Modelling human spatio-temporal behaviour: A challenge for location-based services

David Mountain and Jonathan Raper
Geographic Information Science Group, Dept Information Science,
City University, Northampton Square, London, EC1V 0HB
{dmm, raper}@soi.city.ac.uk

Abstract As phones become "location-aware" over the coming years there is an opportunity to not only ascertain a user's physical location when they make a query, but to store and analyse a history of previous locations. Given user approval, analysis of the spatio-temporal characteristics of the user's current and recent movement can help improve the response to information requests. The EU funded Hypergeo project (www.hypergeo.org) is developing a prototype tourist information system to provide individuals with personalised, location dependent information on the move via handheld computers and connections to the mobile network. As part of this project the GI Science Group at City University have developed a prototype device that can send a user's location over the GSM network and have accumulated large archives of spatio-temporal data for several individuals. These archives are now being analysed in an attempt to produce automated approaches to summarising individual position histories. This will help to derive higher level information from the spatio-temporal data such as enclosing rectangles of typical movement and estimates of current transportation mode.

1. INTRODUCTION

This paper attempts to describe the direction of research conducted by the Geographic Information Science Group and the development that has been achieved to date. Background is given in the form of details about the project funding the research and other research conducted in this area. The software development and algorithms developed for the project are then discussed. The paper considers the project from the perspective of the spatio-temporal (space and time referenced) data sets that are generated more than the project a tourist information system.

2. HYPERGEO PROJECT BACKGROUND

2.1 Project Background and Objectives

The European Commission Information Society Technologies funded Hypergeo (www.hypergeo.org) aims to develop a prototype tourist information service for the delivery of timely location and user sensitive information taking into account new information technologies, geographic information and communication infrastructures.

Several diverse technologies have been combined to implement Hypergeo as a functioning system. Due to the diverse and multi-disciplinary nature of the challenge facing the Consortium, a distributed architecture has been adopted with different Institutions designing components focusing on specific issues. To permit user-friendly access to information, advanced request analysis and natural language processing has been used (Artificial Intelligence Information Analysis Laboratory). In order to take advantage of existing web resources, mechanisms to access distributed information without strong knowledge of its structure have been developed. Personalisation has been achieved through user profile management and the system is location aware through a geolocation system based on the Global Positioning System (GPS).

The Hypergeo system architecture consists of these discrete components linked over the Internet by various middleware protocols including CORBA and Java Remote Method Invocation (RMI). Each component can operate independently and communicate to other components through message sending and updating individual user profile data stored on a central server.

2.2 Location-aware devices and the spatio-temporal data explosion

One element of motivation for the Hypergeo project is the increasing number of technologies and devices that can be used to discern the location of an individual on the earth's surface. This technological advance has been pushed by legal imperative in the USA (the Federal Communications Commission E911 initiative) and potential commercial exploitation of location-based services in Europe. It is these "location-aware" devices that are heralded as the technological driver behind location-based services. Since the time of the position can also be recorded, a spatial and temporal record can be kept for any user storing where they have been and when they were there. A brief discussion of the relevant two main positioning technologies, satellite and terrestrial based solutions, is presented here.

The most viable satellite solution at present is the US global positioning system (GPS), which provides free access to accurate geospatial data through inexpensive handheld positioning devices. Sources of error include atmospheric effects, internal clock errors and lack of satellite availability; the GPS receiver also requires a line of sight to the satellites. While views of the sky are more likely to be uninterrupted in rural areas, in built up areas "urban canyons" can result in only a small portion of the sky being visible (Mountain 2001, in press). When a fix can be attained accuracy is usually very good, within 30 metres, 95 percent of the time.

Most terrestrial positioning solutions exploit the existing mobile telecommunications infrastructure to gauge user location. This is done routinely to the resolution of the mobile cell (cell ID) for the purpose of call routing. More sophisticated handsets or dedicated servers can increase the accuracy of the positional fix, usually with some degree of network modification (Swedberg, 1999). The most popular solutions to date have involved calculating the transmission time between base stations and mobile handset to triangulate the user's position. Processing occurs either on the handset or a network server. Other solutions have included calculating the angle of arrival of signals from two or more base stations and comparing a device's "multipath fingerprint" (the supposedly unique interference signature from every position in the network) with that measured for known locations in the field (Birchler, 1999). None of the proposed terrestrial solutions to date has provided accuracy equal to satellite based solutions.

The proliferation of these location-aware devices is likely to lead to a massive increase in the volume of spatio-temporal data recording the everyday movements of individuals. Knowing traditional GI systems limited ability to handle the temporal dimension (Langran 1992), new data models may be required to store and analyse this spatio-temporal data explosion.

2.3 The phone's evolution into mobile information platform

Further motivation for the Hypergeo project comes from the mobile telecommunication networks transition from being primarily concerned with the exchange of voice data to allowing the exchange of a variety of information types. In the UK and Europe, existing 2nd generation phones will soon be superceded by "2.5G" (increased bandwidth and "always on" connections) then third generation (3G) after sometime after 2003 depending upon available finance and technological barriers (UMTS Forum).

In parallel with change in the telecommunications industry is a shift in the portability of computers in the form of personal digital assistants (PDAs) or handheld computers. Increasingly these devices have network connections through GSM modems, infrared links and in the near future, radio links to networks. These technological steps have led to a blurring of the distinction between handheld computers and mobile phones (Symbian Consortium).

2.4 City University involvement

The primary contribution of the Geographic Information Science Group, City University, to the Hypergeo project is to develop components to manage the geolocation of users. The first element of this is to develop a practical, working, pan-European geolocation device then transmit and store this individual spatio-temporal data in a logical way. The second element is analysis of the position history looking for trends and summarising the data into information that can model the user's spatio-temporal behaviour. These components are discussed in section 4.

3. RESEARCH BACKGROUND

This section is intended to give a brief overview of previous and present research in the field of spatio-temporal analysis; it is not intended to be a comprehensive account.

3.1 Previous research

Modelling phenomena against a spatio-temporal framework owes much to scientists researching the physics of space and time in the first half of the twentieth century. Minkowski visualised space and time as a light cone; a sharply delineated limiting cone defining a boundary between past and future accessible locations (figure 1). All movement must be contained within this region in space-time since nothing can travel faster than the speed of light (Raper 2000).

The Lund School (led by Hagerstrand) of Sweden adapted these ideas as "Time Geography", a spatio-temporal framework against which model human activity. An individual's movement can be described by a space-time path whose gradient represents velocity across the 2-dimensional surface. A vertical line represents no movement and increasingly horizontally sloped lines show faster velocities. Many of these ideas have been formalised and developed by later researchers such as Parkes and Thrift (1980) and Gollege and Stimpson. (1997).

3.2 Current research

While light cones define the absolute physical limit of accessibility, when modelling human activity the constraining cones (or prisms) tend to be defined by more mundane factors (e.g. walking/transport speed or out-of bounds areas). The prism generated describes the maximum spatial extent, or potential path area (PPA), the individual can access at any given time.

Forer (1998) suggested plotting various spatial entities as solid cubic cells occupying space in spatial and temporal dimensions. This shows the individual as a consumer of time space. Discrete structures such as buildings will be spatially static and will always occupy the same two-dimensional space. Individuals will form a single wandering line that will be constant in size but move spatially through time.

More recently Erwig (1999) identified a key set of generic operations such as attribute value at specific times/locations, max/min and start/stop times for spatio-temporal analysis. Imfeld (2000) has extended radial distance functions for analysing spatio-temporal behaviour in animals; a technique which promises to also be suitable for modelling human behaviour.

4. SOFTWARE DEVELOPMENT

Two components have been developed by the GI Science Group to fulfil the commitment to the project. The first, the Position Tracker, is a hardware and software solution for keeping track of the user in real time. The second, the Location Trends Extractor, is a software only solution for analysis and summary of the user's history of movement.

4.1 Position Tracker (PT)

The Position Tracker component (figure 2) comprises the tracking hardware the user must carry with them to relay their location over the GSM network and the software to receive and interpret this data, then update their user profile with new positional information.

The tracking hardware (or 'rig') was a proprietary solution since the Hypergeo project is a pan-European project and there was no viable mobile positioning technology that would work on any network. The Hypergeo solution is a Garmin GPS device connected to a Nokia 9110 communicator via the serial port. A program on the Nokia phone sends the GPS data over the GSM network to a GSM modem where it is parsed by the Position Tracker software. The resulting time- and space-stamp can then be appended to the position history for the relevant user profile (identified by their mobile number). It is hoped that commercial geolocation technologies will replace the present hardware solutions in the near future.

4.2 Location Trends Extractor (LTE)

The Location Trends Extractor is a software solution that allows visual analysis and summary of the potentially vast point history generated by the Position Tracker. This single software component communicates with a central server on which the user profile is stored. The LTE downloads the point history and other relevant data from the profile and returns summaries of the user behaviour that can be used by other components to tailor the information sent to the user. There are three main areas of development for the LTE. Firstly the initial visualisation of results to help the developer understand the data better. Secondly developing summaries of the user's behaviour that can be of more use in tailoring information sent to the user than the point history itself. Thirdly automating this process so that natural breakpoints in the data are identified and summaries returned for sections of the history.

4.2.1 Visual Analysis

Before any summarisation or automation can be attempted, the developer must nurture an understanding of the data. This scientific visualisation (MacEachren and Kraak 1997) involved plotting the point history against time and space (figure 3). It became clear that the key to summarising the user behaviour lay in identifying breakpoints in the data, such as significant spatial or temporal jumps, or sudden changes in speed or direction. These breakpoints can then be used to demarcate 'episodes' in the user's position history. An episode should contain a relatively homogenous spatio-temporal period, such as a single car journey (see section 5.1).

For this reason the visual interface of the LTE was adapted to allow interactive subset selection, based on spatial, temporal or attribute criteria (see figure 4). Several subset functions may have to be performed to whittle the point history down to a single episode. For example a spatial subset may limit the history to a single town, then a temporal subset may be used to limit the history to a single unbroken period of collection. Finally displaying the subset shaded by direction may reveal that the subset is a 'to and from' journey along the same route; shading the temporal history by direction and performing a final subset operation can deliver a single episode.

4.2.2 Spatio-Temporal Behaviour Summary

Having identified individual sections of the position history, these sections can be summarised in a variety of ways. The algorithms for summarising the data are described in more detail in section 5, spatio-temporal algorithms.

4.2.3 Automation

The final stage of development involves automating the summarisation of user behaviour so the LTE can function unattended as a server-side application. First the large history of points for the user must be broken down into discreet temporal sections, each section representing an episode. Analysis can then be performed upon each episode and the results returned to the server.

Usually for a user, analysis should be confined to the most recent points, since the back history of points will already have been analysed and results stored. The problem then arises of setting tolerances for breakpoints for incoming points. A sudden jump in speed may indicate a change in transport or may be a rogue result; the Global Positioning System can occasionally provide very spurious results and the most extreme inaccuracies have been known to be wrong by thousands of kilometres.

Once the end of a new episode has been identified, summary information can be stored. This should include mean speed, sinuosity and other traditional parameters. Beyond this simple analysis of an episode, two further algorithms have been developed which are discussed in more detail in section 5. An envelope is the minimum enclosing rectangle for a set of points. For an episode, a temporal envelope is defined; that is the spatial limits for a set of points that are defined by a minimum and maximum time. Conjectured activity is an algorithm that estimates the user’s transport type given a set of points.

5. SPATIO-TEMPORAL ALGORITHMS

A number of algorithms are in development to provide summaries of the point history data for tailoring information returned to the user from location-based services. The first algorithm, to derive episodes, breaks the point history into a series of discreet sets based on minimum and maximum temporal limits. The remaining two algorithms, envelopes and conjectured activity, then summarise each episode by defining the enclosing rectangle and likely transport use respectively. The Location Trends Extractor software has been developed specifically to implement and visualise these algorithms.

5.1 Episodes

An episode represents a discreet time period for which the user’s spatio-temporal behaviour was relatively homogenous. The intention is to split up the huge data set that may contain months of data into sets of points that can be analysed more effectively. There are clear reasons for doing this since it attempts to use the natural breakpoints in the data set itself to create order within the position history. An episode may be a single car journey, a walk in the countryside, a spatially static period, such as time spent angling; any period for which the data is relatively stable.

The key to defining episode is finding breakpoints in the data. The main indications of breakpoints are listed below.

5.1.1 Temporal/spatial jumps

A large time period for which no data is collected suggests a breakpoint in the data. If after a period of collection where points were received every 2 minutes there is a period where no points arrive for 1 hour, this would suggests a natural breakpoint. The case is also the same for large spatial jumps, however, if a large spatial jump appears without a large temporal jump (e.g. moving several 100km in a few seconds) then this suggests erroneous data collection. Therefore there is no need to look for large spatial jumps directly since these will be identified by temporal gaps or rapid changes in speed.

Temporal jumps can be expected in the data since the geolocation technology utilised is the global positioning system (section 2.2) which will not work indoors or where a sufficient portion of the sky cannot be seen. The data collected at present only therefore records periods when the user is outside hence time in the home or office is recorded as a temporal jump. Utilising terrestrial positioning techniques which may become available through the mobile phone network in the future may lead to more ubiquitous data collection, where the user’s location is recorded, albeit at a lower spatial resolution, whenever the user’s phone is on and has network coverage.

5.1.2 Rapid changes in direction/speed

An indication that the user may have ceased one activity and begun another is a rapid and sustained change in direction or speed. An example of a breakpoint of this kind would be a change in transport from car to on foot, or a reversal in direction, such as might be seen on a parent’s return home after dropping their children at school.

5.2 Envelopes

Envelopes are minimum enclosing rectangles that specify the spatial and temporal limits for the point set; that is maximum and minimum latitude and longitude and the earliest and latest times of collection. This established technique is a simple but effective way of limiting search space, however in this case the temporal dimension has been added hence the term ”rectangle” was no longer considered adequate. Several different kinds of envelopes have been identified for analysis of the point set and despite their names, all store both the spatial and temporal limits for the points.

5.2.1 Spatial Envelopes

Here, maximum and minimum latitude and longitude define the spatial envelope in the same way as minimum enclosing rectangles. An envelope can be defined by clicking and dragging on the display of spatial points, entering maximum and minimum limits, or taking the extremes of the point set as the envelope. This can be stored as a surrogate for the points themselves to give a rough idea of ‘where’ and ‘when’ that point set occurred.

5.2.2 Temporal Envelopes

A temporal envelope stores the same information as a spatial envelope however it is defined by temporal maxima and minima rather than spatial ones. Limits can be specified interactively by the Location Trends Extractor software by clicking and dragging on the temporal interface. Each episode stores a temporal envelope, the spatial and temporal limits for a set of points by a start and end time.

5.2.3 Map Display Envelopes

An application of an envelope is holding the bounding coordinates for map display for the real time delivery of an orientation map for a mobile user. The map scale needs to be context sensitive and this can be achieved by assessing the user’s speed and direction. A user travelling at very slow speed, possibly associated with ‘on foot’ or ‘low speed motor’ transport behaviour, will require a high resolution map, displaying minor roads and pedestrian routes, centred on their current location. A user travelling at high speed will require a coarser resolution map, centred ahead of their current location; either a direct extrapolation from their current direction or a more complex algorithm based on the transportation network.

5.3 Conjectured Activity

Conjectured activity is an attempt to summarise transport behaviour for a set of points. Acknowledging that it is impossible to cater for every transportation type and to be sure of any prediction that is made, transport types have been generalised to four fuzzy classes, and any set of points will have a membership to each between zero and one. The sum of all four classes for any point set will always be one. The four transport activities are listed below.

On foot – This transport type represents slow, self powered transport and static activity. It is characterised by low speeds and high sinuosity. For this transport type the user is expected to have a small range of movement, but be free to move as easily in one direction as another since they will be less confined by transportation networks.

Low-speed motor – This type represents motorised transport at low speed on minor roads (such as cars in suburbs). Speed will be higher than the on foot class and the route possibly less sinuous. When suggesting services for the user, the range will be greater and the user will still have relative freedom to explore in most directions.

High-speed motor – Here the transportation is fast moving, ground based vehicles moving at high speeds in relatively straight lines (such as vehicles on motorways or trains). The user will be heavily confined by the transportation network so ideally this should be considered when returning results regarding location-based services. Direction also plays a greater part here since if the user is moving rapidly South, there is no need to send results that are based to the North of their present location.

Flight – Here the user is entirely confined by the aviation network. The widest field of searches should be used.

6. CONCLUSIONS

The development of "location-aware" phones could create vast, dynamic spatio-temporal data sets for every individual who carries a phone. This may lead to the development of new data models and forms of analysis to find spatial and temporal trends. Location-based services have thus far tended to concentrate on using a mobile carrier's immediate geographic location in isolation but there is merit in using the position history to tailor the results from requests for information further. The work builds on research into Time Geography.

In the course of this research the GI Science group at City University have developed two software solutions and several new algorithms for drawing meaning from an individual's position history. Much work remains to be done in this area. Future work should focus on moving from analysis of the past history to prediction of where the user may be in the future; considering interaction between users, on a one to one scale and *en masse* such as traffic distribution; privacy issues, finding ways to remove the raw data whilst preserving information that may be of use in tailoring queries and temporal scale, generalising older results whilst preserving detail for the recent history.

REFERENCES

Artificial Intelligence Information Analysis Laboratory, Thessaloniki University, http://poseidon.csd.auth.gr/ org (visited 27 July 2001).

Birchler, M. (1999) "E911 Phase II Location Solution Landscape". Federal Communication Commission E911 site. http://www.fcc.gov/e911/mottutorial.ppt.

Erwig, M., Güting, R.H., Schneider, M., and Vazirgiannis, M. (1999) "Spatio-temporal data types: An approach to modelling and querying moving objects in databases," GeoInformatica 3(3).

The European Commission IST Work Programme http://www.cordis.lu/ist/wp2000.htm (visited 27 July 2001).
Federal Communication Commission E911 site. http://www.fcc.gov/e911/mottutorial.ppt.

Forer, P. (1998) "Geometric approaches to the nexus of time, space and microprocess: Implementing a practical model for mundane socio-spatial systems". In Egenhofer, M.J. and Golledge, R. (Eds) "Spatial and temporal reasoning in geographic information systems", Oxford, OUP, pp171-90.

Golledge, R., G. and Stimson, R., J. (1997) "Spatial behavior: a geographic perspective", Guildford.

Hypergeo main site, http://www.hypergeo.org (visited 27 July 2001).

Imfeld, S. (2000) Time, points and space - towards a better analysis of wildlife data in GIS. Ph.D. dissertation. http://www.geo.unizh.ch/~imfeld/diss/ (last visited 27 July 2001).

Langran, G. (1992) *"Time in Geographic Information Systems"*. Taylor and Francis, London.

MacEachern, A. M. and Kraak, M. J. (1997) "Exploratory Cartographic Visualisation: Advancing the Agenda". Computers and Geosciences, 23 (4), 335-343.

Mannila, H., Toivonen, H. and Verkamo, A., I. (1997) "Discovery of frequent episodes in event sequences", Department of Computer Science, Series of Publications C, Report C-1997-15, University of Helsinki.

Miller, H. (1991) "Modelling accessibility using space-time prism concepts within geographic information systems", International Journal of Geographical Information Systems, 5, 287-303.

Mountain, D. M. and Raper, J. F. (2000) “Designing geolocation services for next generation mobile phone systems”, AGI 2000 Conference Proceedings.

Mountain, D. M. and Raper, J. F. (2001) “Spatio-temporal representations of individual human movement for personalising location-based services”, GISRUK 2001 Conference Proceedings.

Mountain, D. M. (2001) “Information anywhere: User friendly location-based services”, GI News Guide to LBS, Navigation and mobile commerce. April/May 2001, pp10-22.

Mountain, D. M. and Raper, J. F. (2001) Positioning techniques for location-based services: characteristics and limitations (in press).

Parkes, D. N. and Thrift, N. J. (1980) Times, Spaces and places. Chichester, John Wiley.

Raper, J., F. (2000) “Multidimensional Geographic Information Science”, Taylor and Francis.
Symbian site, http://www.symbian.com (visited 27 July 2001)

Swedberg, G. (1999) “Ericsson’s mobile location solution”. Ericsson Review No.4,

Theriault, M., Claramunt, C., And Villeneuve, P., Y. (1999) “A Spatio-temporal Taxonomy for the Representation of Spatial Set Behaviours”, Spatio-Temporal Database Management, International Workshop STDBM'99, Proceedings, p1-18, Springer.

UMTS Forum, http://umts-forum.org (visited 27 July 2001).

Figure 1. Minkowski's light cone

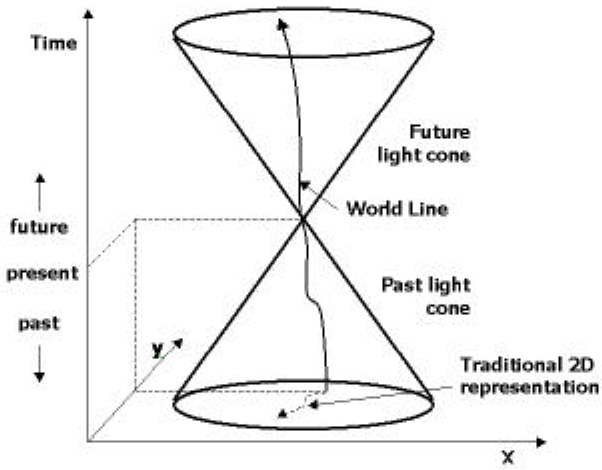

Figure 2. Position Tracker Interfaces

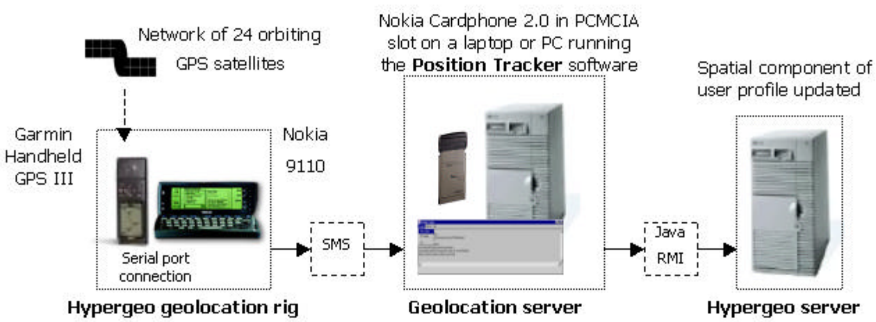

Figure 3. The Location Trends Extractor Interface
There are three main GUIs. A message board for user assistance (top), a spatial view (left) and a temporal view (right)

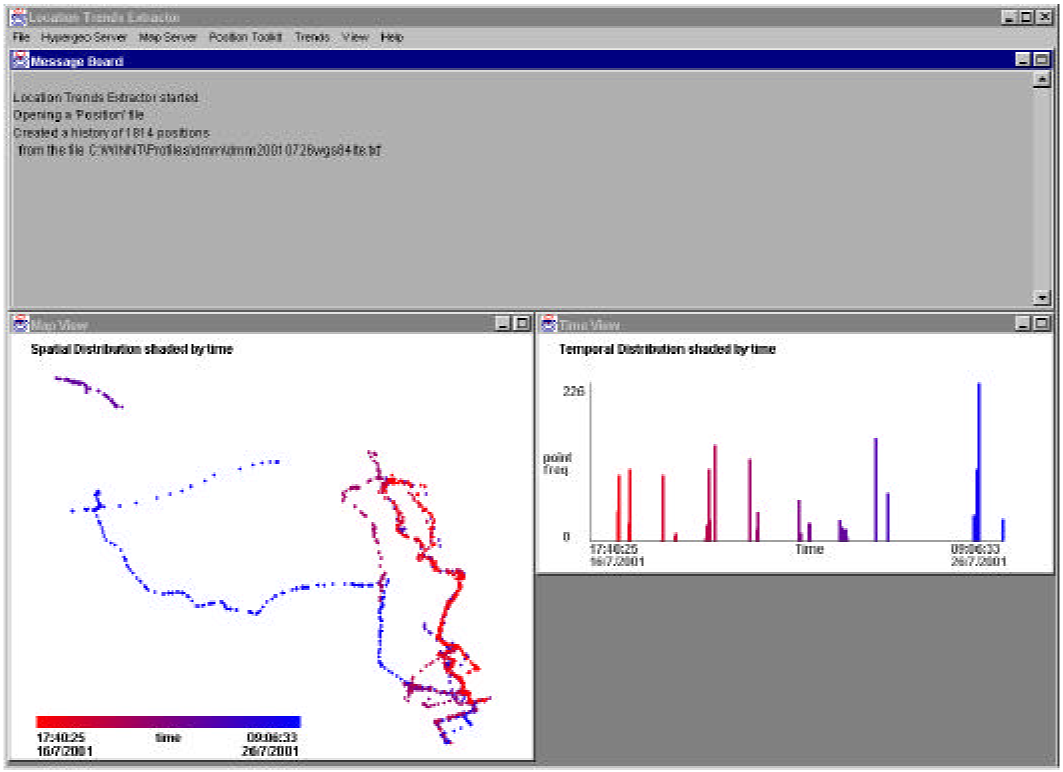

Figure 4. Successive subsetting reducing a complex spatio-temporal history into a single episode