

It's a long way to Monte Carlo: Probabilistic display in GPS navigation

John Williamson¹
¹University of Glasgow
Glasgow, UK
jhw@dcs.gla.ac.uk

Steven Strachan²
²Hamilton Institute
NUI Ireland, Maynooth
steven.strachan@nuim.ie

Roderick Murray-Smith^{1,2}
rod@dcs.gla.ac.uk

ABSTRACT

We present a mobile, GPS-based multimodal navigation system, equipped with inertial control that allows users to explore and navigate through an augmented physical space, incorporating and displaying the uncertainty resulting from inaccurate sensing and unknown user intentions. The system propagates uncertainty appropriately via Monte Carlo sampling and predicts at a user-controllable time horizon. Control of the Monte Carlo exploration is entirely tilt-based. The system output is displayed both visually and in audio. Audio is rendered via granular synthesis to accurately display the probability of the user reaching targets in the space. We also demonstrate the use of uncertain prediction in a trajectory following task, where a section of music is modulated according to the changing predictions of user position with respect to the target trajectory. We show that appropriate display of the full distribution of potential future users positions with respect to sites-of-interest can improve the quality of interaction over a simplistic interpretation of the sensed data.

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous, H.5.2 User Interfaces: Auditory (non-speech) feedback, H.5.2 User Interfaces: Input devices and strategies, G.3 Probability and Statistics: Probabilistic algorithms (including Monte Carlo)

General Terms: Human Factors

Keywords: GPS, Navigation, Uncertainty, Monte Carlo, Feedback, Audio, Probabilistic display

1. MOTIVATION

Problems associated with the usability of mobile devices have been extensively documented. Holland *et al* [3], for example, found that mobile devices tended to demand too much visual attention in their use which may be a potential hazard in some situations. Using a mobile device whilst 'on the move' is a particular problem since it demands significant increase in user attention and cognitive load.

Mobile GPS systems are relatively new and are increasingly being integrated into standard mobile devices such as handheld gaming machines and phones. GPS can be unreliable since there are

frequent problems with spatial resolution, latency and signal shadowing, which may all be detrimental to navigation systems. This, coupled with the user's lack of knowledge of an area in which they are navigating, may render a system unusable.

It is for this reason that we seek to introduce probabilistic, multimodal displays where the user is engaged in continuous negotiation with the system. In this paper, we apply these ideas to the GPS navigation problem on a mobile device augmented with inertial sensing. We demonstrate a probabilistic approach to navigation using a combination of GPS and general inertial sensing. The incorporation of techniques from control and probability theory will allow us to embrace the omnipresent uncertainty and provide a more flexible and usable system. By introducing goal-focused predictive displays to an interface, with appropriate calculation and display of the outcomes, control of the system can be improved. Smith [9] gives a rigorous explanation of the importance of maintaining uncertainty in nonlinear prediction problems. This is not just an observation of interest to technical systems. There is significant, well-controlled experimental evidence (for example, the work of Körding and Wolpert in [5]) that correct display of uncertainty leads to regularised control behaviour in human motor control, in reaching actions and targeting actions. If this can be generalised to broader interaction scenarios then it suggests that uncertain displays have the potential to 'smooth out' the interaction process and make use of a system less frustrating.

The introduction of inertial sensing gives an intuitive (and single-handed) way of interacting with the exploration process. It changes the navigation problem from a simple map-based display into a highly interactive system where the user can actively probe and scan for information, with a display that is focused *on the predicted goals of the user*. The modelling of future user behaviour as a dynamic system gives an effective method for prediction and display which can be controlled via the inertial sensing in a natural, flowing manner. We call this dynamic process *negotiated interaction*, and in this paper we use the example of location-based interaction to give some insight into the potential of the approach. The techniques can be applied in many interaction contexts; GPS navigation is just one particular situation where uncertain display and dynamical modelling are of immediate and obvious value. The notion of active exploration of a dynamic system via controlled sampling to predict a full distribution over the possible eventualities can be used to enhance any interactive system where uncertainty is present.

2. LOCATION-AWARE AUDIO

Location-aware audio systems are not new and many standard GPS come with some form of audio feedback. Work in this area includes that of Holland *et al*. [3] who describe a prototype spatial audio user interface for a mobile GPS. Their interface is designed

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MobileHCI'06, September 12–15, 2006, Helsinki, Finland.
Copyright 2006 ACM 1-59593-390-5/06/0009 ...\$5.00.

to allow mobile users to perform navigation tasks while their eyes, hands or general attention are otherwise engaged. They found with the use of a spatial, non-speech Geiger counter style sound and a prototype ‘back-pack’ based system that very simple and computationally inexpensive spatial mappings are surprisingly effective for helping users to find specific locations. Warren *et al.* [12] have recently shown, with a VRML simulation of a physical environment, that embedding navigation cues within a user’s own music is a promising new approach to the problem of supporting mobile users in navigation tasks. Their initial experimental work suggested that users would be able to follow routes by simply keeping track of the volume and perceived direction of a music source. [2] presents a more complete system for audio trajectory following with a modulated music approach.

Probabilistic approaches to location-aware audio remain largely uninvestigated and we believe that a probabilistic approach could significantly increase the usability and acceptance of mobile GPS systems.

3. CONTROL AND UNCERTAINTY

The system described in this paper differs significantly from many similar navigation and browsing applications in that it explicitly uses display of uncertainty to assist the user. This display fully represents all estimated uncertainty in the prediction of which targets the user may be interested in, and where and how they are likely to move to them. This provides the interactor with clear, direct feedback when entropy (the measure of the ‘spread’ of a distribution; see [6]) is low, and appropriately diffuse feedback as predicted entropy rises.

Displaying the distribution of potential states in this way gives a user sufficient information to make reasonable choices of actions (from a decision-theoretic point of view). Approximating this density with a single best estimate, as is often done in navigation software, gives the user unreasonable confidence in the accuracy of the system and prevents the interactor from choosing an optimal strategy for dealing with the true state of the world. Unrealistically precise feedback in the presence makes smooth, stable control hard; this ‘jumpy’ interaction is familiar to users of conventional GPS devices where the postulated location may sporadically shift in an implausible manner. Showing the true density means the user can appropriately simplify their behaviour when limited information is available.

Importantly, our framework gives the user the power to control the inference parameters as they are interacting, scanning the space and then homing in on regions of interest; for example, by varying the prediction time horizon as a function of a control signal.

4. GPSTUNES

gpsTunes [11] is a system designed to guide a user to a desired target by varying the perceived direction of their current song. So, for example, if a user enters an area with which they are not familiar and they wish to locate a desired area or building, they may inform the system of where they wish to go by clicking a map of that area, which will then alter the audibility and bearing of the music being played depending on how close or far this target is. As they move closer to the target, the audibility of the music will increase as the route to their target becomes less uncertain and they will be notified of their arrival by an additional pulsing sound played over the current track.

4.1 Equipment

The equipment used in the GPS navigation prototype system consists of an HP iPAQ 5550 equipped with a MESH [4, 8] 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 HMC1053 magnetometers and a vibrotactile transducer, used for feedback purposes. The GPS unit is a Trimble Lassen Sq module for mobile devices, and is also built-in as part of MESH (see Figure 1). This module provides us with a 9m resolution with up to 6m resolution around 50% of the time it is used. It also provides us with velocity resolution of 0.06m/s and an 18m altitude resolution.



Figure 1: Left: Mesh device alone and attached to an HP5550 Pocket PC. Right: The MESH circuit board showing the main components related to the navigation task

A standard orthogonal inertial sensor arrangement is used with the sensitive axis of the respective inertial sensors mounted coincident with the principle device axes providing us with direct measures of lateral accelerations, turn rates and magnetic field strength as well as the current GPS latitude and longitude.

4.2 Monte Carlo Propagation for Browsing

The major novel feature of the described system is the browsing interface which facilitates active probing of the locality. This is achieved by projecting possible paths into the future from some location along a given heading. Of course, since the sensed state is noisy, and any prediction introduces further uncertainty, the eventual outcomes form a density over the area being explored.

Ideally, an estimate of the user’s potential future locations would be represented as a probability density function over the navigable space, taking into account likely movement areas, sensor noise and obstructions. This function, however, is normally extremely complex for non-trivial landscapes, and no solution of simple form is available. Instead, it is possible to approximate using a set of samples drawn from the density; this is known as Monte Carlo sampling. It is much more straightforward to draw such approximating samples than it is to directly evaluate it, and the technique lends itself well to the subsequent display of the probabilistic information in a particulate form, such as granular synthesis. Details of Monte Carlo methods can be found in Chapter 29 of [6]. For example a visual display might consist of a point cloud overlaid on a map; goal-directed auditory analogues of this process are described later in the paper.

For the browsing task, a simple algorithm for sampling future possible trajectories is as follows:

- Draw samples $x^0 \dots x^S$ from a distribution ϵ around the current state. This distribution represents the sensor uncertainty at the initial position (e.g. from the shadow maps described later).
- For each step t until some horizon T :

- $x_t^s = x_{t-1}^s + h + l(x_t^s) + \sigma(x_t^s)$ where $\sigma(x_t^s)$ represents the model noise at the new point x_t^s (Gaussian, in our examples), and $l(x_t^s)$ represents the derivative of the likelihood map at that point. h is heading specified by the user. $\sigma(x_t^s)$ can be a constant value or a more complex function; e.g. from a map indicating the resolution or quality of the likelihood map.

- Display the samples x_T^s

This is somewhat similar to the *Hamiltonian* (or *hybrid*) Monte Carlo sampling process; [6], Ch. 30 has further details.

In our implementation, the inertial sensing platform is used to control this scanning, obtaining a heading from the magnetometers to produce h and controlling t via vertical tilt, as measured by an accelerometer. Physical location is estimated via the GPS. Further sensor fusion to increase location accuracy is not performed at this stage.

Intuitively, this process can be imagined as a beam of particles flowing out from around the initial state, probing into likely destinations in the direction the device is being held.

4.2.1 Likelihood Maps

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 pathways on the map. In this case, the map is generated by hand from an existing map, but such likelihood maps can also, for example, be derived automatically from digital photogrammetry maps.

The propagation algorithm can be modified to take account of this map simply by removing particles at a rate inversely proportional to their likelihood given their position. However, for increased computation throughput, our implementation instead modifies the dynamics of the particles such that they are deflected away from regions which are less likely, causing the samples to “flow” across the surface (i.e. by following the derivatives of the map). This produces a browsing system that channels Monte Carlo samples towards regions of increased likelihood, following traversable paths and avoiding obstacles in a natural manner.

It is obviously simple to extend this technique to multiple likelihood maps which can be combined based on context variables. Figure 2 shows an example, where suitable likelihood maps for walking and cycling behaviour are shown. A relatively simple context detection method can then estimate the probabilities of these possible alternatives, and combine these maps to produce a single output map incorporating context information. It is important to note that this includes the full uncertainty in the context variables, integrating the maps over the various outcomes weighted by their likelihood.

4.2.2 A Priori Sensor Uncertainty Maps

One further problem with the naïve propagation algorithm is that it takes no account of the varying uncertainty in sensor measurements, especially the spatially varying uncertainty arising from shadowing and reflection artifacts in GPS fixation. Such maps can be constructed ahead of time given knowledge of the geometry of potential occlusions (for example see [10]). We constructed simple




Figure 2: Particles flowing around the campus likelihood map. Higher red values indicate lower probability. The particle paths are illustrated in blue; the samples at the time horizon are highlighted as bright green dots. From top to bottom: (1) shows particles on likelihood map which is a model for walking behaviour and (2) shows the effect of a more constrained map which models a user on a bike, where particles tend to flow along available paths. (3) and (4) show the effect of the GPS shadow map on the propagation; (3) is a point outside of shadow, while (4) is a nearby point with heavy shadowing. The increased dispersion is obvious.




Figure 3: Estimated GPS shadows for the test campus region. Shadows are computed via a raytracing algorithm, based on satellite azimuth/elevation and an estimated height map for buildings in the area. Darker regions have less satellite coverage.

static occlusion maps for use in our platform with a raytracing technique based on currently locked satellite positions. The combined sensor uncertainty map resulting from estimated shadow positions for our test region is shown in Figure 3.

This map is included in the sampling algorithm by modulating the diffusion parameter ϵ at each time step by the calculated sensor uncertainty at the point. The total sensor uncertainty will be a combination of the map input and accuracy reading produced by the GPS device itself.

The accuracy of a GPS fix is also computed in the sensor hardware in real-time. This includes the number of satellites which have locks and more specific data giving the fix quality and the “horizontal dilution of precision”. This gives a scaling factor for the current uncertainty from 1–50. These are combined with the *a priori* sensor maps to obtain a certainty estimate for the current location.

It would theoretically be possible to improve the accuracy of these maps by comparing GPS readings with the likelihood maps described in the previous section; readings suggesting positions of low likelihood decrease confidence in the current veracity of the sensors. Additionally, we have assumed simple Gaussian diffusion in our spread model, which while a reasonable approximation, could be improved by diffusing particles proportional to the likelihood at their *new* positions (effectively Metropolis-Hastings sampling [6]).

4.2.3 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 prediction time horizon. The interactor can use this to probe further into the possible futures or bring the particle beam in close to examine nearby objects. In particular, this allows the user to experience how the uncertainty in the potential goal space changes. It provides an answer to the question: do all possible movements in

this direction inevitably converge to some likely goal or do they spread out rapidly to a multitude of potential targets? This feedback directly informs the user as to how much effort they will have to expend in scanning the space in the future.

In our implementation 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). This gives the user natural control of the prediction horizon.

There are significant advantages to non-visual display in a mobile context where visual attention is likely to be directed towards tasks not involving the device display. As such, the utility of the navigation system is greatly enhanced by the accurate presentation of the feedback in audio. Two basic prototypes have been implemented to demonstrate these ideas. These are designed to aid the user in target acquisition and trajectory following tasks.

4.3 Granular Synthesis – Target Acquisition

In the target acquisition case, a granular synthesis technique is used to display the output samples. Granular synthesis for probabilistic display is described in more detail in [14].

Each particle is displayed as a short audio “impact sound” drawn from some a selection of waveforms (each goal has one set of distinct source waves). These sounds are drawn from samples of a number of real, physical impacts (e.g. wood, glass, water, etc.) and vary in timbre. In the Monte Carlo case described here, each grain is associated with a sample, and the likelihood of activation with a particular waveform is given by the proximity of the sample to the goal in the location space. More precisely, we define a distribution f_i around each goal i . This set of distributions is used to transform the physical space into the goal space, and the probability of activating a sample grain is given by this distribution. The goal densities are Gaussian in the target acquisition prototype. The particles can be thought of as impacting the target densities; the value of the target map at which they impact modulates the volume of their presentation.

This produces a continuously changing auditory texture which represents the total distribution of particles in the goal space. The sound has a flowing impression which varies from sharply defined audio at low uncertainty to a vaguer mixture of sounds at increased entropy.

5. MUSIC MODULATION – TRAJECTORY FOLLOWING

In this prototype (which is somewhat similar to the *gpsTunes* system) a trajectory is associated with a musical waveform. Monte Carlo sampling is, exactly as in the target acquisition task, used to predict likely positions of the user at future time points. In this case 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 path at the 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, $v = \sum_0^S \tau(x_t^s)$, where τ is 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. This gives the user some sense of the changing probabilities without completely destroying the musical qualities of the source. Wandering off the path produces echoing and muddy sounding output; sticking closely to the path produces clean, crisp sound.

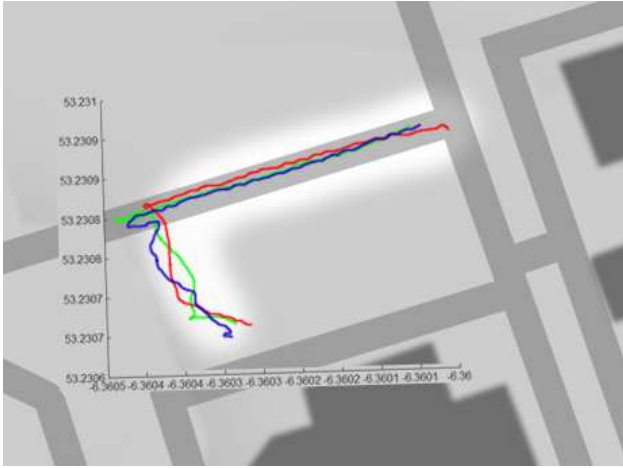


Figure 4: Traces of user behaviour following a trajectory displayed in sound only. The trajectory density is highlighted as a lighter background in this figure. The music is clear and audible while the user follows the trajectory, but reverberation is applied and intensified as their predicted future position strays from the reference density.

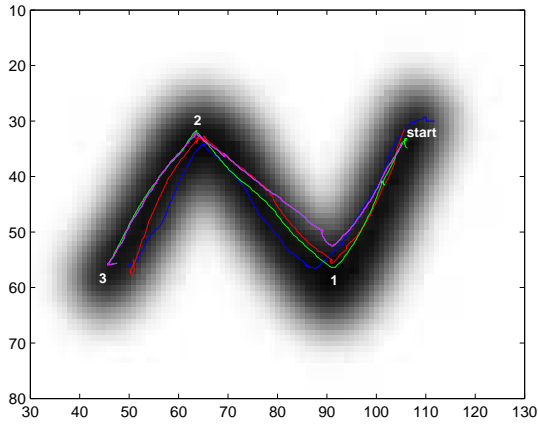


Figure 5: Traces of user behaviour following another trajectory, displayed in sound only. The trajectory density is highlighted as a darker background. Numbers highlight the corners in the trajectory.

An example showing three separate runs with a user following the trajectory with only audio feedback in a campus environment is visible in Figure 4. Figure 5 shows an example with a different route, in a comparatively featureless area (an open field), without visual distractions. The application was started at the same point each time, at the beginning of the trajectory. The music playing at this point is clear and audible whilst the user is pointing along the trajectory, which at the beginning is indicated by a long footpath. The user then walks along the path ahead with clear uninterrupted audio. At a certain point in the path the audio begins to reverb and a change of direction is required. The user scans the trajectory by rotating round to find one direction with clear audible music and heads off in this direction to the end of the trajectory.

Figure 6 shows the trajectory density, along with the path a participant followed, the heading faced and the predicted position. The scanning behaviour at the corners is apparent from the variability of

the heading shown in this diagram. The variability of the pitch angle (and hence the time horizon) also noticeably increases at points of high curvature as the user explores the space. Between corners the pitch angle remains stable. Figure 7 shows the heading time series corresponding to this trajectory, along with the pitch angle, which determines the time horizon.




Figure 6: The path and heading of a user as a trajectory is navigated. The orientation of the device (and thus the direction of the predictions) is shown as blue arrows extending from the current position. The red dots indicate the predicted positions, which are used for the sonification. Numbers correspond to those on Figure 5

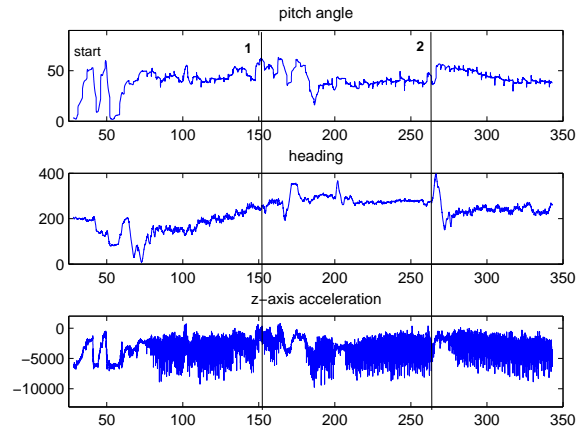


Figure 7: Time series corresponding to the trajectory in Figure 6. The heading and pitch of the device are shown at the top. A pitch angle of 50 degrees corresponds to a prediction of $\sim 50\text{m}$. At the bottom, the acceleration of the device along one axis is plotted, showing the footsteps during navigation. Numbers correspond to those on Figure 5.

5.1 Modality Scheduling

In the current prototype version, audio and vibrotactile feedback are used for the display of future possibilities. Extending this to a fully multimodal system which schedules modalities according to their attentional load and the interest the user shows in targets is possible in this model. Adjusting the target probabilities to (leakily) integrate over a period of time, each modality can be assigned to

a threshold in this integration process. This progression would lead from vibrotactile (most passive) for initial sensing, through audio (more detail, and more attentional load) to a visual display for high fidelity but high bandwidth display, when the user has exhibited significant consistent interest in one or more targets. This directly links the level of interest the user shows in a goal to the display bandwidth used for that goal, without requiring any explicit control on the part of the user. If the joint dynamics of information source and user continue to intertwine, the display of the mobile device becomes progressively higher bandwidth – an ongoing negotiation process (although the ‘negotiations’ can be happening at a range of timescales from almost instant, to over a period of minutes, depending on how the evidence accumulates).




Figure 8: A screenshot of the application in use on the PocketPC device. Predictions are shown in red on this display.

6. ORIENTATION TRACKING

An experiment was conducted to study the effect of the uncertain audio display in a more controlled environment. In this setup, participants stood still and had to acquire targets arranged around them by turning the device, which measured orientation with magnetometers. The participants were not required to move. Acquisition occurred when their measured heading remained sufficiently close to the heading of the targets for a period of time. No GPS signal was used; however positional noise was simulated to produce the effect of a poor quality GPS fix.

Four cases were examined in the experiment: mean display, without additional noise; mean display, with additional noise; uncertain display, without additional noise; and uncertain display with additional noise. However, only the with noise cases are used in the following analysis.

6.1 Experimental Details

Five targets were laid out in the space, each of which had to be acquired three times for each condition (fifteen acquisitions per condition). Acquisition was considered to occur when participants maintained the heading measured by the device within a funnel of 14.03 degrees for 5.4 seconds. Leaving this zone caused the count-down timer to pause until the participant re-entered the capture zone. The targets were arranged in an arc of from approximately $-\pi/2$ to $\pi/2$, at a distance of 71 meters. Target positions were fixed

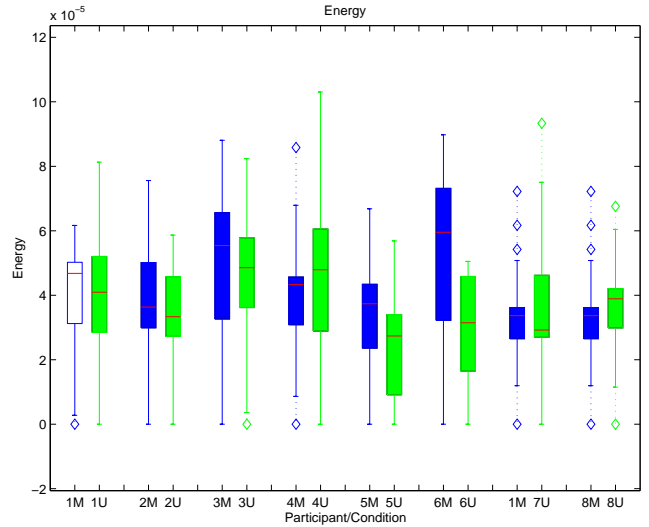


Figure 9: Low frequency energy in the heading signal for each condition, in the portions of acquisition after 63% of the error has been removed. Boxplot shows distribution of energy for each acquisition (blue=mean, green=uncertain). Energy required is reduced for six of the eight subjects with the uncertain display. Energy is computed as $\frac{1}{T} \sqrt{\sum_{t=0}^T \left(\frac{dx}{dt}\right)^2}$.

throughout the trial. Sporadic noise (Gaussian distributed with 4m standard deviation) was used to shift the position of targets in the noisy cases. Noise occurred as steps updated every three seconds, resulting in a square wave like pattern similar to that of true GPS noise.

Heading data was filtered with a low-pass filter, with -3dB rolloff at 8Hz before being displayed and recorded. This eliminates most of the tremor signal (10–12Hz) from the sensed signals, and thus removes significant irrelevant noise from the feedback. The heading data and acquisition times were recorded. The experiment was within-subjects with a counterbalanced presentation order. Eight participants took part in the experiment.

6.2 Results

The mean case, for both the noise and no-noise cases, generally requires more time for acquisition than the uncertain display. Figure 10 gives a boxplot of the time for each acquisition (per participant), illustrating the distribution of timing in each of the cases. Figure 9 shows the energy of the heading signal for each condition and participant (in the low frequency 0.1–2 Hertz band) after 63% of the error has been removed (63% is conventionally used in control theory). There appears to be some reduction in scanning activity in this band for some participants, although the acquisition criterion may have led to successful capture even without significant feedback, leading to anomalous cases where less energy had to be expended. Figure 11 shows a boxplot of the variance of the error between target heading and device heading. There is a significant reduction in the uncertain case as compared to the noisy case. Large deviations from target are less likely when the uncertain display is employed.

Figure 12 shows a typical time series from one participant for the mean and uncertain (with noise) conditions. There is noticeably more searching activity in the mean case, where the participant overshoots the target and has to search back. Figure 13 shows a histogram of error (for the same participant) in the region after the error has been reduced by 63%, for both the mean noise and

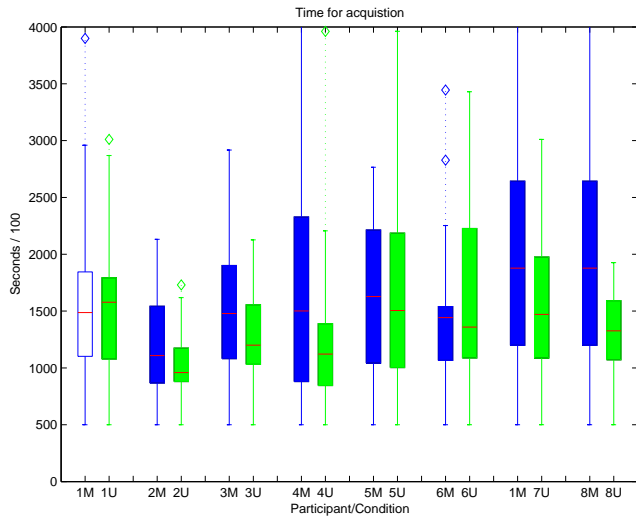


Figure 10: Boxplot showing the target acquisition times in each case. Time to complete is reduced in seven out of the eight cases with the uncertain display.




Figure 11: Boxplot showing the variance of the error during acquisition, for each condition. There is a visible reduction in the variability of the error with the uncertain display; large deviations are less common.

uncertain noise cases. The mean noise cases leads to a distribution of error with heavier tails (more variation during the final stages of acquisition). This is compatible with the variance of error plots in Figure 11. Figure 14 shows a phase plane plot of acquisitions for the same participant. The character of the acquisition is quite clearly altered between the cases, with more variability in the mean case.

6.2.1 Discussion

The results support the hypothesis that the uncertain display requires less effort and results in more stable behaviour. However, the results would have almost certainly been stronger had the selection mechanism been less susceptible to “random” selections. The capture zone for acquisition was over-generous in this experiment, under-penalizing the mean case.

Subjective comments from participants suggest that they felt the targets were larger in the uncertain case than the mean case. They




Figure 12: Heading time series in the mean and uncertain noisy cases for one participant (3). More scanning behaviour is visible in the mean case.




Figure 13: Histogram of error in the mean and uncertain noisy cases for one participant (3). Larger deviations in error are more common in the mean case.

also apparently felt less “in control” in the uncertain case, despite performing better under these conditions. This is unsurprising given the unfamiliarity with ambiguous displays. Participants also noted no change in difficulty between the uncertain case where noise was applied and where no noise was applied; however they noted that the mean case was significantly harder when the artificial noise was applied.

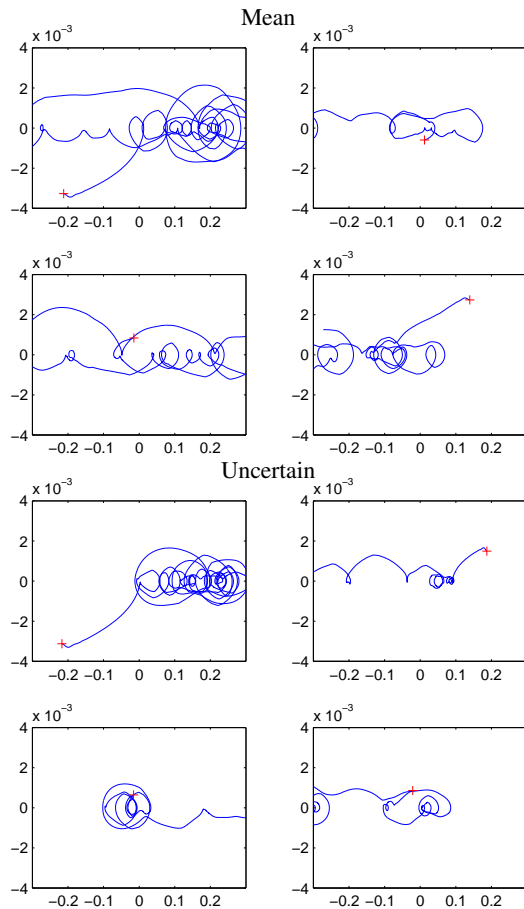


Figure 14: Phase plane plots (error (x axis) against derivative of error (y axis)) of the first four target acquisitions for one participant (3). Top four plots show acquisitions in the mean case, bottom four relate to the uncertain case. Red “+” symbols mark the end point of the trajectory. The acquisition is noticeably noisier in the mean case.

7. CONCLUSIONS AND FUTURE WORK

This paper has demonstrated that the probabilistic, negotiated interaction techniques can be applied effectively to the mobile GPS navigation problem. The Monte 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 evidence of regularised behavior when appropriate display of uncertainty is applied (see the variance plots in Figures 11 and 13). This display ameliorates the effects of noise and inaccuracy of the GPS system; it does not lead to irregular and unnatural movements, as naïve navigation software can do. The navigation systems also demonstrates how interactive sonification of the exploration process can produce a navigation system which can be used eyes-free, and allows the user to bring their sensorimotor systems into the interaction with an augmented environment. It facilitates ‘active perception’ of the environment, drawing information out of the environment as the user needs it.

7.1 Outlook

There is enormous scope for extension of the ideas presented in this paper. Introduction of active selection into the interface (via the pointing-without-a-pointer technique [13]) can be used to permit active focusing of the predictions onto specific goals. This tech-

nique allows simultaneous world-space and goal-space interaction. More sophisticated modelling of user and GPS sensor behaviour can be used to improve the quality and accuracy of the predictive system. The audio and vibrotactile feedback is, at the moment, relatively simple, and there are many degrees-of-freedom in the display which could be used creatively. The model-based sonification techniques of Hermann *et. al.* [1] can be used to display more complex information about the predictions. Very rich information about the environment can then be displayed; for example, aspects of virtual annotations placed in the world, based on probabilistic language models of their content.

Entities which might attract an interactor could include physical ones, such as places or other people; physical-digital artefacts like routes for navigation, or traces left by other users, and purely digital resources such as in-context information resources or messages. In the *negotiated interaction* approach, the user explores the possibilities by directly engaging with these entities, probing and playing with them, and gaining more of a sense of the context during these interactions.

In [7], McCullough writes about the need to simultaneously engage both a human’s brain and hands, that media have to be dense enough to give the impression of a universe of possibilities. We believe that the continuous, stochastic simulation approach used in this paper leads naturally to the sort of organic, rich interaction desired by McCullough, but with a sound technical foundation.

Acknowledgements

We are grateful for support from: HEA project Body Space, SFI grant 00/PI.1/C067, SFI BRG project, Continuous Gestural Interaction with Mobile Devices, EPSRC project Audioclouds, IST Programme of the European Commission, under PASCAL Network of Excellence, IST 2002-506778. This publication only reflects the views of the authors.

8. REFERENCES

- [1] T. Hermann and Ritter H. Listen to your data: Model-based-sonification for data-analysis. In *ISIMADE '99*, 1999.
- [2] M. Jones, S. Jones, G. Bradley, and G. Holmes Navigation by Music: an Initial Prototype and Evaluation *Proceedings of the International Symposium on Intelligent Environments* 95–102, 2006.
- [3] S. Holland, D. R. Morse, and H. Gedenryd. AudioGPS: Spatial audio navigation with a minimal attention interface. *Personal Ubiquitous Comput.*, 6(4):253–259, 2002.
- [4] S. Hughes, I. Oakley, and S. O’Modhrain. MESH: Supporting mobile multi-modal interfaces. In *ACM UIST'04*. ACM, 2004.
- [5] K. P. Körding and D Wolpert. Bayesian integration in sensorimotor learning. *Nature*, 427:244–247, 2004.
- [6] David J. C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, 2003.
- [7] Malcolm McCullough. *Abstracting Craft: The Practiced Digital Hand*. The MIT Press, 1998.
- [8] I. Oakley, J. Ångeslevä, S. Hughes, and S. O’Modhrain. Tilt and feel: Scrolling with vibrotactile display. In *EuroHaptics 2004*, 2004.
- [9] L. Smith. The maintenance of uncertainty. In *Proc International School of Physics ‘Enrico Fermi’, Course CXXXIII*, pages 177–246. Societa Italiana di Fisica, Bologna, Italy., 1997.
- [10] A. Steed. Supporting mobile applications with real-time visualisation of GPS availability. In *MobileHCI 2004*, number 3160 in Lecture Notes in Computer Science, pages 373–377. Springer-Verlag, 2004.
- [11] S. Strachan, P. Eslambolchilar, R. Murray-Smith, S. Hughes, and S. O’Modhrain. gpsTunes: controlling navigation via audio feedback. In *MobileHCI '05: Proceedings of the 7th international conference on Human computer interaction with mobile devices & services*, pages 275–278, New York, NY, USA, 2005. ACM Press.
- [12] N. Warren, M. Jones, S. Jones, and D. Bainbridge. Navigation via continuously adapted music. In *CHI '05: CHI '05 extended abstracts on Human factors in computing systems*, pages 1849–1852, New York, NY, USA, 2005. ACM Press.
- [13] J. Williamson and R. Murray-Smith. Pointing without a pointer. In *ACM SIG CHI*, pages 1407–1410. ACM, 2004.
- [14] J. Williamson and R. Murray-Smith. Sonification of probabilistic feedback through granular synthesis. *IEEE Multimedia*, 12(2):45–52, 2005.