Minds for Mobile Agents

Andrew Heathcote[1,2] (andrew.heathcote@newcastle.edu.au) Charlotte Tanis[2]
Jonne Zomerdijk[2] Tessa Blanken[2] Dora Matzke[2] Denny Borsboom[2]

[1]School of Psychology, University of Newcastle [2]Department of Psychology, University of Amsterdam

Objectives

Effective training through simulation of multi-agent interactions requires autonomous agents guided by realistically complex goals and mental models. We develop a framework that can build, parametrize, and simulate systems of such agents.

Background

Although there is a rich literature on the measurement and modeling of human movements in dense crowds with simple goals (e.g., moving to an exit, e.g., Helbing et al., 2000) the associated "**people as homogenous particles**" approach is questionable in more common scenarios where agents with **diverse** characteristics follow individual plans requiring navigation through a **complicated series of goals**.

We instantiate the same **operational level** functionality (relatively automatic abilities enabling movement towards goals while avoiding obstacles) as social-force models (Campanella et al., 2014) in a utility-maximizing discrete-choice framework that models step decisions (Robin et al., 2009). We also add **strategic level** route finding (Larmet, 2019) and path planning (Hahsler & Hornik, 2007) abilities, allowing agents to plan and re-plan individualized series of goals.

Operational Model

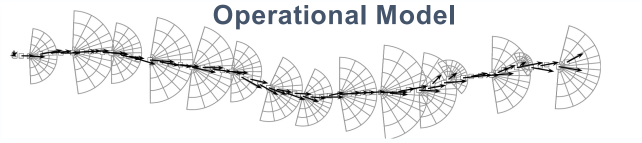

Random utility (U) guides k=33 (11 direction "cones" x 3 velocity "rings": slow/constant/speed up) step choices (r_U = randomness).

$$U_{ik=}\left(U_{ik}^{PS}+U_{ik}^{GA}+U_{ik}^{CA}+U_{ik}^{ID}+U_{ik}^{BA}+U_{ik}^{FL}+U_{ik}^{WB}\right)/r_U$$

Utility is a sum of scaled (*b*) power functions (exponent *a*) of absolute differences (d≥0) between the current position (*meters*) or velocity (i.e., speed, *m/sec*, & angle, *degrees/90*) and **avoidance** or **approach** goals. *Scoring disutility = - utility:*

Repulsion Disutility = b/d^a ***Attraction Disutility = bd^a***

Individual Attraction Utilities

PS = Individuals preferred speed (threshold linear slowing near goals)

GA = Goal angle (tendency to head to goals)

CA = Current angle (tendency to continue ahead, with different weights for left and right sides to account for side preference)

Social Repulsion Utilities

ID = Interpersonal distance (minus infinity at body overlap)

BA = Blocked angle (avoid cones with lots of pedestrians)

Social Attraction Utilities

FL = Follow the leader (promotes lane forming in crowded situations)

WB = Walk beside (social group dependent, also biases FL)

These 7 components produce plausible behavior in complex scenarios (e.g., **shopping in a supermarket**), and the framework can be easily extended to add on new behaviors appropriate for specific contexts (e.g., avoiding visibility to other agents).

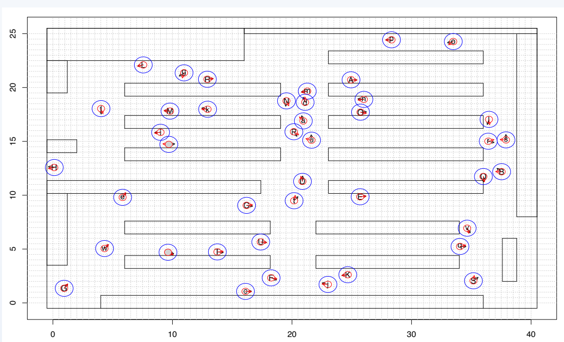

Figure 1. Agents moving around a 40x25m supermarket (black rectangles are shelves). Red circle = 0.6m diameter body; blue circle = 1.5m diameter; red/black arrow = goal/current direction (disparities between arrows occur where pedestrians are avoiding each other); grey body = pause at goal. Agents enter on the bottom left and exit through one of two exits on the middle/top left.

Strategic Model

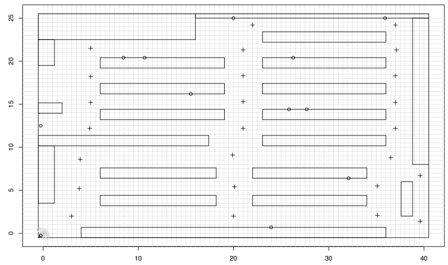

Figure 2. Initial configuration with an agent entering. "+" symbols are waypoints and "o" symbols are the agent's shopping goals. Simulation studies conducted by assigning each agent a randomly selected set of goals to satisfy showed the operational model was competent at satisfying goals while avoiding collisions and avoiding gridlock even at high densities.

Each agent has a sequence of of spatially defined goals of two types:
1) "must visit" goals satisfied when moving within a threshold distance
2) "way-point" goals satisfied when the following goal becomes visible

Route-finding and path optimization (i.e., "traveling salesman") algorithms can be used to initially build sequences ("goal stacks") that satisfy sets of must-visit goals and other constraints such as one-way regions, and random perturbation used to mimic sub-optimal solutions.

Operational factors compromising path plans (e.g., being pushed off course by other pedestrians) can be addressed by re-planning using the same methods or by path re-tracing to the last point where the next goal was visible. Re-planning can also be used with changing conditions (e.g., crowds blocking the planned path at high densities).

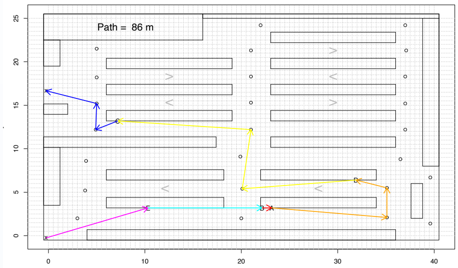

Figure 3. Path plan satisfying 5 must-visit goals (LETTERS) with one-way constraints (grey arrows). Naturalistic paths were obtained by a simple greedy algorithm and allowing way-point goals to be satisfied before the agent gets close to them (e.g., in the orange path the two way-point goals will be satisfied early to produce a smoother trajectory around the end of the aisle).

Conclusions and Future Directions

Step-choices based on utility provide a flexible framework to provide minds for mobile agents, allowing them to operate competently and independently in dynamic multi-agent environments.

The discrete-choice framework enables likelihood-based fitting, providing parameter estimates that illuminate the psychological factors changing behavior in different contexts.

Goal stacks augmented with re-planning provide a flexible means of instantiating dynamic control of complex spatial navigation plans.

The model is being used in projects investigating the effects of space design and movement rules on social distancing and virus spread.

Future work will use position data from movement experiments to quantify individual differences, enabling the model to be calibrated for veridical simulations of complex real-world scenarios.

The same framework is being used to develop cognitive models of strategic decision making in a navel escort simulation task requiring participants to protect a high value target from a submarine.

References

Campanella, M., Hoogendoorn, S., & Daamen, W. (2014). The Nomad model: Theory, developments and applications. *Transportation Research Procedia*, *2*, 462–467.

Hahsler M, & Hornik K. (2007). TSP - Infrastructure for the traveling salesperson problem. *Journal of Statistical Software*, *23*, 1-21.

Helbing, D., Farkas, I., & Vicsek, T. (2000). Simulating dynamical features of escape panic. *Nature*, 407, 487–490.

Larmet V (2019). "cppRouting: Fast implementation of Dijkstra algorithm in R. https://github.com/vlarmet/cppRouting>.

Robin, T., Antonini, G., Bierlaire, M., & Cruz, J. (2009). Specification, estimation and validation of a pedestrian walking behavior model. *Transportation Research Part B*, *43*, 36–56.

Acknowledgements

https://dataversuscorona.com/
https://www.dst.defence.gov.au/partner-with-us/university/rn-uds