Proceedings of the 2009 IEEE International Conference on Systems, Man, and Cybernetics
San Antonio, TX, USA - October 2009

A Cognitive Model of Perceptual Path Planning in a Multi-Robot Control System

David Reitter, Christian Lebiere
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA, USA
reitter@cmu.edu

Michael Lewis, Huadong Wang, Zheng Ma
School of Information Sciences
University of Pittsburgh
Pittsburgh, PA, USA

***Abstract*—We discuss an experiment involving visual path planning for multiple, remote robots in a partially visible building, with a partial 2D map available. Participants in the experiment defined waypoints for each robot to circumnavigate obstacles and explore the building. A cognitively plausible model of visual planning is evaluated using a normalized metric of the fit between model and subject itineraries. We discuss variation in the data and model fit, indicating individual differences in strategies to cope with task demands.**

***Keywords*—Path Planning, Visual Navigation, Robot Control, Cognitive Modeling**

I. Introduction

Human path planning decisions are ubiquitous: whenever we move, be it within rooms, buildings or cities, a path along several waypoints is needed. While robots can solve such tasks and strive for optimality within well-defined constraints, human guidance is often needed in practical applications in order to optimize or even just satisfy a range of ill-defined, task-specific criteria. Understanding this process can guide the design of user interfaces for robot control, to allow human operators to more effectively balance tasks and work more efficiently. In this study, we design a cognitively motivated model to explain and predict visually guided path planning decisions by human operators in a multi-robot control experiment. We propose metrics to evaluate model fit and show that it illustrates differential behavior across subjects, whose strategies do not always scale with the workload.

II. Recent Work

A. Multi-robot Control

Data for this study came from one of the conditions of a multi-robot control experiment investigating how performance scaled with increasing numbers of robots. The study was conducted using USARSim [5], a high-fidelity robotic simulation originally developed for the study of human-robot interaction in urban search and rescue (USAR). USARSim uses a game engine to provide accurate physics and realistic synthetic video needed to replicate an operator's experience in controlling real robots. Validation studies showing agreement for a variety of feature extraction techniques between USARSim images and camera video are reported in [3]. Validation data showing close agreement in detection of walls and associated Hough transforms for a simulated Hokuyo laser range finder are described in [2]. The current UnrealEngine2 integrates MathEngine's Karma physics engine [7] to support high-fidelity rigid body simulation. A series of validation studies have shown close agreement in behavior between USARSim models and real robots being modeled [4, 6, 8, 9, 13].

B. Models of Human Path Planning

[16] present data showing how humans develop spatial plans in a 2D environment that contained a varied number of obstacles. Subject performance was measured by the time to find the path and the number of unnecessary steps taken compared to the optimal path to the goal (errors). The number of turns was found to be a crucial predictor of planning time. The model proposes that subjects choose locally optimal paths (a hill-climbing strategy) and minimize the number of turns; they do not develop a complete plan before committing to initial steps. The model presented in this paper is similar in that the cognitive level also prefers to backtrack locally in order to avoid long-term memory needs. In our visually guided model, local planning results in long, straight lines with few turns.

The integration of cognitive architectures with path finding approaches represents a third field of work. [17,18] aim to blend cognitive and perceptual factors in navigation and planning applications in their work on implementing diagrammatic reasoning in cognitive architectures. Within the context of the ACT-R architecture [26], much work has been done on the problem of spatial planning and representation, including issues of adaptivity in planning [19], and encoding of spatio-temporal stimuli [20]. In this paper, we focus on the case of route planning based on information that is available visually, rather than held in memory.

III. DATA COLLECTION EXPERIMENT

A. *Task*

The path-planning model was evaluated using human data collected in a multi-robot control study. Robots were sent on a search and rescue mission inside a prototypical building with rooms connected by hallways. Obstacles included office furniture, victims and other robots.

MrCS (Multi-robot Control System) is a multi-robot communications and control infrastructure with accompanying user interface developed for experiments in a multi-robot control and RoboCup competition [1] which was used to provide the experimental scenario that required human participants to remotely control a varying number of robots. MrCS provides facilities for starting and controlling robots in the simulation, displaying camera and laser output. Figure 1 shows a screenshot of MrCS. The operator selects the robot to be controlled from the colored thumbnails at the top right of the screen. Robots are tasked by assigning waypoints on a heading-up map on the Mission Panel (bottom left) or through a teleoperation widget (bottom right). The current locations and paths of the robots are shown on the Map Data Viewer (top left).

The "in focus" components of the MrCS interface closely resemble user interfaces used to control experimental USAR robots [11, 12]. The simulated laser range finder model provides 4 meters of coverage. Successive scans are matched to construct the map and localize the robot on it. The resulting map is in a single plane resembling an irregularly sliced pie showing open traversable areas but extending only as far as the laser has scanned. The operator's view through the camera, by contrast, is

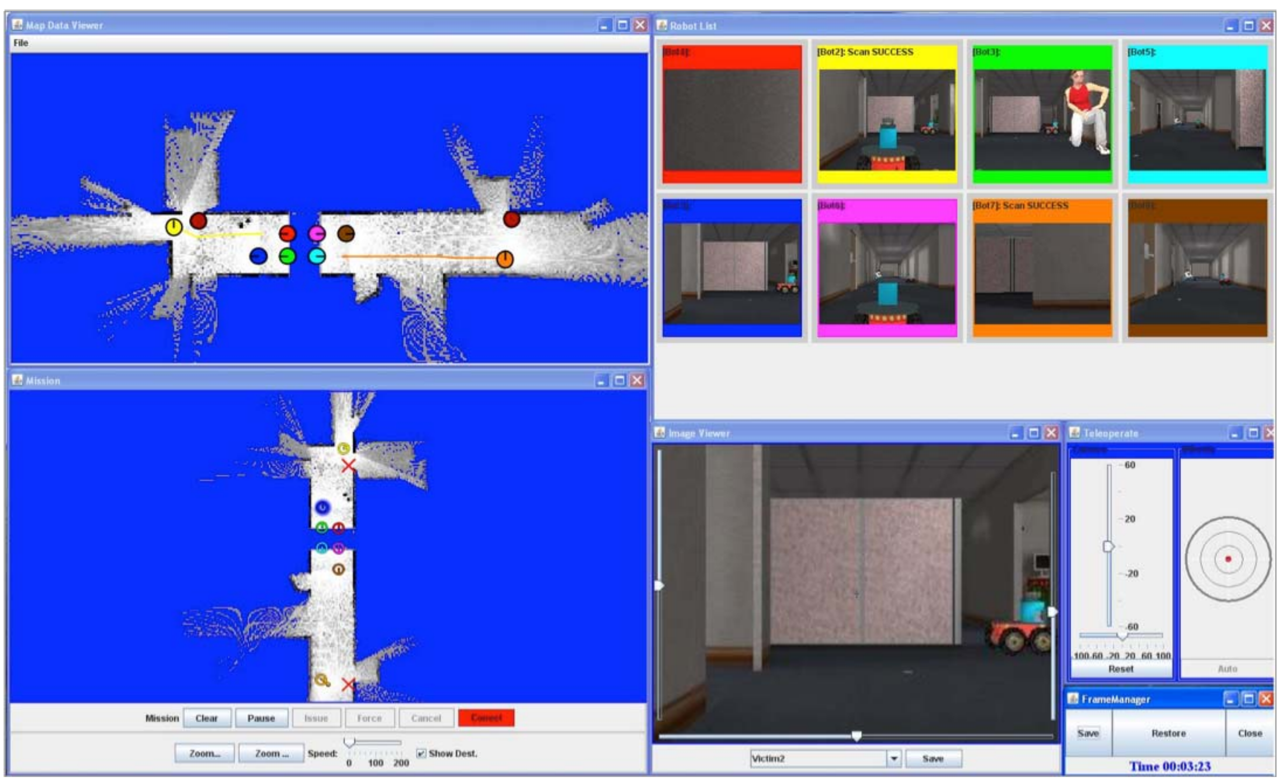

Figure 1: The MrCS graphical user interface.

limited only by obstacles. To explore an area, the operator may select waypoints on the map to specify a path or itinerary for the robot to follow. The robot then drives autonomously between these waypoints detecting and avoiding minor obstacles along the way.

To control multiple robots the operator must perform this sequence repeatedly by selecting waypoints for additional robots while the currently tasked robots are driving between their waypoints. This interface allows operators to control large number of robots in a foraging task (a search and rescue scenario). Other cooperative tasks such as box-pushing are highly demanding and can require an operator's full attention to control only two robots [23]. The demands of controlling multiple robots require operators to continually shift attention among robots, causing them to abruptly switch viewpoints without being able to maintain anchors to landmarks. As a consequence, path planning in this task could be expected to depend almost exclusively on the visual strategy. The resulting path plan entered by the operator as a series of waypoints on the laser-generated map provides an externalized analogue to the segmented paths to immediate goals generated by the visual planning model.

B. *Experiment*

45 subjects from the University of Pittsburgh community participated in the experiment for compensation.

Operators controlled 4, 8, and 12 robots in succession in a between-group repeated measures design. The experiment compared performance between a group performing a search and rescue foraging task, *full-task*, in which participants searched for victims and marked them on a map, to two groups performing only subtasks: participants in the *exploration* condition were asked to explore as large a region as possible but without the requirement to locate and mark victims; in the *perceptual search* condition robots "autonomously" followed trajectories recorded from the exploration condition while operators searched for and marked victims. Major findings of the study [10] were that full-task performance declined substantially between 8 and 12 robots, while perceptual search participants found twice as many victims using 12 robots. On other measures exploration performance was very similar to that of the full-task group with large increases in area explored between 4 and 8 and a flattening from 8 to 12. There were substantially more switches in focus between robots for full-task and exploration groups than for perceptual search. The number of itineraries followed a similar pattern for full-task and exploration and path lengths were nearly identical growing linearly from 4 to 12 robots. Workload ratings on the NASA Task Load Index (TLX, [25]) were high for the full-task and exploration conditions increasing in parallel with the number of robots while perceptual search workload was much lower. A broad summary would conclude that the navigation and path planning required in the exploration condition was the primary contributor to the difficulty and activities required to perform the full task.

Only participant data from the exploration condition were used to model human path-planning behavior in this study (14 subjects). The exploration condition was chosen because navigation and path planning were the sole objective of these participants making their task comparable to the goal-directed planning of the model.

IV. MODEL

A. *Overview*

Among the guiding hypotheses in developing the navigation model was the idea that most local planning problems can be solved visually, with minimal involvement of higher cognition (memory and rule-based algorithms). Memory plays a role when longer itineraries need to be planned without the immediate involvement of visual cues or when significant long-term experience is involved. An integrated model [14] is being developed to cover both aspects using an independently validated cognitive architecture; for the robot-control task, we focus on the visual planning aspects.

The visual system only has access to the part of the visual scene determined by its *visual attention* (cf., the visual representations in [21]). Originally designed to model shifts of visual attention in terms of eye movement saccades, that process in our model corresponds to waypoints designating locations to be connected by straight lines, at which participants set waypoints for the robots. Thus, our visual model represents possible paths to the next waypoint from the point of visual attention to reachable (immediate) goals. Given a first-person perspective, such lines would equal lines of sight; given a two-dimensional (2D) representation given here, possible paths are detected as straight lines that are uninterrupted by walls or other obstacles. The role of the perceptual model component is to identify traversable shapes and select promising ones: the adopted heuristic is to choose the path that brings a robots to a position from which the goal can be reached.

B. *Algorithm*

The visual navigation model always chooses the best route along a line of sight from the current location; the best route is the one that leaves us closest to the goal, without major obstacles in between. Routes that avoid bringing the model to previously visited parts of the territory (using the *visual finsts* mechanism of [22]) are preferred. If a route brings us away from the goal, then we are careful to detect alternative routes along the way: the model inspects the areas left and right at each step, stopping when there is an exit route. Such a way out is likely to be more useful than to retreat further from the target.

Concretely, the model implements the following algorithm:

1. Set visual attention to the start point.
2. Identify straight lines (uninterrupted by obstacles) in all directions from the current point of visual attention.
3. Choose the line that has the least density of walls and unexplored areas in between its end point and the goal (often this will be the line whose end point is the closest to the target).
4. Move visual attention along the chosen line towards the end point. While doing so: If entering a location that has been recently visited (visual finst), identify uninterrupted openings to the left and right, that is, straight lines beginning at the current point of attention and extending orthogonally to the current track. If openings are found, abandon the earlier movement and continue with point 3 (choosing the opening that is closer to the goal). If no openings are found, continue with point 4.
5. Set waypoint.
6. Terminate if goal reached, or if number of steps exceeds the maximum.

Line width (in step 2) is a parameter indicating the assumed width of a robot. Figure 2 shows a model path cutting a corner due to a robot width too low to emulate the subject's choice of waypoints.

The model predicts not only locations of the paths, but also the time required to shift attention from one location to another based on the EMMA model of eye movement [15].

The visual navigation model is conceived as a strategy applicable whenever the goal is deemed easy to reach. Such a goal could also be a sub-goal, with the visual strategy tying in to the cognitive strategy to find micro-solutions and guide exploratory behavior. It should be noted that visual navigation does not require the declarative, explicit storage of a branching point. Any backtracking is visually guided and constrained by visual finsts. As a consequence, the visual model inherits the limitations of visual memory: primarily, there is only a small number of finsts available. Without a memory-based component storing branching points explicitly, the visual model can only identify 5 or so locations as previously visited (if we accept the default assumption of about 5 finsts). The visual model alone may, given a sufficiently complex navigation task, even get caught in a loop, visiting the same locations again and again. The number of finsts chosen here has showed favorable results with our model in the context of maze navigation [14]. The application to 2D navigation in a building required only minimal changes.

C. *Evaluation Metric*

The model was run against subject-produced itineraries; an itinerary is defined as a start and end point (participants chose their end points) and a set of waypoints defined by the participant. Each itinerary was executed autonomously by the robot; it was defined with a static view of the laser-scan map (updated in between planning operations, as robots were scanning their surroundings). Thus, the itinerary provided a natural unit for the evaluation of a path-planning model. The model was given the map view as seen by the participant, as well as a start and end point (higher-level planning, including selecting robots and their target points is not modeled here). Other predictions that may arise from the model concern, primarily, timing. Such data are inherently noisier; reflecting interruptions that may occur unpredictably; also for space reasons we will focus on the planned itineraries.

Itineraries were split (by subjects) into a development (4 subjects) and a test set (10 subjects); results reported here were obtained using the test set only. We exclude data from runs where the model produced unreasonably long itineraries after timing out (due to getting stuck and visiting similar locations), identified as those model itineraries longer than twice the subject's itineraries. We also exclude cases of trivial path planning, where start and end points could be and were connected in a straight line by both subject and model (area near 0). The remaining data set comprised 393 itineraries.

To determine model fit, we propose to measure the area defined in between the two itineraries, that is, the polygon spanned by the concatenation of the itineraries, which share start and end points (see Figure 2). A large area implies a large deviation; a small area means that the itineraries match well; the area is a useful measure for the more complex case of trajectories, where a temporal dimension is added (e.g., [24]). We propose to normalize the area a by regressing out the respective effects of the model itineraries, obtaining a normalized area a_n. A simple approach would use normalization such as

$$a_n = \beta \frac{a}{len(path_M) len(path_S)}$$

A corresponding linear model estimating the size of β failed to account fully for the effect, since residuals (a_n) remained significantly correlated with a. We chose instead the model

$$a_n = \frac{a}{len(path_M)^{\beta_M} len(path_S)^{\beta_S}}$$

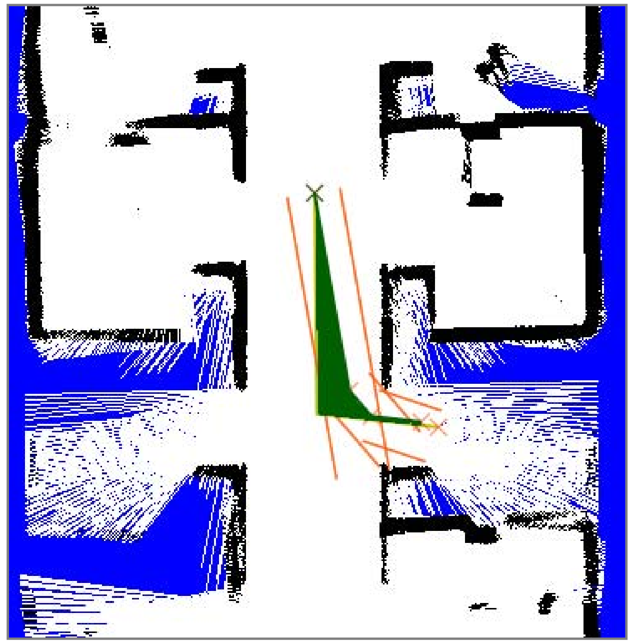

Figure 2: A subject and a model path, leading from a hallway into a room (on the right). Shaded area indicates model fit, straight lines parallel to path indicate width of robot as assumed by model. Markings (cross) show model's possible turn points.

This can be formulated as a linear mixed effects regression model (by subject), estimating β_S=1.55 (p<0.0001) and β_M =-0.08 (not significant). (We retain β_M to account for the general case, in which the model's and subject's path lengths may be less collinear.) The residuals a_n (plus random effects) resulting from this regression model showed no more correlation with either path length; variance in the normalized path divergence area a_n should now be attributed to model performance and participant behavior under controlled experimental conditions.

D. Results

The itineraries produced by model and subject correlated well across subjects (Pearson r=0.97 for subject means). This is, to a large extent, owed to the fact that models were given the same start and end points. Figure 3 shows the varying mean itinerary length over all subjects, indicating that most subjects chose path lengths around 6m.

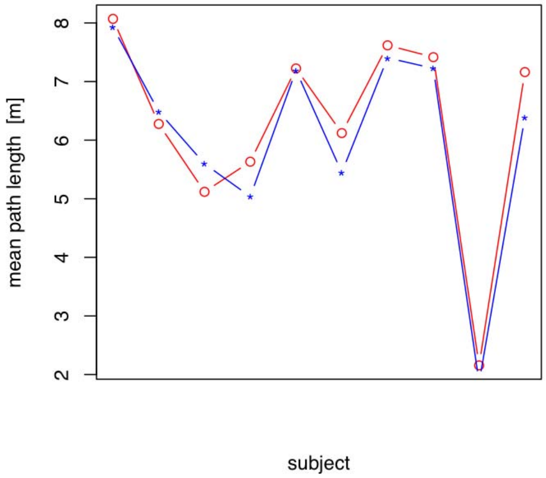

Figure 3: Mean path length of the itineraries for each over the subjects in the evaluation set. Subjects: blue *. Model: red o.

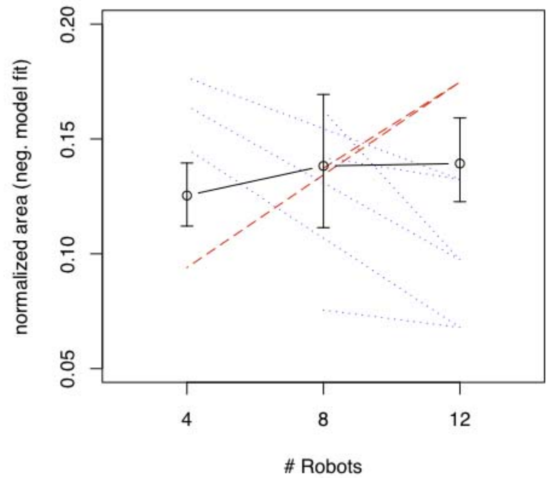

Figure 4: Normalized area indicating neg. model fit over conditions (4,8 and 12 robots), i.e. smaller areas indicate better model fit. Mean (solid black), and means for each subject in the evaluation set (dotted blue: decreasing, dashed red: increasing with # robots).

Comparative analysis (Figure 4) of the normalized area shows good overall model fit with relatively small normalized areas, and also across conditions with 4 and 8 robots. Normalized areas could not be shown to be significantly larger (which would indicate worse model fit) for 12 than for 4 robots (p=0.2, by subject). However, this effect seems to be primarily driven by just three subjects. When splitting subjects into two *strategy* groups, model fit degrades in the 12-robot condition for the larger subject group ($p<0.02$, by subject). Thus, it appears that subjects employ different strategies when task demands increase; which other strategy is chosen varies across subjects.

Itinerary length increased with the number of robots; specifically, participants put in longer itineraries in the 12-robot (μ=1.44m) condition than in the 4-robot condition (μ=1.78m, $p<0.0001$).

V. Discussion

Our relatively simple, perceptual model accounts for much of the path planning found in our data set. For most subjects, the visual path-planning model appears to be more accurate for planning situations under more demand (work load per se was not modeled). A qualitative analysis of cases where the model fits itineraries poorly suggests two common patterns. *First,* participants treat unexplored areas (blue/grey in the map) differently, either cutting through some of these areas, or avoiding them until a laser scan is obtained. The model expects unexplored areas to be accessible (any point's prior probability to be a wall is low), but it lacks inference about the structure of the building. *Second,* participants often chose to traverse corridors and doors in the center rather than cutting corners and merely optimizing distance. We interpret this as a risk minimization strategy to avoid contact with the wall. In the context of our model, this could be explained as an acquired but ultimately implictly internalized strategy rather than part of an explicitly managed set of constraints for the task. Optimizing a robot width parameter improved the fit during model development.

The analysis of the model fit suggests that subjects adapt their strategies as they are given a higher workload. Overall, their itineraries increase in length, although model fit does not significantly decrease with the number of robots even after the effect of path length has been regressed out, but more work is needed to confirm scalability across more subjects. Subjects vary in how they react to the increased workload. The strategy that subjects switch to when under stress appears to be characterized by longer itineraries in the 12-robot condition. This is a reasonable strategy, as it maximizes the time in which each robot will act autonomously.

Further work will have to show whether this strategy change is in the self-chosen start and end points (which would require integrating this visual model with a higher-level planning and allocation model) or due to the way the itineraries are planned. More work is needed to integrate further planning constraints and better awareness of hypothesized structural features of the environment (such as walls). Work is also needed to establish better baselines for human path planning behavior, i.e. to obtain variance information. For this, repeated measurements designs would be needed that exert full control over the problems that are given to subjects. We plan to run such experiments in a high-fidelity, first-person perspective environment, which will require the combination of memory-based and visual planning.

VI. CONCLUSIONS

We have implemented a model of human spatial path planning and evaluated it in the context of a multi-robot control task. Spatial planning in a visually fully grounded situation varies with workload, but also draws on subject-specific strategies. However, even a simple, perceptual model can account for a range of data. If such a model can be validated further, then it suggests that humans avoid in most cases the more costly and time-consuming global planning route in favor of a visually grounded local path-planning algorithm.

REFERENCES

[1] S. Balakirsky, S. Carpin, A. Kleiner, M. Lewis, A. Visser, J. Wang and V. Zipara. Toward hetereogeneous robot teams for disaster mitigation: Results and performance metrics from RoboCup Rescue, Journal of Field Robotics, 24 (11-12), 943-967, 2007.

[2] S. Carpin, J. Wang, M. Lewis, A. Birk and A. Jacoff. High fidelity tools for rescue robotics: Results and perspectives, Robocup 2005 Symposium, Osaka, Japan, 2005.

[3] S. Carpin, T. Stoyanov, Y. Nevatia, M. Lewis and J. Wang. Quantitative assessments of USARSim accuracy". Proceedings of PerMIS 2006, Gaithersburg, MD 2006.

[4] S. Carpin, M. Lewis, J. Wang, S. Balakirsky, C. Scrapper. (2006b). Bridging the gap between simulation and reality in urban search and rescue. Robocup 2006: Robot Soccer World Cup X, Springer, LNAI, 2006.

[5] M. Lewis, J. Wang, and S. Hughes, USARsim : Simulation for the Study of Human-Robot Interaction, Journal of Cognitive Engineering and Decision Making, 1(1), 98-120, 2007.

[6] M. Lewis, S. Hughes, J. Wang, M. Koes, and S. Carpin. Validating USARsim for use in HRI research, Proceedings of the 49th Annual Meeting of the Human Factors and Ergonomics Society, pp. 457-461, Orlando, FL, 2005.

[7] Mathengine, MathEngine Karma User Guide, http:// udn.epicgames.com/Two/rsrc/Two/KarmaReference/KarmaUserGuide.pdf , accessed March 26, 2009.

[8] C. Pepper, S. Balakirsky, and C. Scrapper. Robot Simulation Physics Validation, Proceedings of PerMIS'07, Washington, D.C., 2007.

[9] B. Taylor, S. Balakirsky, E. Messina and R. Quinn. Design and Validation of a Whegs Robot in USARSim, Proceedings of PerMIS'07, Washington, D.C., 2007.

[10] H. Wang, M. Lewis, P. Velagapudi, P. Scerri, and K. Sycara, How search and its subtasks scale in N robots, 2009 Human-Robot Interaction Conference, ACM.

[11] H. Yanco and J. Drury, "Where am I?" Acquiring situation awareness using a remote robot platform. In Proceedings of the IEEE Conference on Systems, Man, and Cybernetics, 2004.

[12] H. Yanco, M. Baker, R. Casey, B. Keyes, P. Thoren, J. Drury, D. Few, D., C. Nielsen and D. Bruemmer, "Analysis of human-robot interaction for urban search and rescue, Proceedings of PERMIS, 2006.

[13] M. Zaratti, M. Fratarcangeli and L. Iocchi. A 3D Simulator of Multiple Legged Robots based on USARSim. Robocup 2006: Robot Soccer World Cup X, Springer, LNAI, 2006.

[14] D. Reitter and C. Lebiere. A subsymbolic and visual model of spatial path planning. In Proc. Behavior Representation in Modeling and Simulation (BRIMS), Provo, UT, 2009.

[15] D. D. Salvucci. An integrated model of eye movements and visual encoding. Cognitive Systems Research, 1:201–220, 2001.

[16] D. Fum and F. del Missier. Climbing the mazes: A cognitive model of spatial planning. In Proceedings of the Third International Conference on Cognitive Modeling, pp. 126–133, Veenendaal, The Netherlands, 2000. Universal Press.

[17] B. Chandrasekaran. Multimodal cognitive architecture: Making perception more central to intelligent behavior. In Proceedings of the AAAI National Conference on Artificial Intelligence, pp. 1508–1512, Boston, MA, 2006.

[18] H. A. Dye. A diagrammatic reasoning: Route planning on maps with ACT-R. In Proceedings of the Eigth International Conference on Cognitive Modeling, Ann Arbor, MI, 2007.

[19] W.-T. Fu. An ACT-R adaptive planner in a simple map-navigation task. In F. Detje, D. Doerner, and H. Schaub, editors, Proceedings of the Fifth International Conference on Cognitive Modeling, pages 99–104, Bamberg, Germany, 2003.

[20] T. Johnson, H. Wang, J. Zhang, and Y. Wang. A model of spatio-temporal coding of memory for multidimensional stimuli. In Proceedings of the 24th Annual Meeting of the Cognitive Science Society, Fairfax, VA, 2002.

[21] J. Glasgow and D. Papadias. Computational imagery. In P. Thagard, editor, Mind Readings. MIT Press, Cambridge, MA, 1998.

[22] Z. W. Pylyshyn. The role of location indexes in spatial perception: A sketch of the finst spatial-index model. Cognition, 32:65–97, 1989.

[23] J. Wang. Human Control of Cooperating Robots. PhD dissertation, University of Pittsburgh, 2007.

[24] N. Pelekis, I. Kopanakis, I. Ntoutsi, G. Marketos, G. Andrienko and Y. Theodoridis. Similarity Search in Trajectory Databases. In: Proceedings of the 14th IEEE International Symposium on Temporal Representation and Reasoning (TIME 2007), Alicante, Spain, 2007.

[25] S. G. Hart and L. E. Stavenland, "Development of NASA-TLX (task load index): Results of empirical and theoretical research," in Human Mental Workload, P. A. Hancock and N. Meshkati, Eds., pages 139-183. Elsevier, 1998.

[26] J.R. Anderson. How can the human mind occur in the physical universe? New York, NY: Oxford University Press, 2007.