TRACKING AND ANALYSIS OF THE MOVEMENT OF PEDESTRIANS

Sarah Clayton
Dr Neil Urquhart
Prof Jon Kerridge
Napier University

ABSTRACT

Pedestrian movement is unconstrained. For this reason it is not amenable to mathematical modelling in the same way as road traffic. Individual pedestrians are notoriously difficult to monitor at a microscopic level. This has lead to a lack of primary data that can be used to develop reliable models.

Although video surveillance is cheap to install and operate, video data is extremely expensive to process for this purpose. An alternative approach is to use passive infrared detectors that are able to track individuals unobtrusively. This paper describes the use of a low cost infrared sensor for use in tracking pedestrians. The sensor itself, manufactured by a British company, is designed to count people crossing an arbitrary datum line. However, with the development of additional software, the functionality of these sensors can be extended beyond their original design specification. This allows the trajectories of individual pedestrians to be tracked.

Although the field of view of each sensor is relatively small ($4m^2$), trajectories can be matched in software as they leave the field of view of one sensor and enter that of another in real time. In this way, sensors can be deployed in an array and individual pedestrians can be tracked across a wide area. The sensors have been deployed in a busy indoor corridor, adjacent to the entrance to a busy café.

The data collected is then to be used to develop a model of the way in which pedestrians move through a space. Common behaviours, such as obstacle avoidance, are captured. These data can then be used as the basis of simulations that use artificial intelligence techniques in a multi-agent software environment.

1. INTRODUCTION

The study of pedestrian dynamics is concerned both with pedestrians in normal situations and in evacuation situations. While the latter has always been subject to intense study, governments across the world are now placing pedestrians at the heart of their transport policies. The aim is not only to facilitate pedestrian movement at transport interfaces, such as rail stations, but also to encourage this as a mode of transport in its own right. An example is the development of the London Pedestrian Model. This is the result of collaboration between Transport for London (TfL) and the Central London Partnership (CLP), who commissioned the consultancy, Intelligent Space Partnership (ISP), 'to develop a model of total walking volumes for every node on the street network in Central London' (Desyllas et al., 2003). It arose from the Mayor's Transport Strategy (GLA, 2001), where a key aim was 'to make London one of the most walking friendly cities for pedestrians by 2015.'

Figure 1: Visibility analysis for Central London

Travel on foot is ubiquitous, and encompasses all human activities. However, our understanding of this form of transport hasn't advanced much since the work of Fruin (1971) on levels of service in the 1970s. The reasons for this are two fold. Pedestrian movement is unconstrained, unlike road traffic, with a great deal of variability in acceleration and direction. This is very difficult to model mathematically. The second difficult is in being able to gather data about people's behaviour as they travel through a space. Although a few studies have published models based on such data, the data sets used are small.

These difficulties are, however, increasingly being overcome due to the decreasing cost of computing resources and the emergence of new technologies allowing the gathering of primary data. These technologies range from the ubiquitous, such as video surveillance, to extremely sophisticated and expensive laser based scanners.

Macroscopic models are concerned with aggregate behaviours, and predictors for these. In the case of the London Pedestrian Model, mentioned above, metrics such as pavement width, proximity to London Underground stations and visibility analysis across the whole area of study are used to predict pedestrian behaviour and densities.

Microscopic models are at the other end of the spectrum, in terms of levels of detail, and are concerned with behaviour at the individual rather than aggregate level. Behaviour models are instantiated in autonomous agents that can act independently in the simulated environment.

Many of these simulations exist, each based on different models of pedestrian behaviour. The most popular model is the social forces model developed by Helbing (1991). Helbing drew comparisons between fluid dynamics and pedestrian flow, and established a number of partial differential equations describing the behaviour of agents in his model, that depend on attractive and repelling social forces, such as desire or dislike of proximity to others. Helbing's model has been widely adopted in part because it is straightforward to implement, but also because simulations are able to produce common emergent behaviour such as lane formation. However, there is some debate about the relevance of classical physics to human behaviour. This model has no basis in observed human behaviour.

Discrete choice models have also been employed in pedestrian simulations. Commonly used in economics, these are employed where two or more distinct choices are available. The assumption is of the maximisation of utility. Antonini et al (2004). employed this as the basis of their model, and found a high correlation between their agents' behaviour and gathered data. However, the data sets used were extremely small: 36 trajectories from a digital video sequence taken from Lausanne Metro station (Antonini et al., 2004), and 190 trajectories collected in Sendai, Japan (Teknomo, 2002).

As the examples above show, the lack of primary data is an extremely limiting factor in developing models of pedestrian behaviour. The reasons for this are the difficulties in collecting this data. Historically, this has been done by processing video data. However, as Willis et al. (2004) found, although this is very cheap to collect, it is extremely expensive and laborious to process, despite the very rich data produced.

In some instances, very low tech approaches have been taken. In one study, undertaken by Daamen and Hoogendoorn (2004), individual users of a train station were followed as they alighted or boarded trains and their walking speeds timed. At the opposite end of the spectrum, the Pedestrian Accessibility and Movement Environment LAboratory (PAMELA) at University College London employs side scanning laser, originally developed for the motor industry, to track pedestrians in a laboratory setting equipped with a reconfigurable environment.

All these approaches require a trade off between the equipment costs in collecting data, and the cost of processing it. While video data is cheap to collect, it is extremely time consuming and labour intensive to process. At the same time, side scanning lasers are extremely expensive and can only be used in a laboratory environment, despite returning high quality data.

2. DATA COLLECTION

The aim of current research at Napier is the development of a cheap, straightforward and unattended means of collecting data. At the heart of this effort is a passive infra-red sensor developed by a British company. The Irisys MkII sensor has a number of characteristics that make it suitable for this purpose, despite the fact that this is beyond its design specification. Normally deployed in retail environments, its purpose is as a people counter, counting the number of people that cross an arbitrary datum line.

This device has a number of capabilities that make it useful in this context. Its built in digital signal processor (DSP) performs a great deal of calculations on the information it receives, and returns the coordinates of targets it picks up as simple Cartesian coordinates. It is able to track and identify each individual target across its field of view and sends this information, at 30 millisecond intervals, down a serial connection to a PC. This frees us from the need to do any image processing in order to obtain this information. Placed on the ceiling of the area being observed, people moving beneath the sensor appear as ellipses. Algorithms on the sensors are able to discern these from the background, and return the centroid of the ellipse as well as other characteristics.

2.1. OPTICAL DISTORTION

However, as stated above, this use is beyond the design specification of the sensor, which means that a number of difficulties need to be overcome. The lenses of infra red devices are generally made of germanium, which is an expensive raw material. The quality of the lenses themselves is generally very poor. They introduce quite severe optical distortion to the data returned. Although this is common problem with lenses and the equation governing this type of distortion is well known and understood, finding means to correct it is a recurring theme in machine vision literature. The equation cannot be reversed analytically, and numeric means must be employed.

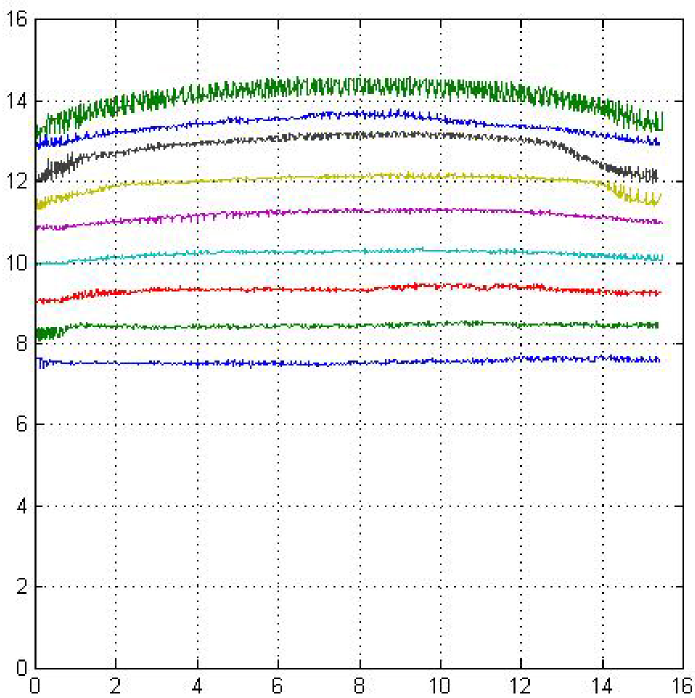

Figure 2: Effects of radial distortion

Figure 2 shows the effects of radial distortion. In contrast with the ‘ground truth’ where each of the above lines were ten centimetres apart, the lines in the data returned become more compressed the further they are from the centre of the lens. They also become increasingly curved, and suffer an increasing amount of noise in the data.

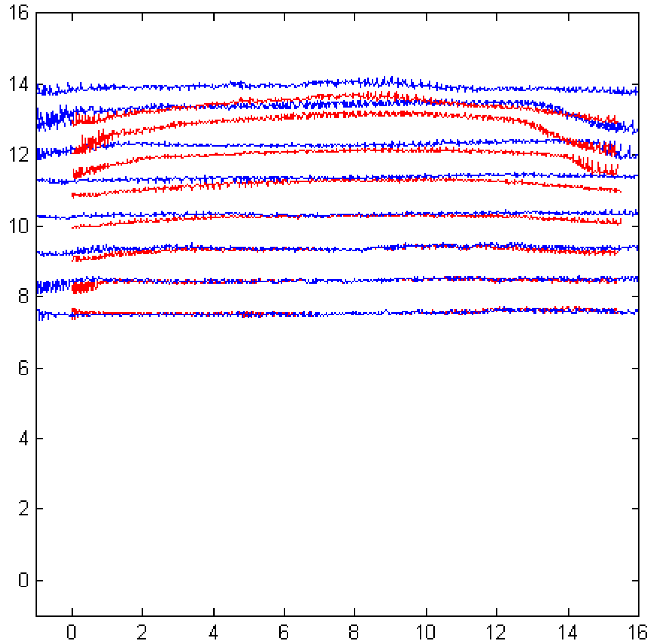

Figure 3: Correction of radial distortion

Figure 3 shows the same data after it has been corrected. Although there is still distortion at the edges, most of it has been mitigated. Most of the lines are now equidistant and straight.

The method for finding a solution to this problem had to ensure first of all that the lines recorded were absolutely straight. The infra red source also had to be moving; as the Irisys MkII sensor picks up changes in temperature compared to the background, any stationary infra red source would very quickly vanish from the returned data set. Calibration diagrams are often used to correct radial distortion in visible spectrum lenses, but this method wasn't suitable in this context.

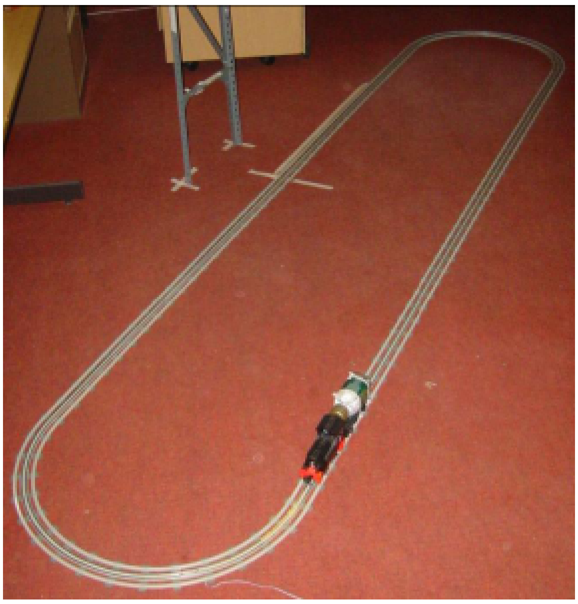
Figure 4: Test track

In order to ensure both a straight line and a moving target, a test rig was built to house the sensor, and a two metre length of Lego railway track was placed beneath it, orthogonal to the y-axis of the sensor's own coordinate system (Figure 4). An infra-red source was provided by a bottle filled with warm water and this was pulled along the trace by a Lego locomotive (Figure 5).

Figure 5: Lego locomotive with infra-red source

The collected data was then processed using Matlab, which produced the necessary correction parameter (Figure 3).

2.2. TARGET MATCHING

The second limitation of the Irisys Mk II sensor is the relatively small field of view, approximately 4m^2. This constraint means that in order to observe any significant amount of space, the sensors must be deployed in an array. The site used for experiments at Napier University is the corridor adjacent to a busy café at the Merchiston campus. This corridor leads off the campus foyer and is also the route to a number of teaching rooms.

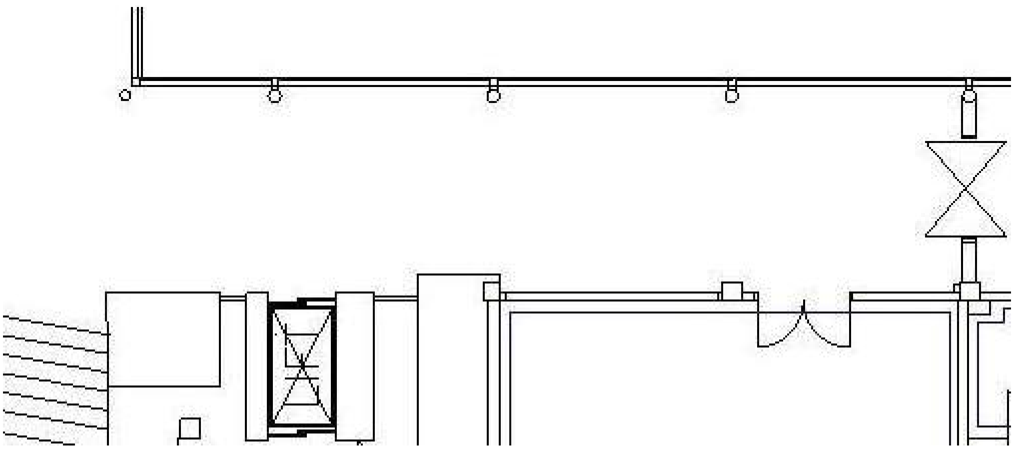

Figure 6: Floor plan of Apex corridor

In Figure 6 above, the café is situated to the right of the lift, through the double doors. The foyer is to the far left. The teaching rooms are reached through the double doors in the far right. Six sensors are deployed along the corridor, which is approximately 18 metres long and 4 metres wide.

Processing data from sensors in real-time is challenging, and requires the use parallel processing by the receiving software. This has been achieved using a Java API called JCSP that implements the principles of Communicating Sequential Processes (CSP) in the Java programming language. CSP is a theory of parallel processing that can guarantee that many of the problems inherent in parallel processing, such as deadlock and race conditions, do not occur. This makes parallel software relatively easy to produce, and extremely robust once it is running.

A prototype system was developed as an exhibit for the Edinburgh Science Festival, where children were encouraged to walk their names underneath two infra-red sensors mounted above them on a specially built rig. This was an extremely popular exhibit, and remarkably, given the traffic that past through during the festival, there were no software failures. The software scaled up extremely well from two to six sensors when an extended version was deployed in the corridor on the Napier campus.

Two computers are required to gather and process the data returned by the sensors. The first is dedicated to reading the serial communications from the sensors and parsing the raw data, this is then transmitted down an Ethernet connection to a computer that performs matching. In this way, a clear view of the movement of pedestrians in the area is built up.

An example of the information that is received is shown in Figure 7. Both targets originated in the foyer. Target 34 is proceeding to the teaching rooms whereas target 35 is entering the café.

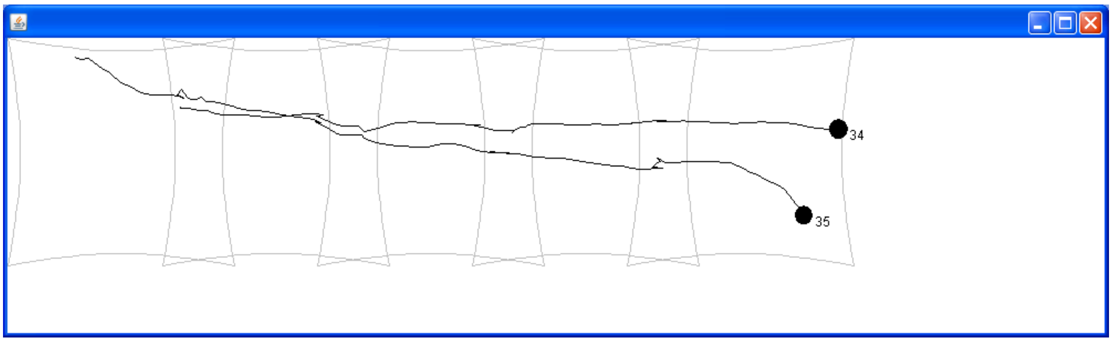

Figure 7: Data recorded in Apex corridor

Although matching the targets in this example worked well, there is still a great deal of noise in the overlap areas, which can be smoothed out using various means, such as splines, during post-processing.

The corridor is also equipped with a network camera that records activity as video. This is used to disambiguate data that cannot be explained by the infra-red traces alone. The example in Figure 8 shows a target (83) moving in a circular fashion. This turned out to be someone walking in circles while they were talking into their mobile phone.

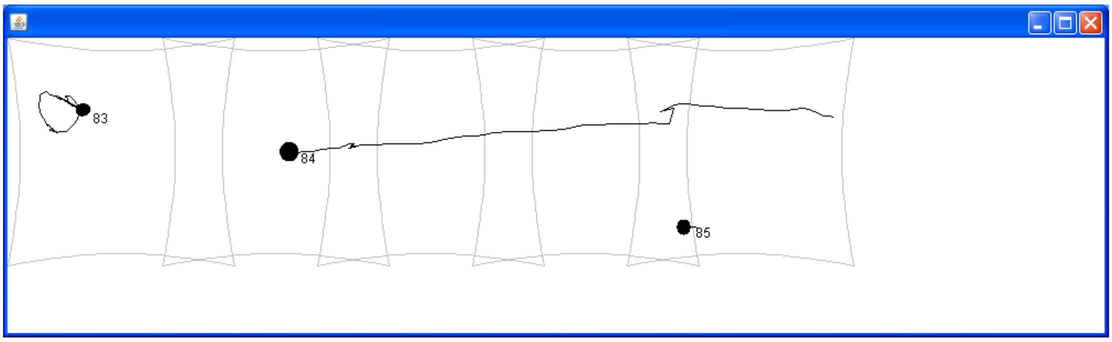

Figure 8: Loitering target

With this infrastructure in place, there is the potential for data gathering is unlimited. In contrast to the studies mentioned above, the relied on very small data sets to validate their models, the data analysis phase of this research will be able to call on much more substantive primary data.

3. DATA ANALYSIS

A common characteristic of research in the field of pedestrian behaviour is the imposition of a model explaining pedestrian behaviour based on little or no empirical proof. Even where there is some empirical evidence to support these models, it is unlikely that they entirely explain human behaviour.

The approach taken in this research is to make no assumptions about the way people behave. The model that will result from this work will be derived from the data itself, rather than from application of theories derived from, for example, classical physics or economic theory. Machine learning is a well established branch of artificial intelligence, and is often used in complex environments where the factors affecting behaviour are unknown.

Learning Classifier Systems (LCS), first proposed by John Holland in 1976, are an ideal means of implementing machine based learning in this research. LCS have often been used in robotics as well as agent-based systems. They have these characteristics: each LCS has sensors, with which it senses its environment and effectors with which it acts on its environment. Learning is carried out by an in built genetic algorithm that optimises responses to the environment according to the reward it receives. This reward is in proportion to how 'correct' responses are. In this case, this will be how closely the LCS responds to the environment compared with the example data is presented with.

The data gathered will be used to train the LCS in how to respond in common situations, such as overtaking, maintaining group cohesion, avoiding obstacles and other people, etc. Over time, the LCS evolve to respond to situations in exactly the same manner as people do. As Eiben and Smith (2003) put it, LCS embody mechanisms "via which the system can learn generalisations about the environment." These generalisations will then form the core of a model of pedestrian behaviour entirely derived from empirical data.

4. CONCLUSION

Pedestrian behaviour is an oft neglected and obscure part of transport research, despite its ubiquity. Government policy throughout the world is moving more and more towards encouraging this form of transport, in its own right as well as part of encouraging the use of public transport. However, studies in this area suffer many technical challenges in tracking pedestrians as well as simulating them.

In this paper, a novel and cost effective approach to tracking pedestrians is described. This is providing vast amounts of empirical data for further analysis. The proposed framework for analysing this data is also described. The intention is to not only automate the collection of data but also the discovery of general behavioural rules people employ when walking through an area.

REFERENCES

J. Desyllas, E. Duxbury, J. Ward, and A. Smith, "Pedestrian demand modelling of large cities: an applied example from London," Centre for Advanced Spatial Analysis (UCL), London, UK., Tech. Rep. 62, 2003.

The Mayor's Transport Strategy. Greater London Authority, 2001.

J. Fruin, "Pedestrian Planning and Design", Metropolitan Association of Urban Designers and Environmental Planners, 1971.

D. Helbing, "A mathematical model for the behavior of pedestrians," Behavioral Science, vol. 36, pp. 298–310, 1991.

G. Antonini, M. Bierlaire, and M. Weber, "Discrete choice models of pedestrian walking behavior," Transportation Research Part B, August 2004.

A. Willis, N. Gjersoe, C. Havard, J. M. Kerridge, and R. Kukla, "Human movement behaviour in urban spaces: implications for the design and modelling of effective pedestrian environments," Environment and Planning B: Planning and Design, vol. 31, no. 6, pp. 805–828, November 2004.

K. Teknomo, "Microscopic pedestrian flow characteristics: Development of an image processing data collection and simulation model," Ph.D. dissertation, Tohoku University Japan, Sendai, March 2002.

W. Daamen and S. Hoogendoorn, "Pedestrian traffic flow operations on a platform: observations and comparison with simulation tool SimPed," in Computers in railways IX (Congress Proceedings of CompRail 2004), J. Allen, C. Brebbia, R. Hill, G. Sciutto, and S. Sone, Eds., May 2004, pp. 125–134.

A. Eiben and J. Smith, Introduction to Evolutionary Computing, ser. Natural Computing Series. Springer-Verlag, 2003, ch. 7, pp. 115–128.