

Show me the way to Monte Carlo: density-based trajectory navigation

Steven Strachan¹
Hamilton Institute
NUI Maynooth¹
steven.strachan@nuim.ie

John Williamson²
Dept. of Computing Science
University of Glasgow²
jhw@dcs.gla.ac.uk

Roderick Murray-Smith^{1,2}
Dept. of Computing Science
University of Glasgow²
rod@dcs.gla.ac.uk

ABSTRACT

We demonstrate the use of uncertain prediction in a system for pedestrian navigation via audio with a combination of Global Positioning System data, a music player, inertial sensing, magnetic bearing data and Monte Carlo sampling for a density following task, where a listener's music is modulated according to the changing predictions of user position with respect to a target density, in this case a trajectory or path. We show that this system enables eyes-free navigation around set trajectories or paths unfamiliar to the user and demonstrate that the system may be used effectively for varying trajectory width and context.

Author Keywords

GPS, Navigation, Uncertainty, Monte Carlo, Feedback, Audio, Tracking, Probabilistic Display.

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous

INTRODUCTION

In this paper we describe the first implementation of a completely hand-held system, which utilises a combination of Global Positioning System (GPS) data, a music player, inertial sensing, magnetic bearing data and Monte Carlo sampling and modulates a listener's music in order to guide them through a density to a desired physical location. In this instance along a set trajectory or path. Car navigation systems are now widespread and well-designed but there are openings in developments in pedestrian navigation systems. What if users wish to allow their phones to *control* them, and guide them from a starting position, through their desired trajectory to a goal point, subject to constraints? This system offers a new general mechanism for providing highly interactive context-aware applications. The densities could represent $P(C_i|x)$ - the probability of context state C_i given the current state vector x . By treating our system as a separate density layer in any application it is possible to provide different functionalities. For example, densities could be used to represent differing contexts such as local socio-economic levels or crime rates, areas of interest to tourists or various Geographic Information Systems data.

LOCATION-AWARE AUDIO

Many standard GPS come with some form of audio feedback. The majority of car navigation systems use combi-

ned audio and visual feedback, with direct voice commands being the most popular mechanism for influencing drivers. Speech-based feedback is also popular in pedestrian based GPS projects, especially for visually-impaired users. The use of non-speech based audio cues includes, e.g. Loomis *et al* [5], who describe an experiment where users are guided along a route of predefined way-points using their back-pack based system, which uses spatialised audio in order to convey information about the surrounding virtual environment. Other work in the area includes that of Holland *et al*. [2] who describe their prototype spatial audio user interface, *AudioGPS*.

Using music as the mechanism for guiding users has also been previously investigated. Work on music based guidance includes *gpsTunes* [8], where initial testing of a prototypical system had shown that it was possible to allow users to navigate in the real world, using a combined Audio/GPS player to aid navigation. Similarly Etter *et al* [1] describe, *Melodious Walkabout*, which again utilises a users music to guide them to their desired location. A study was conducted which concluded that it was possible, after users had gained some initial experience, to guide people by adapting their own music in a spatial way. Jones *et al* [3,4,9] have conducted a similar study with their *OnTrack* system. They show that it is possible to guide a user through a number of audio beacons to a desired location using continuously adapted music.

One fundamental difference between previous systems and our own is that whereas those systems tend to provide a user with an audio cue from a distant beacon or way-point and let the user take any route they desire to that location, we are attempting to provide continuous controlled audio feedback which can provide flexibly varying constraints. The system we describe in this paper is significantly different from other similar navigation applications in that it explicitly uses display of uncertainty in order to assist the user. Probabilistic approaches to the display of uncertain information in interaction remain largely uninvestigated but we introduced an approach with a location-aware audio feedback example in [10]. The system uses a probabilistic approach to presenting auditory information about the user's current state. We found that appropriate display of uncertainty could increase the usability and acceptance of mobile GPS systems. The use of inertial sensing in our system makes it fully interactive, with tilting and panning used to allow the user to actively probe the local environment in an tightly-coupled,

embodied manner. Also, where previous systems used audio as their only form of feedback, we also introduce vibrotactile feedback for a fully handheld, multi-modal interface. The inclusion of this kind of system in a mobile phone would be advantageous if, for example, a user could be guided to where they wished to go whilst simultaneously taking a call, utilising only the vibrotactile feedback.

TRAJECTORY-TUNES

As an example application of our density probing system we developed a trajectory following application which follows naturally from the *gpsTunes* system described in [8, 10]. It is designed to guide a user along a desired trajectory using audio and vibrotactile feedback via their mobile device, be that a PDA, music player or mobile phone. If a user is traversing from one point to another in an area with which they are not familiar, there may be an optimal trajectory to that point or a trajectory which avoids any potential hazards. In this situation it is up to our system to guide the user through this preferred path.

The desired trajectory is represented by a density layered on top of a map of the local area, as in figure 3-(5). Monte Carlo propagation is then used for browsing this density map, which allows us to actively probe the locality by simulating possible paths into the future from the estimated location, along the current heading, enabling us to predict likely positions of the user at future time points. If the user, at a predicted future position, has strayed from the correct trajectory, this information may be fed back to the user giving them the chance to take control and change their behaviour.

Hardware

The equipment used consists of an HP iPAQ 5550 running windowsCE equipped with a MESH [7] inertial navigation system (INS) backpack consisting of 3 Analog Devices $\pm 2g$ dual-axis ADXL202JE accelerometers, 3 Analog Devices $\pm 300deg/s$ Single chip gyroscopes, 3 Honeywell devices HMC1053 magnetometers and a vibrotactile device. Our



Figure 1. Left: MESH pack & attached to a PocketPC & open.

GPS is a Trimble Lassen Sq module, produced for mobile devices, and is also built-in as part of MESH (see figure 1). In this system, apart from utilising the GPS for positioning, we have also used the accelerometers to calculate pitch and roll, the magnetometers in conjunction with the accelerometers to achieve tilt-compensated heading and the vibrotactile unit for feedback.

Monte Carlo Propagation for Browsing

The major novel feature of the described system is the browsing interface which facilitates active probing of the locality




Figure 2. Particle beam flowing around the campus likelihood map. The particle paths are illustrated in blue; the samples at the time horizon are highlighted as bright green dots.

as described in [10]. This is achieved by projecting possible paths into the future from some location along a given heading using Monte Carlo sampling. Details of Monte Carlo methods can be found in Chapter 29 of [6]. A straightforward propagation of particles through the search space would lead to a fairly simple distribution of points at the time horizon, which would be unlikely to model likely possible user destinations effectively. It is extremely unlikely, for example, that the user will be inside a solid wall at any point in the future. To represent these varying positional likelihoods we use a simple likelihood map, giving a probability p of being in a particular position (as measured by the sensors) in the mapped area. An example of such a map is shown in Figure 2; in this example the buildings have very low likelihood and there is increased likelihood around paths.

Feedback

This system utilises both audio and vibrotactile feedback. The audio feedback consists of a distortion of the musical waveform. The distortion takes the form of a reverb effect which is modulated by the likelihood of the user being on the correct path at the Monte Carlo time horizon. This is computed by summing the values of the likelihood map at the Monte Carlo sample points to estimate the overall probability of being on the path at the horizon, $v = \sum_0^S t(x_i^s)$, where t is the trajectory probability density function. This value is used to modulate the reverb parameters such that a low probability of remaining on the trajectory results in increased reverberation, producing echoing and muddy sounding output, and this is also mapped into vibrotactile feedback so that a low probability of remaining on the trajectory results in a ‘stronger’ vibrotactile feedback. Sticking closely to the path produces clean, crisp sound with no vibrotactile feedback.

Scanning Via Variable Time Horizon Prediction

One of the most important ways in which the user can interact with the navigation system is via the direct manipulation of the Monte Carlo prediction time horizon. The interactor can use this to probe further into the possible futures or bring the particle beam in close to examine the nearby area. In particular, this allows the user to experience how the uncertainty in the potential goal space changes. In our implementation the time horizon is controlled via vertical tilt (sensed by the accelerometers in the MESH hardware), by analogy to artillery fire. Higher tilt levels project the particles further into the space, with correspondingly greater uncertainty (depending on the model). A tilt back looks into the future and a tilt forward brings us back to the present. Our system was cali-

brated to have a 20m look-ahead, so at maximum extension the user will be hearing feedback from 20m ahead.

EXPERIMENT: TRAJECTORY FOLLOWING

An initial probe was conducted to demonstrate that the system could guide users to a desired location, and also to examine the effects of varying trajectory width, and the presence of visual distracters.

Method

In total 6 participants took part in the experiment, aged between 20 and 29. All participants had used a mobile phone or PDA before but only 2 had any experience with GPS navigation. Five trajectories were used in total with four of them taking the same form but with varying width. The main trajectory used (trajectories 1-3 and 5 in figure 3) represent a well known path on the university campus as shown in figure 3(5). This trajectory was then translated over to a wide-open, featureless playing field and given three different widths. Users were first required to traverse trajectories 1-3, presented in a counter-balanced order. The 4th trajectory presented was a simple N-Shape which was also placed over the open playing field. The final trajectory presented was the identical shape to trajectory 1 but this time it was placed back over the main campus, over paths and round buildings. Before the experiment began participants were given a 5 minute description of the system followed by a practice run, to gain a feel for using the system over a relatively simple trajectory. Heading data was recorded, along with timestamps, 3-axis accelerations, latitude, longitude, ground speed, pitch angle of the device and total uncertainty.

Results

The principal result from this experiment is that it is possible for this system to guide users to a set location, with no user failing to reach the end point of any trajectory. A number of different strategies were employed by the users. Some users were very active in probing the locality, taking a cautious and careful approach, as in figure 6(a). This figure shows a quiver plot where the blue dots represent the user's current position, the direction of the cyan arrows represents the heading direction, the length of the cyan arrows represents the tilt of the device and the red dots represent the current Monte Carlo prediction location. If these predictions are located on the white area, negative feedback is produced, if they are located on the black area there is no feedback. Other users were relatively inactive in scanning for the most part, but became very active when it was required, employing a 'straight-ahead' approach while receiving no feedback and only scanning when they began to move off of the correct path to find another good direction. This led to a zig-zagging or bouncing behaviour as shown in figure 6(b). One other interesting behaviour observed is when the user 'clings' to the edge of the trajectory, as in figure 6(c). They move along the path keeping touch with the edge, using the feedback from it as a guide, reassuring themselves.

The inclusion of trajectory 4 allowed us to examine the effects of adding tight corners where a real understanding of the trajectory space is required in order to successfully




Figure 4. Completion times for six participants over all five trajectories.

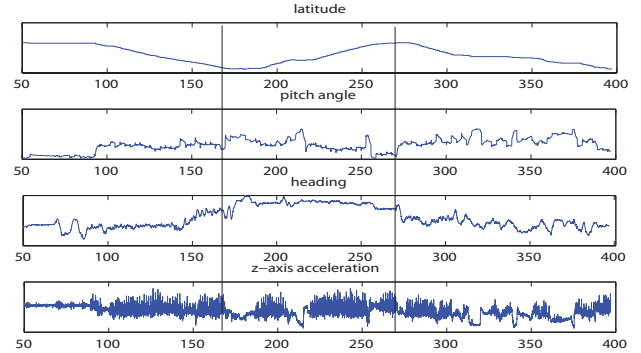


Figure 5. A comparison of the effect of a corner on trajectory 4. Vertical lines indicate the location of the corner.

complete the course and from the figure for completion time. From observation it was clear that users had most trouble with this trajectory. Figure 6(d) shows an example from one user who traversed this N-shaped trajectory successfully. When the user reaches the corners of the trajectory a lot more probing activity is observed in the quiver plot, since at this point the user is required to fully exploit the degrees of freedom, in order to recover the trajectory. Figure 7(right) shows the tilt and walking activity for the same example. Observe from the z -axis accelerometer data, that at the corner points in the latitude plot the user stops, then there is a burst of activity in the pitch angle, where the user is attempting to look-ahead, and a shift in the heading to the correct direction.

Looking at figure 5 showing the completion times for each participant, we see that the participants generally finished more quickly on the widest trajectory (trajectory 3) although comments from the users suggested that some of them found the lack of feedback and relative freedom in this case slightly disconcerting. Figure 7 shows the boxplot for scanning energy, defined as $\sqrt{\sum_{t=0}^T (\frac{dx}{dt})^2}$, where x is the heading signal. This shows that users tended to scan less for the widest trajectory number 3 and most for the narrowest trajectory number 2. This is intuitive as we would expect users to react to and increase scanning immediately after feedback and in the case of trajectory 2 they are generally receiving more feedback than in the wider trajectory number 3. Interestingly, we see that the completion time for trajectory 5, through the campus, is significantly lower than for all other trajectories, including its equivalent trajectory 1, on the open playing field. So, while we have shown that in a featureless environment like a playing field, people were able to follow the path,



Figure 3. All trajectories used in the experiment. Trajectory 5 is located in a campus setting with all other trajectories located on the playing field in the top left corner. Trajectory 2 was 9m wide and was the narrowest, trajectory 1 was 18m wide and trajectory 3 was 36m wide.

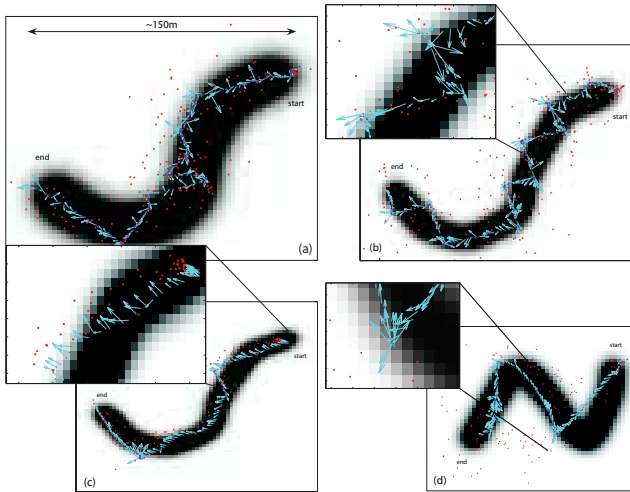


Figure 6. (a) A 'cautious' approach to trajectory 3 (b) A 'straight-ahead' approach to trajectory 1. (c) A 'clinging' approach to trajectory 1. (d) A user traversing the 'N'-shaped trajectory number 4.




Figure 7. Left: Plot of scanning energy for each participant on each trajectory. Blue circles indicate the mean value for that trajectory. Right: Plot of completion times for each participant on each trajectory.

their performance improves significantly when the audio and vibrotactile constraints from the PocketPC are coupled with the natural constraints of the environment (paths, buildings etc).

Some users also commented that they found the vibrotactile feedback more useful than the audio feedback, although there was no difference in the way the feedbacks were triggered. This could be due to the on/off nature of the vibrotactile feedback (on if they were straying off of the path and off if they were ok) whereas the audio feedback was part of the music they were listening to. It may have been difficult then to perceive small reverberations in the sound compared to small vibrational pulses.

CONCLUSIONS

This paper has demonstrated that probabilistic, multimodal, handheld interaction techniques can be applied effectively to allow users to explore general density functions in space, with an example of pedestrian GPS navigation. The Mon-

te Carlo sampling method provides an effective way of integrating probabilistic models into practical interfaces, and displaying the results in a multimodal fashion.

The field trials have shown that it is possible to guide users to a desired location over a set trajectory or path and a number of interesting behaviours have been observed. Interactive sonification of the exploration process produced a navigation system which may be used eyes-free, where the user brings their sensorimotor systems into the interaction with an augmented environment. It is clear from this initial data that it is possible for users to navigate to the final location in a featureless environment like a playing field, using audio and vibrotactile feedback alone. Their performance and confidence improves significantly when the audio and vibrotactile constraints from the system are coupled with the natural constraints of the environment, suggesting that the system is promising for a range of realistic use cases.

Acknowledgements

We are grateful for support from: HEA project Body Space, SFI grant 00/PI.1/C067, IRCSET BRG SC/2003/271, Continuous Gestural Interaction with Mobile devices, EPSRC project Audioclouds, IST Programme of the European Commission, under PASCAL Network of Excellence, IST 2002-506778.

REFERENCES

1. Etter, R. Melodious walkabout - implicit navigation with contextualized personal audio contents. *Adjunct Proceedings of the Third International Conference on Pervasive Computing*, vol. 191(204).
2. Holland, S., Morse, D. R., and Gedenryd, H. AudioGPS: Spatial audio navigation with a minimal attention interface. *Personal and Ubiquitous Computing* 6, 4 (2002), 253–259.
3. Jones, M., and Jones, S. The music is the message. *interactions* 13, 4 (2006), 24–27.
4. Jones, M., Jones, S., Bradley, G., and Holmes, G. Navigation by music: an initial prototype and evaluation. *Proceedings of the International Symposium on Intelligent Environments*. Microsoft Research (ISBN: 1-59971-529-5) (2006), 1849–1852.
5. Loomis, J. M., Golledge, R. G., and Klatzky, R. Navigation system for the blind: Auditory display modes and guidance. *Presence: Teleoperators and Virtual Environments*, 7, 193-203., 1998.
6. MacKay, D. J. C. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
7. Oakley, I., Ångeslevä, J., Hughes, S., and O'Modhrain, S. Tilt and feel: Scrolling with vibrotactile display. *EuroHaptics 2004*. 316–323.
8. Strachan S. and Eslambolchilar P. and Murray-Smith R. and Hughes S. and O'Modhrain S. gpsTunes - controlling navigation via audio feedback. *Mobile Human-Computer Interaction 05*.
9. Warren, N., Jones, M., Jones, S., and Bainbridge, D. Navigation via continuously adapted music. *CHI '05: CHI '05 extended abstracts on Human factors in computing systems*. ACM Press (New York, NY, USA, 2005), 1849–1852.
10. Williamson, J., Strachan, S., and Murray-Smith, R. It's a long way to Monte Carlo: probabilistic display in GPS navigation. *MobileHCI '06*. ACM Press (New York, NY, USA, 2006), 89–96.