

An Interest-Based Framework for Modelling of Human Gross Motion in a Dynamic Environment

David K. Wood, Steven Scheduling

Australian Centre for Field Robotics, University of Sydney

Sydney, NSW 2006, Australia

{D.Wood, Scheduling}@acfr.usyd.edu.au

Abstract

This paper investigates a framework for modelling human gross motion in a dynamic environment. The concept of 'interest' as a model parameter is introduced, and a structure for an interest tracking system is developed. A number of motion primitives are derived, based on observation of the available datasets. Results are presented using data collected from the Fish-Bird auto-kinetic artwork.

1. Introduction

THIS paper proposes an approach to the question 'How do people move through a space containing a selection of both obstacles and objects of interest?'. This is a question which has been posed in both the engineering and psychology communities many times over a number of years. Fajen and Warren [Fajen and Warren 2003], [Warren and Fajen. 2004] designed a model which can account for human motion in environments with clearly defined obstacles and a single goal. Other work [Vasquez et al., 2005; Bruce and Gordon, 2004; Panangadan et al., 2004] has investigated the motion of people through real environments, but these have not investigated the issue of dynamic points of interest. The long term problem which is under investigation is that of a general model of human motion which can account for multiple behaviour types and multiple potential objects of interest. In this paper a method of modelling interest is presented as a first step towards this goal.

It should be clearly stated here that interest for the scope of this paper is not necessarily the same 'interest' that a person is exhibiting towards an object. The definition of interest used herein is 'an observable parameter of human gross motion which has a significant effect on that motion.' It is a construct to assist in modelling which is derived from, but separate to, human interest modelling as practiced by psychologists, where interest is defined as 'the power of attracting or holding

one's attention' [WordNet, 2007]. The primitives from which the interest model is constructed represent a selection of types of motion which humans have been witnessed to exhibit in the available data.

This interest-based architecture should allow for the construction of a model which is reasonably accurate in any space where the points of interest are known. Work such as [Bruce and Gordon, 2004], where the motion of people through sensor blind spots is modelled, is useful in a tracking field, however it requires large amounts of training to be fully effective. The objective of this work is to at least provide a solid starting point for learning algorithms, thus significantly shortening training times, or ideally to be an accurate, off-the-shelf model of human motion through a dynamic environment.

There are 3 main fields of application for a model as discussed above. Firstly this model has applications in short-time-horizon prediction. By improving the quality of predictions of human motion, then a number of additional applications become available. For example, if an intelligent billboard were to detect a drop-off in the level of interest being exhibited towards it by people walking past, it could switch to a different ad which was known to be more attention-getting to increase its attractiveness to the audience.

Secondly, it can be used to improve the performance of tracking solutions. This new model will allow search areas to be reduced when scanning for a target correlation when predicting. It will also allow for improved target disambiguation by reducing the size of these search areas, decreasing overlaps. This will be useful in a number of applications. One of these is security, where tracks which display interest towards a secure area could identify that target for closer scrutiny. It could also be used in a manner similar to that in [Panangadan et al., 2004], where a model is learned of people's behaviour patterns while in a space, and people who do not conform to that model are identified.

Finally there are potential applications where the model is used to close a control loop. One such example is in crowd control, where relative interest levels of

dynamic environment elements, such as electronic billboards, could be varied to influence the motion of people through a given space. An example of this application is in museums with dynamic exhibits, where the behaviour of exhibits could be modified in order to induce an even spread of people around the museum. This would minimise crowding at individual exhibits, maximising patron comfort and increasing the number of patrons which could be accommodated.

This paper discusses the theoretical framework for an interest model, and the project from which the data has been collected. Results of the interest tracker on the collected logs are shown, and the potential for this interest model to be of use for the motion modelling problem is discussed.

2. Test-Bed Project

2.1. The Project

The Fish-bird project is an artwork, produced over 3 years by Mari Velonaki and her collaborators at the Australian Centre for Field Robotics (ACFR) [Velonaki et al., 2005]. The work involves two self-mobile wheelchairs (Fig 1) which inhabit a space into which the audience enter and interact with the chairs. All objects in the space are tracked using a combination of ceiling-mounted cameras and a pair of SICK laser scanners. The behaviour of the chairs is determined by a script running on a control PC, which sends commands to the chairs over a pair of Bluetooth wireless links. The chairs are also equipped to communicate with the audience using a small thermal printer designed to print sales receipts. The script instructs the chairs to 'speak' at appropriate times, with statements selected from a database of potential phrases. These printouts are in a cursive script, to encourage the audience to accept the chairs as entities in their own right rather than mechanical constructs.



Figure 1: The Fish-Bird artwork in Wood St galleries, Pittsburgh, USA.

One of the key drivers behind the design of the system was to make all the technology completely concealed, in order to facilitate the audience personifying the chairs. As such, the interactions between the audience and the chairs should be completely interest-driven, as any potentially techno-phobic reactions are minimised.

2.2. The Logged Data

The data being logged for this work consists of a structure which completely describes the observed state of the exhibition space at a point in time. It includes the filtered position estimates of the chairs and any people in the space. It also contains information on the behaviour which the chairs are executing, and information about progress through the driving script. One of the unique aspects of this data is the fact that the people are completely unaware of the observation. This is important for the artwork as well as the research, as the reactions of the people to the chairs will differ if they believe they are being watched. From a research point of view, poses something of an issue, as it severely limits the amount of data which can be gathered. It would be considered an invasion of privacy to record video feeds without the knowledge of the participants, and so the only data which is logged is that of the final tracker outputs. This makes the data completely anonymous, as no information beyond position is recorded from the tracker. It is also not possible to check aspects such as possible points of interest against the video feeds. This information does however pose a unique dataset, as it is extremely rare for such large quantities of information to be collected on an unaware audience.

The logs being used have been collected over two separate exhibitions of the work, a 3-week exhibition at the Art-Space gallery in Sydney, Australia in late 2005, and a 3-month exhibition at the Wood St. Galleries, Pittsburgh, USA in early 2006. For the times when the museum was open to the public, the world state was sampled at a rate of 5 Hz, resulting in over 3 Gb of data from the Wood St exhibition.

2.3. Special features for Interest modelling

One of the greatest advantages in using the wheelchairs as test subjects is that it is possible to vary their intrinsic parameters to change the audience's perception of them. For example, a person is more likely to exhibit interest towards that chair. This interest is also likely to be influenced by the movement of the chairs, as it has been demonstrated that interest grabbing potential is linked to the unexpectedness of the observed event [Horstmann, 2005]. This has a parallel to work done in perception, where measures such as the Shannon entropy which indicate the unexpectedness of a particular distribution are used to identify easily trackable or otherwise informative features in an environment [Suresh Kumar et al., 2004]. This characteristic will influence the movement of the audience, making it a facet of the psychological interest

models that this interest model should replicate.

3. Interest Model Definition

3.1. Background

The definition of interest used in the introduction was derived from two different directions, each of which has a bearing on the problem as posed. Firstly, this parameter interest is of use in a modelling sense, as it enables some assumptions to be made about the final structure of the model which shall result in a significant reduction in complexity. Secondly, the interest parameter was derived from existing psychology and engineering research on the motion of people through a space [Fajen and Warren 2003; Warren and Fajen. 2004; Vasquez et al., 2005; Panangadan et al., 2004].

Engineering Background

The majority of the engineering support for this particular definition of interest is on the grounds of computational efficiency and transfer learning across different spaces. By introducing the intermediate state of interest into the model, it is possible to make the gross motion of the person conditionally independent from the potentially dynamic elements of the environment inducing that interest. This structure will allow the easy transfer of a heuristic or learned model of human motion between environments, without a need to re-examine the interest-caused motion in the new environment. As a result, this modelling architecture will be significantly more portable between different environments, and as such more easily applicable to a wider range of problems. This is a significant advance over some of the existing work such as that of Vasquez et. al. [Vasquez et al., 2005] where the structure of the model is learned specifically for the application space.

The conditional independence between the motion of the person and the behaviour of the environment becomes more significant when the application of learning techniques such as Dynamic Bayesian Networks (DBNs) [Theeuwes et al., 2004] is considered. By introducing conditional independence into the definition of interest, this has a direct translation into the structure of any probabilistic modelling structure which could be used for higher-level modelling tasks.

Psychology Background

In general, the work done in Psychology on this field has focussed on visual experiments using eye-tracking to determine the “Attentional Capture” properties of various stimuli. Of particular interest in this context is work on the influence of distractions on a simple search task by Theeuwes [Theeuwes et al., 2004]. This work shows that attention-grabbing behaviour in the environment is significantly attention grabbing, even when a person is being asked to perform a specific task. This work shows the strength of attention capture on a person’s behaviour, where attentional capture will strongly influence a person’s behaviour, regardless of prior instructions or

intentions. This is a useful result for this modelling work, as it suggests that the interest-based interactions will be detectable, regardless of a person’s prior intentions.

3.2. Types of Behaviour

The use of primitives is an established technique for the modelling of human motion [Jenkins and Matarić, 2002]. As such, a number of possible gross motion primitives have been investigated, along with how they show interest in an object in the environment. To date, only two of the following primitives have been implemented, however they have all been observed to occur in the Fish-Bird dataset. One aspect of the interest modelling which is of note in this situation is establishing a reduced set of primitives which any given object will be able to induce in a person. For example, behaviours such as the “following” primitive will not be enacted towards stationary targets. The investigation of different types of points of interest, and appropriate motion primitives for each one is an area of potential further work.

Direct Approach Primitive

This is the first primitive to be examined, and its existence is based on the results of work by Fajen and Warren [Fajen and Warren 2003], where it is shown that a person’s trajectory will rapidly take them into a heading towards the object of interest. This primitive can easily be used to generate an interest estimate, by examining how the heading of a person compares to bearings from that person to potential targets. An example of this is shown in figure 2.




Figure 2: The Approach Primitive. The person (circle) is walking towards target 1 (cross). This indicates that interest in 1 is greater than interest in either of 2 or 3.

Close Proximity Primitive

This primitive is of use in an environment where potential points of interest are static, resulting in a behaviour pattern in which a person approaches the target and stops to inspect it. By using range as an interest indicator, the interest of a person in an object can be estimated, even when that person is stationary, and thus the estimate of their heading is poor. Although not found in the literature, this primitive is based on the assumption that the object of interest of the person is the object to which they have the nearest physical proximity. This assumption has not yet

been verified, however based on empirical observations of the data it appears to be present. Figure 3 shows behaviour consistent with this primitive.




Figure 3: The Close Proximity Primitive. The person (circle) is closest to target 1 (cross). This indicates that interest in 1 is greater than interest in either of 2 or 3.

Following Primitive

This gross motion primitive is best examined by investigating the rate of change of range to a potential point of interest. This primitive is more complex, as low or zero rates of change of range will indicate if a person is performing either of two interest displaying actions. Firstly, matching the motion of a dynamic target is an extremely strong interest indicator. This gives a range rate of approximately zero. This form of motion is shown in figure 4.

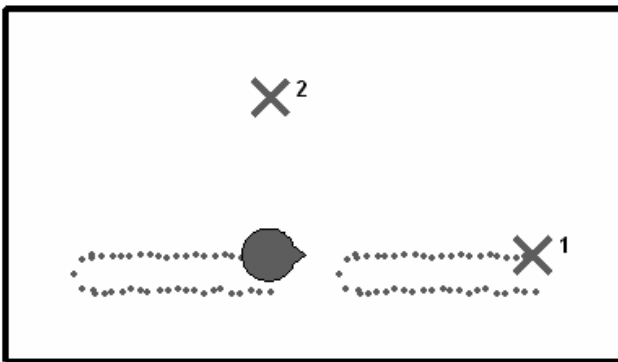


Figure 4: The Follow Primitive. The person (circle) is matching the motion of target 1 (cross). Interest in target 1 is greater than interest in target 2 despite the shorter range to target 2.

There is also the case where a person is circling an object of interest without actually approaching it. Again, this is a strong indicator of interest, with a rate change of approximately zero towards the object of interest. This motion type is shown in figure 5.




Figure 5: The Follow Primitive, the person (circle) is circling target 1 (cross). Interest in target 1 is greater than interest in target 2.

There is also a third motion type which can be detected with this method, a large negative range rate indicates that the person is approaching the point of interest (Fig 2). This re-enforces the bearing-based primitive, which should also show strong interest in the target in this situation. This primitive has been observed in the data collected, however it has not yet been robustly implemented.

These 3 primitives can account for the bulk of behaviours observed in the Fish-Bird dataset. When other objects of interest are investigated, it may be necessary to generate additional primitives for the motion types observed.

3.3. Interest Estimation through Heading

As the position observations are outputs of a tracking filter, then the observations contain both position and variance for all the objects being tracked. In order to use this information most effectively, any transformation to calculate interest parameters should propagate the variance as far forward as possible. Given the non-linearity of the function used to calculate heading/bearing, then an appropriate technique is the use of the unscented transform [Julier et al., 2000]. This technique is useful due to its ability to rapidly propagate a Gaussian estimate through an arbitrary non-linear transform. Given a function

$$Y = f(X_1, X_2, \dots, X_n) \quad (1)$$

which maps an arbitrary combination of n random input variables, the unscented transform approximates each distribution of mean and variance as a set of sigma points given by:

$$\sigma = \left\{ \bar{X} - \sqrt{\Sigma}; \bar{X}; \bar{X} + \sqrt{\Sigma} \right\} \quad (2)$$

The mean and variance of this set of points is identical to that of the initial distribution, so the mean and variance of the transformed sigma points form an approximation to the true post-transform distribution. This transform is significantly faster than Monte-Carlo methods, as only a small number of points are transformed, and allows for good approximations to be obtained as long as the transform of the original distribution is close to Gaussian. In this case, the input distribution is treated as a 4D Gaussian with the following terms:

$$X = \begin{bmatrix} x_1 \\ y_1 \\ x_2 \\ y_2 \end{bmatrix} \quad \Sigma = \begin{bmatrix} \sigma_{1xx}^2 & \sigma_{1xy}^2 & 0 & 0 \\ \sigma_{1yx}^2 & \sigma_{1yy}^2 & 0 & 0 \\ 0 & 0 & \sigma_{2xx}^2 & \sigma_{2xy}^2 \\ 0 & 0 & \sigma_{2yx}^2 & \sigma_{2yy}^2 \end{bmatrix} \quad (3)$$

Where (x_1, y_1) and (x_2, y_2) are the position of the two targets between which a bearing is being estimated. The sigma points are then transformed through

$$\theta = \tan^{-1} \left(\frac{y_2 - y_1}{x_2 - x_1} \right) \quad (4)$$

to produce a single bearing estimate given by $\bar{\theta}$ and σ_{θ} .

This same transform is used both to estimate the heading of the person between any two time-steps in the data, and to estimate the bearing from the person to any potential points of interest in the room.

The heading estimates are also subject to a threshold condition for distribution approximation. Due to the fact that the transform is moving to a closed space in bearing, any approximation in which $\sigma_{\theta} \geq \pi$ is replaced by a uniform distribution over the range $-\pi \leq \theta \leq \pi$ due to issues regarding the effect the tails of the Gaussian have as they wrap around the distribution. As the heading estimates are Kalman filtered [Bar-Shalom and Fortmann, 1988] to smooth them, this term is only invoked extremely rarely, when the person has been stationary for an extremely long period of time.

The bearing estimates are never subject to this condition in this dataset, as any objects which are close enough together to produce bearings with extreme variances will already have been blobbed into a single track. In future, however, an improved vision-based tracker will allow tracking of objects in close proximity to each other, and this rule may have to be further investigated.

The interest estimate from heading in any given target is then calculated according to the Mahalanobis distance between the heading estimate for the target as a distribution, and the bearing estimate to the point of interest as a single point. This provides a distance, with the relative levels of interest determined by the variance in the heading. This gives appropriate behaviour, as it becomes more peaked as the confidence in the heading grows, and more uniform as confidence in the heading decreases. The actual interest parameter is given as:

$$I_R = 1 - \text{Norm}(\text{Dist_Mahalanobis}(\theta, \psi)) \quad (5)$$

where θ is the heading of the person, and ψ is the bearing to the target. The Mahalanobis distance is normalised and then inverted to ensure that a small distance is treated as a high interest.

Two alternative measures were also considered. These were the Kullback-Leibler divergence [Kullback and Leibler, 1951] between the heading and the bearing distributions, and the Euclidian distance between two Gaussians as proposed by Grover [Grover, 2006]. These have the disadvantage that they indicate a measure which

is significantly more dependant on the variance in the bearing and heading than the Mahalanobis distance. As such, it is possible with these two measures for a bearing with high mean difference from the heading, and large variance to have a smaller distance measure from the heading than one with a closer mean value, but smaller variance.

3.4. Calculation of Interest through Range

The estimation of interest through range is significantly less straight-forward than estimation through heading. The transform from position to range is again non-linear, and is calculated through the unscented transform where

$$R = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2} \quad (6)$$

using the same 4D base distribution as was used for the heading calculations. This is chosen over the Euclidian distance between Gaussians discussed above, as the Euclidian distance will change with variations in the covariances of the position measurements, and it is required that the metric utilised has a stable mean for a constant separation of means between the two points.

The unscented transform gives the required estimates of \bar{R} and σ_R . The Interest due to range is then calculated as:

$$I_R = \frac{1}{R} \quad (7)$$

In this case the non-linearity of the range to interest conversion is useful, as close proximity is a strong interest indicator. In the future, additional filtering and calculations will be implemented to allow for estimation of the third primitive as a direct follow-on from the range readings.

3.5. Fusion of interest results

The output of both interest calculation methods is a vector of distance measures with as many elements, n , as there are points of interest in the room. In order to effectively fuse these two pieces of information, they are first normalized, so that the interest levels sum to one. If we assume that the object of interest is contained within these vectors, then this allows us to easily scale the various disparate measurements into a set of more easily comparable values. If the true point of interest lies outside the set of objects in the interest vector, then any falsely high interest readings will be of a form that approximates the primitives of the person's motion regardless, so the difference for the model is purely academic.

Once these two interest distributions have been normalized, they can be treated as a PDF of interest over each of the objects in the room. The final interest level is then taken as a fusion of these two PDFs. These PDFs are combined using Bayesian fusion [Durrant-Whyte, 2005] as given by

$$I_{final} = \alpha \cdot I_{Heading} \cdot I_{Range} \quad (8)$$

where α is a re-normalisation parameter. This method behaves in an appropriate manner, weighting the resultant interest towards the most informative observation. This

technique also allows for the easy addition of more terms into the fusion, as well as the potential for transformations of observations using a sensor model. This would allow weighting of observations to favour the primitives which are more appropriate for any given object of interest.

4. Tracker Results

The results from the Pittsburgh data confirm the majority of the assumptions and design criteria discussed above. The fusion of interest observations from the primitives is also investigated. In this dataset, the points of interest used are the two wheelchairs, the exit to the room, and any other people which may be in the room at the time. This demonstrates the capabilities of the system to handle spatially static and spatially dynamic points of interest, as well as points of interest which may only be present for a short period of time.

4.1. Interest Tracker

The two primitives for the interest tracker that have been implemented are closely tied to the motion type. Figure 6 shows the results of the interest tracker running on person #287335, shown in the available video¹. In this video, the length of the red lines towards the targets from the represent interest due to Range, green represents interest due to heading, and blue represents the final interest estimate. As is clear from the video, the person walks towards the Blue chair, showing high interest in that chair, then turns and walks back towards the exit. The appearance of person #287336 also appears as a spike in interest, as that person appears to the tracker close to person #287335

As expected for this motion primitive, the measure of interest due to heading is generally more informative so the final interest estimate tends to track that. It should be noted that at the start of the trajectory, the interest due to range is strongly aligned towards the door, which is reflected in the final interest estimate.

Figure 7 shows the interest tracker results from a different person. Initially the person is stationary giving a large heading variance, and thus the range interest estimate showing interest in the red chair propagates strongly through to the final interest estimate. At 6.5 seconds, the track starts moving, yielding an improved heading estimate, and thus allowing the interest estimate through heading to contribute more strongly to the final estimate.

4.2. Real-Time capability

Another important requirement for this work is real-time portability. As the shown in Table I, even when running under MATLAB, the interest tracker can perform all the

required operations significantly faster than real-time. This is encouraging, as the successful implementation of this technique at this speed implies that the complete system will be able to run in a computationally efficient manner. This will allow for simple integration with the existing Fish-Bird project system.

TABLE I
MATLAB PROFILER RESULTS

Length of Data File (Seconds)	Number of Observations of People	Time to Process (Seconds)	Real-Time Processor Utilisation @ 5Hz
3598	82509	827.1	23.0%
3598	51933	380.7	10.5%
3599	50669	437.1	12.1%

Profiler results running on a P4 3.00 GHz with 512 Mb of RAM

It should be noted that the 3 files shown above are the files with the largest number of observations from the entire 3-month set of logs. As such, the processor times shown are for peak load periods, rather than an average dataset. Secondly, implementing the code in C++ for integration with the project should result in a significant performance increase, allowing for higher sample rates without an associated increase in processor costs. It should also be stated that 40-50% of the computational cost in this code was the performing of the unscented transform, and so optimisation of the technique used to calculate the headings and bearings should allow a significant reduction in the computational cost.

4.3. Conclusions

The interest tracker itself is performing to the specifications as described earlier. The distributions produced match expectations, given the design of the primitives. This tracker is capable of handling a dynamic number of dynamic objects within the space. This is demonstrated through the fact that other people in the room are considered to be potential interest targets, in addition to the 2 chairs, and a manually defined area encapsulating the entrance/exit to the room.

As presented, this work has a number of shortcomings. Firstly, the lack of learning techniques to adapt to variations between the pre-defined primitives and the true behaviour of the people within the space. Secondly it is not capable of correlating between actions in the environment and the resultant change in interest levels exhibited. Work to overcome these limitations has proceeded using DBNs, where the inputs to the system can be enumerated, and the relationships between these inputs and the behaviour of people within the space to be more easily examined. The prediction of interest also becomes possible once this system is fully implemented, which will allow a more powerful predictive and modelling tool.

The data-structures used in this training system are of manually-tagged interest data, so that comparisons can also be performed between the interest as learned by the system and interest as seen by a human trainer. This will allow further confirmation of the ties between the interest

¹ - Video available at http://www.jugglethis.net/ACRA2007/D_Wood_ACRA-07.avi

Note: Video cannot be streamed from this location, it must be saved locally to play

being modelled and the psychological definition of interest.

The work presented herein demonstrates the existence of interest as postulated, and confirms that this parameter is detectable in the datasets collected. This will lead to more accurate modelling of the motion of people through spaces in which there are dynamic elements.

Acknowledgments

This work is supported by the ARC Centre of Excellence programme, funded by the Australian Research Council (ARC) and the New South Wales State Government. The authors would also like to acknowledge the help of Mari Velonaki, James Hudson, Stefan Williams and David Rye from the Fish-Bird project.

References

- [Fajen and Warren 2003] B. Fajen and W. Warren, "Behavioural Dynamics of Steering, Obstacle Avoidance and Route Selection," *Journal of Experimental Psychology: Human Perception and Performance*, vol. 29, pp. 343–362, 2003.
- [Warren and Fajen. 2004] W. H. Warren and B. R. Fajen, "Behavioural Dynamics of Human Locomotion," *Ecological Psychology*, vol. 61, pp. 61-66, 2004.
- [Vasquez et al., 2005] D. Vasquez, T. Fraichard, O. Aycard, and C. Laugier, "Intentional Motion On-line Learning and Prediction," presented at International Conference on Field and Service Robotics (FSR), 2005.
- [Bruce and Gordon, 2004] A. Bruce and G. Gordon, "Better motion prediction for people-tracking," Proceedings of the IEEE International Conference on Robotics & Automation (ICRA), 2004.
- [Panangadan et al., 2004] A. Panangadan, M. Mataric, and G. Sukhatme, "Detecting Anomalous Human Interactions using Laser Range-finders," presented at IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2004.
- [WordNet, 2007] *WordNet Search - 3.0*, <http://wordnet.princeton.edu/perl/webwn>, last accessed 5/09/2007
- [Velonaki et al., 2005] M. Velonaki, D.Rye, S. Scheduling and S. Williams. *Fish-Bird Circle B - Movement C*, Nicholas Tsoutas (curator), Artspace Visual Arts Center, Sydney, Australia, September – October 2005
- [Horstmann, 2005] G. Horstmann, "Attentional Capture by an Unannounced Colour Singleton Depends on Expectation Discrepancy," *Journal of Experimental Psychology / Human Perception & Performance*, vol. 31, pp. 1039-1060, 2005.
- [Suresh Kumar et al., 2004] F. R. Suresh Kumar, B. Upcroft and H. Durrant-Whyte, *Statistical Models for Perception of Natural Features*, 2004
- [Jenkins and Matarić, 2002] O. C. Jenkins and M. J. Matarić, "Deriving Action and Behaviour Primitives from Human Motion Data", presented at IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2002.
- [Julier et al., 2000] S. Julier, J. Uhlmann, and H. F. Durrant-Whyte, "A New Method for the Nonlinear Transformation of Means and Covariances in Filters and Estimators," *IEEE Transactions on Automatic Control*, vol. 45, pp. 477-482, 2000.
- [Bar-Shalom and Fortmann, 1988] Y. Bar-Shalom and T. E. Fortmann. *Tracking and Data Association*. Academic Press, 1988.
- [Kullback and Leibler, 1951] S. Kullback and R. A. Leibler. On information and sufficiency. *Annals of Mathematical Statistics*, vol. 22(1), pp. 79–86, March 1951.
- [Grover, 2006] R. Grover, "A Model for Sensing and Reasoning in Autonomous Systems" Ph.D. Thesis, Australian Centre for Field Robotics, University of Sydney, Australia, 2006
- [Durrant-Whyte, 2005] Hugh Durrant-Whyte, Course notes: *Introduction to Estimation and Data Fusion* [Internal Publication, Centre of Excellence for Autonomous Systems, Sydney, Australia], 2005
- [Cowell et al., 1999] Robert G. Cowell, A. Philip Dawid, Steffen L. Lauritzen, David J. Spiegelhalter. "Probabilistic Networks and Expert Systems", Springer-Verlag, New York, 1999
- [Theeuwes et al., 2004] J. Theeuwes, A. F. Kramer, A. Kingstone, "Attentional Capture Modulates Perceptual Sensitivity", *Psychonomic Bulletin & Review*, vol. 11 (3), pp. 551-554, 2004

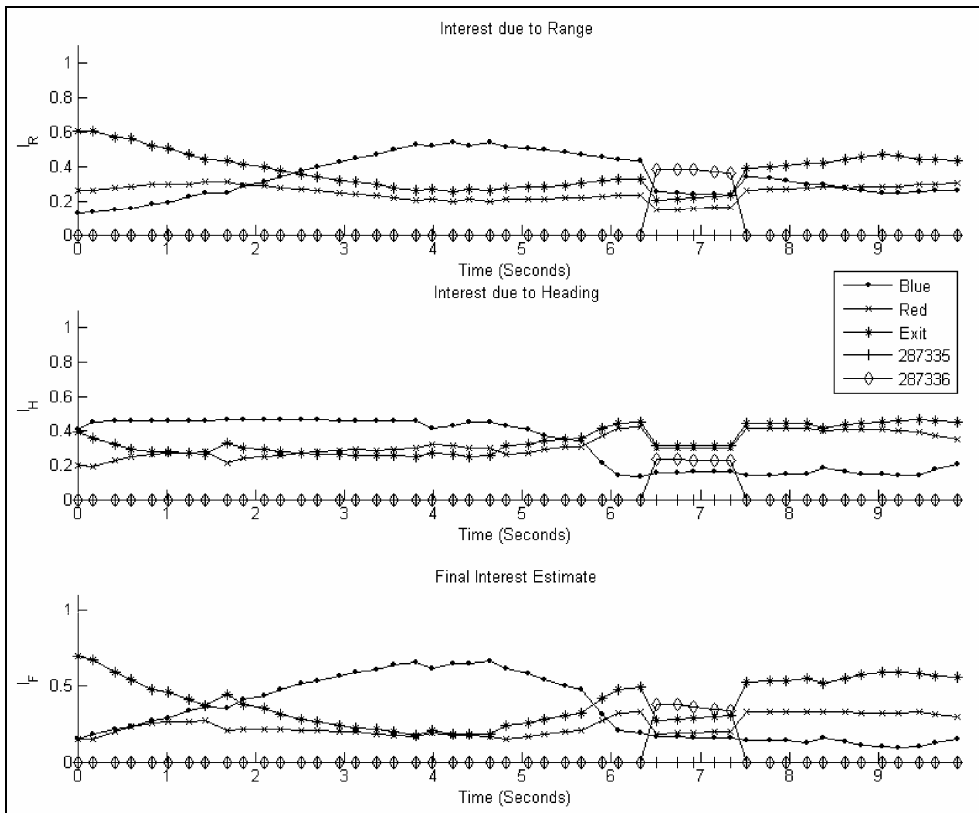


Figure 6: The Interest Tracker Results on a track conforming mostly to the heading-based motion primitive.

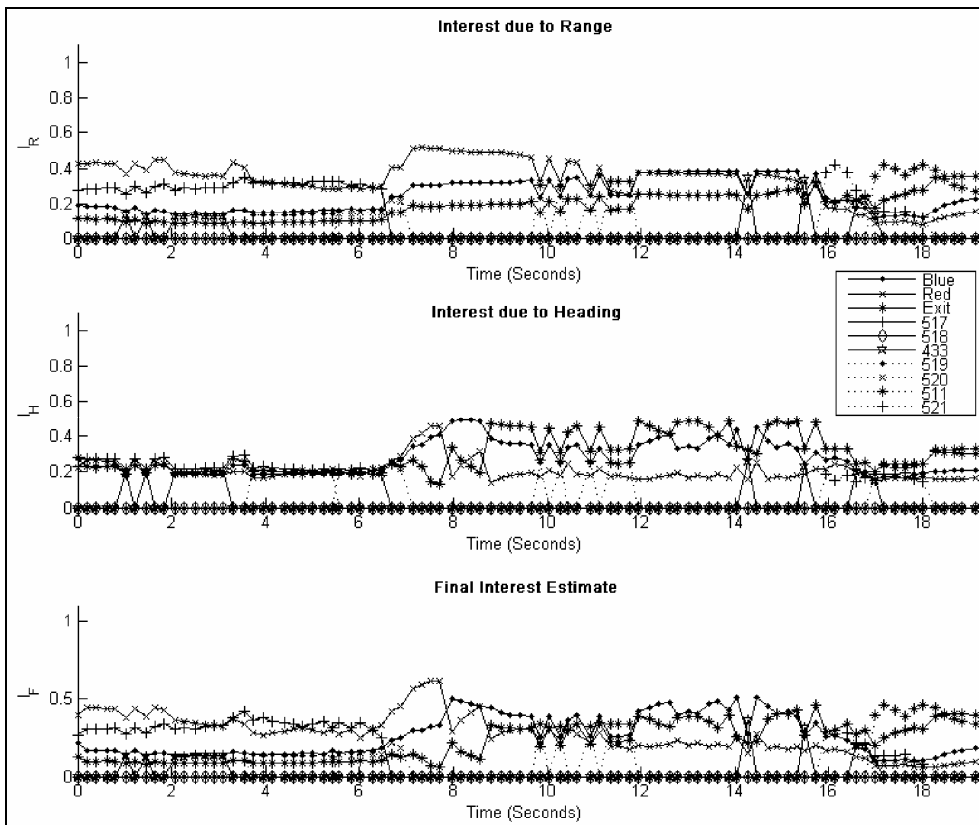


Figure 7: The Interest Tracker Results on a track conforming to multiple motion primitives.