Artificial Intelligence for Human-Robot Interaction
Papers from the AAAI 2015 Fall Symposium

Missteps in Robot Social Navigation

Andrew Sutcliffe and **Joelle Pineau**
School of Computer Science, McGill University,
McConnell Engineering Bldg. Room 318,
3480 University, Montreal, Qc, Canada, H3A 0E9

Neil Tenenholtz
Vecna Technologies
36 Cambridgepark Dr,
Cambridge, MA 02140

Introduction

Assessing the quality of robot social navigation is a challenging problem fraught with human obstacles. From preconceived notions to perspective or point of view, evaluations can differ from person to person. Most work in the field of robot navigation is focused on creating algorithms that produce efficient robot trajectories. We posit that the evaluation of trajectories in a social context is essential and distinct to trajectory generation. In this work we recorded a manually driven powered wheelchair through different scenarios and asked expert evaluators to assess the quality of the powered wheelchair's movement. These evaluations were then compared to post-experiment assessments from trajectory generation algorithms and social navigation concepts. Our results show that it is possible to build a simple model to predict expert evaluators' responses. Unfortunately, there is no clear consensus amongst these experts on what quality behaviour is. This suggests that while current navigation algorithms offer strong heuristics for the generation of smooth trajectories in well-defined environments, their efficacy in evaluating social navigation is less obvious. We believe that more emphasis must be put on dynamic and reactive navigation algorithms as any heuristic approach will be limited due to variance in people's behaviours and expectations.

Methods

A protocol was devised to gather expert evaluations of robot navigation in a hallway. In this case the experiment was run with an intelligent powered wheelchair, though the driving was done manually. Subjects were asked to observe pairs of wheelchair trajectories from video recordings and to select the video representing better behaviour. Trajectories were evaluated using heuristics optimized for social navigation these heuristics are detailed in the following section. Across pairs of trajectories, differences in the computed heuristic were then compared to individual votes.

Heuristics

Three navigation heuristics, which find optimal trajectories by minimizing a cost incurred by the robot moving through space, were considered. These heuristics were chosen because they focus on features that have been shown to have an impact on social navigation. They are: the model predictive equilibrium point control (MPEPC) (Park, Johnson, and Kuipers 2012), predictability (Dragan, Lee, and Srinivasa 2013), and proxemics (Lu, Allan, and Smart 2013). Each heuristic is a linear combination of its features, with weights adjusted for specific robotic platforms. A new custom heuristic was defined as a combination of all features present in the different considered functions. These five factors are: linear acceleration, linear speed, angular acceleration, angular speed and proxemic distance. All theses factors are robot centric.

Scenarios

Six scenarios were constructed: three people oncoming (S-3-0-0), two oncoming and one passing (S-2-0-1), one oncoming, one passing and one side-by-side (S-1-1-1), one on coming and one passing with one not moving facing the wall (S-1-0-1*), turning a corner with one oncoming and one passing (L-1-0-1) and, finally turning the corner with one oncoming and one side-by-side (L-1-1-0). Figure 1 illustrates the starting locations and goals of each of these scenarios.

Protocol

Eight experienced roboticists participated as evaluation experts. Evaluations were conducted through a Google Survey embedded with private Youtube videos. Subjects were asked to evaluate pairs of different videos from the same scenario. Subjects evaluated all possible unordered pairings, presented in a randomized order, by selecting which video in the pair showed better wheelchair behaviour or indicate if there was no significant difference. There were six videos each for the first two scenarios and five videos each for the other four scenarios. In total, each subject evaluated 70 video pairs across the six scenarios. Videos were filmed in a hallway with the camera located in a corner at the intersection with the ceiling, allowing for an unobstructed view of the scene. Three actors were recruited for these videos and did not change throughout the six scenarios. A confederate drove the powered wheelchair under the instruction to vary his behaviour to produce a diverse range of trajectories.

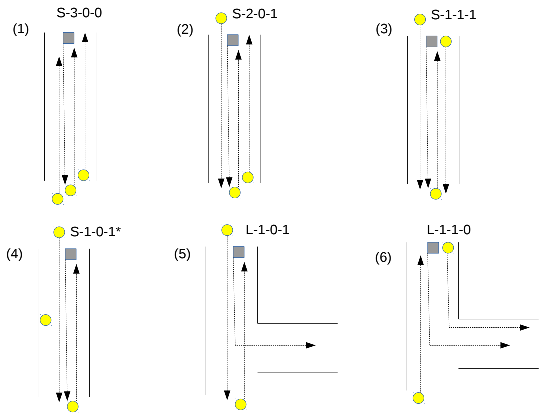

Figure 1: *6 scenarios evaluated. Circles are people while the square is the driver in the powered wheelchair. Arrows indicate the intentions of the actors and of the driver.*

Analytical Testing

The data was analysed through logistic regression models which retained the assumption of a linear combination of features. Each expert vote on a pair was taken as an individual data point. Instances where experts expressed no preference between the videos were ignored, which resulted in with 531 individual data points. Additionally, a score was produced for each pair using a weighed combination of the heuristics described above. For each pair, the scores were compared, producing a signed difference between the members of a pair. Since the heuristics chosen compute cost, the resulting sign of the difference indicated a preference computed by the automatic approach. The weights of the scoring function were trained using a logistic regression model to try and best fit expert opinions.

A first model was trained using all features. Subsequent models were trained by leaving out one of the features in order to assess the impact of each feature on the significance of different features using analysis of variance. In order to compare the results to the agreement amongst experts, 10 fold cross-validation was performed. Votes by different experts on the same pair were grouped together to avoid splitting them between the test set and the train set. Cross-validation produces estimates for the accuracy of our models which were then compared to the agreement amongst expert evaluators. This agreement was simply computed as the average over all pairs of the size of the majority vote on each pair expressed.

Results

Analysis of variance on the different models lead to the conclusion that the three most significant features were: angular speed, angular acceleration and proxemic distance. The average majority amongst expert votes on each pair was of 77.4%. The average accuracy over 10-fold cross-validation of the top models are reported in Table 1.

Test set	Accuracy	95% confidence interval
Full Model	69.3%	(+/- 23.5%)
Angular + Proxemic	72.5%	(+/- 25.0%)

Table 1: Accuracy of logistic regression.

Discussion

Angular acceleration, angular speed and proxemic distance were identified as predictors for expert evaluations. Given that on average there was a 77.4%, 22.6% split on average amongst the evaluators, a predictive power of 72.5% is acceptable. The improvement in accuracy in the model considering only angular acceleration and speed as well as proxemic distance may be due to over-fitting on the part of the full model. Unfortunately, when taken into context, the lack of solid consensus amongst expert evaluators, comes to light. Not only do experts not agree on what quality behaviour is, but there is convincing evidence in literature that the assessment of robot motion may depend on a variety of factors such as situational factors (Komatsu, Kurosawa, and Yamada 2011; Forlizzi 2008) or robot appearance (Takayama and Pantofaru 2009).

The problem then becomes one of identifying situations and using the appropriate behaviour given the situation. Similar work is being done through the PlanIt project (Jain, Das, and Saxena 2014), which has the goal of collecting large amounts of data through crowd sourcing in order to learn human preference functions in social contexts. Currently, robots are quite novel and therefore expectations from one person to the next may vary widely. Consequently, pedestrians will exhibit a range of reactions to robotic navigation. Therefore, identifying and reacting to social missteps in real-time will be essential to efficient and high-quality robot navigation.

Conclusion

As demonstrated in this work, traditional heuristics associated with social robot navigation can predict human assessments of robot trajectories in specific scenarios. However, human assessments tend to be noisy and therefore traditional approaches might find themselves limited in scope. As robots are deployed in more and more social environments, the number of different social scenarios will grow rapidly. To overcome this hurdle, navigation algorithms must become reactive, detecting social missteps and responding to them appropriately. Although some have begun tackling these ideas (Morales et al. 2012; Park and Kuipers 2013) through algorithms for walking along side someone, we believe that further work is necessary. Further work should aim at creating algorithms that integrate novel sources of feedback such as human pose detection and affect and speech recognition, to achieve significant, meaningful improvements aimed towards more interactive algorithms.

Acknowledgements

This work is supported by the Scholarship BMP Innovation FQRNT-NSERC and Vecna Technologies Inc.

References

Dragan, A. D.; Lee, K. C.; and Srinivasa, S. S. 2013. Legibility and predictability of robot motion. *2013 8th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* 1:301–308.

Forlizzi, J. 2008. Robots in Organizations : The Role of Workflow , Social , and Environmental Factors in Human-Robot Interaction.

Gulati, S., and Kuipers, B. 2008. High performance control for graceful motion of an intelligent wheelchair. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on*, 3932–3938. IEEE.

Hall, E. T. 1969. *The hidden dimension*, volume 1990. Anchor Books New York.

Jain, A.; Das, D.; and Saxena, A. 2014. Planit: A crowd-sourced approach for learning to plan paths from large scale preference feedback. *arXiv preprint arXiv:1406.2616*.

Komatsu, T.; Kurosawa, R.; and Yamada, S. 2011. How Does the Difference Between Users Expectations and Perceptions About a Robotic Agent Affect Their Behavior? *International Journal of Social Robotics* 4(2):109–116.

Lu, D. V.; Allan, D. B.; and Smart, W. D. 2013. Tuning cost functions for social navigation. In *Social Robotics*. Springer. 442–451.

Morales, Y.; Satake, S.; Huq, R.; Glas, D.; Kanda, T.; and Hagita, N. 2012. How Do People Walk Side-By-Side ? Using A Computational Model Of Human Behavior For A Social Robot. *Proceedings of the Seventh annual ACMIEEE International Conference on HumanRobot Interaction HRI 12* 301–308.

Park, J. J., and Kuipers, B. 2013. Autonomous person pacing and following with model predictive equilibrium point control. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, 1060–1067. IEEE.

Park, J. J.; Johnson, C.; and Kuipers, B. 2012. Robot navigation with model predictive equilibrium point control. In *IROS*, 4945–4952.

Takayama, L., and Pantofaru, C. 2009. Influences on proxemic behaviors in human-robot interaction. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, 5495–5502. IEEE.

Yoda, M., and Shiota, Y. 1997. The mobile robot which passes a man. *Proceedings 6th IEEE International Workshop on Robot and Human Communication. RO-MAN'97 SENDAI* 112–117.

Zucker, M.; Ratliff, N.; Dragan, A. D.; Pivtoraiko, M.; Klingensmith, M.; Dellin, C. M.; Bagnell, J. A.; and Srinivasa, S. S. 2013. Chomp: Covariant hamiltonian optimization for motion planning. *The International Journal of Robotics Research* 32(9-10):1164–1193.