

“© 2011 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

Learning Navigational Maps by Observing Human Motion Patterns

Simon T. O’Callaghan, Surya P. N. Singh, Alen Alempijevic and Fabio T. Ramos

Abstract—The study of human motion patterns is crucial for the development of social robots that share the environment with humans. This paper presents a methodology to allow a robot to navigate in a complex environment by observing pedestrian positional traces. A continuous probabilistic function is determined using Gaussian process learning and used to infer the direction a robot should take in different parts of the environment. The approach learns and filters noise in the data producing a smooth underlying function that yields more natural movements. Our method combines prior conventional planning strategies with most probable trajectories followed by people in a principled statistical manner, and adapts itself online as more observations become available. The use of learning methods are automatic and require minimal tuning as compared to potential function or spline function regression. This approach is demonstrated and tested against a cluttered office and open forum environments using laser and vision sensing modalities. It yields interactive paths that are similar to the expected people behaviour (e.g., across a street at a pedestrian crossing) without any *a priori* knowledge of the environment or explicit programming.

I. INTRODUCTION

As robots start to share the environment with humans, the study of human motion patterns becomes increasingly important. In the robotics community, the social interaction between humans and robots have been object of numerous studies [1], [2], [3]. It presents significant challenges and has proved vexing for numerous reasons – not the least that such interactions are often qualitative and difficult to measure.

However, human locomotion is incredibly informative. Considering the points people transit not only informs how to navigate, but also improves interaction. Consider, for example, navigating around spilled coffee. Trying to sense this is incredibly difficult, yet following the cues of people walking around it is relatively much easier. That is, understanding and generalising peoples’ behaviour is paramount to defining better decision making rules and improve safety of humans and robots. This, however, can be difficult as the variability of human motion patterns is immense even in structured environments such as offices and corridors. Also challenging is estimating and modelling the uncertainty in these motions from noisy sensors, which are generally subject to occlusions and misdetections.

In this paper we study how robots can learn a potential function for an environment by observing the motion pattern of people and consider its use for navigation. Given a set

of traces from pedestrians walking in a populated area, we seek to learn a function that maps any arbitrary location to a direction of travel – we call this the navigational map. This navigational map reflects people’s behaviour in a statistical manner and helps the robot to understand the environment. For example, people (almost) always avoid obstacles when walking. By learning how humans walk in a closed area, a robot can avoid obstacles without relying solely on perception. Equally, assumed social boundaries such as personal workspaces or the area between a television and a viewer often mean that the shortest path is not always the optimum.

Traditionally, integrating such behaviour into the path planner is challenging as the definitions of preferred spaces are mostly qualitative, hence making metrics (and sensing) difficult [4]. However, learning the motion pattern can lead to trajectories that incorporate a very abstract level of reasoning on the environment without explicitly understanding the underlying principles influencing them.

The process of observing and tracking pedestrian movements has received considerable attention in the robotics community and includes approaches based on a variety of sensing modalities including monocular vision [5], [6], laser [4], and inertial measurement [7]. Human navigation models, such as those proposed by [8], [9] and [10], are primarily focused on steering and obstacle avoidance. The suggested continuous curvature models complement this work by providing a more informed initial navigation map. While the process of extracting and tracking people is non-trivial (particularly when multiple people are moving), this work is concerned with using these trajectories for informing robot navigation, preferably in as natural a manner as it is presumed that this begets more interaction.

We rely on a machine learning technique to develop a model of how people traverse the environment. Using traditional path planning techniques as a prior, the resulting navigational map emulates human-like motion trajectories in a sound, statistical manner. In order to learn this map from pedestrians’ traces, we explore the benefits of Bayesian learning and, in particular, a popular regression technique known as a Gaussian process [11].

A key aspect of this paper is the manner in which the problem is formulated. By modelling human motions as a deviation from a prior path plan, it becomes possible to infer preferred areas (a subsequently areas of repulsion) all within the Bayesian statistical framework.

This paper is organised as follows. Section II introduces the problem formulation and describes the algorithms developed and illustrate them with a pedestrian example. Section

S. T. O’Callaghan, S. P. N. Singh and F. T. Ramos are with the Australian Centre for Field Robotics, University of Sydney, J04, NSW 2006 Australia. A. Alempijevic is with the Mechatronics & Intelligent Systems Group, University of Technology Sydney, NSW 2007. This work is supported in part by the Australian Research Council (ARC) Centre of Excellence program. Correspondence should be addressed to s.ocallaghan@cas.edu.au

III details the theory behind the approach and Section IV presents experimental results using real datasets. Finally, Section V provides conclusions.

II. PROBLEM FORMULATION

Intuitively, the proposed approach seeks to produce a navigational map, a function which maps a location to a normalized velocity in a manner that incorporates the motion patterns of people. This would enable the trajectory taken by a person travelling to a destination to be estimated at any location allowing the robot to navigate in a human-like manner. Essentially, this is achieved by learning a continuous function that describes how people deviate from some prior belief on the path to a destination. This deviation function, Ψ , is combined with the prior navigational map, H , to produce a navigational map that captures the general trends of how humans traverse the environment.

A. The Prior Map

The prior map, H , can potentially be derived from a wide variety of navigational techniques described in the literature, e.g. potential fields - for an overview, see Ch. 4 of [12] or sampling-based algorithms [13]. It can be viewed as a method of incorporating expert knowledge about the environment, such as perhaps the location of obstacles, into the posterior map. In this paper, a naïve prior is adopted to illustrate the ability of the proposed technique to capture social navigation trends in the map despite very little prior understanding of the environment. Essentially, the prior direction at any location, \mathbf{x}_* , on the map is the arctangent of the line joining that location to the destination, \mathbf{x}_D . More formally:

$$H(\mathbf{x}_*) = \arctan(\nabla(\mathbf{x}_D - \mathbf{x}_*)). \quad (1)$$

B. Learning the Motion Patterns of People

Generally speaking, our technique can be divided into two main segments. Initially, the outputted paths of a people tracking algorithm such as [14], [15] are used in an offline learning phase to estimate Ψ for a given destination. Subsequently, an online algorithm uses the generated function to traverse the environment in a way that reflects the tendencies of people.

For the remainder of this section, data obtained by simulating the path traces of pedestrians using a zebra-crossing, Fig. 1, is employed as an illustrative example to help detail the steps of our approach and highlight some of its advantages.

1) *Offline Phase: Learning the Deviation Function:* The procedure for the offline learning phase is detailed in Alg. 1.

Each navigational map is conditioned on a specific destination. Consequently, only observed path traces leading to the goal are relevant to the problem. A KD-tree ball search can be used to identify path traces ending at, or passing through, the objective. Fig. 1(a) shows the traces that would be chosen, represented in blue, when generating a navigational map for the indicated destination. Fig. 1(b)

Algorithm 1 Learn Deviation Function, $\Psi(\theta)$

Input: H :- Prior Map, \mathbf{x}_D :- Destination, Path Traces

Output: θ :- Parameters of Ψ

- 1: Identify Path Traces leading to \mathbf{x}_D
 - 2: Determine Tangent of identified Path Traces at each observation location, X .
 - 3: Calculate target vector, \mathbf{y} :
 $\Psi(X) = \text{Tangent}(X) - H(X)$
 - 4: Train Gaussian process using training data (X, \mathbf{y}) to learn the parameters, θ , of the underlying Deviation Function.
-

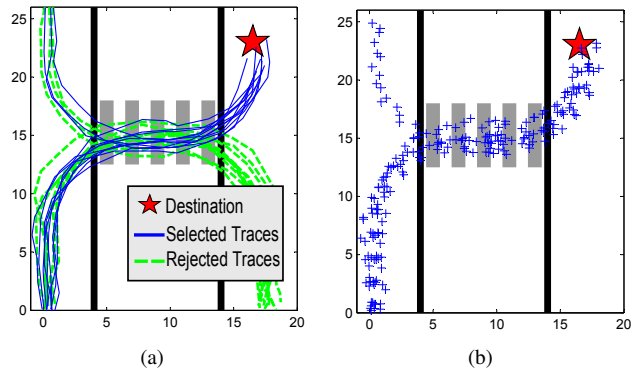


Fig. 1. Plan view of zebra-crossing with simulated pedestrians using it to get from one sidewalk to another. (a) Pedestrian's path traces. (b) The observed locations of the pedestrians used to generate the traces

reveals the observed locations of the pedestrians, X , that were used to create the selected traces.

The training data's target vector, \mathbf{y} , used for learning the underlying deviation function is obtained from the selected raw data by initially approximating the trajectory of the pedestrians at each observed location and subtracting it from some pre-determined prior.

An important advantage to our approach is that the form of the prior map can be entirely user-specified. It can range from trivial (no map) to very naïve solutions such as in this example where the prior's trajectory is always directed towards the destination or it could also take on a more complex form. Fig. 2(a) compares the prior and the path trace trajectory at every observed location of the pedestrians. From this, training data for Ψ , such as that shown in Fig. 2(b), can be determined using Step 3 of Alg. 1.

A Gaussian process, GP, is then trained to estimate a non-parametric probabilistic model of Ψ . The predictive mean and variance functions of our example's deviation model were sampled and are shown in Fig. 3. The predictive mean represents the model's estimate of how much the human motion pattern deviates from the expected direction of navigation (the prior). The associated variance quantifies the model's confidence in each prediction.

2) *Online Phase: Navigating to the Destination:* Once a model for Ψ has been trained, it becomes possible to predict the motion pattern of people at any location, \mathbf{x}_* , in the environment and to obtain an associated variance on each

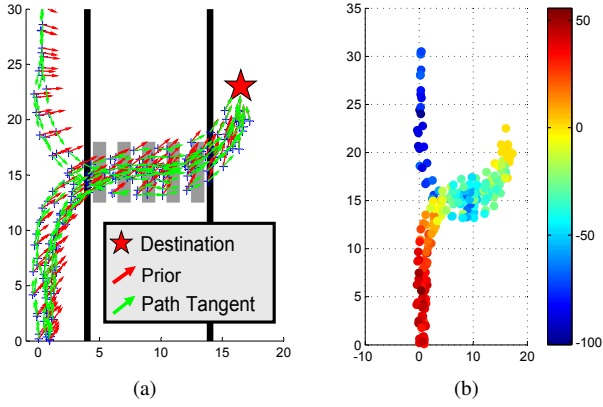


Fig. 2. (a) Comparison between prior and trace trajectory. (b) Training data passed to GP: location (metres) versus deviation angle (degrees)




Fig. 3. (a) Grid samples from the predictive mean, shown in degrees, and (b) variance functions of Ψ .

prediction. Alg. 2 describes how such a model can be used to navigate a robot towards the destination in a human-like manner.

Algorithm 2 Navigate to Destination, \mathbf{x}_D

Input: H , \mathbf{x}_D , θ , \mathbf{x}_*

- 1: **while** $\mathbf{x}_* \neq \mathbf{x}_D$ **do**
 - 2: $\text{Direction}(\mathbf{x}_*) = H(\mathbf{x}_*) + \Psi(\mathbf{x}_*, \theta)$
 - 3: Move forward along $\text{Direction}(\mathbf{x}_*)$
 - 4: **end while**
-

The robot queries the probabilistic model with its current location and receives an estimate of the deviation angle. This is then added to the original prior to produce a prediction of the direction that a human would take in that position. The robot moves along the trajectory given by the navigational map until the destination is reached.

Figure 4 compares the prior navigational map with the posterior map produced after incorporating the information gathered from observing how people move through the environment. The social context of the zebra-crossing is reflected in the posterior map and it has the effect of channelling the flow from left to right through the designated crossing zone.

Figure 5 demonstrates the routes that robots with various starting locations would take to reach the destination using Alg. 2 on the trained navigational map. The resulting paths




Fig. 4. Comparison between prior (a) and our proposed navigational map after being trained using observations of human motion patterns in the region (b).

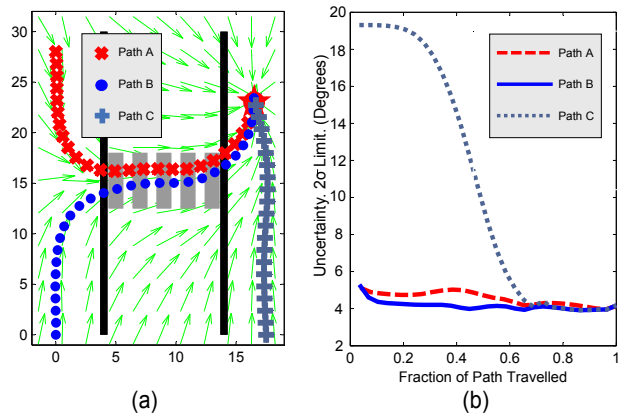


Fig. 5. (a) Routes taken to destination using the posterior navigational map for different starting locations. (b) 2σ boundary of trajectory prediction along each route measured in degrees (right).

are in keeping with expected social context – they navigate as we might expect a person would take in the same situation.

C. Incorporating Predictive Uncertainty into the Platform’s Actions

An important output of our probabilistic model is the predictive variance. This measure of uncertainty manifests itself in the navigational map as a force that returns the trajectories to those of the prior map in areas where there is insufficient observations to make a reliable prediction on the motion pattern of people. Examining the region below the zebra-crossing in Fig. 5 where no pedestrians are observed, the model is correctly uncertain about the value of Ψ here and thus the trajectories of the navigational map return to the prior.

The variance could also be used to influence the level of caution that the robot exercises as it travels towards the destination. Areas of high uncertainty generally correspond to regions where few or no (pedestrian) observations. Regions of high variance could be used as an indication that the robot should adjust its speed appropriately or increase the computational power dedicated to perceiving its surroundings. Figure

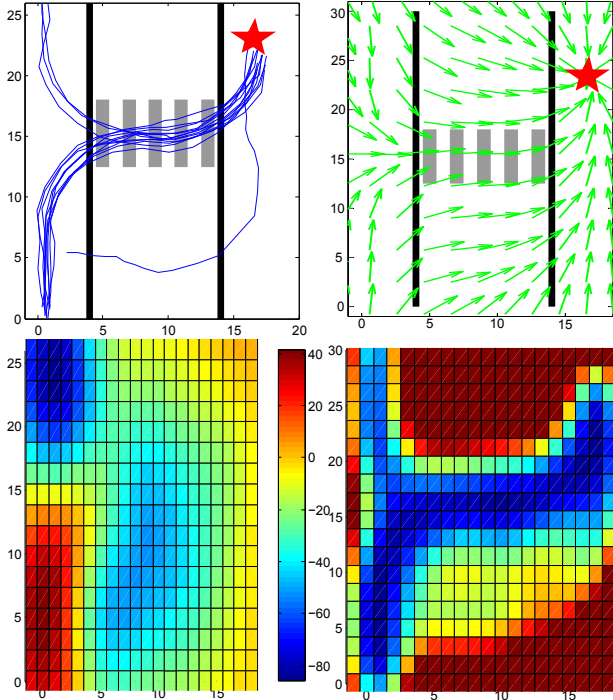


Fig. 6. Series of images showing the changes to the navigational map when a new observation is made of a pedestrian crossing the road without using the zebra-crossing. Clockwise from top left: Observed path traces. Posterior navigational map. Predictive variance function for Ψ . Predictive mean function for Ψ .

5(b) compares the 2σ boundaries of the three sample routes. As expected, Path C initially has the highest uncertainty as the platform moves through a region where no pedestrians were observed.

A variation to the proposed approach would be to use the variance as a repulsive force in the navigational map. This would have the effect of causing the robot to move towards regions where it has observed pedestrians and is more certain about the value of Ψ .

D. Integrating New Observations into the Model

After the training phase, incorporating additional observations into the model can be done online. The theory behind this procedure is detailed in Section III. Fig. 6 illustrates the effects on the outputs of our method when the observation of a jaywalker's path trace is integrated into the map.

Interestingly, the trajectories of the navigational map below the zebra-crossing adapt to reflect the behaviour of a person walking in that region (i.e., to get to the other side of the road as directly as possible) compared to the trajectories in Fig. 5. Importantly, however, the variance in this region, although lower than areas where no observations were made, is still higher than the section over the zebra-crossing and the route preferred by the majority of pedestrians.

III. LEARNING MECHANISM

The proposed method is based upon the Gaussian process' ability to predict $p(\Psi|\mathbf{x})$, where Ψ is the deviation taken by people from the prior trajectory at \mathbf{x} , a physical location

within the environment. The GP is used to fit a likelihood function to the training data $\{\mathbf{x}_i, y_i\}_{i=1 \rightarrow N}$ where N is the number of training points and y_i , the training output or target data, corresponds to the angle between the prior and the tangent to the path taken by the person at a specific location. The resulting continuous function can then be used to interpolate between data points allowing predictions to be made on the human motion patterns over the entire region using the well understood Bayesian statistical framework.

A. Gaussian Process Fundamentals

Several important characteristics of the Gaussian process makes it well suited to our intended application.

- Continuous models: GPs do not require a discretised representation of an environment. Similarly, they are able to predict the motion pattern of people at arbitrary locations.
- Predictive Variance: Unlike most other regression models, GPs provide uncertainty estimates for predictions at any set of locations. This uncertainty takes into account the local density of observations and their noise level.
- Marginal Likelihood: The parameters of the model are marginalised during the training. This eliminates the need for any hand-tuning.
- Flexibility: GPs are non-parametric regression models and can thus approximate an extremely wide range of motion models.

The Gaussian process itself can be viewed as a distribution over an infinite number of possible functions thus performing inference takes place directly in the space of functions. By assuming that all deviation angles, indexed by their corresponding location in the environment, are jointly Gaussian, we obtain

$$\Psi(\mathbf{x}_*) = \mathcal{N}(\mu, \sigma^2), \quad (2)$$

where

$$\mu = k^\top(\mathbf{x}_*, X)^\top [K(X, X) + \sigma_n^2 I]^{-1} \mathbf{y}, \quad (3)$$

$$\sigma^2 = k(\mathbf{x}_*, \mathbf{x}_*) - k(\mathbf{x}_*, X) [K(X, X) + \sigma_n^2 I]^{-1} k(X, \mathbf{x}_*). \quad (4)$$

Here, \mathbf{x}_* refers to a query or test location, X the training inputs, σ_n^2 the variance of the global noise and K is the covariance matrix. The elements of the covariance matrix $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ are defined depending on a covariance function k parameterised by hyperparameters. In this application, the hyperparameters' and σ_n 's optimal values for the datasets are derived by maximising the log marginal likelihood using a simulated annealing algorithm followed by quasi-Newton gradient ascent.

By setting $\mathbb{E}[\Psi(\mathbf{x}_*)] = 0$ during training and inference, the predicted deviation will naturally revert back to the prior in areas lacking sufficient observations. An extensive explanation and derivation of the Gaussian process can be found in [11].

B. Training the Covariance Function

A number of commonly used covariance functions were tested on a variety of datasets. Given the nature of the process being modelled, it is unsurprising that stationary covariance functions performed best, most notably the Matérn class [16]:

$$k_{Matern}(\mathbf{x}, \mathbf{x}_*) = \sigma_f \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}r}{l}\right)^\nu K_\nu\left(\frac{\sqrt{2\nu}r}{l}\right) \quad (5)$$

where r is the Euclidean distance between inputs \mathbf{x} and \mathbf{x}_* while the trainable hyperparameters l and σ_f represent the lengthscale and amplitude, respectively. K_ν is a modified Bessel function and ν is shape parameter that regulates the smoothness of the interpolation. As special cases, $\nu \rightarrow \infty$ equates to infinitely differentiable squared exponential covariance function and $\nu \rightarrow \frac{1}{2}$ corresponds to the sharper exponential covariance function. Cross validation led to the conclusion that $\nu = \frac{5}{2}$ produced the best results.

The hyperparameters loosely correspond to the style of motion adopted by the pedestrians in the environment. Consequently, they appear to generalise well to similar scenarios such as a different flow in an office block or another area of a park provided an identical method for deriving the prior is used.

C. Storing and Updating the Inverse Covariance Matrix

A traditional drawback to the Gaussian process is its long computational time which makes it unsuitable for many online operations where realtime performance speeds are required. The GP's runtime complexity of $\mathcal{O}(N^3)$ is primarily as a result of the requirement to invert the covariance matrix, K , in equations 3 and 4.

Fortunately, the elements of K are independent of query points. Consequently, its inverse can be precomputed and stored for use in Alg. 2 thus eliminating the computational bottleneck during online use. Additional observation locations can be subsumed into K^{-1} using the matrix inversion lemma and submatrix inversion as detailed in [17].

IV. RESULTS

The proposed approach was tested on a number of different datasets. Here we present the results from two such evaluations, namely the UTS RobotAssist Project [18] and the Edinburgh Informatics Forum Pedestrian Database [6].

A. UTS RobotAssist Dataset

The experimental setup involved two stationary SICK LMS-200 laser rangefinders, shown in blue in 7(a), positioned to cover a combined area of approximately 50 m² of a typical office environment. The path traces of passing employees were extracted using a Hidden Markov Model people tracking algorithm [19]. Trajectories leading to a chosen destination, indicated by the red star, were selected as inputs to our algorithm, Fig. 7(c).

Once again, the naïve potential function described in II-A was used as the prior map, Fig. 7(b). The target vector, Fig. 7(d), was obtained using Line 3 of Alg. 1, enabling Ψ to be

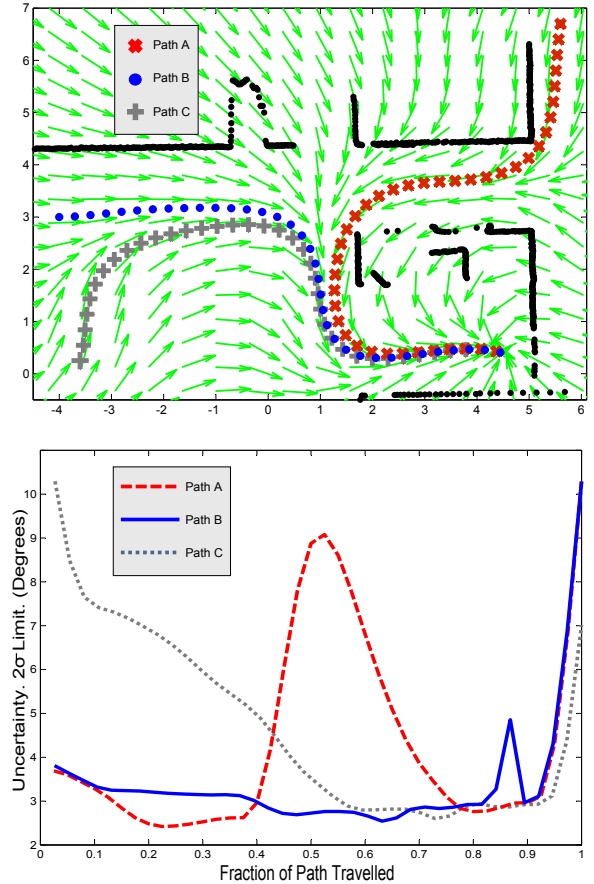


Fig. 8. Sparsely sampled quiver plot of the posterior navigational map with examples of resulting routes to the destination taken from three different starting locations (top). 2σ boundary (degrees) vs. fraction of route travelled for each sample route (bottom).

learnt through the GP framework. Samples from the resulting predictive mean and 2σ functions are shown in Fig. 7(e) & (f).

The posterior navigational map, represented as a quiver plot in Fig. 8, merges the observations of path traces into the prior map to generate more socially-informed routes to the destination. Three typical starting locations were chosen as inputs to Alg. 2 and their resulting paths were plotted.

An important advantage to incorporating the motions of people into the navigational map is the ability to indirectly sense obstacles that may be occluded or even undetectable to the robot's sensors. The large table on which one of the rangefinder's is placed in Fig. 7(a) close to the center of the image is below the sensor's plane and hence does not appear in the laser returns plotted in Fig. 7(b) & (c). Observing that people avoid this area leads to trajectories like Path C which also circumvent the obstacle without ever requiring to reason about the boundaries or type of the obstruction. Although in this case the obstacle is quite tangible, our approach would just as easily handle more abstract obstructions such as a restricted area due to a slippery surface or paths that may impinge on the work areas of other people.

A key feature of the Gaussian process is its capacity to

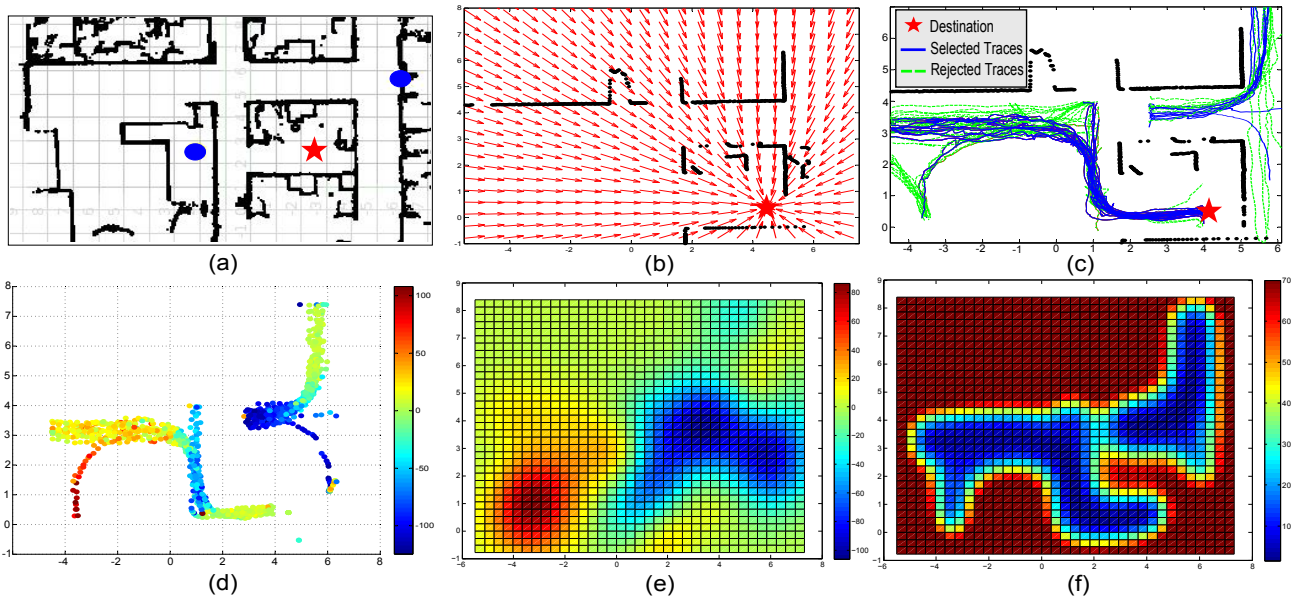


Fig. 7. Top row: (a) Occupancy grid of office area - positions of laser rangefinders marked in blue. (b) Prior navigational map - determined using Equation 1. (c) Path traces of pedestrians detected by lasers. All axes are in measured in metres. Bottom row: (d) Training data passed to GP - target vector is measured in degrees. (e) Samples from the predictive mean function (middle) and (f) 2σ boundary function of Ψ .

infer the most likely value of the Ψ at any point in the map based on the trends of observed motion patterns in the region. The benefit of this is clearly illustrated in the region occluded to both rangefinders - around $\mathbf{x}_* = [2, 3.5]m$ in Fig. 7(c). Despite a lack of path traces in this area, the GP correctly predicts the most likely trajectory of pedestrians in the vicinity. Crucially, an associated variance is also produced for each prediction and can be used as an indicator of the caution that should be exercised in each particular region. The second graph in Fig. 8 illustrates the 2σ boundary for each of the sample paths. The area occluded from both sensors is made apparent by the large peak in uncertainty approximately halfway along the curve of Path A.

Similarly, the variance associated with the planned trajectory of Path C is initially quite high due to a low density of observations in the region as only 2 pedestrians were observed here. However as the trajectory brings the robot into the more populated corridors, the uncertainty on the predicted value of the deviation from the prior falls to within a 2σ boundary of approximately 3 degrees.

B. Edinburgh Informatics Forum Pedestrian Database

The database consists of a large set of detected pedestrians walking through the Informatics Forum at the University of Edinburgh, Fig. 9(a), and has been gathering data on a daily basis since August 2009. The method of data acquisition involves a fixed camera suspended 23 metres above the floor that records the location of tracked targets on a frame-by-frame basis. The tracked trajectories of a few hundred detected targets are shown in Fig. 9(b).

The posterior navigational map shown in Fig. 9(c) was generated from data gathered on September 10th 2009. The

destination that this particular map was conditioned on is indicated by the red star in the lower portion of Fig. 9 (b) & (c). For illustrative purposes, four sample trajectories have been plotted using different starting locations. One area where people deviate strongly from the prior prediction is at the bottom of the image where a staircase blocks a direct line of travel. As Path A demonstrates, this deviation is also learnt by our algorithm and incorporated into the navigational map. Another interesting pattern captured by the proposed technique is the elongated S-shaped trajectory which was seen to be adopted by many pedestrians when moving from the top left of the image to the destination, as indicated by Path B.

The forum itself is a wide open area which differs considerably from the more confined spaces of the previous dataset. This fact is reflected in the smoother trajectories of the pedestrians and consequently in the longer lengthscales learnt by the covariance function. An important outcome resulting from this smoother motion pattern is that the GP is more confident at inferring the value of the deviation function in regions far from observed traces compared similar distances in the office dataset. Examining the uncertainty contours in Fig. 9(c), it can be seen that the variance function rises much more gradually to its maximum value in areas lacking in observations relative to Fig. 7(f).

C. Cross-Validation of Results

A fundamental requirement of the proposed algorithm is its ability to accurately model human-like trajectories in different environments. The capacity of our technique to emulate such trajectories was tested using cross-validation on both datasets. Table IV-C summarises the results of our

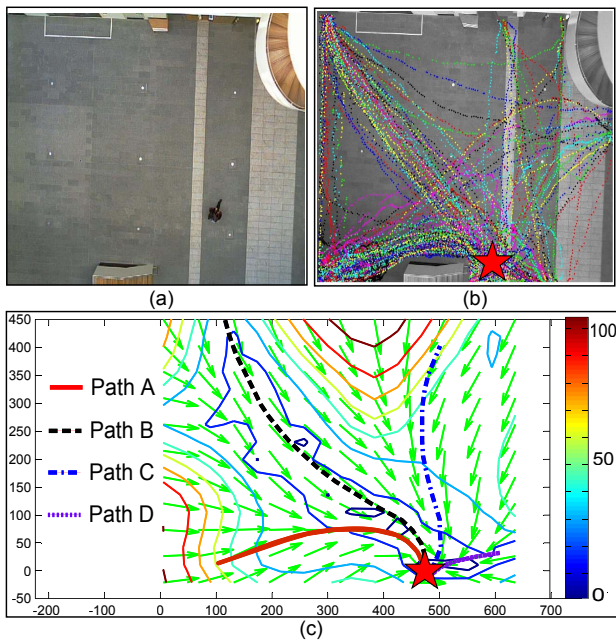


Fig. 9. (a) Plan view of forum. (b) Subset of tracked trajectories. (c) Posterior navigation map (quiver plot) with uncertainty contours and their associated colour bar superimposed. 4 sample trajectories from typical starting points are also included. Uncertainty is measured in degrees (2σ boundary). The axes of the graph correspond to pixel numbers in the image. Each pixel has a width of 24.7 mm.

TABLE I
CROSS-VALIDATION RESULTS

	σ Limit	2σ Limit
UTS RobotAssist	75.4	95.3
Edinburgh Informatics Forum	61.3	84.8

analysis. For the UTS RobotAssist dataset, the observed trajectories were divided equally into a training set and a validation set. Path traces from five different days were used in testing the algorithms model learnt for the Edinburgh Informatics Forum dataset.

The figures in the σ and 2σ columns represent the percentage of test points that fell within one and two standard deviations of the prediction, respectively. Indicatively, the models learnt for both datasets explain the validation cases quite well, particularly the RobotAssist dataset. This is most likely due to the fact that its narrow walkways confined the range of motion hence reducing outliers. It was found that the majority of outliers in the Edinburgh dataset leading to the slightly lower than expected percentage of traces falling within the 2σ limit was as a result of pedestrians who tended to wander aimlessly around the open area before choosing a destination.

V. CONCLUSIONS

In this paper we have introduced a powerful tool for incorporating the motion patterns of people into the trajectories of robots. We have demonstrated its ability to address several important challenges currently facing robotic

navigation such as perceiving obstacles that are traditionally difficult to observe and motion planning in a manner that conforms to social assumptions. The proposed approach is illustrated and tested on a variety of datasets which demonstrate the algorithm’s capacity to encapsulate social context in navigation.

Handling this problem within the framework of the Gaussian process avoids the necessity to discretise the world or the resulting trajectories. Crucially, an associated predictive variance exists along each trajectory which can be used to dictate the the level of confidence the robot have in the model for each region of the map. An additional benefit of adopting a Bayesian approach is the ability to learn the sensor noise levels and characteristics of the underlying function in tandem through the optimisation of the marginal likelihood function without the requirement of hand-tuning the model’s parameters.

We believe that this work is an initial step towards integrating conventional decision making algorithms and path planning with the complex decision making processes in humans and their social behaviours.

ACKNOWLEDGMENTS

We thank the RobotAssist project and The University of Technology Sydney for providing the office dataset used in this paper. This work has been supported by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government. S. Singh additionally acknowledges the support of the Rio Tinto Centre for Mine Automation.

REFERENCES

- [1] M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun, “Learning motion patterns of people for compliant robot motion,” *International Journal of Robotics Research*, vol. 24, pp. 31–48, 2005.
- [2] W. Burgard, A. Cremers, D. Fox, D. Hähnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun, “Experiences with an interactive museum tour-guide robot,” *Artificial Intelligence*, vol. 114, no. 1-2, pp. 3–55, 2000.
- [3] B. Jensen, “Motion tracking for human-robot interaction,” Ph.D. dissertation, École Polytechnique Fédérale de Lausanne, 2004.
- [4] J. Müller, C. Stachniss, K. Arras, and W. Burgard, “Socially inspired motion planning for mobile robots in populated environments,” in *Proc. of International Conference on Cognitive Systems*, 2008.
- [5] O. Masoud and N. P. Papanikolopoulos, “A novel method for tracking and counting pedestrians in real-time using a single camera,” *IEEE Transactions on Vehicular Technology*, vol. 50, no. 5, pp. 1267–1278, Sep. 2001.
- [6] B. Majecka, “Statistical models of pedestrian behaviour in the Forum,” Master’s thesis, School of Informatics, University of Edinburgh, 2009. [Online]. Available: <http://homepages.inf.ed.ac.uk/rbf/FORUMTRACKING/>
- [7] E. Foxlin, “Pedestrian tracking with shoe-mounted inertial sensors,” *IEEE Computer Graphics and Applications*, vol. 25, pp. 38–46, 2005.
- [8] B. Fajen and W. Warren, “Behavioral dynamics of steering, obstacle avoidance, and route selection,” *Journal of Experimental Psychology: Human Perception and Performance*, vol. 29, no. 2, pp. 343–361, 2003.
- [9] G. Arechavaleta, J.-P. Laumond, H. Hicheur, and A. Berthoz, “An optimality principle governing human walking,” *Robotics, IEEE Transactions on*, vol. 24, no. 1, pp. 5–14, Feb. 2008.
- [10] W. Huang, B. Fajen, J. Fink, and W. Warren, “Visual navigation and obstacle avoidance using a steering potential function,” *Robotics and Autonomous Systems*, vol. 54, no. 4, pp. 288–299, 2006.
- [11] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning*. MIT Press, 2006.

- [12] H. Choset, K. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. Kavraki, and S. Thrun, *Principles of Robot Motion*. MIT Press, 2005.
- [13] D. Hsu, R. Kindel, J.-C. Latombe, and S. Rock, "Randomized Kinodynamic Motion Planning with Moving Obstacles," *The International Journal of Robotics Research*, vol. 21, no. 3, pp. 233–255, 2002.
- [14] C. R. Wren, A. Azarbayejani, T. Darrell, and A. P. Pentland, "Real-time tracking of the human body," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 19, pp. 780–785, 1997.
- [15] J. K. Aggarwal and Q. Cai, "Human motion analysis: a review," *Comput. Vis. Image Underst.*, vol. 73, no. 3, pp. 428–440, 1999.
- [16] M. Stein, *Interpolation of Spatial Data: Some Theory for Kriging*. New York: Springer, 1999.
- [17] A. Singh, F. Ramos, H. D. Whyte, and W. J. Kaiser, "Modeling and decision making in spatio-temporal processes for environmental surveillance." in *ICRA*. IEEE, 2010, pp. 5490–5497.
- [18] A. Alempijevic, S. R. Kodagoda, and G. Dissanayake, "Mutual information based data association," in *Proceedings of the Fifth International Conference on Intelligent Sensors, Sensor Networks and Information Processing*, Dec. 2009.
- [19] S. Sehestedt, S. Kodagoda, A. Alempijevic, and G. Dissanayake, "Efficient learning of motion patterns for robots," in *Australasian Conference on Robotics and Automation*, 2009.