Bayesian Human Motion Intentionality Prediction in Urban Environments

Gonzalo Ferrer[a],[*], Alberto Sanfeliu[a]

[a]*Institut de Robòtica i Informàtica, UPC-CSIC. Llorens Artigas 4-6, 08028 Barcelona, Spain.*

Abstract

Human motion prediction in indoor and outdoor scenarios is a key issue towards human robot interaction and intelligent robot navigation in general. In the present work, we propose a new human motion intentionality indicator, denominated Bayesian Human Motion Intentionality Prediction (BHMIP), which is a geometric-based long-term predictor. Two variants of the Bayesian approach are proposed, the Sliding Window BHMIP and the Time Decay BHMIP. The main advantages of the proposed methods are: a simple formulation, easily scalable, portability to unknown environments with small learning effort, low computational complexity, and they outperform other state of the art approaches. The system only requires training to obtain the set of destinations, which are salient positions people normally walk to, that configure a scene. A comparison of the BHMIP is done with other well known methods for long-term prediction using the Edinburgh Informatics Forum pedestrian database and the Freiburg People Tracker database.

Keywords: Human Motion Prediction; Pattern Recognition; Crowd Analysis;

1. Introduction

The development of social mobile robots that interact with humans and perform tasks together in everyday environments for guiding or accompany tasks [10][9] , either domestic or public spaces, requires the design of new navigation tools that take into account the human motion intentionality. Similarly, the

---

[*]Corresponding author. Tel: +34 93 4010863 Fax: +34 93 4015750
  Email addresses: `gferrer@iri.upc.edu` (Gonzalo Ferrer), `sanfeliu@iri.upc.edu` (Alberto Sanfeliu)

development of more intelligent tracking and surveillance systems boils down to the development of new tools, for instance, a long-term prediction as proposed in the present paper. We define the term "human intentionality long-term prediction" as a tool to forecast intentions or motivations that drive human behaviors. We specifically use this term whenever a person has the intention to walk towards a destination, then he/she is in the path to go to a destination that can be far away from the actual human spatial position.

There are a number of real world applications where motion intentionality can be applied. For example, human motion scene surveillance, where we can detect abnormal trajectories making use of prediction information, and thus identifying all those erratic trajectories; human tracking, to obtain a robust tracking of human motion; robot navigation in crowded human environments, in order to adjust robot navigation to be aware of human motion; or social interaction, to improve the social interaction between robots and humans. All these issues require prediction information, preferably a long-term prediction.

In general, the analysis of future scenarios is done using short-term prediction, which is a propagation of the current state into a certain time horizon, but it has several limitations, specially appreciable on more complex environments. In the present paper we propose a new long-term human motion intentionality predictor based on geometry criteria, called Bayesian Human Motion Intentionality Predictor (BHMIP). The BHMIP has been validated in outdoor urban environments by using two well-known databases. Furthermore, the proposed method has been implemented in a real robot experiment, using the Tibi robot [16], while developing a robot companion task (in Fig 8 is depicted a robot companion experiment). This concrete robot experiment demonstrated to be largely enhanced after using prediction information while in general, the use of forecasting tools like the BHMIP may improve a countless number of robot applications, specially in urban or outdoor environments. In this work, we assume that we have a prior knowledge of the scene destinations, which have been obtained through an automatic procedure, for example by clustering the locations where people enter or leave a location (we will later discuss this issue).

2. Related Work

Concerning the wide variety of human motion predictors in the literature, two major human motion predictors (HMP) groups can be distinguished: a geometric-based group and a place dependent-based group. For the latter group, we have to learn the prediction model for each one of the environments where the HMP is used. The geometric-based group does not always need to learn the human motion intentionality (HMI) for each specific environment, although training is also required in one way or another.

Bennewitz *et al.* [1], propose a place-dependent method in which they analyze a collection of people's motion behaviors in an indoor environment by a clustering technique that uses the Expectation-Maximization algorithm. Once they learn the classes of motion trajectories, they use these primitive trajectories as patterns for human motion trajectory association, and thus, inferring HMI. One of the main disadvantages of the place-dependent methods, as we will discuss later, is the lack of flexibility on abnormal observations, that is, all those erratic trajectories due to a person stopping or changing its destination. Our algorithm is able to quickly adapt to changes in intentionality and perform successfully.

Vasquez *et al.* [17] cluster different motion patterns by a dissimilarity measure which allows the use of pairwise clustering algorithms in order to group observed trajectories into patterns. Chen *et al.* [3] propose a clustering method and three different prediction strategies based on the quality of the matching.

Using heuristics and geometric criteria, Foka *et al.* [8] propose a geometric-based method for human motion prediction that uses human motion intentionality in terms of goals. Prediction is done by identifying final destinations based on the instantaneous tangent angle in combination with a grid-based probability assignation to all final destinations. It has been used for on-line prediction for robot navigation in dynamic environments. Our approach belongs to this group of HMI predictors, but it outperforms the Foka approach. Another geometrical approach proposed by Ferrer *et al.* [7] predicts future trajectories by minimizing the variance of curvature of forecast paths.

A mixed approach, proposed by Ziebart *et al.* [18] uses both place dependent and geometric criteria. They use a reward function to generate the optimal paths

towards a destination. This method can also be used as a modeling of route preferences as well as for inferring destinations. This method requires intensive place-dependent training, which is an important drawback for the generalization of its use as a predictor.

Dee *et al.* [4] propose a vision-based prediction to infer intentionality, characterized as the combination of obstacles and free space pixels under the field of view of the person. Liu *et al.* [11] propose a method for long term prediction using localization awareness.

3. Problem overview

In this section, we propose a prediction indicator capable of quantifying the human motion intentionality (HMI) implicit on a trajectory with respect to the current position and orientation. This intentionality indicator should capture the probability that a human trajectory reaches a destination point d_m, which is a clear indicator for the inherent intentionality. To achieve this, we define the variable ϕ_{nm}, which is the angle between the current orientation of the target n and the vector to the destination point d_m, a relative measure of the orientation with respect to a destination (see Fig. 1 for clarification).
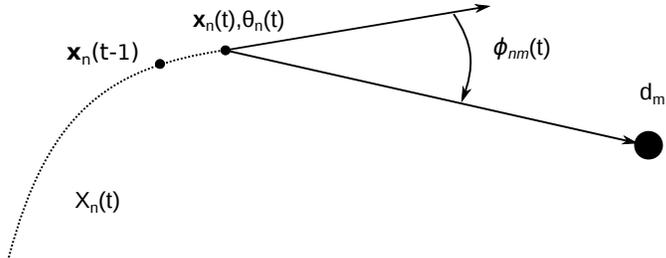

Figure 1: The angle ϕ_{nm} is defined as the angle described by the orientation vector of the target n at time t and the $\mathbf{x_n}(t) \rightarrow d_m$ vector.

In brief, we present a basic formulation necessary to analyze real trajectories. Let

$$X_n(t) = \{\mathbf{x_n}(1), \mathbf{x_n}(2), \ldots, \mathbf{x_n}(t)\} \tag{1}$$

be a set of T positions (people detections) where each point $\mathbf{x_n}(t) = [x(t), y(t)]_n$ is the position at time t of the nth trajectory with respect to the world reference

frame. Additionally, we define the orientation $\theta(t)_n$ as a function of the current position and the previous position, with respect to a global reference frame.

Moreover, we define a set of destination positions $D = \{d_1, d_2, \ldots, d_M\}$, that represents positions in a scene that persons can go to. The existence of destinations is a requirement for most of the existing HMI methods. The set of destinations points must be known in advance, and there are different methods to compute them. In our case we use a Expectation-Maximization method. Other works, like [12] and [18], approximate these destinations as short-term propagations of the current state.

As it can be seen in Fig. 1, $\phi_{nm}(t)$ is the angle defined by the first derivative of the current trajectory and the $\mathbf{x_n}(t) \rightarrow d_m$ vector. By doing this, $\phi_{nm}(t)$ becomes a measure relative to a destination, while $\theta(t)$ is a global measure of the target orientation. This difference will allow us to obtain a good characterization of the human motion intentionality.

Applying the kernel density estimation method [15] to real HMI databases, we have verified that there exists a high similarity of the ϕ angle *pdf* and a Gaussian function, or a Von Mises distribution if we would want to take into account the periodicity of the variable.

The scheme depicted in Fig. 2 corresponds to the graphical model that describes the basis of our algorithm. The relation of the destination d_m at time t is given by the relation of the positions and this structure determines the calculation of the probabilities used in the next section.

Although it is out of the scope of the present paper, a real implementation requires some filtering of the detections in order to eliminate noise and outliers. We will simply take into account that a previous filter exists, in general, but for the understanding of our approach is not a mandatory to consider it.

4. Bayesian Human Motion Intentionality Predictor

The problem of estimating the best destination is reduced to a sequential data classification, where the decision of choosing a destination is taken at each instant of time while the human is walking. Our technique is inspired in a complete Bayesian framework, in order to classify the motion intentionality. We will begin our analysis using an infinite window, that takes into account from
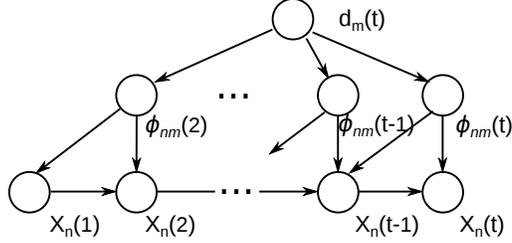
Figure 2: Graphical model of our classifier, called Bayesian Human Motion Intentionality Predictor.

the first observed position of the trajectory to the current one. We will discuss later the importance of the observed positions of a trajectory, depending on its time elapsed since its observation, that is, if more recent observed positions of a trajectory are more significant to the destination classification than older ones.

The method proposed in this paper, denominated Bayesian Human Motion Intentionality Prediction (BHMIP), is a geometric-based method which uses a Bayesian classifier to compute the best prediction to a given destination position, for each position $\mathbf{x_n}(t)$ of the trajectory $X_n = \{\mathbf{x_n}(1), \mathbf{x_n}(2), \ldots, \mathbf{x_n}(t)\}$. We model the probability

$$P(\mathbf{x_n}(t)|\mathbf{x_n}(t-1), d_m) = \mathcal{N}(\phi; 0, \sigma_\phi^2) \tag{2}$$

as a Gaussian function. In Fig. 3 is depicted an example of this probability function to two destinations centered at the position $\mathbf{x_n}(t)$.
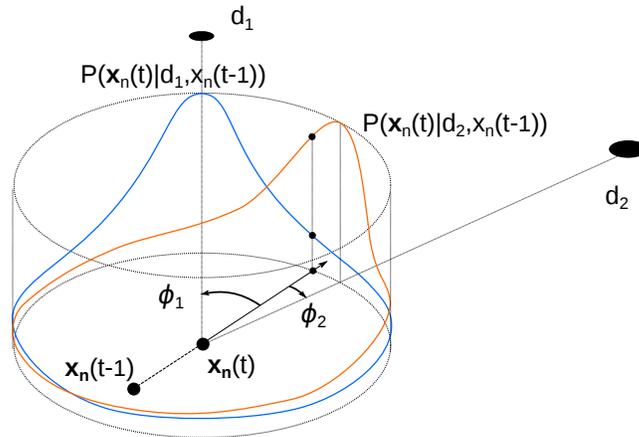

Figure 3: Different probability functions shifted depending on their respective destinations.

4.1. Naive BHMIP

The method is simple, it only requires an initial learning of the positions deployed as destinations.

We formulate the BHMIP in the following manner: if the *posterior* probability to a specific destination d_m is greater than the *posterior* probability to go to another destination d_j, that is

$$d^* = d_m \quad \text{if} \quad P(d_m|X_n(t)) > P(d_j|X_n(t)) \quad \forall m \neq j \tag{3}$$

for all $j = 1, ..., M$ destinations, then d_m will be the best destination d^* describing the current human motion intention.

The joint probability $P(X_n(t)|d_m)$ can be obtained easily as follows using the chain rule:

$$
\begin{aligned}
P(X_n(t)|d_m) &= P(\mathbf{x_n}(1), \mathbf{x_n}(2), \ldots, \mathbf{x_n}(t)|d_m) \\
&= P(\mathbf{x_n}(t)|d_m, \mathbf{x_n}(t{-}1), \ldots, \mathbf{x_n}(1)) \cdot \\
&\quad P(\mathbf{x_n}(t{-}1)|d_m, \mathbf{x_n}(t{-}2), \ldots, \mathbf{x_n}(1)) \cdot \\
&\quad \vdots \\
&\quad P(\mathbf{x_n}(1)|d_m)
\end{aligned}
\tag{4}
$$

For each trajectory $X_n(t)$ we have only considered dependence of position in between consecutive positions as the ϕ_{nm} variable is function of two points at instants t and $t-1$. Consequently, the Eq. (4) can be rewritten more compactly as:

$$P(X_n(t)|d_m) = P(\mathbf{x_n}(1)|d_m) \prod_{\tau=2}^{t} P(\mathbf{x_n}(\tau)|d_m, \mathbf{x_n}(\tau{-}1)) \tag{5}$$

Using the Bayes theorem we can compute the *posterior* probability of the destination d_m, given the current and previous positions of the trajectory $X_n(t)$:

$$P(d_m|X_n(t)) = \frac{P(X_n(t)|d_m)P(d_m)}{P(X_n(t))} \tag{6}$$

where $P(d_m)$ is the prior probability to reach the destination d_m. By replacing Eq. (5) into the Eq. (6), we obtain a compact formulation of the BHMIP.

4.2. Sliding Window BHMIP

Another unknown arises: what if the intentionality of a person changes in the middle of a walk? As we are evaluating partial trajectories $X_m(t)$ and not the full observed trajectory X_m, we don't know the true final destination. Inspired by the Sliding Window method [5] for sequential classification, we define a length or time interval of past positions $\mathbf{x_n}(t)$, discarding the previous person's positions, as a measure of the flexibility to intentionality changes. Thus, the Eq. 5, is rewritten as:

$$P(X_n(t)|d_m) = \prod_{\tau=t-w}^{t} P(\mathbf{x_n}(\tau)|d_m) \qquad (7)$$

This solution, although its impact is minor on typical trajectories, has proved to work fine specially in abnormal trajectories, where an unexpected behavior of a person is observed.

4.3. Time Decay BHMIP

We have incorporated an additional feature in the BHMIP to weight in a different way the contributions of the past positions to our proposed method. Intuitively, we believe that the more recent positions observed are more determinant than older ones.

Based on the Sliding Window BHMIP approach, we propose a variation that makes use of a non-constant window that degrades as a function of time:

$$P(X_n(t)|d_m) = \prod_{\tau=t-w}^{t} P(\mathbf{x_n}(\tau)|d_m)^{g(t-\tau)} \qquad (8)$$

The domain of the function $g(t-\tau)$ is the elapsed time since the observation $\mathbf{x_n}(\tau)$ took place and it is positive by definition. The output of the g function is in the interval $[1, 0]$, being 1 for the current position, the higher weighted probability, and 0 for the most distant in time. For this reason, we propose to use the exponential function,

$$g(t) = e^{-\frac{t}{\sigma_w}} \qquad (9)$$

that satisfies all the above mentioned requirements. The performance of the Time Decay approach (BHMIP-D) will be discussed in the following section.

8

5. Experiments

We have used two different databases to validate the results presented in this paper. The Edinburgh Informatics Forum Pedestrian database [14] is a set of human trajectories at the Informatics Forum, the main building of the School of Informatics at the University of Edinburgh (see Fig. 4 left). The data covers several months of observations which have resulted in about 1000 observed trajectories each working day, and at present has more than 92.000 observed trajectories. The sample rate is 9 positions per second. A set of eight final destinations are identified.

The Freiburg People Tracker [13] ground truth consists of a set of 162 person tracks at the city center (see Fig. 4 right), resulting in more than 10000 manually labeled positions in order to validate their multi hypothesis people tracker. The sample rate is approximately 50 positions per second.

A total of 1280 trajectories from the Edinburgh database were studied, each trajectory consisting of an average of 100 positions. The trajectories are divided into two sets, one set for training and the other for testing. As the Freiburg database has less person tracks available, 30% of the trajectories were used for training purposes and 70% of the trajectories were used for testing, each trajectory consisting of an average of 500 positions.

5.1. BHMIP parameter learning

One of the main advantages of the geometrical, or place-independent methods for prediction is that a set of primitive trajectories is not required. However, they are not excluded from training: the set of destinations requires to be defined somehow. One simple solution might be a manual setting of the trajectories into salient positions such as stairs, elevators, doors, etc.. The solution adopted in this work, altought, has been a clusterization of positions into destinations.

A general view of the Edinburgh Informatics Forum is depicted in Fig. 4 left. The scene is taken from a camera fixed at 23m above the floor. For the purpose of the present paper, 8 final destinations d_m have been estimated using the Expectation-Maximization algorithm [2] of the starting and ending positions of each trajectory. Some of these final destinations coincide with entry/exit points. The same procedure has been carried out for the Freiburg

database (Fig. 4 right), where all the initial and final points of each trajectory are depicted and the clusterization of the destinations is shown, as well as the covariances on those positions.
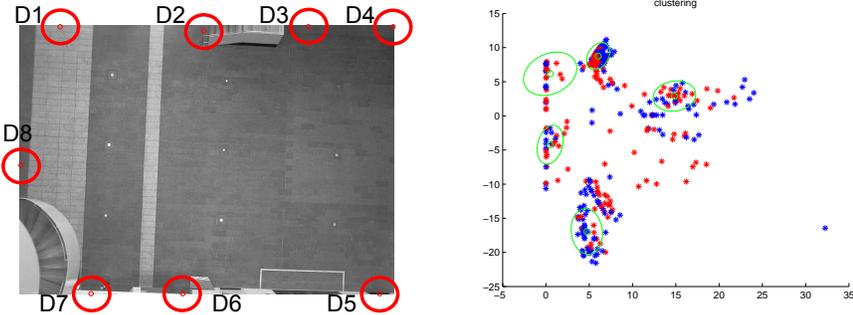

Figure 4: On the left, the map of the Edinburgh Informatics Forum. Eight different destinations are drawn as red circles. On the right, the Freiburg database environment, where are shown the clusters of the destinations. In both cases a EM algorithm has been used to obtain the destinations.

As discussed in Sec. 4, the two BHMIP methods proposed make use of past information. The width w and σ_w of the window are parameters of the system, and the performance of the predictor depends on these values.
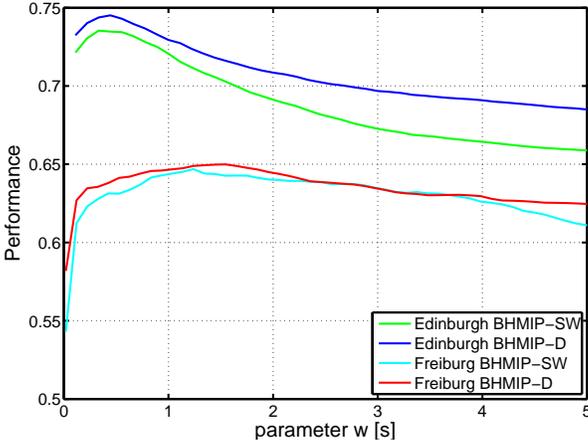

Figure 5: Performance versus the parameter w and σ_w (in seconds), using the Edinburgh and the Freiburg database.

We have evaluated the accuracy of prediction by doing the following: we provide an observation of the trajectory corresponding to the interval $[0, t]$ from the beginning $t = 0$ to the current time t. We incrementally augment the current

10

time t until it reaches the end of the trajectory $t = T$. For every observation provided of a given trajectory, we calculate the most likely destination according to each method and then check if the prediction is or not correct. Fig. 5 draws the performance of each method depending on the value w and σ_w (the window width), which is measured in seconds.

For the Edinburgh database results, there is a clear optimum for the Sliding Window BHMIP at $w = 0.444s$ and for the Time Decay BHMIP at $\sigma_w = 0.5s$. The results for the Freiburg database are not clear, the best performance is for $w = 1.5s$, however the performance presents an almost flat behavior for $w > 2s$. For this reason, we have chosen a time width w and σ_w roughly around $0.5s$ for both databases, which corresponds to the reaction time of the system to make good predictions.

5.2. BHMIP testing

In this subsection, we analyze the proposed methods, the Sliding Window Bayesian Human Motion Intentionality Predictor (Sec. 4.2) and the Time Decay BHMIP (Sec. 4.3). The w and σ_w parameters are the ones obtained previously. In addition, a comparison with other methods is shown. We will stress on the advantages of our approach with respect to the state of the art methods.

A comparison with other approaches is difficult since no other methods use destination points as they have been defined in this paper. The geometrical-based method we have implemented for comparison is the approach proposed by Foka *et al.* [8]. Originally, this method treats all the cells on a grid map, corresponding to obstacles, as possible destinations and computes the best cell as the destination a person aims to. One limitation of the method appears when the destination is not an obstacle. For this reason, we have made the following modifications in the method: the interesting cells of the map are translated into the same destination points defined in our approach, and by doing this, we are not constraining the position of the destinations to grid cells. The other implemented method is the approach proposed by Bennewitz *et al.* [1] as one of the most representative prediction methods of the place-dependent group. A set of primitives is obtained for both databases applying the Expectation-Maximization algorithm. With the purpose of comparing with the BHMIP, we have assumed a correct prediction if the corresponding primitive associated

11

to the current trajectory finishes at the desired destination. We have used a different sampling rate for the two databases, however the four methods use the same sampling rate when testing each database.
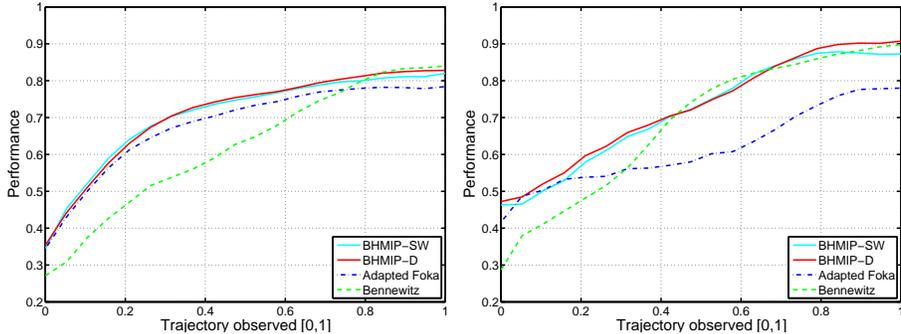

Figure 6: On the left, performance of the algorithm against other state of the art methods, using the Edinburgh database. On the right, the performance on the Freiburg database. The horizontal axis is the percentage of time to complete the trajectories, where T is the total time to complete the trajectories.

A comparison of the results, for the four methods, is shown in Fig. 6. The figure on the left shows the performance in the Edinburgh database. The geometrical methods (BHMIP-SW, BHMIP-D and Adapted Foka) present a better performance at early stages of the evaluation than the place-dependent (Bennewitz), although at long observations it outperforms the geometrical, scoring almost 84% of success rate.

In Fig. 6 right is depicted the performance using the Freiburg database. As can be seen, the Bennewitz method could not obtain such a good performance after short observations (t small), but it improves the accuracy after considering longer observations, same as before. However, the performance of the place-dependent varied for each scenario. At the Freiburg city center (right), where 5 destinations and 20 primitive trajectories are clearly defined, the method obtains almost 90% of accuracy when the complete trajectory is considered and a mean of 68.47% (See Table 1). On the other hand, the performance obtained for the complete trajectory at the Edinburgh scenario (which consists of 8 destinations and more than 50 primitive trajectories) is 84% but its average performance is under 62.07%. The reason for such a performance at the Edinburgh Forum is due to the fact that the trajectories are in an open area with less structural

constraints, therefore more unexpected trajectories might occur. Nevertheless a fully observed trajectory can't be considered as a prediction problem but a classification problem and thus, full trajectory observations are not so relevant.

Considering the previous reasoning, the BHMIP is a method well suited for real-time prediction. It can provide rapidly an accurate prediction of people trajectories. Also, the geometrical-based methods are more indicated in open scenarios, where prediction is done in areas consisting of many destination points.
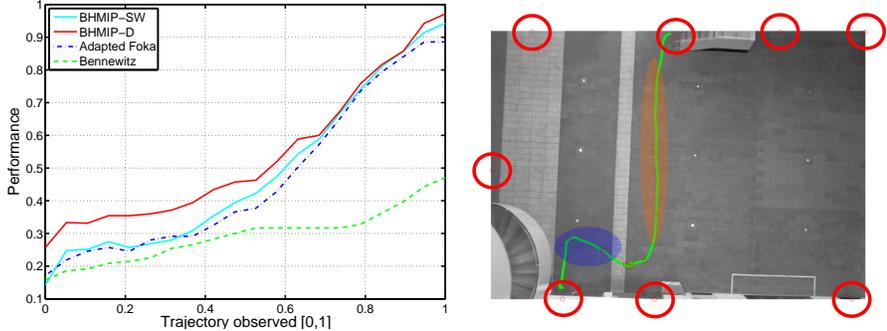

Figure 7: On the left, performance of the algorithm using a selected set of abnormal trajectories in the Edinburgh database. On the right, an example of an abnormal trajectory, where a change in the destinations occurs in the middle of the path.

Observing the behavior of the human trajectories, we have seen that there are people that walk from an origin to a destination without stopping, maintaining a constant speed. Other people stop one or several times and change their velocity. There are also people that modify their trajectory several times and then stop in some places before reaching their destination. All these abnormal cases are specially interesting since they represent a challenge for any prediction method. That is the reason why we have selected a set of specially challenging trajectories to classify. An example of an abnormal trajectory can be seen in Fig. 7 right. All these trajectories have in common that they are abnormal and predominantly appear changes in intentionality.

The results are shown in Fig. 7 left. All the methods fail to determine initially the correct destination. It is not surprising since we have set the dataset to contain trajectories with changes in destinations and other abnormalities. Nonetheless, the geometrical approaches behave in general better at

13

Table 1: Overall performance of the prediction methods using different databases.

	Edinburgh	Freiburg	"Abnormal"
BHMIP-SW	70.81%	71.59%	49.00%
BHMIP-D	71.10%	72.74%	54.29%
Adapted Foka	68.22%	61.93%	47.00%
Bennewitz	62.07%	68.47%	29.29%

late stages of the intentionality prediction. They are capable of quickly recover from changes in destinations while the place-dependent approach presents more problems to adapt.

Table 1 shows a summary of the overall performances of the evaluated methods, each percentage is the mean of the performance of every experiment carried out before. Furthermore, we can observe that both BHMIP methods present a better performance in all situations, being the BHMIP-D slightly better and more adaptable to abnormal trajectories than the Sliding Window BHMIP.

5.3. Experiments in a robotic task

The main motivation to develop the BHMIP is, of course, a direct application to robotic tasks. More concretely, the presented work was developed as a key tool for a robot companion experiment [6]. The main task of the robot in this experiment is accompanying a person. The robot companion predicts the human target's destination and uses this information to anticipate where the human will be and positioning the robot in the best place. In addition, the companion task is carried out successfully under the presence of obstacles and moving persons (see Fig. 8). We have implemented a people detector based on laser information and a particle filter for the person tracking system. Under the framework of the robot companion task, we observed that the consideration of human motion prediction enhances the performance of the method: a set of 45 real-life experiments with different volunteers have been carried out and each participant filled out a questionnaire. The measurement was a simple rating on a Likert scale between 1 to 7. For the evaluation score, repeated ANOVA measurements were conducted. Three different scores were examined: "Robot's

Intelligence", "Human-Like Motion" and "Level of confidence". For each score examined, the average of the results showed an increase of 38.89%, 52.78% and 43.9% respectively, compared to not using the prediction information, while accompanying the volunteer.

Figure 8: Example of a real robot experiment: Left image corresponds to the robot camera. On center image Dabo is accompanying a person to her desired goal while navigating in a crowded environment. The right image corresponds to the robot GUI, built on ROS architecture. Green cylinders correspond to persons' detections and the orange cylinder corresponds to the target to be accompanied.

Real trajectories are subject to disturbances, as a part of the real world settings. The BHMIP handled correctly most of the uncertainties and quickly adapted to changes in intentionality.

We have presented just a simple example where the human prediction enhances a certain utility, but the use of forecasting tools like the BHMIP may improve a countless number of robot applications, specially in urban or outdoor environments.

6. Conclusions and Future Work

In the present paper we have presented a novel and accurate human motion intentionality indicator, denominated Bayesian Human Motion Intentionality Prediction (BHMIP), which is a geometric-based long-term predictor. We have proposed two variants of the algorithm: the BHMIP-SW (sliding window) and the BHMIP-D (time decay). The performance of both is successful, but the BHMIP-D behaves better under challenging trajectories, as demonstrated in the abnormal database.

We have presented a simple formulation, a low computational complexity and it outperforms other state of the art approaches. The system requires

minimal training, since it is place independent and only the set of the scene destinations must be obtained, although it might be obtained using geometrical or manual methods as well.

A comparison of the BHMIP is done with other well known methods for long-time prediction using the Edinburgh Informatics Forum pedestrian and the Freiburg People Tracker databases. Additionally, experiments in a real scenario are carried out including a set of volunteers walking in the presence of a mobile two-wheeled robot, to validate the overall performance of the BHMIP.

In the future work we will obtain a propagation model of the predicted human motion in order to complete the forecast information that can be extracted from persons' observations.

7. Acknowledgements

[1] M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun. Learning motion patterns of people for compliant robot motion. *The International Journal of Robotics Research*, 24(1):31, 2005.

[2] C.M. Bishop et al. *Pattern recognition and machine learning*, volume 4. Springer New York, 2006.

[3] Z. Chen, L. Wang, and N.H.C. Yung. Adaptive Human Motion Analysis and Prediction. *Pattern Recognition*, 44(12):2902–2914, 2011.

[4] H. Dee and D. Hogg. Detecting inexplicable behaviour. *British Machine Vision Conference*, 477:486, 2004.

[5] T.G. Dietterich. Machine learning for sequential data: A review. In *Structural, syntactic, and statistical pattern recognition*, pages 15–30. Springer, 2002.

[6] G. Ferrer, A. Garrell, M. Villamizar, I. Huerta, and A. Sanfeliu. Robot interactive learning through human assistance. In *Multimodal Interaction in Image and Video Applications*, pages 185–203. Springer, 2013.

[7] G. Ferrer and A. Sanfeliu. Comparative analysis of human motion trajectory prediction using minimum variance curvature. In *Proc. of the 6th Int. Conf. on HRI*, pages 135–136, Lausanne, Switzerland, 2011. ACM.

[8] AF Foka and PE Trahanias. Probabilistic Autonomous Robot Navigation in Dynamic Environments with Human Motion Prediction. *International Journal of Social Robotics*, 2(1):79–94, 2010.

[9] Nourbakhsh I. Dautenhahn K. Fong, T. A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42:143–166, 2003.

[10] A. Garrell and Sanfeliu A. Cooperative Social Robots to Accompany Groups of People. *The International Journal of Robotics Research*, 31(13):1675–1701, 2012.

[11] X Liu and H Karimi. Location awareness through trajectory prediction. *Computers, Environment and Urban Systems*, 30(6):741–756, November 2006.

[12] M. Luber, G. Diego Tipaldi, and K.O. Arras. Place-dependent people tracking. *The International Journal of Robotics Research*, 30(3):280–293, 2011.

[13] M. Luber, G.D. Tipaldi, and K.O. Arras. Place-dependent people tracking. *International Journal of Robotics Research*, 30(3), March 2011.

[14] B. Majecka. *Statistical Models of Pedestrian Behaviour in the Forum*. Msc dissertation, School of Informatics, University of Edinburgh, 2009.

[15] B.W. Silverman. *Density estimation for statistics and data analysis*, volume 26. Chapman & Hall/CRC, 1986.

[16] E. Trulls, A. Corominas Murtra, J. Pérez-Ibarz, G. Ferrer, D. Vasquez, J.M. Mirats-Tur, and A. Sanfeliu. Autonomous navigation for mobile service robots in urban pedestrian environments. *Journal of Field Robotics*, 28(3):329–354, 2011.

[17] D. Vasquez, T. Fraichard, and C. Laugier. Markov Models: An Incremental Tool for Learning and Predicting Human and Vehicle. *The International Journal of Robotics Research*, 28(11):1486–1506, 2009.

[18] B.D. Ziebart, N. Ratliff, G. Gallagher, C. Mertz, K. Peterson, J.A. Bagnell, M. Hebert, A.K. Dey, and S. Srinivasa. Planning-based prediction for pedestrians. In *Intelligent Robots and Systems. IEEE/RSJ International Conference on*, pages 3931–3936. IEEE, 2009.