A Neurocomputational Approach to Modeling Human Behavior in Simulated Unmanned Aerial Search Tasks

Brandon S. Perelman (BPerelman@gm.slc.edu)
Department of Cognitive and Learning Sciences, 1400 Townsend Drive
Houghton, MI 49931 USA

Shane T. Mueller (ShaneM@mtu.edu)
Department of Cognitive and Learning Sciences, 1400 Townsend Drive
Houghton, MI 49931 USA

Abstract

In the present study, we used a synthetic task environment representing an aerial wilderness search and rescue task to test and validate neurocomputational models of the human hippocampus. Participants completed two tasks: a search task, in which they searched a wooded area for multiple targets, and a route choice task, in which they searched high and low density probability regions, parametrically varied by relative distance, for a single target. The decision task was designed to examine participant evaluation of cost versus expected reward. We compared participant behavior during this task with (1) an optimal model, which makes decisions based on a cost to reward ratio, and (2) a neurocomputational model based on hippocampal architecture, which generates activation from probability weighted reward sites. We discuss initial results using this paradigm and the bases by which we will model performance during unmanned aerial search for multiple targets.

Keywords: UAV; wilderness search and rescue; hippocampus, neural network; navigation; synthetic task environment.

Unmanned Aerial Vehicles in Wilderness Search and Rescue

Wilderness search and rescue (WiSAR) operations consist of time sensitive searches for missing persons lost in rugged terrain. Due to the potentially large size of the search areas, WiSAR operations often benefit greatly from aerial assets, such as rotary and fixed wing aircraft (Hoekstra, M. [West Michigan Search and Rescue], personal communication, July 25, 2012). Comparatively low operational cost requirements and other domain specific factors have catalyzed research on the feasibility of unmanned aerial vehicles (UAVs) for WiSAR operations (e.g., Adams et al., 2007; Adams et al., 2009; Goodrich et al., 2008; Goodrich, Morse, Engh, Cooper, & Adams, 2009).

One goal of this research is to minimize the number of personnel required for UAV operations. This is accomplished via interface improvements (e.g., Cooper & Goodrich, 2008) and efforts to optimize the means by which operators conduct their search, often based loosely on Bayesian approaches to modeling missing person behavior (e.g. Lin & Goodrich, 2009; 2010). To examine human search behavior using a UAV, we created a synthetic task environment (STE; Perelman & Mueller, 2013), using the Psychology Experiment Building Language, approximating the cognitive requirements of WiSAR using UAVs.

The present study represents an effort to apply neural models of the hippocampus, primarily used to study rodent navigation, to human tasks. First, such an effort offers to inform our understanding of the unmanned aerial task. Neurocomputational models permit the study of human factors memory effects in this domain. Second, the United States Air Force, as one of the primary end users of UAV technology, has expressed interest regarding the possible use of transcranial magnetic stimulation, a technique informed by research in network dynamics, to improve pilot performance in UAV tasks (e.g. AFRL, 2007). Third, biologically-inspired neural networks form a useful bridge between theory and behavior. Finally, a goal of UAV research is to permit a single operator to control multiple UAVs. The present STE and modeling afford a foundation toward studying multiple unmanned aerial agent navigation.

Assessing Human Behavior during Unmanned Aerial Search

STEs are task-centric simulations (Cook & Shope, 2004) that sacrifice physical fidelity for a great deal of control and flexibility in the experimental examination of cognitive phenomena. Many of the existing UAV flight simulators (e.g., Aviones) replicate the instrumentation involved in UAV control. A lower fidelity STE allows us to collect data from undergraduate populations without the training necessitated by more complex environments. This STE is designed to permit the study of memory-related problems, such as those noted by Cooper & Goodrich (2008), and decision making in the unmanned aerial systems domain.

The present STE (Figure 1) is intended to approximate the cognitive task requirements of WiSAR using UAVs. Participants control the UAV continuously by indicating a destination point on the north up topographical map (left). Once the UAV reaches the destination point (red target), it will orbit until a new destination is selected. Simultaneously, participants search for targets (blue tents) using the track up satellite image representing the view through the UAV's sensor package (i.e., camera window, right). Flight duration is governed by a parameterized fuel level depicted as a numeral and a gauge (top right).

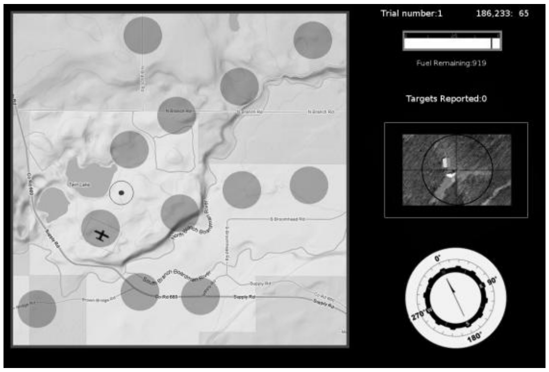

Figure 1. The STE during the multiple target search task.

Experimental Method

The STE is currently capable of collecting data in two paradigms: a multiple target search task, and a probability density decision task. Two experiments were conducted to examine participants' behavior in these tasks. The two experiments were identical with the exception of the number of participants (30 in the first experiment and 12 in the second) and the parameters of the decision test. Participants were recruited from the Michigan Technological University undergraduate participant pool. We will only report basic results from the multiple target search task, and focus in greater detail on the decision test, discussing results from both experiments together.

Multiple Target Search Task

Method

In each of the five trials in this task, participants searched for six targets amongst the 12 possible target locations. This layout was designed as an analogue to the use of probability maps in WiSAR (i.e., heat maps of probable missing person locations; Ferguson, 2008; Lin & Goodrich, 2010; Perkins, Roberts, & Feeney, 2003). Participants were given sufficient fuel to cover roughly 40 % of the search area. Fuel was limited to encourage participants to use maximally efficient search strategies. Parameters of this task were identical during both the first and second experiment.

This task requires that participants solve a version of the Euclidian traveling salesman problem (E-TSP; Graham, Joshi, & Pizlo, 2000). Though the traveling salesman problem (including E-TSP) is considered computationally intractable, humans produce near optimal solutions to this problem (e.g., Graham et al., 2000; MacGregor & Ormerod, 1996). Currently, our models have the ability to solve this task in a ground based scenario (cf. Mueller, Perelman, & Simpkins, in press), but they have not yet been validated for aerial search dynamics. Consequently, we will primarily focus on modeling and examining the decision task. However, we will report some basic behavioral results from this more complex search task as well.

Results

Participants exhibited an effect of learning whereby performance improved in terms of percent map explored, $F(4, 22) = 4.60$, $p = .008$, and targets flown over, $F(4, 26) = 4.32$, $p = .008$. Participants identified roughly half of the targets on each trial ($M = 2.97$, $SD = 1.32$), and remembered more than half of these ($M = 1.88$, $SD = 1.45$). Generally, targets that were identified later in the trial were more likely to be remembered, but the great degree of variance in the number of targets that were identified precludes meaningful quantitative analysis (results not shown).

Importantly, the search task requires making numerous small scale decisions about potential cost-reward tradeoffs. Thus, our initial models will focus on a simpler decision test that examines how these tradeoffs are evaluated.

Decision Test

Method

The second task, a probability density decision task (hereafter referred to as the decision test; Figure 2), was created to test participants' evaluation of cost versus reward for regions of varied probability density. Participants were asked to search a rectangular high density probability region, depicted as an oasis, which could be exhaustively searched by flying over a single point, and a long linear low density probability region, depicted as a road, which would require a greater period of time to search exhaustively. Participants searched for a single target, and the probability of the target spawning in either region was randomized but equal. Each of tests was conducted in a random order. Target location was chosen independently on each trial, appearing on the road or oasis with equal likelihood. The critical dependent variable of interest was whether the participant searched the road or oasis first. In all cases, this could be determined unambiguously using simple rule based analyses. Since the placement parameters of each reward feature (i.e., the road and the oasis) in Experiment 2 were based observed behavior in Experiment 1, we will discuss these results separately.

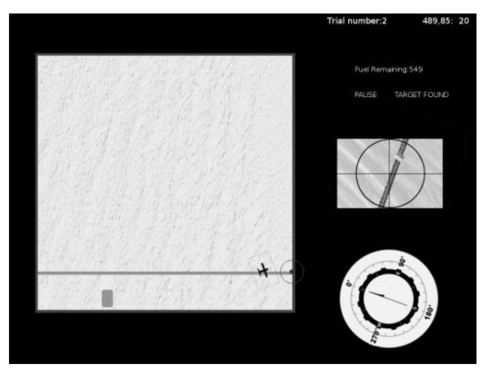

Figure 2. The STE configuration during the decision test. The UAV is currently flying over the road.

Experiment 1 Results

30 participants completed nine trials in a 3 (road: close, medium, far) x 3 (oasis: close, medium, far) design, with the oasis and road distances from the starting location parametrically varied and presented in a random order (see Figure 3, panel A). Each panel in the figure shows the ratio of participants who visited the road versus the oasis. Participants exhibited a clear bias toward the oasis in most conditions, even when the oasis was at its farthest and the road was at its closest (making it reasonable to follow the road to the oasis). Logistic regression determined that both road (p = .036, OR = .59) and oasis (p = .012, OR = 1.90) location reliably predicted participants' navigation decisions in each trial of this task, X^2 (1, 270) = 11.40, p = .003. Data regarding participants' preferences are available in Table 1. To understand this tradeoff better, we conducted a second experiment in which we varied the road and oasis distances based upon behavior observed in this experiment.

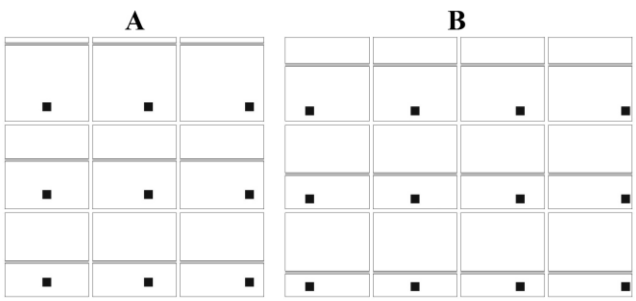

Figure 3. Feature locations in Experiment 1 (Panel A) and Experiment 2 (Panel B). Road and oasis are depicted as a horizontal bar and square, respectively. The UAV started each trial in the lower left hand corner.

Table 1: Proportion of Participants choosing each Feature in Experiment 1 (Road / Oasis)

	Oasis Close	Oasis Med.	Oasis Far
Road Far	0 / 30	2 / 28	4 / 26
Road Med.	1 / 29	5 / 25	4 / 26
Road Close	3 / 27	5 / 25	7 / 23

Experiment 2 Results

In order to investigate more deeply participants' evaluations of cost versus reward at various distances, the above experiment was repeated, using 12 participants, with the distances varied in 12 randomly presented trials according to a 3 (road: close, medium, far) x 4 (oasis: very close, close, medium, far) design as shown in Figure 3, panel B. Generally, the distances associated with each condition were adjusted so that participants' responses would follow an expected gradient. Participants seemed to exhibit a strong preference toward exploring the high density probability region (i.e., the oasis). The majority of participants favored exploring the road over the oasis in only two conditions, where the road was at its closest and the oasis at its farthest (see Table 2).

Logistic regression found that road (p < .001, OR = .34) and oasis (p = .001, OR = 1.90) locations reliably predicted navigation decisions in each trial of this task, X^2 (1, 144) = 30.56, p < .001.

Table 2: Proportion of Participants choosing each Feature in Experiment 2 (Road / Oasis)

	Oasis Very Close	Oasis Close	Oasis Med	Oasis Far
Road Far	0 / 12	1 / 11	3 / 9	4 / 8
Road Med.	0 / 12	2 / 10	5 / 7	6 / 6
Road Close	6 / 6	6 / 6	7 / 5	8 / 4

Modeling Participants' Behavior in the Decision Test

The results of the decision test showed that participants had a strong preference to search the high density location, even when it might be better to take a longer route. We will first examine how a model of pathfinding based upon neural architecture makes this tradeoff. Next, we will examine the optimal model's preferences to understand the types of information human participants use to make their decisions.

Neurocomputational Model

The neurocomputational model used in the present study is based upon a simple spreading activation model, described in greater detail by Mueller et al. (in press), that represents aspects of the hippocampal architecture and interconnectivity in *cornu ammonis* layers 1 (CA1) and 3 (CA3) (see Samsonovich & Ascoli, 2005). In this biologically-inspired model of the hippocampus, the CA3 cell layer encodes possible locations in the environment. The interconnections amongst these cells, by their weighted associations, form a cognitive map of the environment. The environment here is represented by a 25 x 25 grid space. In animals, hippocampal place cell ensembles selectively fire at a high rate when the organism is in a specific location in the environment (O'Keefe & Nadel, 1978). In this model, each grid cell (in this context, grid cell refers to cells within a grid space, distinct from entorhinal cortical grid cells) represents a place cell ensemble. With exploration, the connection weights amongst the grid cells (i.e., CA3 place cells) grow to represent known adjacent locations between which travel is possible. This allows experiential formation of a cognitive map of the environment.

The model also includes a CA1 layer, which contains nodes that represent contexts comprised of sets of past experiences. In that sense, CA1 cells represent the aggregate features of discrete contexts to which an organism may like to return. For every location in the environment, the CA1-

CA3 connection associates that goal context with a specific location in the cognitive map (i.e., grid space), which is defined by the CA3-CA3 interconnections. When a context is selected as a destination, the CA1 cell activates its associated CA3 cell, and activation diffuses throughout the CA3 network creating a “goal scent” emanating from the goal location. Activating multiple goal locations simultaneously results in a topographical landscape of values, ranging from 0 to 1, which an agent can navigate using simple hill climbing (Figure 4) to locations of perceived reward (i.e., goal sites). This mechanism permits the agent to navigate to the goal location in the CA3 layer associated with the goal context in the CA1 layer. Although this algorithm is fairly greedy, neighboring goal sites will reinforce one another leading to a small boost in preference for contiguous goal areas.

For the purpose of modeling the present task (i.e., an aerial search task), we have retained characteristics of the model that Mueller et al. (in press) applied to a human ground-based search task. In that model, the CA3 cells were assumed to be interconnected as the participants were already familiar with the environment. In the present model, we also assume that all of the CA3 cells are interconnected as, for the UAV, there are no boundaries to movement (i.e., direct routes are possible between any two locations in the environment), and the UAV operator can see the entire map without having to learn routes. Therefore, decisions are driven solely by a map of the potential reward locations, (i.e., oasis and road locations) weighted by estimates of the reward probability.

Within this 25 x 25 grid, the road occupied 1 x 25 cells and the oasis 3 x 3 cells. Since the probabilities of finding the target in either feature location were equal, the probability of finding a target in each road cell was .05 / 25 = .02, and the probability of finding a target in each oasis cell was .05 / 9 = .056, and the activation in each of these grid cells was parameterized appropriately. Therefore, the intensity of the initial activation driving the spreading activation (i.e., goal scent) in each grid cell is dependent that grid cell’s probability of containing the target.

Results (see Figure 4) show that the model has a strong preference for the road, unless the oasis location is especially close. This is because the model’s decisions are driven heavily by the initial strongest contact with distance-discounted reward values, and the difference in probability between individual road locations and the oasis cells contributes little to the decision. This is in stark contrast to our behavioral results. Consequently, we computed an optimal path model, to understand whether humans were on average better or worse than the model.

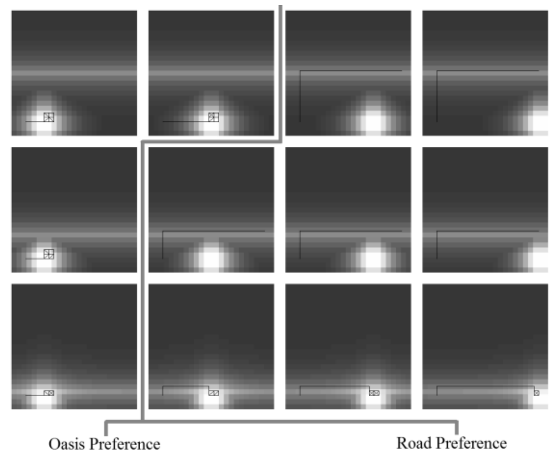

Figure 4. Neurocomputational model results for the decision test parameterized to fit feature locations in Experiment 2. Activation is displayed as a heat map and the agent’s path, originating in the lower left hand corner, is shown in black. Agent generally prefers exploring the road first.

Optimal Model

In the decision test, there is some ambiguity regarding the nature of an optimal decision. If, for instance, one had a limited but unknown amount of fuel or time (e.g., if nightfall was approaching), it might be optimal to search the highest density regions before fuel runs out. Alternately, if one assumes ample fuel and time, then perhaps the optimal route is the most efficient. In the present task, an optimal model is concerned with the complete search path, including the distance and time needed if the target is not found on the first leg of the path. Consequently, we computed expected search time for two search paths on each map: one that visited the oasis first, followed by the road, and a second that visited the road, then the oasis. In each case, the route that produced the smallest expected exhaustive search time was chosen as the optimal path.

The results of the optimal model are shown in Figure 5, along with the results of the neurocomputational model and the average trajectory of our participants. The optimal model is shown with a hashed red line as a ratio of costs between the two possible routes, with a 1:1 cost ratio represented as a diagonal line. Here, the best decision depends on whether the optimal ratio is above or below the diagonal. In nearly every case, the majority of participants are consistent with the direction of the optimal model. Thus, although individual participants did not always behave optimally, they were more efficient in aggregate than the neurocomputational model. Furthermore, there is a very strong correspondence between the optimal cost ratio and the proportion of participants choosing each option, as seen in Table 3.

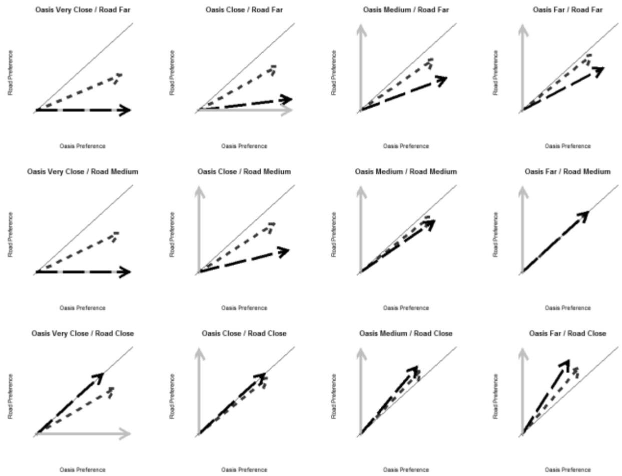

Figure 5. Arrow plots for each of the 12 conditions in the decision test, Experiment 2. The angle of the arrow indicates the proportion of participants choosing the road versus the oasis (lined arrow), the cost ratio produced by the optimal model (short dashed arrow), or the choice made by the neurocomputational model (solid arrow), compared with no preference (45 degree line). 90 and 0 degrees indicate absolute preference for the road or oasis, respectively. Panels from left to right, and bottom to top, indicate increasing oasis and road distances, respectively. As the oasis and road distances increase within series, agents' preferences for that feature decreases.

Across both experiments there were extremely strong relationships between the optimal model ratio and the proportion of participants making each choice. Relationships between the data and the neurocomputational model are weaker but still significant (see Table 3).

Table 3: Regressions amongst Agents

	Optimal : Participants	Optimal : NC Model	Participants : NC Model
Exp. 1	R^2 = .83, p < .001	R^2 = .55, p = .023	R^2 = .69, p = .005
Exp. 2	R^2 = .79, p < .001	R^2 = .65, p = .001	R^2 = .36, p = .039

This may indicate a decision process with noise and is reminiscent of the literature on probability matching in signal detection theory (see Mueller & Weidemann, 2008), where most participants compute something close to the optimal model, but are subject to small deviations that influence them one way or the other. We attribute this disparity amongst agents' performance to the fact that the optimal model and humans have access to the entire map space, whereas the neurocomputational model relies upon local decisions using diffuse information.

Alternative Neurocomputational Model

Our initial computational model uses source activation driven directly by the correct probabilistic search environment. This assumption produces behavior different from participants and the optimal model. It is possible to parameterize the neurocomputational model to produce behavior more similar to the data and optimal decision, by reducing the source activation for the road. To do this, the road source activation needs to be reduced by a factor of 50. Although it is possible to fix the model in this particular environment, it is unlikely that this would translate generally to other configurations of reward, distance, and probability.

Discussion

The results of the present study indicate that the majority of participants made decisions that were optimal and, in aggregate, these decisions were related to the likelihood ratio of the complete alternative search paths. Our initial neurocomputational model, driven by simple activation dynamics, failed to approximate this probability-matched ratio when appropriately parameterized, and only succeed when the lower probability, but closer, path was discounted heavily. This suggests that the simple activation hill climbing model is also inappropriate for multi-goal search (as described in the search task here and in Mueller et al., in press). This failure likely stems from the fact that the neurocomputational model does not assess the global reward associated with an entire search plan, but rather is biased by first contact with search probability.

Global versus Local Solutions to Search

Although a local decision model can perform reasonably well in search tasks, these findings suggest that mechanisms for representing and computing the overall value of a planned search must be considered. This is also the primary insight of work on human solutions to the E-TSP problem, in which local and simple heuristic solutions tend to fail, in that they are both less accurate than humans and less efficient.

These results suggest several possible additions to the model that would enable better performance. First, previous research has shown that a foveating pyramid architecture would enable better global solutions to multiple object pathfinding (Pizlo et al., 2006). Using such a scheme, activation dynamics could be used to independently solve routing problems at different scales. Thus, at a coarse scale, the agent may use activation driven by higher level nodes, representing target clusters, to determine which part of the map to visit first. Such a hierarchical scheme may be important for many types of search problems.

A second possibility (compatible with the hierarchical scheme) is to implement means for spatial planning and mental simulation. Such a model would require means for representing, generating, and evaluating a plan. A spatial plan could be represented in a hippocampal model, via a short term memory buffer, for a sequence of spatial goals (Jensen & Lisman, 1996). Plans could be generated via mentally simulating the spatial context change using the hierarchical method discussed earlier. Finally, evaluation of the plan could build on the basic local spreading activation mechanisms we used in the current model, integrated over simulated steps. It should be noted that low level planning has already been explored in models of hippocampal activity, and aspects of phase precession within theta rhythms have been hypothesized as a mechanism for evaluating possible alternative goal states (Samsonovich & Ascoli, 2005; Koene, Gorchetchnikov, Cannon, & Hasselmo, 2003).

Summary

In this paper, we examined the ability of a simple goal activated spatial map to guide search behavior. Results showed that although the model was capable of guiding search, it did not capture basic human behavioral patterns, which indicated preference for reward probability over the weight of initial contact with activation from distant rewards. Human decisions, in aggregate, were highly related to the expected utility ratio between alternate search plans, suggesting humans were evaluating complete search plans and making decisions based on the entire flight path. Results suggest that the basic spatial map model must incorporate capabilities to represent and evaluate more complete search plans in order to perform more optimally, and in order to provide an accurate account of human performance.

References

Adams, J. A., Cooper, J. L., Goodrich, M. A., Humphrey, C., Quigley, M., Buss, B. G., & Morse, B. S. (2007). Camera-equipped mini UAVs for wilderness search support: Task analysis and lessons from field trials. *Journal of Field Robotics, 25,* 1-2.

Adams, J. A., Humphrey, C. M., Goodrich, M. A., Cooper, J. L., Morse, B. S., Engh, C., & Rasmussen, N. (2009). Cognitive task analysis for developing unmanned aerial vehicle wilderness search support. *Journal of Cognitive Engineering and Decision Making, 3,* 1-26.

Cooke, N. J. & Shope, S. M. (2004). Designing a synthetic task environment. In Schiflett, S. G., Elliott, L. R., Salas, E., & Coovert, M. D. (Eds.), *Scaled Worlds: Development, Validation, and Application* (pp. 263-278). Surry, England: Ashgate.

Ferguson, D. (2008). GIS for wilderness search and rescue. In *Proceedings of the ESRI Federal User Conference*. Washington, D. C.

Goodrich, M. A., & Cooper, J. L. (2008). Towards combining UAV and sensor operator roles in UAV-enabled visual search. In *Proceedings of the 3rd ACM/IEEE International Conference on Human-Robot Interaction*. Amsterdam, the Netherlands.

Goodrich, M. A., Morse, B. S., Engh, C., Cooper, J. L., & Adams, J. A. (2009). Towards using UAVs in wilderness search and rescue: Lessons from field trials. *Interaction Studies*, *10*, 453-478.

Graham, S. M., Joshi, A., & Pizlo, Z. (2000). The traveling salesman problem: A hierarchical model. *Memory & Cognition, 28,* 1191-1204.

Jensen, O. & Lisman, J. E. (1996). Hippocampal CA3 region predicts memory sequences: Accounting for the past precession of place cells. *Learning & Memory, 3,* 279-287.

Koene, R. A., Gorchetchnikov, A., Cannon, R. C., & Hasselmo, M. E. (2003). Modeling goal-directed spatial navigation in the rat based on physiological data from the hippocampal formation. *Neural Networks, 16*, 577-584.

Lin, L. & Goodrich, M. A. (2009). UAV intelligent path planning for wilderness search and rescue. In *IEEE International Conference on Intelligent Robots and Systems*. St. Louis, MO.

Lin, L. & Goodrich, M. A. (2010). A Bayesian approach to modeling lost person behaviors based on terrain features in wilderness search and rescue. *Computational and Mathematical Organization Theory, 16*, 300-323.

MacGregor, J. N. & Ormerod, T. C. (1996). Human performance in the traveling salesman problem. *Perception & Psychophysics, 58,* 527-539.

Mueller, S. T., & Weidemann, C. T. (2008). Decision noise: An explanation for observed violations of signal detection theory. *Psychonomic Bulletin & Review, 15*, 465-494.

Mueller, S. T., Perelman, B. S., & Simpkins, B (in press). Pathfinding in the cognitive map: Network models of mechanisms for search and planning.

O'Keefe, J., & Nadel, L. (1978). *The hippocampus as a cognitive map*. New York, NY: Oxford University Press.

Pizlo, Z., Stefanov, E., Saalweachter, J., Li, Z., Haxhimusa, Y, & Kropatsch, W. G. (2006). Traveling salesman problem: a foveating pyramid model. *The Journal of Problem Solving*, *1*, 83-101.

Perelman, B. S. & Mueller, S. T. (2013). A Synthetic Task Environment for experimentally testing human factors effects in unmanned aerial search. [Psychology Experiment Building Language program]. Available from https://sites.google.com/a/mtu.edu/aerialste/.

Perkins, D., Roberts, P., & Feeney, G. (2003). *Missing person behavior: An aid to the search manager*. Northcumberland, UK: Centre for Search Research.

Samsonovich, A. V., & Ascoli, G. A. (2005). A simple neural network model of the hippocampus suggesting its pathfinding role in episodic memory retrieval. *Learning & Memory*, *12*, 193-208.

Air Force Research Lab. (2007). *Enhancing warfighter cognitive abilities with transcranial magnetic stimulation: A feasibility analysis* (Report No. AFRL-HE-WP-TR-2007-0095). WPAFB, OH: AFRL.