thrun, S., Bradski, G., and Olukotun, K. (2006). The information-form data association filter. In *Proceedings of the 2005 Conference on Neural Information Processing Systems (NIPS)*. MIT Press.
- Viterbi, A. J. (1967). Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory*, 13(2).
- Zhao, H. and Shibasaki, R. (2005). A novel system for tracking pedestrians using multiple single-row laser-range scanners. *IEEE Transactions on Systems, Man and Cybernetics — Part A: Systems and Humans*, 35(2):283–291.