

Towards more train paths through early passenger intention inference

Peter Colborne-Veel, Nathan Kirchner, Alen Alempijevic

Centre for Autonomous Systems, University of Technology, Sydney, Sydney, Australia

Email for correspondence: peter.j.colborne-veel@student.uts.edu.au

Abstract

In public train stations, the designed way finding tends to induce individuals to conform to specific egress patterns. Whilst this is desirable for a number of reasons, it can cumulate into congestion at specific points in the station. Which, in turn, can increase dwell time; for example, loading and unloading time increases with concentrations of people trying to load/unload onto the same carriage. Clearly, an influencing strategy that is more responsive to the current station situation could have advantages.

Our prior research studies in Perth Station demonstrated the feasibility of reliably and predictably influencing passengers egress patterns in real time during operations. This capability suggests the possibility of active counterbalancing of the egress-alternatives while maintaining way finding. However, the prerequisite for such capability is the availability of knowledge of passenger's intention at a point in their journey where viable egress-alternatives to their destination exist.

This work details an approach towards an early (in the passenger journey) passenger intention inference system necessary to enable active egress-alternative influencing. Our contextually grounded approach infers intention through reasoning upon observed system and passenger cues in conjunction with a-priori knowledge of how train stations are used. The empirical validation of our intention inference system, which was conducted with data acquired during operations on a platform in Brisbane's Central train station in Queensland, is presented and discussed. The findings are then employed to argue the feasibility of an influencing system to reduce passenger congestion and the potential service impacts.

1. Introduction

Congestion in train stations is a major operational issue and one that is set to become more pronounced with population growth (Thompson, K., 2012; Wang, B. and Legaspi, J, 2012; Li, Z, and Hensher, D., 2013). A build-up of congestion at different points throughout a station is able to increase the dwell time of trains which in turn limits the number of train paths possible (Puong, A., 2000). This problem is not isolated to a single train either and a delay to a single train can propagate and cause problems for the rest of the train system (Higgins, A and Kozan, E, 1998). It follows that any reduction in this congestion could reasonably lead to an improvement in train dwell time and train paths.


The congestion that leads to reduced dwell time is in part caused by individual passenger egress decisions, e.g. passengers who are trying to exit the platform can run into passengers who are trying to board the train, slowing down both. The ability to counterbalance passenger egress with alternatives could negate some of the causes of congestion and therefore lead to operational benefits. The existing methods of counterbalancing are

limited to signage and announcements. These approaches are limited because the operator must wait until the congestion is apparent before knowing what counterbalancing to apply.

Our prior research studies in Perth Station demonstrated the feasibility of reliably and predictably influencing passengers egress patterns in real time during operations (Caraian, S., Kirchner, N. and Colborne-Veel, P., 2015). This capability suggests the possibility of active counterbalancing of the egress-alternatives while maintaining way finding, a responsive passenger influencing system. However, the prerequisite to being able counter balance egress alternatives for passengers is the ability to infer where a passenger intends on going. We need to know this because passengers with different intentions will require different and sometimes opposite influence strategies, e.g. passengers who want to go to the concourse will need to go to the exit points where as passengers who want to board a train could be encouraged to move away from the exit point. Consider Fig.1, the application space for such intention inference introduces five constraints that would need to be overcome. More elaborate explanations are provided in the proceeding sections but in general these constraints are,

- 1) Would need to infer a person's intention with enough time to encourage the passenger to make alternative egress decisions.
- 2) Would need to produce passenger insights early enough within a sensor field of view so that influencing can be both applied and verified (closed loop influencing).
- 3) Would need to operate in areas where only sparsely distributed person detection sensors are possible, e.g. in areas without infrastructure to mount onto, etc.
- 4) Would need to work within privacy constraints set by application space.
- 5) Would need to work on demographics of people who are typically non-compliant with traditional influencing strategies.

Figure 1: The need for intention inference to counter balance egress alternatives



A major problem that is anticipated for intention inference within these constraints is the ambiguous nature of human movements within disjoint sensor field of views and narrow windows of time. As a motivating example, if a person wanders around a platform they may enter a sensors field of view in a way that looks like they are heading towards the exit of the platform, when in fact, they have no intention of leaving. Other such ambiguities exist throughout a passenger's journey and these would need to be overcome for intention inference to be feasible. One possibility would be to contextually ground passenger cues by looking at concurrent system events, e.g. if someone is walking towards the exit but a train hasn't arrived at the platform for some time, then perhaps they are less likely to actually be leaving, and we could exploit this to differentiate between two ambiguous person cues.

The method section and experimental component of this paper are devoted to outlining a preliminary exploration of intention inference that uses passenger and system cues in ways to reduce ambiguity within this application space. The method is evaluated using real data from Brisbane's Central station. The particular intention inference scenario is a limited case of what may be possible with future work. The results of this study are then used to argue for the feasibility of a responsive influencing system and the potential operational benefits that this may enable.

2 Previous work on intention inference

The field of intention inference is divided into two branches, continuous intentions such as, predicting the position of a tennis player's racket and discrete intentions, such as predicting which staircase a person will use. For this paper we are interested in discrete intentions. To the best of our knowledge no system specifically developed for intention inference within the train station application space exist, however many systems have been developed for other applications (Kelley, R., et al 2008; Kooij, J., et al, 2014; Patel, M., Valls Miro, J. and Dissanayake, G., 2010; Mandalia, H., Salvucci, D., 2005). Typically, current intention inference methods use machine learning algorithms adapted for the task. Three promising approaches to intention inference that we found were Hidden Markov Models (HMMs), Support Vector Machines (SVMs) and Gaussian Process Models (GPs).

The HMM approach has been used extensively for the problem of inferring discrete intentions (Kelley, R., et al 2008; Kooij, J., et al, 2014; Patel, M., Valls Miro, J. and Dissanayake, G., 2010). HMM is a subset of a broader range of algorithms known as Dynamic Bayesian Network or DBN. The drawback of these approaches is that a model needs to be constructed before they are able to learn anything from data.

Support Vector Machines are considered a reliable machine learning tool for classification and have found many applications including being used for discrete intention inference. The closest existing research to train passenger intention inference that we found was for the problem of inferring car driver intentions, specifically, inferring when a lane change will occur (Mandalia, H., Salvucci, D., 2005). In this work, an SVM was trained with driver data from Japan and this approach was shown to outperform alternatives for this application. They were able to use cues from the car such as steering angle and cues from the environment such as distance to next car.

Gaussian Process models are a stochastic alternative to HMM and SVM approaches and have a benefit of being able to model the uncertainty for any given prediction as well as the prediction itself. Gaussian Process models can be used for both regression (continuous intentions) and classification (discrete intentions) and are able to learn a model from data and features alone. Typically prior work has found GP models to be useful for continuous intention inference such as inferring where in Cartesian space a player is trying to hit a table tennis ball. These models have also been used for the problem of socially acceptable path planning for robots. In that case the GP was trained to produce a probability heat map of how people use the space (Alempijevic, A., Fitch, R. and Kitchner, N). The advantage of the GP in this work was that it was able to capture complex pedestrian phenomena such as not walking in front of a television without needing to be trained explicitly. The disadvantage of GP models is because they are nonparametric they are computationally expensive.

For our work we have selected the SVM approach owing primarily to its prior success for discrete intention inference. A HMM was decided against because of the need to build a model and the fact that data was available for an SVM to learn the model from the data. A GP approach is interesting because of the way that it models uncertainty. We decided against its use for this work because an SVM was suitable, however, a GP could have advantages for a more complete inference system, e.g. could model uncertainty.


3 Constraints on intention inference for this application

3.1 The need to infer intention before decision point

Intention in this work refers to a passenger's overarching objective while moving through a train station, e.g. they want to exit the station. A decision point is a point in time in a passenger's journey when/where they select one of multiple options to achieve the same intention, e.g. choosing between a staircase and an elevator, illustrated in Fig.2.

A responsive influencing system would need to be able to influence a person's decision before the decision point has been reached in order for influence to be feasible. Trying to influence a change in decision after the decision point may be unproductive because the passenger will already consider themselves to have past that point in their journey and could be reluctant to recant their decision.

Figure 2: Decision point – choosing between stairs and elevator



The likelihood of influencing a passenger's decision is contingent on the probability of successful intention inference and that of successful influencing. Neither of which are likely to be 1, and so a question remains as to how we could maximise success. One possibility would be attempting to infer intention and influence at multiple points throughout their journey. This seems feasible as multiple decision points are known to exist, e.g. choosing how to enter the station, choosing the ticket gate, choosing the staircase to enter the platform, choosing the platform boarding position, etc.

It follows that one constraint on the intention inference system would have to be that it could infer a passenger's intention before a decision point has been reached. Furthermore that such a system could be implemented at successive decision points to maximise the likelihood of success.


3.2 Narrow window of time for closed loop influencing

A key feature that a responsive system has over existing influencing methods is that it is closed loop and targeted towards individuals. Closed loop is an engineering term that means the system is able to influence passenger's egress, verify that the desired effect was

achieved and modulate the influencing otherwise. Existing methods of influence using signage can be thought of as open loop meaning that they can influence passengers egress but they have no way of verifying the success of the influencing cannot respond and change their influence accordingly. Operator announcements could be considered either open or closed loop depending on how aware the operator is of the situation, but this method only considers crowds and not individuals.

Where as a responsive system could perform closed loop influencing on individuals. One requirement of this is that the intention inference system needs to be able to infer a person's intention and apply influencing within a field of view of a sensor so that the effect can be monitored and modulated within the same sensor if needed, illustrated in Fig.3.

Figure 3: The advantage of closed loop influencing



The effective sensor range of our system is 8 meters and closed loop influencing was shown to be possible in our previous work given 2-3 meters of distance to influence a person (Caraiian, S., Kirchner, N. and Colborne-Veel, P., 2015). With an average walking speed of 1.4m/s this leaves only the first 1 second of a person's path to infer their intention and apply closed loop influencing within a single sensor. Potentially, it could be possible to use two sensors, where one inferred their intention and the other applied closed loop influencing, but this expends significantly more hardware and infrastructure, along with adding complexities.

3.3 Sparsely distributed sensors

Complete sensor coverage throughout a person's journey is unlikely due to the already mentioned increased hardware and infrastructure expenditure, and increased complexities. It follows that an intention inference system for this application could not require complete sensor coverage. This introduces a problem of ambiguity within sensor field of views which must be overcome by an intention inference system. Within a limited field of view, two passengers with opposing intentions may have very similar looking paths, e.g. a passenger who is leaving will move towards the exit and a passenger who is wandering around a platform may also move in the direction of the exit.

One possibility would be to contextually ground passenger cues by looking at concurrent system events, e.g. if someone is walking towards the exit but a train hasn't arrived at the platform for some time, then perhaps they are less likely to actually be leaving, and we could exploit this to differentiate between two ambiguous person cues. Using system events to contextually ground passenger cues is an approach that will be explored in the experimental component of this paper.

3.4 Operating within privacy restrictions

A key feature of the responsive influencing system being developed is that it respects passengers' privacy. This requires that any information used by our system must not be able to associate to private information outside of the system. Our previously devised and validated Sensor Hardware Platform (SHP - Kirchner, N, et al. 2014), which is used in this work, employs this under pinning data sensitivity.

3.5 Applicable to non-compliant demographics

In order for responsive influencing to be most effective it needs to be deployable to influence passengers who are typically non-compliant with traditional methods. This means that the intention inference component of this system cannot necessarily depend users to opt in, such as a mobile application, or a similar level of compliance and instead should be general enough that it could potentially be usable on all passengers. This is not to say that a mobile application could not be integrated into an overall system and could help with intention inference of some passengers, but that is would not be solely sufficient.

4 Intention inference for this application

4.1 Sensing passenger and system cues

To overcome the ambiguity present within a narrow window of time of a person's path, it may be feasible to contextually ground passenger cues using system cues. The requirements to do this would be to detect time matched passenger and system cues. For passenger cues we would need the ability to gather time stamped, position and velocity traces of individuals who enter the sensors field of view so that their path could be constructed. For system cues we would need the ability to gather time stamped observations of system events, e.g. the time since a train arrived.

The task of generating person cues can be broken down into identifying individuals and tracking them within the sensor region. In this work we have used a previously developed system for person detection and tracking system (Kirchner, N., et al. 2014) and manually coded these individuals with intention labels. Manual coding alone was used for generating system cues. The system cue of particular interest was the time since the last train arrived.

4.1.1 Passenger Cues

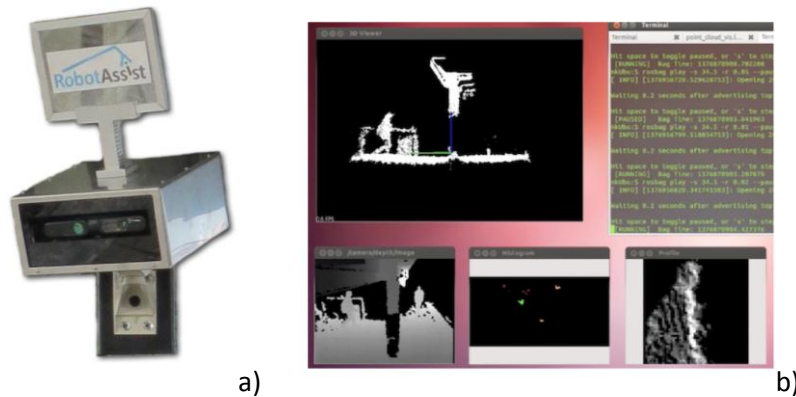
Our previously developed person detection and tracking system [Hordern and Kirchner, 2010, Kirchner et al., 2012, Kirchner et al., 2014] was implemented on our SHP. Our SHP, which has been devised, developed and empirically evaluated, is shown in Fig. 3. Our SHP has been demonstrated to be capable of robust people detection, tracking, and counting system in train stations [Kirchner et al., 2014].

4.1.2 Train event detection

The problem of sensing system cues such as train events is currently being investigated by another researcher. For this paper, we were interested in train arrival events and without any existing system that could be used for this task, the events were manually coded by from the sensor output by noting the time when defined train events were observed. The sensors that

were used for this task were time synchronised to ensure consistency between system cues and person cues within the two sensors.

Figure 3: Our SHP for robust people awareness – shown in a) & b)



4.1.3 Passenger intention labelling

To train and evaluate a supervised machine learning algorithm, it is necessary to provide labelled examples of each intention. A known limitation of the person tracking system used for this work is the fact that because it depends heavily on a velocity model for tracking a person when there are disjoint sensor regions, the velocity model cannot correctly associate traces across the different sensors. Concurrent work is being directed at this problem. This presented some difficulty for constructing the ground truth for the intention inference system because we needed to associate the traces of people moving within one sensor with people leaving to the concourse in the sensor mounted at the exit. We were able to label the persons intention using a manual method where we looked at the 3D image of a person in one sensor and manually compared it with the images of people exiting. The result was a set of tracks labelled with the person's intention.

4.1.4 Feature selection for intention inference

To investigate how passenger cues and system cues could be used for an intention inference system, we selected three cues, passenger direction, passenger speed and the time since the last train arrived. The first two are simple passenger cues and the second is one example of a system cue among many others that are likely to exist.

For finding the direction we looked at the angle of each passengers mean velocity vector within the first 1 second of their path. Speed was found by looking at the magnitude of their mean velocity vector within the first 1 second. It was important to limit the information we used to only the first 1 second of path because this is the necessary window of time that could allow for closed loop influencing to occur; as previously detailed. Also important is the fact that using a mean over 1 second of data (30 positional traces approximately) reduces the noise of the reading helping to reduce the undesired effect of erroneous traces.

The system cue that was selected is the time, in seconds, since the last train arrived. In this case we considered the time since a train arrived on a single platform. By arrived we mean that the time is 0 the moment when a train reaches the platform, stops and the doors open.

4.1.5 Algorithm training and evaluation

The approach that was taken to infer a person's intention was to train a support vector machine (SVM) using the passenger cues in combination with the system cue. This algorithm is a part of the field of machine learning and essentially learns relationships from example data to predict labels for new data.

With a data set acquired using the method outlined above, we could then train and evaluate the SVM algorithm and analyse the results. To ensure statistical validity we randomly partitioned the given data set into two equally sized sets, one for training and the other for testing. We then swapped the set that was used for training and the set that was used for testing around. This was performed 100 times over and the mean accuracy and variance was generated. Each feature was evaluated independently and then all of the possible combinations were tested. A success was recorded if the inferred intention matched the manually labelled intention $\{-1, +1\}$ with -1 meaning the person was not going to concourse and +1 meaning that they were, a failure was recorded for an incorrect inference.

5 Empirical evaluation

To evaluate the feasibility of intention inference for our application we conducted a study with two focal hypotheses.

H1 – It is feasible to infer a train passenger's intention at a decision point and within a window of time (1s) and space that could allow for closed loop influencing

H2 – The feasibility of this intention inference can be improved by incorporating system cues together with passenger cues to reduce ambiguity

5.1 Inferring whether a passenger intends on leaving to concourse or not

In our study we decided to focus on the problem of inferring whether a person intends on leaving the platform or not. The ability to infer if a person is trying to leave is of interest to a responsive influencing system because the passenger must make a decision about how to leave the platform, e.g. they may need to decide whether to exit using a particular set of stairs or elevator. By inferring if a person intends on leaving early during their departure from the platform we could maximise the potential for using responsive influencing to try and influence the passenger to make decisions that will reduce congestion overall, e.g. we could encourage them to use a more distant staircase if we know that they will run into people trying to enter the platform using the closest staircase.

This problem is also of interest because it presents a scenario where ambiguity is likely to be present. Arbitrary social phenomena such as wandering around while talking on a phone mean that, for a particular window of time, a person's path could appear to be heading for the exit but they are in fact wandering around. An intention inference system within this application space would need to be able to differentiate between these cases where the person's path is ambiguous. Something that we looked at for this study was whether we could contextually ground a person's path observation so that the two cases could be differentiated. Considering the problem of inferring whether a person is trying to leave vs. not leave we decided that a possible system cue would be the time since the last train arrived to the platform. Intuitively, we suspected that during the period immediately after a train's

arrival experiment, the likelihood of seeing someone walk towards the exit and actually exit would be greater than after a prolonged time since the last train arrived.

5.2 Experiment setup

Our SHPs were mounted along a platform in Brisbane's Central station where they produced positional traces and velocity readings for each individual who passed through the SHPs field of view. These paths were manually associated with the person's intention by using a sensor mounted at the exit to check whether people left. An image of the sensor setup at the train station is provided in Fig.4 and of the raw sensor output from the two sensors in Fig.5.

Figure 4: Sensors mounted at train station



Figure 5: Raw sensor depth image used for inference (LEFT), labelling (RIGHT)



5.3 Experiment results

Our SVM was trained with three features, the first two being person cues of direction and speed, the third being a system cue time since train arrival, which in this case was time since train arrival on platform 5 given the specific setup of the sensor. Each feature was tested for the two cases, firstly in isolation and then all combinations were tested. The mean accuracy and variance for each was recorded and is presented in the Tab.1 and Tab.2. To calculate these results, the dataset was randomly partitioned into a training set and a testing set for an SVM classifier. The accuracy was measured by the ratio of correct classifications to the total

number of examples in the testing set. This process was repeated 100 times with different random partitions and the average accuracy and variance were recorded below. In Tab.1 The value listed under C is the average success rate or accuracy of the classifier when correctly inferring that a person is going to the concourse. The value listed under !C is the average accuracy of the classifier when correctly inferring that a person is not going to the concourse. In Tab.2 the variance of the success rate for the 100 different partitions is provided.

Table 1: Classification Accuracy

Classification Accuracy	C	!C	Overall
Direction	0.44	0.91	0.68
Speed	0.66	0.82	0.74
Train Time	0.72	0.65	0.69
Direction + Speed	0.65	0.82	0.73
Direction + Train Time	0.70	0.72	0.71
Speed + Train Time	0.75	0.77	0.76
Direction + Speed + Train Time	0.71	0.81	0.76

Table 2: Classification Variance

Classification Variance	C	!C	Overall
Direction	0.02	0.02	0.01
Speed	0.02	0.01	0.01
Train Time	0.02	0.02	0.01
Direction + Speed	0.03	0.01	0.01
Direction + Train Time	0.03	0.03	0.01
Speed + Train Time	0.02	0.02	0.01
Direction + Speed + Train Time	0.03	0.02	0.01

As can be seen from these results, we found that direction alone was a poor feature to distinguish between the two cases. This was anticipated by the authors given the ambiguous nature of human movements within a narrow window of time. People orient themselves in particular directions for all kinds of arbitrary social reasons and this leads to direction within a small window of time being ambiguous. Generally, including direction made the variance of the inference system worse and the accuracy not any better. This finding reaffirmed our suspicions that arbitrary human phenomena such as walking around randomly while on the phone etc. would lead to ambiguous paths. These results highlight the importance of contextually grounding passenger cues.

Unexpectedly, the degree to which speed was successful at distinguishing between the two cases was notable. In our study we found that using speed alone had an accuracy of (0.66 variance 0.02 for people leaving to concourse, 0.82 variance 0.01 for people not leaving). Prior to commencing this study it was not obvious that speed would prove a valuable distinguishing feature. These results suggest that passengers who alight from trains may move statistically faster than passengers who are walking around waiting for trains. Intuitively, the reason for this could be that passengers who alight from trains have a reward for moving quickly, namely, they will reach their destination sooner, whereas, passengers

who are wandering about have no reward for moving more quickly because they are wasting time until the train arrives.

The third feature that was used was time since the last train arrived, a system cue that we wanted to use as a test to see if it was possible to contextually ground person cues. In the specific experimental setup the sensor was mounted so that it could see passengers who alighted from trains on platform 5 but not on platform 6. For this reason we used the time since the last train arrived on platform 5. From our results we found that this feature had reasonable predictive ability alone for each case (0.72 variance 0.02, 0.65 variance 0.02 respectfully) and importantly it combined with the person cues to achieve better intention inference than was possible without the combination (0.75, 0.02 variance, 0.77, 0.02 variance). These results suggest that if someone has alighted from a train you are more likely to see them within the first minute since the train arrived than you are to see someone who is not going to the concourse. This result is interesting because it shows that by contextually grounding the person cues (direction and speed) by including a system cue (time since train arrival) we could improve intention inference efficacy. Intuitively if a person is moving quickly, towards the exit and a train recently departed, they are more likely to intend on leaving the platform.

6 Discussion

6.1 Feasibility of responsive influencing system

Given what we have learned about the opportunities and limitations for intention inference from our study, in this section we would like to cover the feasibility of an influencing system overall and the potential service impacts that may be possible. The major hurdle for intention inference for this application was the ambiguity present in passenger paths within small windows of time. In our study, we showed how little predictive value direction had within the field of view of a sensor. This was anticipated by the authors and can be understood as the consequence of arbitrary social phenomena that exists within this space. Our results showed that by incorporating a system cue together with the passenger cues we were able to improve the inference accuracy better than was possible using the features alone.

In particular, we focused on the problem of inferring whether or not a passenger intends on exiting the train station or not. Given that we can infer this information, it is worthwhile taking a moment to explain how this could actually be used by an influencing system. Suppose that there were sensors distributed throughout a platform such that the relative crowding in each region was known (e.g. south exit, north exit, platform 1 north etc.). When a sensor inferred that a person intended on leaving the platform, the corresponding influencing device as described in (Caraiian, S., Kirchner, N. and Colborne-Veel, P., 2015) could classify them as an alighting passenger and indicate that they should head towards a particular exit, otherwise, it could classify them as a boarding passenger and indicate that they should move away from the exits. Such a setup could overcome congestion caused by boarding passengers trying to board near platform exits. Moreover, the reduction in congestion could allow for faster and more consistent boarding and alighting of passengers. It is reasonable that the dwell time of the trains would decrease as a result of this and hence allow for more train paths.

More generally, our results from this study suggested that some of the ambiguity in the passenger cues could be overcome by contextually grounding passenger cues. This contributes to the feasibility of an influencing system overall because it opens up the possibility of intention inference at other decision points where this ambiguity exists. To get an idea of what may be possible for intention inference at other decision points, we can spend a moment speculating about what other system cues may be exploitable. People tend to look at information screens to see when their train will arrive. In many stations such as Central and Circular Key in Sydney, information screens are available while the person is on the concourse. It may be possible to exploit this fact and create a component to look at a person's head orientation to see if they are looking at a particular screen for some time. This could give some indication that the person is about to enter the platform associated with that screen or we could look at a ticket machine for an indication that the person is likely to enter the station or the time of day, etc. It follows that many system cues could potentially be exploitable at earlier stages of the passengers journey, giving rise to the possibility of multiple decision points to influence.

At no point during our method did we require passengers to opt in to our inference system, such as by signing up to a mobile app or use any sensitive passenger information, such as individual Opal card data. We are not suggesting that an optional mobile app can't be used for intention inference but only that it would be insufficient in and of itself. We used only the first 1 second of a person's path to make it so that the intention inference could be applied to a closed loop influencing system and that it could be inferred early. Adhering to these constraints and achieving 76% inference accuracy gives an indication that early intention inference is more feasible than not by incorporating system cues to contextually ground passing cues. From the proceeding discussion it is reasonable to say that a responsive influencing system is more feasible than not for use within this application space.

6.2 Operational benefits

Minimising train dwell time is a key operational objective because of its relationship to the number of train paths possible. A major cause of prolonged dwell time is congestion and a potential operational advantage of using a responsive system could be a better distribution of passengers leading to faster boarding and alighting of trains. Passenger behaviours that lead to congestion, e.g. alighting passengers colliding with boarding passengers, could be discouraged by early intervention in the passengers way point selection, e.g. we could encourage boarding passengers to stay away from exits. Because we have shown that it may be feasible to use influence at multiple decision points early during the journey, there would be multiple opportunities to try and prevent congestion from occurring. This is a clear operational benefit over existing methods that must wait until the congestion is apparent.

7 Conclusion and Future Work

In this work we have explored the feasibility of an intention inference system that could be used for the application of informing a responsive influencing system. We were able to identify key constraints that a complete intention inference system would have and explored how these could be overcome. The key problem that we have identified is the problem of ambiguity in passenger positional traces within the necessary narrow windows of time. We

were able to conduct a study to show that some of this ambiguity could be overcome by contextually grounding person cues with system cue. Through our discussion we outlined the feasibility, and potential operational benefits, of an influencing system.

Although an initial step towards an intention inference system has been made there is still a lot that needs to be improved. In particular, our study was limited to only two intentions and although the study suggests that it may be possible to exploit system cues at other decision points, the extent to which they could be exploited remains an open question.

Future work is to conduct more probing experiments that would be both longer in duration and also could look at other decision points. An interesting question is whether there could be higher order passenger cues observed by looking at passengers path in preceding sensor field of views, our study only looked at one sensors field of view. Another interesting line of inquiry would be to explore the available system cues and to see how these other system events could be used to reduce ambiguity at different decision points.

Acknowledgments

The IP contained here within is Patent Pending [Kirchner, 2012a, Kirchner, 2012b, Kirchner, 2015]. This work is supported by UTS' Centre for Autonomous Systems and RobotAssist (www.robotassist.org). Further information on our activities is available at www.nkHRI.com

References

- Alempijevic, A., Fitch, R. and Kirchner, N. (2013). Bootstrapping Navigation and Path Planning Using Human Positional Traces, IEEE International Conference on Robotics and Automation (ICRA), Karlsruhe, Germany.
- [Hordern and Kirchner, 2010] Hordern, D. and Kirchner, N. (2010). Robust and Efficient People Detection with 3-D Range Data using Shape Matching. In *Proceedings of the 2010 Australasian Conference on Robotics and Automation*.
- Kelley, R., et al (2008), Understanding Human Intentions via Hidden Markov Models in Autonomous Mobile Robots, Proceedings of the Third Annual ACM/IEEE International Conference on Human-Robot Interaction.
- Kooij, J., et al, (2014), Context-Based Pedestrian Path Prediction, European Conference on Computer Vision 2014.
- Mandalia, H., Salvucci, D., (2005), Using support vector machines for Lane-Change Detection, Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting.
- Patel, M., Valls Miro, J. and Dissanayake, G., (2010), Dynamic Bayesian Networks for Learning Interactions between Assistive Robotic Walker and Human Users, Advances in Artificial Intelligence Lecture notes in computer science, Vol. 6359. pp 333-340.
- Caraian, S., Kirchner, N. and Colborne-Veel, P., (2015), Moderating a robot's ability to influence people through its level of socio-contextual interactivity, Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction.

Kirchner, N., et al. (2014), A robust people detection, tracking and counting system, Proceedings of the Australasian Conference on Robotics and Automation, Melbourne Australia.

[Kirchner et al., 2012] Kirchner, N., Alempijevic, A., and Virgona, A. (2012). Head-to-shoulder signature for person recognition. In Proceedings - IEEE International Conference on Robotics and Automation, pages 1226–1231.

Wang, B. and Legaspi, J, (2012), Developing a train crowding economic costing model and estimating passenger crowding cost of Sydney CityRail network, Australasian Transport Research Forum 2012 Proceedings, Perth, Australia.

Puong, A., (2000), Dwell Time Model and Analysis for the MBTA Red Line, MIT's Open Courseware Project.

Thompson, K., (2012), A socioeconomic study of carriage and platform crowding in the Australian railway industry, CRC for Rail Innovation.

Li, Z, and Hensher, D., (2013), Crowding in Public Transport: A Review of Objective and Subjective Measures, *Journal of Public Transportation*, Vol. 16, No.2, 2013.

Higgins, A and Kozan, E, (1998), *Modeling Train Delays in Urban Networks*, Transportation Science, Vol. 32, Issue 4, November 1998