

A NATURALISTIC COMPUTATIONAL MODEL OF HUMAN BEHAVIOR IN
NAVIGATION AND SEARCH TASKS

By

Brandon S. Perelman

A DISSERTATION

Submitted in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

In Applied Cognitive Science and Human Factors

MICHIGAN TECHNOLOGICAL UNIVERSITY

2015

© 2015 Brandon Perelman

UMI Number: 3706310

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI 3706310

Published by ProQuest LLC (2015). Copyright in the Dissertation held by the Author.

Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against unauthorized copying under Title 17, United States Code



ProQuest LLC.
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106 - 1346

This dissertation has been approved in partial fulfillment of the requirements for the Degree of
DOCTOR OF PHILOSOPHY in Applied Cognitive Science and Human Factors.

Department of Cognitive and Learning Sciences

Dissertation Advisor: *Shane T. Mueller*

Committee Member: *Edward T. Cokely*

Committee Member: *Myounghoon "Philart" Jeon*

Committee Member: *Christopher W. Myers*

Department Chair: *Susan L. Amato-Henderson*

Table of Contents

Preface	6
Abstract.....	7
Executive Summary	8
Defining the Problem Space.....	8
Generating a Naturalistic Computational Model	9
Introduction.....	13
Modeling Spatial Problem Solving.....	16
Chapter 1: Literature Review	18
1.1 Lines of Research in Spatial Problem Solving.....	18
1.1.1 Pathfinding	20
1.1.2 Wayfinding.....	21
1.1.3 Orienteering.....	21
1.1.4 Combinatorial Optimization Problems	22
1.1.5 Navigation vs. Search	24
1.2 Function and Purpose of the Hippocampus.....	26
1.2.1 Problem Solving: Spatial and Visual Memory.....	26
1.2.2 Navigation	29
1.3 Hippocampus Neuroanatomy.....	30
1.3.1 Afferent Structures	30
1.3.2 Hippocampal Formation – Gross Neuroanatomy.....	31
1.3.3 Hippocampal Formation – Cells and Functions.....	32
1.3.4 Network Dynamics in the Hippocampus	33
1.4 Hippocampus-based Neural Network Models.....	34
1.4.1 Early Models – Sequence Prediction	34
1.4.2 Hippocampus-based Navigation Models	35
1.4.3 Modeling Human Behavior	37
1.5 Computational Approaches to Modeling Problem Solving.....	38
1.5.1 E-TSP: Problem Solving in the Third Person	39
1.5.2 Wayfinding and Navigation: Problem Solving in the First Person.....	42
1.5.2.1 Levels of Encoding – Fine & Coarse Categorization and Errors in Representation	44
1.5.2.2 Superordinate Distortions – Evidence of Heuristics?	46
Chapter 2: New Analysis of Prior Empirical Research.....	47
2.1 Criteria Selection.....	48
2.1.1 Perelman & Mueller 2014 Results Interpretation.....	49
2.1.2 Perelman & Mueller 2014 Discussion	50
2.2 Multiple Optimization Criteria in TSP Solutions.....	53
2.2.1 Results – Estimated Time to Find / Time-Discounted Reward Intrusion	54

2.2.2	Results – Distance to First Contact Intrusion.....	55
2.3	General Discussion	56
Chapter 3:	Empirical Research.....	57
3.1	Rationale & Hypotheses	57
3.2	Study 1 – Subject Matter Expert Interviews	59
3.2.1	Study 1 – Method	59
3.2.1.1	Wilderness Search and Rescue.....	63
3.2.1.2	Geospatial Intelligence Analysis.....	64
3.2.1.3	Bail Bonds Enforcement	65
3.2.1.4	Search in Military Operations	66
3.2.2	Study 1 – Results	68
3.2.2.1	Identifying Target Locations versus Search Implementation	68
3.2.2.2	Offloading and Organizing Search Information	71
3.2.2.3	Solving the Probability Space	74
3.2.2.4	Trust in Automation versus Training	75
3.2.3	Study 1 – Conclusions	76
3.3	Study 2 – Laboratory Investigation of Planning and Memory Representations	77
3.3.1	Study 2 – Method	80
3.3.2	Study 2 – Data Collection and Analysis.....	86
3.3.3	Study 2 – Hypotheses	86
3.3.4	Study 2 – Experiment 1 Results.....	88
3.3.5	Study 2 – Experiment 2 Results.....	92
3.3.6	Study 2 – Discussion	97
3.4	Study 3 – Naturalistic Investigation of Planning and Memory Representations	100
3.4.1	Study 3 – Method	101
3.4.2	Study 3 – Hypotheses	104
3.4.3	Study 3 – Task 1 Results.....	104
3.4.4	Study 3 – Task 2 Results.....	108
3.4.5	Study 3 – Discussion	109
Chapter 4:	Modeling Search Behavior	111
4.1	Prior Efforts Modeling Navigation and Search.....	112
4.1.1	Modeling Planning – Spreading Activation Model.....	117
4.1.2	Modeling Planning – Hierarchical Spreading Activation Model	118
4.1.3	Modeling Planning – Linear Plan Model – Random Evolution	122
4.2	Modeling Planning using Mixed Optimization Criteria.....	125
4.2.1	Creating a Plan and Plan Evaluation.....	126
4.2.1.1	Parameters – Log Number of Segments.....	128
4.2.1.2	Parameters – Goodness of Fit.....	128

4.2.1.3 Parameters – Path Length.....	129
4.2.1.4 Parameters – Time-Discounted Path Length.....	129
4.2.1.5 Parameters – Average Angle.....	130
4.2.2 Fitting a Plan.....	130
4.2.3 Following a Plan – Local Decision-Making.....	132
4.2.4 MCM Results.....	133
4.2.4.1 Observed Differences in Planning and Optimization	134
4.2.4.1.1 Problem 2 – Lightbulb.....	139
4.2.4.1.2 Problem 4 – Z.....	139
4.2.4.1.3 Problem 5 – Z with Clusters	140
4.2.4.1.4 Problem 7 – V with Clusters	141
4.2.4.2 Simulations & Model Fitness.....	142
4.2.5 MCM Discussion.....	145
Chapter 5: Using the Model – Recommendations for Implementation.....	149
5.1 Integrating Planning and Navigation	150
5.2 Integrating Planning and Memory	152
5.3 Potential Domains of Application	154
5.3.1 Emergency Egress	155
5.3.2 Military and Law Enforcement Operations.....	157
5.3.3 Urban Design.....	159
5.4 Trust in Automated Plan Generation – a Proposed Framework for Application.....	160
5.4.1 Automatic Route Generation	161
5.4.2 Experiment 1: Measuring Alternative Optimization Criteria in Search Plans	162
5.4.3 Experiment 2: Trust in Automation – Path Congruency	163
5.4.4 Expected Implications and Discussion.....	164
References.....	166

Preface

Chapter 2 of this dissertation contains novel analysis of data collected by Perelman and Mueller (2014), presented in Section 2.1, and Mueller, Perelman, Tan, and Thanasuan (*manuscript under review*), presented in Section 2.2. These individuals contributed to the research efforts in the following ways. Section 2.1 contains a re-visitation of data originally collected by Brandon S. Perelman and his advisor, Shane T. Mueller. Shane T. Mueller provided supervisory input on the experiment design, the spatial problem stimuli used in that study, and data analysis. Brandon S. Perelman programmed the experimental tasks, designed the experiment, collected the data from participants during the study, analyzed the data, and wrote Section 2.1. For the present dissertation, participant solutions were subjected to a number of analyses (Pathmapping analysis, multidimensional scaling, and Gaussian mixture modeling) which were conducted collaboratively by Brandon S. Perelman and Shane T. Mueller.

Shane T. Mueller, Brandon S. Perelman, Yin-Yin Tan and Kejkaew Thanasuan all contributed to the research effort that produced the data described in Section 2.2. Yin-Yin Tan and Kejkaew Thanasuan collected data from 16 of the participants. Shane T. Mueller designed the experiment, conducted the analyses, and wrote the manuscript submitted as Mueller, Perelman, Tan, and Thanauan (*manuscript under review*). Brandon S. Perelman collected data from the remainder of the participants. For the present dissertation, Brandon S. Perelman independently analyzed the data, and wrote Section 2.2.

Abstract

Planning, navigation, and search are fundamental human cognitive abilities central to spatial problem solving in search and rescue, law enforcement, and military operations. Despite a wealth of literature concerning naturalistic spatial problem solving in animals, literature on naturalistic spatial problem solving in humans is comparatively lacking and generally conducted by separate camps among which there is little crosstalk. Addressing this deficiency will allow us to predict spatial decision making in operational environments, and understand the factors leading to those decisions. The present dissertation is comprised of two related efforts, (1) a set of empirical research studies intended to identify characteristics of planning, execution, and memory in naturalistic spatial problem solving tasks, and (2) a computational modeling effort to develop a model of naturalistic spatial problem solving. The results of the behavioral studies indicate that problem space hierarchical representations are linear in shape, and that human solutions are produced according to multiple optimization criteria. The Mixed Criteria Model presented in this dissertation accounts for global and local human performance in a traditional and naturalistic Traveling Salesman Problem. The results of the empirical and modeling efforts hold implications for basic and applied science in domains such as problem solving, operations research, human-computer interaction, and artificial intelligence.

Executive Summary

The spatial problem solving tasks of search and navigation have been investigated in many basic and applied communities. For example, these tasks are often described as *pathfinding* in the basic psychology literature, whereas in the applied literature they are referred to as *wayfinding* or *routing*. While the underlying research questions are similar, the terminology, methods, and approaches to modeling used by all of these communities vary widely, and there is very little fruitful intercommunication among these communities. Modelers in these communities have yet to produce functional agent-level accounts of human behavior that are useful in applied settings. One major deficiency of prior modeling efforts is an inability to account for multiple optimization criteria; the models generally attempt to solve the problem according to a single success criterion, while the human data suggests that multiple criteria influence spatial problem solving. The purpose of this dissertation is to create a computational model of human navigation and search that uses strategic planning and bottom-up heuristic pathfinding processes to approximate human behavior.

Defining the Problem Space

The problem is most chiefly one of scope – the basic science community does not typically produce models applicable to human behavior, and the applied science community does not produce accounts of spatial problem solving that are sufficient to explain individual agents' behavior. While basic science has focused largely on the minutiae of the neurobiological processes underlying spatial problem solving (i.e.,

pathfinding), it has essentially ignored higher level cognitive components that humans use to solve these tasks. The end results of this effort are models that, with a very high degree of biological fidelity, are capable of making accurate predictions of *non-human* (typically rodent) behavior. In contrast, applied approaches to modeling search and navigation (i.e., wayfinding) have focused largely on group dynamics. So while these approaches approximate aggregate human behavior fairly well, they fail to account for the behavior of *individuals*. This is critical, as these models are often applied in domains in which lives are at stake. Some approaches to modeling human behavior account for higher cognitive human capabilities in spatial problem solving, however these models typically address combinatorial optimization problems like the Traveling Salesman Problem (TSP) rather than spatial navigation and search. In short, to date, no modeling efforts have filled this capability gap.

The goal of this dissertation is to merge aspects of models of pathfinding, routing, and wayfinding to create a naturalistic computational model that (1) provides a strong account of human behavior in planning and execution in navigation and search, and memory related to search, (2) uses heuristic processing to remain computationally inexpensive and viable for multi-agent simulations, and (3) drives agent behavior to create useful predictions to real-world problems.

Generating a Naturalistic Computational Model

Creating such a model requires integrating a number of aspects from models developed in disparate fields. Pathfinding models offer simple, biologically plausible

mechanisms for driving local decision making in navigation, and for encoding and retrieving context and sequence memories crucial for navigation and spatial memory. However, prior research in modeling soldiers' behaviors during search operations (Mueller, Perelman, & Simpkins, 2013) indicates that applying these mechanisms to account for human behavior in spatial tasks requires higher level planning.

While specific mechanisms for constructing these plans have yet to be implemented in published human pathfinding models, models of spatial problem solving use high level planning to solve combinatorial optimization problems (e.g., Pizlo, Stefanov, Saalweachter, Haxhimusa, & Kropatsch, 2006). These models of problem solving involve subdividing the problem space hierarchically into either clusters (e.g., Pizlo et al., 2006) or linear representations (e.g., MacGregor, Ormerod, & Chronicle, 2000). These models are biologically inspired but not necessarily plausible, or useful in multi-agent simulations, because they use computationally expensive brute force computation at certain levels of processing. These models represent computational approaches to problem solving, and are not intended to model biological *processes*. By substituting computationally inexpensive biologically plausible pathfinding algorithms for these brute force computations, we can provide a reasonable heuristic account of how humans form and execute spatial plans.

Finally, wayfinding models emphasize the importance of environmental features in human spatial navigation. Pathfinding models assume that cognitive maps (i.e., an agent's spatial conceptualization of the environment) are created through exploration, ignoring the importance of environmental context salience in the cognitive map

construction process. In essence, the probability of encoding any two environmental contexts are equal. In contrast to that approach, wayfinding models assume that spatial navigation is nearly entirely facilitated by environmental cues, though few have provisions for allowing the agent to encode and retrieve cues. That is, agents in most wayfinding models follow signs to specific locations, such as exits, but are incapable of forming their own cue-context associations that would allow them to navigate (i.e., “turn left at the room with the large potted plant in the corner.”). The wayfinding approach will be used to inform a model with a capability to weigh the probability of encoding environmental contexts based on environmental cues, and to use that information to inform plan construction and execution during spatial navigation (c.f., “handrails” in orienteering) and the agent’s memory for where it searched and what it found.

In summary, these fields in spatial problem solving all have something to offer in constructing a computational model that is useful for making predictions about human behavior (see Figure 2). Before the relevant components contributed by each of these fields can be integrated, however, empirical studies are needed to elucidate specific details of this process. Specifically, the routing literature offers two distinct types of solutions for modeling high level planning - cluster versus linear representations. The empirical task lies in disambiguating these two possibilities found in prior studies (Perelman & Mueller, 2013a). Studies 2 and 3 of this dissertation achieve this disambiguation, and identify differences in plan construction and execution as it is done in third-person routing tasks (Study 2), and in first-person spatial navigation in a naturalistic task similar to a hasty search conducted during search and rescue operations.

This information on high level planning in spatial navigation and search tasks determines which of two candidate modeling approaches better represent plans in these tasks – cluster or linear planning dynamics.

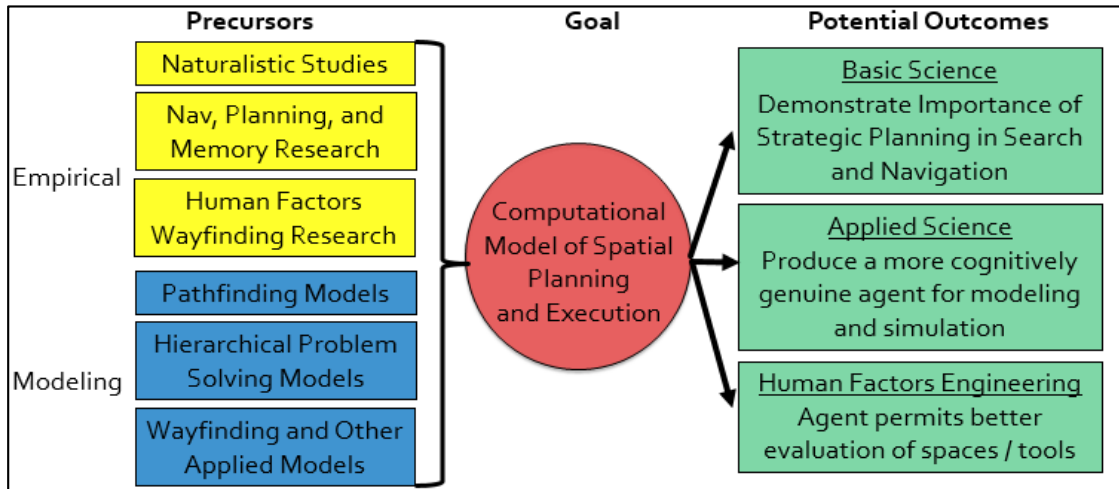


Figure 1. Prior and current empirical and modeling research intended to inform the computational modeling effort, and potential uses for the model. For potential methods for applying this model to other domains, see Chapter 5.

The end result of this effort is a naturalistic computational model capable of approximating human considerations in planning and executing spatial navigation and search routes. The final section of this dissertation contains design recommendations for integrating the model in service of other domains, such as multi-agent egress simulations and search and cordon operations in the military and law enforcement. The goal of this modeling effort is to effectively approximate human performance in the task used by Perelman and Mueller (2014) to distinguish optimization strategies for problem solving in a path planning task.

Introduction

Both basic and applied researchers have attempted to understand, or create algorithms for controlling, an agent's navigation through an environment. In the basic psychology community, this process is often called *pathfinding*. This research primarily seeks to understand how organisms form and use cognitive maps (Tolman, 1948) by identifying the neural, cognitive, and physical processes by which pathfinding is accomplished. In the applied human factors literature, a similar problem is often called *wayfinding* (Lynch, 1960), and this research seeks to understand environmental cues that direct the flow of humans through an environment. Human performance in these tasks shares many similarities with other types of spatial and sequential problem solving, for example the TSP (e.g., MacGregor & Ormerod, 1996) and the Tower of Hanoi (e.g., Samsonovich & Ascoli, 2005). One component tying together all of these problems is planning, or mental simulation.

A key assumption of this dissertation is that solving complex problems requires planning to subdivide the problem space into elements that can each be solved in turn, reducing the agent's cognitive burden and allowing it provide near-optimal solutions to otherwise intractably complex problems. Many hierarchical models of problem solving achieve this subdivision deterministically through a process of continually subdividing the problem space into increasingly fine components, whereas the present argument states that this subdivision requires only two levels, the higher level containing the plan representation, and the low level at which the detailed plan is executed. In both cases, breaking up the problem space is designed overcome memory limitations, and is

peripherally similar to what Miller (1956) first described as *chunking* in immediate memory.

Carrying this to a greater extreme, in many applied spatial problem solving domains, wherever possible the information is offloaded from the human entirely. As I will discuss in Chapter 3, GPS logging tools allow wilderness search and rescue (WiSAR) operators to view where they have searched without having to remember the details (SME 1, May 15, 2014), and command and control elements track the movement of units in military operations (SME 5, March 7, 2015). Likewise, bail bonds enforcers construct databases to organize the information they use to constrain their searches (SME 2, May 13, 2014).

The theory described herein proposes that a mechanism (or mechanisms) used for pathfinding and spatial planning decomposes general spatial constructs into sequentially ordered memory representations, a process necessitated by memory limitations, through which the agent can plan a trajectory. This is different from chunking, which refers to grouping items based on semantic or perceptual similarities. This thesis posits that the representations used in the aforementioned tasks are hierarchically subdivided into smaller components, the characteristics of which will be revealed by sequence errors in the appropriate tasks. One assumption of this argument is that the spatial representations used in planning are identical to those remembered after completing a spatial task. This argument is based on empirical evidence indicating similarities between spatial memory and list memory (Perelman & Mueller, 2013a), and the responsibility of the same structure, the hippocampus, for both episodic memory and mental simulation in spatial

navigation (Lyon & Gunzelmann, 2011; Lyon, Gunzelmann, & Gluck, 2008; Samsonovich & Ascoli, 2005; Takahashi, 2013). The empirical research in this dissertation will characterize these representations in generalizable spatial problem solving tasks.

It is important to draw a distinction between the actual search process (i.e., probing locations in the environment to find one or more targets, a process optimized by minimizing the estimated time to find, or the average distance among those points), which I propose relies on sequence memory, and foraging behaviors that are more closely linked to semantic memory processes (e.g., Hills, Todd, & Goldstone, 2008). The present thesis is reconciled with that work by drawing a distinction between the cognitive mechanisms used for semantic search and spatial foraging, and the cognitive mechanisms recruited for sequence memory and navigation (i.e., choosing the best route among a set of points, a process optimized by minimizing path length) using mental simulation. While this manuscript will not empirically investigate foraging, certain tasks in applied settings that relate to the present research are closely tied to semantic search and foraging (see Chapter 3, Section 2, below).

For the purpose of this dissertation, a distinct definition of the term *optimization* differentiates it from the term as typically used in mathematical contexts, such as operations research. Traditionally, optimization refers to finding the best solution from all feasible solutions, maximizing or minimizing a reward or cost function. For this dissertation, this definition should be expanded to include the phrase, "...or the best solution produced after a process of satisfaction according to specific criteria." Therefore,

optimization criteria refers to linearly-related parameters by which a given solution may be evaluated. And, the result of an optimization process according to a specific optimization criterion may not, in fact, be the global optimal solution according to that criterion. While this approach is similar to that used to model decision making according to multiple optimization criteria (c.f., Emmerich & Deutz, 2006), the terminology used herein should not be viewed as interchangeable with the terminology as used by researchers in those fields (e.g., operations research and economics).

Modeling Spatial Problem Solving

Models have value in predicting human search trajectories and memory effects, investigating effects of environmental characteristics on search and memory (e.g., when should search professionals trust their “gut instincts” for generating search plans versus using planning software? When should search professionals trust their memory for where they searched, versus when is it necessary to meticulously record actual search trajectories using a system like GPS?), and utility in building heuristic search algorithms for machine cognition. The explicit purpose of the present modeling endeavor is to construct a biologically plausible cognitive agent capable of realistic spatial problem solving across multiple domains.

While predicting human behavior in operational (i.e., military, law enforcement, and first responder) environments is useful, an ecologically valid cognitive agent holds the most utility for civilian applications. Predicting human navigation behavior based on the *learnability* of an environment is useful in civil engineering applications such as

public transportation system design and emergency egress modeling, domains currently serviced predominantly by flow models. Because the pathfinding mechanisms described herein are computationally inexpensive, they can be used in multi-agent simulations. While agent-based simulations exist for modeling emergency egress, intelligence in these models is often implemented in the form of information use (e.g., Viswanathan, Lees, & Ayd, 2012) rather than experiential learning, therefore, very few (e.g., Hajibabai, Delavar, Malek, & Frank, 2007) take into account differences between egress routes selected by experts and novices (i.e., individuals who frequent a building who know it very well versus those who do not). Legion, a popular Egress model, takes only destination, speed, separation from other agents, and reaction time as input parameters (Pan, Dauber, Han, & Law, 2006), and thus makes no attempt to model agent behavior based on knowledge of the environment, psychological responses to stress, or other factors. Modeling these differences is important for predicting egress flow in environments where the ratio of regular workers and new visitors is relatively constant, such as hospitals.

In sum, current models are inadequate for describing human navigation and search in the real world. Many pathfinding models, created in service of neuropsychology, are intended to approximate the neural bases of rodent behavior, while many biologically-inspired human models are designed to solve combinatorial optimization problems. Conversely, models used for solving real-world problems fail to incorporate important cognitive aspects that impact agent behavior. This modeling effort will produce an intelligent human-like agent capable of solving these types of problems.

The utility of this model lies in its generalizability to many domains in which spatial navigation is critical, such as egress modeling, urban planning, and instructional design for military and law enforcement personnel engaged in cordon and search operations.

During the course of this dissertation proposal I will provide (1) a literature review of human spatial problem solving performance, neural correlates, and modeling approaches, (2) new analysis of prior empirical research, (3) a series of interviews, laboratory studies, and a naturalistic study investigating human performance in spatial problem solving, (4) a discussion of prior computational approaches to solving this problem, and (5) a model of planning and plan execution that is capable of solving the types of naturalistic problems for which current models of spatial problem solving are inadequate. The unique contribution of this dissertation is a computational model of planning and execution in spatial problem solving, applicable across multiple domains, that bridges the gap between biologically plausible pathfinding models, and ecologically valid wayfinding models.

Chapter 1: Literature Review

1.1 Lines of Research in Spatial Problem Solving

While basic neuroscientists who study spatial problem solving are primarily concerned with comparative psychology approaches (i.e., studying animal behavior), scientists with an eye on application have investigated navigation in humans in several main domains or lines of research (see Table 1). The first, pathfinding, is primarily studied in the animal literature; the human literature is comparatively underdeveloped.

The second, wayfinding, as it is called in the human factors literature, is concerned with investigating environmental cues (i.e., landmarks) that facilitate learnability. Third, sports psychologists often study orienteering from the expert / novice perspective. Finally, route planning has been investigated extensively by researchers interested in visual memory.

Table 1. Relevant domains and lines of research examining spatial problem solving.

Domain	Characteristics	Agent Goals	What to optimize?	Example Contributors
Pathfinding (functional neuroanatomy and modeling literature)	Basic science approach (comparative psychology and neuroscience) Experiential learning Focus on the <i>agent</i>	Localization Learn environment & encode objects in the environment Mental Simulation	Path length / distance “Effort” and efficiency	Hasselmo Samsonovich Rolls McNaughton Burgess Levy O’Keefe Gunzelmann
Wayfinding (human factors approach)	Applied science approach (e.g., emergency egress models) Rule-based flow through environment Focus on the <i>environment</i> (e.g., cue salience, learnability, landmarks, signage)	Localization Navigate to goal locations in the environment (often at first-pass through environment)	Signage salience (location and type) and environment layout Environmental flow	Caduff Timpf Gaisbauer Frank Rüetschi Klöpfer
Traveling Salesman Problem	Intractably complex problems to which humans generate near-optimal solutions Performance modeled using various types of models (nearest-neighbor, cluster-based models, and trajectory / convex hull models)	Generate a solution among all of the target locations having the shortest path length	Path length (shortest)	Pizlo Graham Tenbrink MacGregor Stefanov Kong Schunn
Naturalistic / Applied Traveling Salesman Problems	“Applied” version of TSP. Goal is to hit targets as early as possible. Applied to domains like orienteering, robot navigation, wilderness search and rescue, etc. Used in situations where the distance that can be traveled is limited and possibly uncertain, traveling requires effort (distance-reward tradeoff), and where rewards dynamically degrade or change	Generate a solution that maximizes the number of targets hit in the minimum distance.	Average distance between nodes Estimated Time to Find	Blum Charikar Tenbrink Albers Goel Lin Goodrich Adams Waharte Trigoni
Orienteering	Uncertain goal locations (controls) Requires that humans use “handrails” to permit localization Orienteer must consult the map, monitor the terrain, and move simultaneously. Requires heuristics (start-forward, end-back).	Navigate through waypoints (controls) to a destination. Maintain efficient routes. Maintain localization. Set control flags	Path efficiency (fastest time) Speed / difficulty tradeoff “Control flag placement	Blum Eccles Mottet Saury Vicente Wang

Table 1 Continued

<p>Spatial Memory</p>	<p>Participants' memory for locations and what is contained within them.</p> <p>Localization and path planning research.</p> <p>Generally, human spatial memory believed to be hierarchical – coarse and fine grained structure. Geographical errors are one example of this.</p>	<p>Reconstruct a path, or landmarks along a path</p> <p>Recall or reconstruct a top-down map</p> <p>Distance estimates</p>	<p>Correct routes / sequences</p> <p>Recall for relative geographic locations</p> <p>Distance estimates</p>	<p>Wiener Mallot McNamara Collett Wehner Stevens & Coupe Hirtle & Jonides</p>
<p>Foraging</p>	<p>Animals search for resources distributed in patches.</p> <p>Mechanisms sometimes compared to semantic foraging in human memory retrieval</p>	<p>Minimize energy expenditure (i.e., maximize reward versus time)</p>	<p>Energy expenditure</p> <p>Time spent in a particular patch</p>	<p>Hills Todd Jones Charnov Stephens</p>
<p>Machine Navigation</p>	<p>How to most efficiently get from Point A to B</p> <p>How to cover the search space the best</p> <p>Biomimicry approach: colony based (bees, bacteria, ants, etc.)</p>	<p>Robot must navigate from one location to another</p> <p>Many robots must coordinate movement / search together</p>	<p>Path efficiency (ETA)</p> <p>Estimated Time to Find</p> <p>Coverage</p>	<p>Van den Berg Lin Wong Bourgault Verscheure Zhang Evers Nilsson</p>

1.1.1 Pathfinding

Pathfinding describes an agent's ability to find the shortest or least effortful path among locations in a problem space (Mueller et al., 2013). Pathfinding can be further subdivided into two distinct fields – comparative psychology research that seeks to understand how an agent navigates through an environment using mental simulation, and machine cognition or artificial intelligence, in which the goal is to provide mechanisms for driving a robot. Early approaches to pathfinding in the AI community include Dijkstra's (1959) algorithm, and the A* algorithm developed for the robot *Shakey* (Hart, Nilsson, & Raphael, 1968). Empirically, while pathfinding is mainly studied in animals, however Mueller et al. (2013) investigated and modeled pathfinding in dismounted infantry conducting cordon and search operations in an urban environment using models of varying complexity. This work indicated that simple models similar to those used to approximate rodent behavior can be applied to certain human tasks.

1.1.2 Wayfinding

Wayfinding is human navigation as studied in the human factors community, and it is typically studied from the points of view of landmark navigation and information processing (Montello, 1993). As such, wayfinding depends heavily on environmental features such as landmarks and context salience (Caduff & Timpf, 2008). The goal of wayfinding research is to inform design, such as in city planning (Gaisbauer & Frank, 2008) and public transportation networks (Rüetschi, 2007; Timpf, Rüetschi, & Caduff, 2005). Designs are typically improved by creating intuitive flow (Rüetschi, 2007; Timpf, Rüetschi, & Caduff, 2005) or improving the salience, accessibility, and legibility of signage (Hajibabai et al., 2007). Wayfinding can be understood as the study of how the environment facilitates navigation, and is typically studied for its applications in civil engineering, specifically public transportation system engineering and emergency egress modeling.

1.1.3 Orienteering

Orienteering is studied from a sports psychology perspective, with the goal of improving athlete performance. Orienteers face a problem similar to that studied in the wayfinding literature. However, the key differences are that landmarks (i.e., environmental cues) are typically natural rather than man-made, and that a topographical map is provided to participants to facilitate planning and localization. In the wayfinding

domain, the burden lays on the human factors engineers and designers to create an environment that is easily navigable. In orienteering, the burden lies on the agent (i.e., the competitor), creating a challenge that fosters competition.

In orienteering, participants compete for the fastest time using a north-up proprietary orienteering map to navigate among waypoints called *controls* (Eccles, Walsh, & Ingledew, 2002). Some styles of this game involve the participants locating control flags, whereas others require the participants to place control flags accurately (Mottet & Saury, 2013). The maps used in orienteering depict five levels of vegetation density, which serve as a measure of how easily a competitor can traverse a location. Therefore, orienteers must consider speed vs. distance and speed vs. effort tradeoffs during route planning. Orienteering is regarded as cognitively demanding because orienteers must simultaneously move (often run), pay attention to local terrain, and consult their maps (Eccles et al., 2002). Because many orienteering competitions use a fixed time limit, it is possible that competitors will not reach all of the controls. For this reason, orienteering is described as a Discounted-Reward Traveling Salesman Problem (Blum et al., 2007) and is mathematically similar, in terms of optimization strategy, to more general search problems such as wilderness search and rescue. This problem paradigm can be applied to other route planning tasks as well.

1.1.4 Combinatorial Optimization Problems

The Traveling Salesman Problem (TSP) belongs to a family of problems known as combinatorial optimization problems. The TSP, formulated in the 1930s (Schrijver,

2005), requires that the solver create the shortest path among a series of points (cities). A main goal of TSP researchers is to understand how humans employ heuristics to solve a computationally intractable problem. TSP is regarded as NP-hard or computationally intractable, because a “brute force” (i.e., exhaustive) approach to solving it requires generating $(n - 1)! / 2$ separate trajectories. A number of heuristic algorithms have been devised for solving this particular routing problem (see below). TSP-like problems such as the Vehicle Routing Problem (Dantzig & Ramser, 1959) are often used in operations and management and other applied settings.

One particular characteristic of routing problems that separates them from pathfinding is perspective. Pathfinding is often studied in domains where the agent perceives the environment in the first person, whereas routing problems require the agent to solve a problem from an overhead, third person perspective. Evidence that optimization occurs independently of problem space perception (i.e., in the first or third person) contradicts traditional models of spatial problem solving, and necessitates a novel approach comparing data from first and third person implementations of the same task (see Chapter 3, Studies 2 and 3, below).

Though there is little crosstalk between researchers investigating routing and pathfinding problems, one study in the comparative psychology literature attempted to tie together TSP-solving dynamics and pathfinding in the rat. In a novel experiment, de Jong, Gereke, Martin, & Fellous (2011) tasked rats with completing a TSP, substituting food pellets for cities. The results of this study indicated that rats learned to optimize their solutions (i.e., minimize their path lengths) over repeated exposures to problems

consisting of food pellets. The results of this study indicate that perhaps optimization occurs separately from problem space perception. This is important as many models of TSP-solving in humans assume that optimization occurs mainly during perception (see below). Additionally, this study indicates that TSP-solving through pathfinding, in the first person, is optimized in a similar way to traditional TSP-solving, accomplished in the third person from an overhead view.

1.1.5 Navigation vs. Search

At this point, it is critical to draw a distinction between two spatial problem solving behaviors: search and navigation. In navigation tasks, the agent must solve a problem that bears many resemblances to a TSP. In these problems, (1) the agent is tasked with visiting every location, (2) the rewards (i.e., the probability of finding a target) associated with all of the locations are equal and static, and (3) optimal performance was assumed to minimize overall path length. Prior attempts at modeling human search behavior (e.g., Mueller et al., 2013; Perelman & Mueller, 2013b) have set the problem space with TSP-like (i.e., navigation-goaled) characteristics. Note, however, that unlike a traditional TSP, the agents in these models did not have the return to base requirement.

Nevertheless, search problems are quite different from navigation. Specifically, (1) the agent does not need to visit every location, and can end its search as soon as it satisfies the search criteria (e.g., finding one or more hidden objects), (2) the rewards associated with each location may be different (i.e., the probability of finding the target in

one location may be more than the next). Furthermore, if the target is mobile (as is the case in many applied settings, such as wilderness search and rescue), then the reward associated with each possible target location is dynamic. Finally, (3) the agent typically wants to optimize its trajectory to minimize the estimated time to find the target. For these reasons, search tasks have many characteristics in common with orienteering.

In tasks such as aerial intelligence, surveillance, target acquisition and reconnaissance, in military applications, and aerial search in civil search and rescue operations, experts construct probability maps to define the search space. However, once this space is complete, the procedure for plotting a route through that probability space is far less formulaic, being executed “off-hand” by either the incident commander (in civil operations) or the pilot (in military and intelligence operations; see Chapter 3, Section 2).

Routing a path through a probability space is functionally similar to E-TSP. The difference lies in constraints within the problem space – E-TSP focuses on optimizing for the shortest path length, whereas real-world search and navigation tasks might incorporate other constraints. For example, a shopping trip involves visiting a number of destinations, the order of which should be determined by environmental constraints such as food spoilage (it is better to buy food later so it will not sit in a hot car), weight (if you must carry all of the items, then it is better to buy the heavier items last), and other constraints such as store hours. Applied (e.g., Evers, Dollevoet, Barros, & Mansuur, 2011) and naturalistic (Perelman & Mueller, 2014; Ragni & Wiener, 2012; Tenbrink & Seifert, 2011) studies have investigated human performance in these types of search tasks, but the

literature in this area is notably sparse. One purpose of this dissertation is to further understand and model human behavior in this type of task.

1.2 Function and Purpose of the Hippocampus

Traditionally, human problem solving (Newell & Simon, 1972) and spatial navigation (Marr, 1971) have been treated primarily as a hippocampus-based behavior due to the requirement for mental simulation (i.e., exploring and comparing possible future moves, and selecting from among them; Samsonovich & Ascoli, 2005). In 1957, Scoville and Milner published the results of follow-up studies on psychosurgery patients, one of whom is well known in the neuropsychology literature as H.M. Following bilateral medial temporal lobe resections, which destroyed roughly two thirds of the hippocampus, the amygdala, and other adjacent structures, H.M. and another patient who had received a similar surgical intervention exhibited profound memory deficits. Comparing these observations with those of patients receiving similar but less radical surgeries, who also displayed memory deficits, Scoville and Milner (1957) concluded that the hippocampal formation is vital for memory formation.

1.2.1 Problem Solving: Spatial and Visual Memory

Serial memory is believed to play a role in many processes such as language learning, planning, social behaviors, and motor skills (Hurlstone, Hitch, & Baddeley, 2014). The body of literature concerning serial memory in the verbal domain is rather extensive compared to the visual and spatial domains, about which comparatively little is known. To inform our understanding of memory in those domains it is necessary to draw

analogues from the verbal memory literature. For an in-depth discussion of memory characteristics, refer to Baddeley, Eysenck, and Anderson (2009). In this section I will discuss the spatial and visual memory literature in the context of problem solving.

Problem solving often requires mentally simulating future possible problem states and planning based upon those simulations. For example, the Tower of London task (Shallice, 1982), which is often used to test for cognitive deficits in planning, requires that the subject rearrange a set of disks. The Tower of London necessitates mentally simulating the transitional states between the current and goal configurations.

Samsonovich and Ascoli (2005) modeled human performance in a similar task, the Tower of Hanoi. While these problems are excellent diagnostic tools, they tell us little about how humans represent objects in memory, how they can use these representations to construct plans, and why constructing a plan might be necessary in the first place.

One characteristic of human memory that has been clear since the earliest days of memory research is that it has limits. Problem solving often requires that people mentally simulate series of “moves”, as is the case in the aforementioned tower problems and in Rubik’s Cube solving, for example. This is similar to the problem described by Miller (1956) in which comparisons among stimuli are limited by the subject’s ability to store representations of those stimuli in memory. To fully understand the utility of planning in problem solving, we should take the example of a very difficult problem: the TSP.

The Euclidean Traveling Salesman Problem (E-TSP) is described as an NP-hard computationally intractable problem. Despite the difficulty of the problem, human subjects generate near-optimal heuristic solutions to this problem (MacGregor &

Ormerod, 1996). Furthermore, these solutions may be optimizing criteria that are irrelevant in E-TSP, but useful in spatial problem solving (see Chapter 2).

Most computational accounts of TSP-solving (see below) assume that humans are able to produce these near-optimal solutions by splitting the problem into more manageable pieces (e.g., clusters of points or a higher level linear plan) in two or more hierarchically arranged layers. I assume that this process is equivalent to the chunking mechanism described by Miller (1956) which allows humans to increase the amount of information available in memory. Recent research (Pizlo & Stefanov, 2013) indicates that relatively large TSP problems are solvable even when the agent is afforded a limited working memory capacity.

In the visual and spatial domains, the exact mechanism by which memories are created and represented remains unclear. Visual and spatial memory representations share certain characteristics with verbal working memory representations. Object representations exhibit degradation in memory (Osugi & Takeda, 2013). Investigations (Zhang & Luck, 2008) suggest that these object representations are not resources that can be scored according to resolution-number tradeoffs. Rather, visual working memory permits storage of a discrete number of fixed-resolution representations. In visual and spatial recall, humans exhibit primacy and recency effects (Parmentier, Andres, Elford, & Jones, 2006). Furthermore, errors in these tasks are often similar to the transpositions seen in list memory research; subjects are more likely to swap serial positions of closely grouped stimuli during serial recall than they are of more distant stimuli, an effect known as locality constraint (Hurlstone et al., 2014). This effect has been observed in an abstract

dot task (Parmentier et al., 2006), the Corsi block task (data not yet published), and a simulated aerial search task (Perelman & Mueller, 2013a). Effects of serial position on recall have also been noted in a spatial visualization task (Lyon et al., 2008). If this general memory mechanism underlies all sequence completion problems, then we should expect similar errors in spatial navigation tasks.

1.2.2 Navigation

Navigation can be conceptualized as the retrieval and reordering of past-encoded spatial memories into a logical sequence that begins with the current location and ends with the goal destination. This sequence then provides a basis for mental simulation from the agent's current location to its destination. Tolman (1948) first labeled the abstracted environments, through which mental simulation would allow an organism to navigate, "cognitive maps." O'Keefe & Dostrovsky (1971) provided neural evidence for cognitive maps, later organized into a comprehensive theory by O'Keefe & Nadel (1978), through their single cell recordings of freely moving rats. Place cells, as they are called in rodents, are spatially sensitive hippocampal pyramidal cells. Subsequent work has identified other spatially sensitive cells in rodents, such as boundary vector cells (O'Keefe & Burgess, 1996) and head directions cells (Taube, Muller, & Ranck, 1990), spatial view cells (Rolls, Robertson, & Georges-François, 1997) in humans, and grid cells (Hafting, Fyhn, Molden, Moser, & Moser, 2005) which are common to many mammals, including rodents and humans. Environmental representations from all of these spatially sensitive hippocampal

place cells comprise the “cognitive map” that the organism uses to navigate through the environment.

1.3 Hippocampus Neuroanatomy

In order to accomplish its function in encoding and retrieving information, the hippocampus must receive information from afferent structures. The nature of these connections gives insight into information representations relevant to the present research.

1.3.1 Afferent Structures

Environment perception necessary for navigation is facilitated via a combination of idiothetic (e.g., motor feedback, proprioception, and vestibular input) and allothetic (e.g. vision, olfaction, etc.) information that allows environmental landmark perception and path integration (Whitlock, Sutherland, Witter, Moser, & Moser, 2008). The role of the visual system is to provide spatial (e.g., topographic relationships) and nonspatial (e.g., object recognition) information via the dorsal and ventral streams, respectively, to the hippocampus by way of the perforant path, with the peri- and postrhinal cortices forming the bridges between the neocortex and the entorhinal cortex (EC; Wang, Gao, & Burkhalter, 2011). While visual input is exceedingly useful for navigation and localization, it may not be a requirement. In humans, verbalized spatial information (i.e., auditory presentations of direction words, such as “left,” or “right”) produces spatial memory effects equal to those produced during virtual self-motion in an environment

devoid of landmarks (Lyon & Gunzelmann, 2011). In terms of neural evidence, rats that are deafened and blindfolded still show spatially selective place cell activation when traversing an environment (Hill & Best, 1980), indicating that, at least in rodents, place cells can be activated using idiothetic information alone.

1.3.2 Hippocampal Formation - Gross Neuroanatomy

Most afferent information enters the hippocampal formation through the perforant path. The perforant path provides a connection between the hippocampus and the EC, which may be subdivided into the highly interconnected medial EC (MEC) and lateral EC (LEC) regions. The LEC receives input mainly from the perirhinal cortex, but also receives input from the olfactory and insular cortices as well as the amygdala.

Functionally, the LEC processes olfactory, visual, and tactile information. The MEC receives input primarily from the postrhinal cortex, but also from the presubiculum (i.e., a region that processes information necessary for the function of head direction cells), the occipital lobe, and the retrosplenial cortex (Witter et al., 2000). In rodents, the LEC receives projections from the lateromedial visual field via the temporal cortex (i.e., the ventral stream) while the MEC receives input from the anterolateral visual field via the posterior parietal cortex (i.e., the dorsal stream; Wang et al., 2011). This is consistent with prior observations of high spatial selectivity in the MEC (Hafting et al., 2005) and low spatial selectivity in the LEC (Hargreaves, Rao, Lee, & Knierim, 2005).

Information from the EC is transmitted into *Cornu Ammonis* (CA) hippocampal layers 1 and 3. In addition, the perforant pathway transmits information to the dentate

gyrus (DG), which itself projects to CA3 through unmyelinated axons known as mossy fibers. The DG is a highly plastic structure within the hippocampal formation that is believed to facilitate encoding and consolidation of spatial memories and drive exploration (Saab et al., 2009). Projections from both the LEC and MEC into CA1 are separate, whereas projections into CA3 directly and through the DG converge, forming a combined “object + place” representation (Knierim, Lee, & Hargreaves, 2006).

1.3.3 Hippocampal Formation – Cells and Functions

CA1 and CA3 pyramidal cells have long been recognized as spatially sensitive and important for navigation (O’Keefe & Dostrovsky, 1971). Activity in these areas has been recorded in many environments and under different environmental conditions such as darkness (for a comprehensive review, see Redish, 1997). Importantly, firing in these cells depends upon environmental cues, so rotating cues in the environment causes the spatial fields to rotate as well (Burgess & O’Keefe, 1996). Spatial sensitivity in these cells has been observed in rodents (O’Keefe & Dostrovsky, 1971) and primate species (Rolls & O’Mara, 1995) including humans (Ekstrom et al., 2003), though some research suggests a greater importance of parahippocampal regions in humans (Aguirre, Detre, Alsop, & D’Esposito, 1996).

Neuroanatomical research and computational modeling work have elucidated the roles of these cells in tasks involving memory. Early models, based primarily on connectivity in CA3, focused on the hippocampus’ role in performing cognitive functions

such as sequence prediction (Levy, 1989; Levy, Colbert, and Desmond, 1990; Levy, 1996; Lisman & Jenson, 1996), a task used in both problem solving and navigation.

1.3.4 Network Dynamics in the Hippocampus

The aforementioned models, and subsequent models (e.g., Samsonovich & Ascoli, 2005), assume that the hippocampus works through a phenomenon called phase precession. Theta rhythm is a continuous oscillatory pattern of neural activity. The peak of each theta cycle can be defined as the time point of maximum neural activity within the network (Skaggs, McNaughton, Wilson, and Barnes, 1996). As an organism explores the environment, it moves through the receptive fields of individual place cells, known as place fields. As the organism traverses a specific place field, the phase at which that place cell fires with respect to the theta peak advances. This phase precession can be said to encode the organism's specific location relative to its progression through the place field.

Samsonovich & McNaughton (1997) propose an alternative (or perhaps complimentary) function for this phase precession. This alternative function is based upon two-dimensional characteristics of phase precession that can only be identified via averaging cumulative network activity data iteratively. Specifically, phase precession is modulated by the head direction of the animal and, cumulatively, generates a fan-like pattern at various points in the environment. The proposed alternative function of phase precession, espoused by contemporary models (e.g., Ascoli & Samsonovich, 2013; Mueller et al., 2013), is that the direction of phase precession reflects potential future choices. In the navigation domain, this corresponds to potential future directions. Our

modeling efforts in the problem solving domain assume that phase precession provides strategic options for higher-level planning. The purpose for describing this neuroanatomy is (1) to illustrate the extent to which a layered neural structure can facilitate planning (2) to show the extent to which navigation, problem solving, and memory encoding and retrieval are all considered to be bottom-up hard wired processes. This pattern will be apparent in other models inspired by biology outside of the hippocampus (e.g., Pizlo et al., 2006; see below).

1.4 Hippocampus-Based Neural Network Models

Generally, computational models of the hippocampus use data structures that represent individual brain regions or cell types. That is, they do not attempt to faithfully represent the true complexity of the hippocampal formation and its afferents especially in terms of interconnectivity. Rather, cognitive models make specific assumptions regarding the roles and functions of brain structures, then “fill in the blanks.” For example, NeuroNavigator (Ascoli & Samsonovich, 2013) uses an array of 1,000,000 neurons representing the DG, CA1 and CA3 regions, which is far greater than the ~180,000 neurons estimated in the hippocampus proper, discounting the DG (Akdogan, Unal, & Adiguzel, 2002). Earlier models focused upon network dynamics in CA3.

1.4.1 Early Models – Sequence Prediction

Early computational models of the hippocampus (Levy, 1989; Levy, Colbert, & Desmond, 1990) used mainly CA3 network dynamics to facilitate sequence completion,

treating CA1 primarily as a decoder about which many arbitrary assumptions were necessarily made (Levy, 1996). Due to the computational simplicity of the problems, and computing limitations of the day, networks were rather small (e.g., ~500 neurons; Levy, 1996). Despite their apparent simplicity, these early models were capable of a number of complex sequence prediction operations, including finding short cuts and one-trial learning (Levy, 1996). The latter, specifically, is congruent with the short term synaptic plasticity necessary for forming recency traces in navigation (Mueller et al., 2013).

A number of these early modeling observations remain in many contemporary models. First, asymmetry, especially among CA3-CA3 connections, is viewed as beneficial in solving sequence prediction-like problems (cf., symmetrical networks such as Hopfield networks, which typically stabilize into one of two states; Hopfield, 1982). Second, problem space learning is a Hebbian process of strengthening synaptic weights. Third, the actual path construction algorithms, or sequence predictions, are applied to the CA3 layer, though the CA1 layer plays either a late role, as a decoder, or an early role, as a data structure containing goal states (cf., Mueller et al., 2013).

1.4.2 Hippocampus-based Navigation Models

Navigation entails joining individual perceptual contexts sequentially via mental simulation to create routes by which an organism can move from one location to another. Through a similar mechanism, a cognitive architecture capable of sequence prediction is also capable of spatial navigation. Hippocampus-based models capable of spatial navigation use place cell dynamics to simulate rodent trajectories, from starting locations

to experientially-learned reward sites (i.e., goal locations). In these models, the environment is typically represented using a grid (Burgess, Recce, & O'Keefe, 1994). One of the early models by Burgess and O'Keefe (1996) illustrates the mechanisms common among these models: Through exploration, the simulated agent encounters a reward location to which it may later like to return. A downstream "goal cell" is activated by features contained within that location. In a rat, this might be simply food or water, or perhaps a particularly interesting mating scent trace. A one-shot neural mechanism increases the connection weights between that goal representation and the place cells active when the rat is in that location, binding together *what* (the reward) and *where* (the rat's location where it found the reward). Moving away from this goal location causes a monotonic decrease in place cells with associations to the goal location. In two-dimensional space, we can assume that each goal location has excitability placed in a Gaussian distribution on top of the reward's location. In the Burgess & O'Keefe (1996) model, the agent navigates the topography created by these distributions (i.e., proximities to each goal) via simple hill-climbing.

Many research groups have focused on dynamics specifically intended to model rodent behavior. These include models using boundary vector cell theory (Barry et al., 2006), and the spiking neuron models CATACOMB (Cannon, Hasselmo, & Koene, 2002) and NeuroNavigator (Ascoli & Samsonovich, 2013). By comparison, models of human navigation are relatively sparse.

1.4.3 Modeling Human Behavior

Researchers have modeled aspects of navigation in humans, such as localization, using cognitive architectures designed specifically to approximate human behavior (ACT-R; Lyon & Gunzelmann, 2011; Lyon et al., 2008). However, most attempts to model the full pathfinding and search processes executed by humans have relied upon adaptations of non-hierarchical, two-dimensional rodent models. Results from these studies indicate that rodent models are inadequate to meaningfully predict human behavior. Specifically, these models do not incorporate hierarchical planning or mental simulation capability (e.g., Samsonovich & Ascoli, 2005). Mueller et al. (2013) used an architecture based on a model created by Samsonovich and Ascoli (2005) to approximate search behavior by dismounted infantry soldiers conducting cordon and search operations in an urban environment. The model used in this study assumes that the CA3 layer contains the experientially-learned cognitive map, while the CA1 layer contains the goal representations that are used to activate the appropriate CA3 cells. Mueller et al. (2013) used algorithms of varying complexity, including random walk, novelty-seeking, goal-driven, and exhaustive models. The model that produced reasonable approximations of human behavior required that a high level plan be generated by the experimenter and programmed into the model. The agent did not generate the high level plan itself, and this dissertation proposes that this functionality is necessary for approximating human behavior.

At the execution level, all of these models used similar dynamics to the model originally implemented by Burgess and O'Keefe (1996). Goal locations generated

“activation” which diffused throughout the learned CA3-CA3 interconnections. The agent then used simple hill-climbing to navigate among these locations. While the agent makes decisions based upon diffused information, its decisions are essentially myopic. For this reason, the spreading activation model is essentially a nearest neighbor model that constrains trajectories on the basis of environmental barriers (i.e., walls). The model that produced human-like trajectories placed a goal location in each room of the search space (functionality encoded by experimenters), constituting a higher level planning that could be executed using a relatively simple mechanism. However, this higher level planning appears necessary to approximate human behavior. Further research in this area outside of the infantry domain (see below) has elucidated many of the shortcomings of a two-dimensional spreading activation mechanism for modeling human trajectories. Given the results, and the similarity of this problem to others such as E-TSP, it seems that perhaps a problem-solving model, as opposed to a navigation model, may provide inspiration for modeling human behavior in such tasks.

1.5 Computational Approaches to Modeling Problem Solving

Spatial problems requiring navigation among points in a problem space, such as the Vehicle Routing Problem, TSP, and naturalistic versions of this problem (e.g., Perelman & Mueller, 2014) are combinatorial optimization problems often referred to as NP-hard or NP-complete due to the requirement of brute force to solve them optimally. However, humans solve such problems near-optimally every day during the course of their daily errands, but also in operational environments such as WiSAR and land

navigation. In order to solve such problems, it is necessary to process the problem space at multiple levels (i.e., hierarchically) to break the problem into more manageable pieces. This section contains a review of the literature on spatial problem solving in the first person (i.e., wayfinding and land navigation) and third person (i.e., E-TSP), including evidence for hierarchical spatial representations and approaches to modeling human problem solving in these areas.

1.5.1 E-TSP: Problem Solving in the Third Person

As mentioned, E-TSP is a routing task in which participants must plot the shortest route, returning to the starting location, through a number of nodes, or cities. The deep structure of the task is similar to naturalistic routing problems (Perelman & Mueller, 2014). In contrast with first person navigation, this problem space is different in a number of key ways. First, the entire environment is available to the subject. Second, the subject must solve the problem using vision alone. In E-TSP, but not in naturalistic tasks that closely resemble E-TSP, the only optimization criteria is path length. Furthermore, the nodes among which the participant must plot the route are identical – there are no semantic bases upon which to categorize these points aside from visual characteristics (i.e., Euclidean distance from other points within the problem space). Finally, as previously mentioned, analysis of data collected by Mueller, Perelman, Tan, and Thanasuan (*manuscript under review*) suggests that a characteristic of human behavior in this task is a tendency to bias optimization criteria, such as estimated time to find, even if

the criteria are irrelevant. Therefore, it seems that people are applying multiple heuristics simultaneously.

Computational models for solving E-TSP often involve nearest-neighbor, cluster, region, or trajectory-based hierarchical planning strategies, whereby the agent solves the problem at varying grain sizes. It is important to note that while they are not necessarily inspired by hippocampus neuroanatomy, the hierarchical E-TSP models involve planning mechanisms; by recoding the problem in terms of manageable chunks, the models provide heuristic near-optimal solutions similar to those created by humans.

Nearest-neighbor models provide suboptimal solutions to E-TSPs, but are often used because they are able to provide a different solution for each starting point, creating a range of solutions that better approximates human behavior (MacGregor & Ormerod, 1996). Retrospective verbal report research indicates that humans use this strategy to some degree during TSP solving (Tenbrink & Wiener, 2009). It is worth noting that nearest-neighbor models behave very similarly to spreading activation models where the goal locations are sparse and unable to reinforce one another (spreading activation models allow multiple nearby goals to reinforce one another, making highly dense clusters of goal locations more attractive to the agent given similar distances to first contact). Nearest-neighbor mechanisms provide a connection between path-planning (i.e., with a bird's eye view) and navigation (i.e., with an egocentric view), as many models of navigation use spreading activation.

Some computational models of problem solving targeting E-TSP use a cluster-based hierarchical architecture (e.g., Graham, Joshi, & Pizlo, 2000; Pizlo et al., 2006)

corresponding to traditional metaphors of visual attention as a zoom lens (Eriksen & St. James, 1986), retrospective verbal report (Tenbrink & Wiener, 2009), and observations of human subjects solving E-TSP (Pizlo, Rosenfeld, & Epelboim, 1995) indicating that human subjects evaluate problems at multiple grain sizes, scanning for both global and local attributes. A basic one-dimensional (as each successive approximation is twice as fine as its parent grain) hierarchical model for solving E-TSP solves the problem at increasingly finer resolution using bisection. This architecture forms a pyramid by which the problem is conceptualized at many different grain sizes. Advantageously, hierarchical architectures permit the agent to generate near-optimal heuristic solutions to route-planning problems while reducing the amount of information (analogously, cognitive demand) handled by the agent at any given time. That is to say that, while the entire TSP is computationally intractable, an agent can generate excellent solutions by optimally solving a number of less complex problems. And, importantly, such models exhibit statistically equivalent behaviors to human subjects. Recent modeling efforts (e.g., Pizlo et al., 2006) advance biological plausibility by incorporating non-uniform acuity within the hierarchy, corresponding to photoreceptor configuration in the human eye.

Cluster-based models require assumptions about how humans solve the higher (i.e., between-clusters) and lower (i.e., within-clusters) level problems. Often this is explained by assuming that humans can exhaustively solve the simplified problems. However, retrospective self-report studies indicate that humans sometimes solve TSPs without exhaustive brute force solutions. One empirically supported way to accomplish this is by first forming a higher-level trajectory using Gestalt-style top-down processing

that incorporates strategic optimization criteria, then solve the decomposed problems at a finer resolution (i.e., for individual points) using bottom-up hardwired perceptual systems or nearest-neighbor heuristics (Tenbrink & Wiener, 2009). One way to conceptualize the higher level trajectory is as an evolving convex hull (MacGregor et al., 2000) or a linear plan based upon a trajectory passing through multiple cluster centroids (Kong & Schunn, 2007). In this approach, a convex hull is drawn around the problem as a coarse plan. This plan is then modified to include individual locations. The modeling approach used in the present dissertation is a computational adaptation of this general approach (see Chapter 4, Section 4.2).

1.5.2 Wayfinding and Navigation – Problem Solving in the First Person

First person navigation and wayfinding, such as that used in daily errands, sports (e.g., orienteering), and operational environments (e.g., land navigation), seems to require at least a two-layer structure in which spatial representations can be dissociated into coarse- and fine-grained representations of the environment (Hirtle & Jonides, 1985; Hochmair & Frank, 2000; Huttenlocher et al., 1991; Plumert & Hund, 2001; Stevens & Coupe, 1978; Weng, Jiang, & Qu, 2008; Werner & Diedrichsen, 2002). These separate representations are typically referred to as clusters (e.g., Hirtle & Mascolo, 1986), although their shapes are often recognized as irregular (i.e., non-Gaussian; Stevens & Coupe, 1978).

Evidence for this dissociation may be found in the human literature using tasks that differ greatly from E-TSP. These tasks include estimating the relative directions of

cities (Stevens & Coupe, 1978) or locations within cities (Huttenlocher, Hedges, & Duncan, 1991), or even the relative locations of buried objects in small environments like sandboxes (Huttenlocher, Newcombe, & Sandberg, 1994). In these tasks, humans typically use landmarks for travel vectors and self-localization, but may also draw on visual (e.g., optic flow), olfactory, somatosensory, vestibular, and proprioception cues. In these tasks, the entire environment is likely unavailable to the participant, in which case they may be forced to operate using mental models or beliefs about the problem space (Hochmair & Frank, 2000). There may be semantic reasons to cluster certain locations, for example membership to a superordinate national category (Stevens & Coupe, 1978) or some other kind of association regarding that location (e.g., associating the grocery store, bank, and beverage distributor as your pre-football game errand route. These associations may include experiential and semantic clustering criteria; Hirtle & Jonides, 1985).

This two-level hierarchical structure differs from the hierarchical structures suggested in the largely visual models proposed for solving E-TSP (e.g., Graham et al., 2000; Pizlo et al., 2006) in that the aforementioned E-TSP models require multiple hierarchical levels as compared with only two. Second, E-TSP algorithms depend entirely upon vision for clustering whereas wayfinding does not require vision at all; given opportunity to explore an environment using non-visual modalities, blind and vision-impaired people perform similarly (i.e., exhibit the same types of errors) to healthy controls (Loomis et al., 1993). Furthermore, as mentioned before, clustering (i.e., categorization) in wayfinding may be accomplished using criteria that are entirely

nonspatial (e.g. Hirtle & Jonides, 1985; Stevens & Coupe, 1978). Finally the problem spaces are not necessarily Euclidean, but may also include networks like city streets (Hochmair & Frank, 2000), and may even include optimization criteria such as traffic on those streets (Weng, Jiang, & Qu, 2008).

1.5.2.1 Levels of Encoding – Fine & Coarse Categorization and Errors in Representation

The human wayfinding literature suggests a two-level hierarchical system (e.g. the Category Adjustment model, Huttenlocher, et al., 1991), whereby spatial encoding is accomplished in a fine to coarse fashion in which humans perceive the physical locations of objects in the environment and then categorize them based on some higher order criterion. The neuropsychological literature suggests a similar mechanism for encoding spatial information, and the opposite mechanism for retrieval. Specifically, during encoding afferent information enters the hippocampus mainly unidirectionally from the EC to the DG, to the CA3 layer (encoding relative locations) and finally to CA1 (encoding location contents; Amaral & Lavenex, 2006). In this line of research, error distributions are used to infer characteristics of the representations used in these spatial tasks. Parallels between the present distributions and other distributions, such as those seen in verbal memory (e.g., Hirtle & Jonides, 1985) and episodic memory (e.g., Huttenlocher et al., 1991) are used to further extrapolate on these characteristics.

Errors evidencing spatial representations appear within less than a second from stimulus presentation (Werner & Diedrichsen, 2002), about the duration of iconic memory (Sperling, 1960), indicating that at least some of the distortion found in object

representation occurs during encoding. Generally speaking, spatial representations tend to emphasize memory for the *coarse* categorizations over the fine grained estimates. For example, where locations are clustered according to semantic criteria, within-cluster distance estimates tend to be underestimated, while across-cluster distances are overestimated (Hirtle & Jonides, 1985; Hirtle & Mascolo, 1986). In addition, in memory the relative directions of cities, judgments tend to be distorted by the superordinate geographical units of each city's state (Stevens & Coupe, 1978). Straight paths are generally remembered as shorter (Hochmair & Frank, 2000) and curves are underestimated or omitted (a characteristic which holds implications for our path reconstruction task presented in Perelman & Mueller, 2013a). In addition, directional judgments are often distorted to conform to cardinal or orthogonal directions (i.e., 0, 90, 180, and 270 degrees; Sadalla & Montello, 1989), producing underestimations of wide angles and overestimation of tighter angles (Loomis et al., 1993). Huttenlocher et al. (1991) asked participants to recall the location of a dot within a circle, finding that participants spontaneously subdivided the circle into quadrants, tending to drift the dot's placement toward the prototypical center of each quadrant. These results, taken together, indicate that humans tend to encode memories for specific locations poorly, supplementing these relatively low fidelity fine grained encodings with high fidelity encoded-higher level representations. While empirical studies examine these representations in the context of error, is it possible that these errors reflect heuristics that enable humans to make reasonable judgments given biological information processing constraints?

1.5.2.2 Superordinate Distortions – Evidence of Heuristics?

Heuristics are simple decision-making processes, perhaps biases, which can enable good decision-making in complex environments that are often intractably complex (Cokely, Schooler, & Gigerenzer, 2009). Or, when misapplied, these heuristics can produce errors. Examining the body of literature concerning biases in spatial representations, it is easy to conclude that these biases are errors that, at best, allow researchers to make inferences about memory representations. However, one study's conclusions illustrate how these biases may be viewed as computational heuristics, every bit as clever as those applied to E-TSP. Take, for example, the predicament often faced by protagonists in television and movies – there are two ways to reach the destination; one is through a scary tunnel filled with threats, while the other involves traversing a far safer but longer route. In short, the path through the tunnel involves a high cost over a short distance, whereas the path around involves a low cost over a long distance. Kosslyn, Pick, and Fariello (1974) tested human distance judgments across barriers, and found that distance estimates across barriers were grossly overestimated. Clearly a bias in experimental settings, this bias may represent a heuristic in naturalistic settings for solving cost-reward evaluations for persons faced with a similar task –attempt to traverse the barrier, or to try and go around it?

Within the wayfinding literature, researchers have identified a number of other heuristics used to solve these otherwise computationally intractable problems. Golledge, Jacobson, Kitchin, and Blades (2000) break the problem space into two specific tasks; wayfinding, which they define as selecting segments from possible routes between two

points, and navigation, which consists of the lower level decision making that occurs dynamically to produce the followed course. According to this definition, navigation requires self-localization within the environment, whereas the wayfinding stage appears to accept strategic input considerations. These strategic input considerations constitute strategies or heuristics, among which the authors identify a number of interesting considerations, such as minimizing turns (c.f., the Least-Angle heuristic; Hochmair & Frank, 2000), optimizing for effort (as in orienteering), optimizing for path length, minimizing intersections (c.f. 2-Opt TSP algorithm; Croes, 1958). However, additional heuristics for solving this problem have been identified, most of which are at least at face value similar to the Category-Adjustment Model (Huttenlocher et al., 1991), which states that in cases of uncertainty, recalling distances (and perhaps other spatial information) may be accomplished by taking the remember value for that specific location and weighting it by the category-wise value. This heuristic for spatial encoding permits the judgment of distances with reasonable accuracy as it presumes low fidelity in the fine grain memories of object locations, and high fidelity in the (comparatively fewer) higher level representations.

Chapter 2: New Analysis of Prior Empirical Research

Developments in this present line of research inspired us to revisit data collected in prior studies, specifically the dataset collected by Perelman and Mueller (2014), in which participants attempted to optimize for path length or time-discounted reward in TSP problems designed to disambiguate strategic optimization criteria, and the dataset

collected by Mueller, Perelman, Tan, and Thanasuan (*manuscript under review*), in which participants completed traditional TSP problems in the PEBL implementation of the task in order to form a normative dataset. Novel analysis of the data collected in these two studies has inspired the research and modeling approaches presented below.

2.1 Criteria Selection

Perelman and Mueller (2014) created problems containing spatial target layouts designed to disambiguate two types of solutions, one optimizing for path length, and the other optimizing for estimated time to find. Participants ($n = 28$) in that study completed the problems under two instruction conditions, one designed to encourage path length optimization, and the other designed to encourage participants to minimize the estimated time to find a target hidden among those points, a distance-discounted reward TSP. Perelman and Mueller (2015) further explored this dataset using a pathmapping approach to calculate pairwise divergence among all drawn paths within each problem to create a symmetrical divergence matrix (similar to a distance matrix), projecting this onto a metric space using Kruskal's Non-Metric Multidimensional Scaling (via the isoMDS function of the MASS package; Venables & Ripley, 2002), and using stepwise flexible mixture modeling to identify solutions associated with particular mental models of the problem space. The authors then aggregated these solutions by cluster, and plotted them in an overlaid fashion (see Figure 2).

2.1.1 Perelman & Mueller 2014 Results Interpretation

The results of this study indicated that participants strategically adapted optimization criteria according to instructions, although participants typically produced a variety of solutions in each instruction condition. Refer to Figure 2, Problem 2 (first row); when provided with shortest path length instructions, most participants produced a solution that involved tracing the loop clockwise (Cluster 5), while the remainder of solutions in this group fell into Cluster 4, which consisted mainly of a shoelace-style trace across the bottom portion of the problem space. Conversely, when participants were instructed to minimize estimated time to find, they produced solutions that fell into Clusters 1 and 3, which represent two separate approaches to solving the bottom portion of the problem space in rows, with the remainder falling into Clusters 2 and 4, which contain shoelace solutions and other suboptimal strategies.

For Problem 4 (Figure 2, row 2), when asked to generate shortest path solutions, participants nearly all (27/28) completed the Z-shaped problem by initially traveling left, then solving for the Z in lines. When instructed to minimize estimated time to find, four participants adopted a strategy that emphasized distance to first contact (the mathematically optimal solution to this problem optimizing time-discounted reward). Of the remaining participants, roughly half adopted variations on this strategy, some of which were very poor, and the other produced the same paths that they had when optimizing for path length.

Problem 5 (Figure 2, row 3) contained a similar Z-shaped problem appended with clusters on each of the vertices and the character's tail. When asked to generate shortest

path solutions, participants generally (aside from a few cases that appear to be errors) solved the problem using the route used in Problem 4, Cluster 3, but the local solutions to the points within the clusters varied. The remainder of the participants produced paths for Problem 5 that fell into Cluster 3, the cluster containing the greatest variance in that set. When attempting to minimize estimated time to find, most (18) participants' solutions fell into Clusters 2 and 4, which are very similar to the optimal estimated time to find solution to this problem globally, but differ locally (i.e., on a point-by-point basis).

Finally, Problem 7 consisted of a V with clusters at each point. When instructed to provide shortest path solutions, participants produced responses (Clusters 1 and 2) that are similar to the global optimal shortest path solution, differing mainly in local solutions. Cluster 2 contains paths similar to those in Cluster 1 with the exception that the first move is to the right rather than the left, respectively. When instructed to minimize estimated time to find, participants' solutions generally fell into Cluster 2 (i.e., the same solutions they provided for shortest path) and Cluster 3, the globally optimal solution for this problem minimizing estimated time to find. Paths falling into Cluster 4, representing paths from both instruction sets roughly equally, mostly contain the highly suboptimal shoelace strategy.

2.1.2 Perelman & Mueller 2014 Discussion

The results of this analysis indicate a number of interesting phenomena. First, participants' solutions reflected strategic thinking, rather than automatic application of rigid bottom-up processing to engage the problem space. In Problems 2 and 7, some

participants applied the highly inefficient shoelace strategy. In some cases, solutions reflect a nearest-neighbor strategy that emphasizes distance to first contact (e.g., Problem 4, Cluster 2). Second, given options that are equal in terms of global efficiency, participants seem to apply different reasonable strategies (see Figure 2, Problem 2, Clusters 1, 3, and 5) as well as others that are suboptimal (Clusters 2 and 4). Third, comparing solutions to each of the instruction sets, the differences appear to reflect *global* rather than *local* planning differences. For example, in Problem 4, Cluster 1, participants' solutions followed the same global trend (down, to the vertex of the top angle of the Z, then split off to the tail somewhere down the segment connecting the two vertices), but differed in terms of local solutions given that global plan.

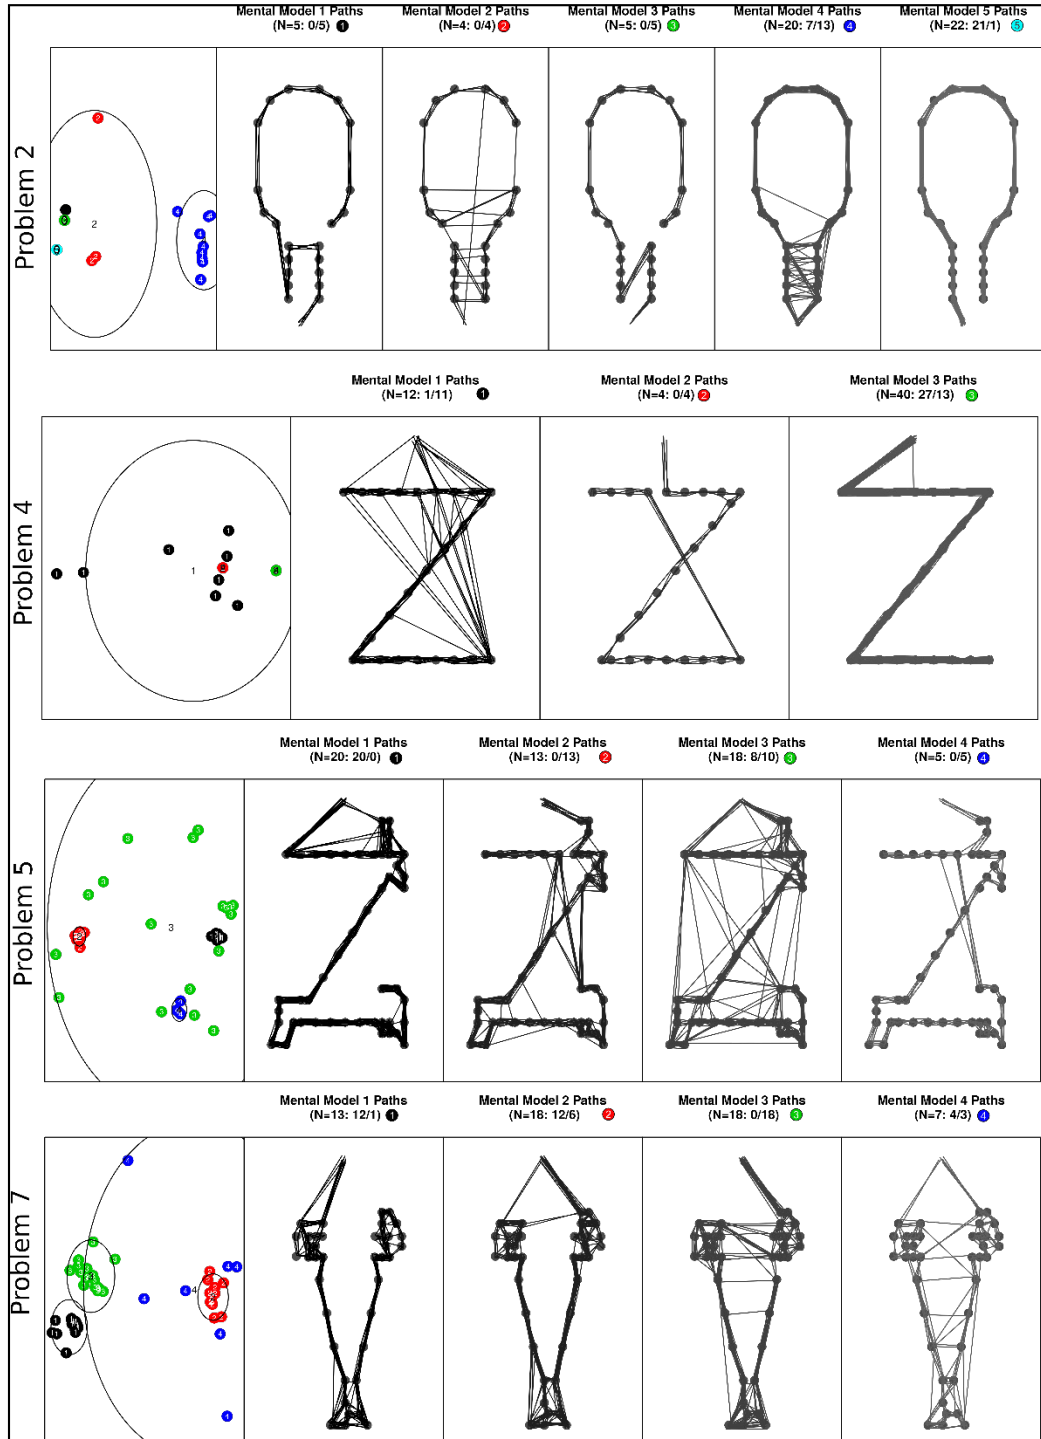


Figure 2. Mixture modeling results for four candidate problems. Leftmost panel in each row shows the isoMDS solution, and remaining panels show clusters of solutions (plotted with jitter). Title indicates the number of solutions in cluster, along with the breakdown between the two instructions (path length / estimated time to find). Details of each solution are discussed in the text.

2.2 Multiple Optimization Criteria in TSP Solutions

Optimal solutions to the TSP minimize the path length in a solution that starts at any point and returns to home. Mueller, Perelman, Tan, and Thanasuan (*manuscript under review*) tested 29 subjects in 50 computerized fixed starting location TSPs consisting of six (five trials), 10, 20, and 30 (15 trials each), to create a normative dataset for the PEBL implementation of the problem. In this section, I reanalyze their data to detect the influence of another optimization criteria, specifically estimated time to find (i.e., time-discounted reward), on their solutions.

This section contains the method for this analysis, which I will refer to as the *reverse solution* analysis paradigm. Since TSP solutions are closed loop, a solution and its reverse (e.g., solving the problem space clockwise vs. counter clockwise) have an equal path length. However, the two solutions may differ in terms of other temporal criteria, such as distance to first contact and estimated time to find (i.e., time-discounted reward). Given that participants are asked to generate solutions on the basis of path length optimization alone, the observations of the two equally-good candidate solutions (the participant-generated solution and its reverse) should be equal between the solution that is better, in terms of that other optimization criteria, and its reverse. To investigate the intrusion of other optimization criteria into TSP solutions, the reverse paradigm solution calculates the efficiency of the reverse solution against its participant provided solution, where a skew away from .5 for a given solution indicates a bias, or the intrusion of that other optimization criteria.

2.2.1 Results – Estimated Time to Find / Time-Discounted Reward Intrusion

As per the reverse solution procedure described above, estimated time to find estimates were calculated for each solution and its reverse as the sum of the cumulative sum of each node distance in the solution multiplied by the number of segments – the order of that segment (i.e., counting down by one from the total number of segments in the solution). This time (or distance, in a TSP) discounted value increases with the average distance required to visit each location, so better (i.e., lower) values are achieved by covering more of the problem space as early as possible in the total solution. Of the 1358 solutions suitable for analysis, 861 (63.40%) were consistent with the form that was superior in terms of estimated time to find, indicating the intrusion of time-discounted reward strategies into traditional TSP solutions. An exact binomial test indicated that this difference was significant from the expected distribution given chance, $p < .001$.

Given that both solutions (i.e., the participant-generated solution and its reverse) are equally good in terms of path length, is using a time-discount optimized strategy impairing or helping participants' ability to generate good path length solutions? In order to investigate this question, I applied the reverse solution paradigm to only the solutions that matched the optimal solution for that problem. If applying time-discount strategies to TSP is helping participants, then we should see a greater skew toward this strategy in the problems that are solved optimally versus problems generally, and some of the worse solutions (those with path lengths 15% longer than optimal). Of the 317 optimal solutions in this dataset (23.34% of all solutions), 206 (64.98%) were also superior in terms of

estimated time to find. Of the 93 solutions containing path lengths at least 15% worse than optimal, 66 (70.97%) favored the solution form with superior estimated time to find.

One possibility for this behavior is that a bias toward ETF-optimization occurs throughout the problem solving process. Another possibility is that the bias is the result of a lack of foresight – that is, the bias appears because participants underweight the segment returning to home at the end of the tour. To investigate this second possibility, I performed the aforementioned analysis with the exception that the last segment in each tour was omitted, eliminating any variance attributable to the last segment's length. Removing the final segment, participants favored the solution form with superior estimated time to find in only 52.28% of trials, an effect that was only approach significance according to a binomial test, $p = .098$. However, a binomial test did reveal a significant effect for the optimal solutions, $p = .018$, with 56.78% of them favoring the form with superior estimated time to find. No such effect was found for the poor solutions, $p = .534$. The results of this analysis indicate that a great portion, but not all, of the variance in the observed bias is attributable to participants' inability to account for the return to home.

2.2.2 Results – Distance to First Contact Intrusion

Distance to first contact is the primary consideration of the nearest-neighbor algorithm. Distance to first contact for each solution and its reverse were calculated according to the reverse solution paradigm described above. Of 1358 total trials, participants selected the solution form favoring distance to first contact in 931 (68.56%)

of those. An exact binomial test found this effect to be significant, $p < .001$. In addition, of the 317 optimal solutions, 213 (67.19%) favored distance to first contact, and 75 (80.65%) of the 93 inefficient (i.e., path length at least 15% greater than optimal) solutions favored distance to first contact. These results suggest that distance to first contact also intruded into participants' solutions.

2.3 General Discussion

Analysis of human performance in a computer implementation of the TSP indicates the intrusion of irrelevant optimization criteria, specifically time-discounted reward and distance to first contact (Section 2.2). This quality is reflected in the mental models with which participants approach global optimization, as seen in Section 2.1. These results, taken together, indicate that computational models of human problem solving in TSP-like naturalistic tasks should be capable of producing similar solutions given the input of multiple, perhaps competing, optimization criteria.

Regarding the analysis in Section 2.2, the proportions of trials favoring time discounted reward and distance to first contact were roughly the same for the all trials and the optimal trials, but the proportions of bias were higher by roughly 10% for the very poor trials, indicating that this bias toward irrelevant optimization criteria may, in some cases, be deleteriously affecting participants' ability to solve TSP solutions optimally. Further research will be required to explore this hypothesis.

Chapter 3: Empirical Research

In naturalistic tasks, the optimization criteria may be defined in terms of very specific goals. In some cases, these goals are unitary. For example, if the task is to plot a route for an unmanned aerial vehicle (UAV) designed to deliver items to consumers, then assuming that there are no issues with air traffic, the problem space appears largely like a traditional E-TSP. However, if that same UAV is tasked with searching for a missing person, the problem space appears more like a time-discounted TSP (Perelman & Mueller, 2014) in which the optimal strategy is to search as many locations as early as possible to maximize the probability of finding the target before it moves, or otherwise perishes. The purpose of this empirical research is to build on prior research in the areas of third and first-person problem solving, and provide evidence for either Gaussian or linear spatial representations in naturalistic problem spaces. This research is designed to inform computational models of human problem solving.

3.1 Rationale & Hypotheses

The purpose of this dissertation is to investigate and model human planning and memory in naturalistic search tasks. The literature contains a clear consensus that memory is required for pathfinding, as it is in essence experiential learning in the spatial domain. However, the ways in which memory is required for solving problems that require planning and mental simulation is less clearly established in the literature. Key assumptions, moving forward, are that (1) near-optimal human performance in problem solving tasks cannot be explained by bottom-up mechanisms alone, and solving such

problems requires top-down planning, and (2) that these top-down plans are a product of hierarchically organized (i.e., at least two-layer) memory structures that allow complex problems to be subdivided into manageable pieces.

It is worth noting that one consequence of these assumptions, which is beyond the empirical scope of this dissertation, is that patients with memory deficits should exhibit degraded performance in solving problems such as the TSP. This assertion is supported by simulated lesion studies using transcranial magnetic stimulation in healthy controls (Cutini, Di Ferdinando, Basso, Bisiacchi, & Zorzi, 2008), as well as subjects suffering closed-head traumatic brain injuries (Basso, Bisiacchi, Cotelli, & Farinello, 2001).

Although the present studies investigate human performance in tasks that are ecologically valid and used in the real world, the investigative substrate of this dissertation is memory, more specifically memory representations involved in, and resulting from, planning during search operations. Three human subject studies will investigate the following research questions:

1. What planning-based issues do professionals face in search and rescue domains? What role does memory play in actual search operations?
2. What are the characteristics of memory representations in search? Are these memory representations based on contextual sequencing? Does the planning stage influence the memory representations either through proactive interference or providing scaffolding for retrieval, or are they simply subject to decay?

3. Does planning and plan execution differ in the first person, as opposed to 2-D topographical environments? Do memory representations differ between the two environments?

3.2 Study 1 – Subject Matter Expert Interviews

During the course of their operations, professionals working in domains such as wilderness search and rescue, bail enforcement, and firefighting create and execute complex plans. These plans permit the operators to execute relatively complex searches by simplifying (i.e., chunking) those complex trajectories into relatively simple representations.

3.2.1 Study 1 – Method

The first study of this dissertation consisted of six interviews with subject matter experts (SMEs), recruited via convenience sampling from the author’s professional network, designed to improve our understanding of planning during search operations in the aforementioned domains. This number of experts follows the recommended guidelines for cognitive task analysis (Clark, Feldon, Van Merriënboer, Yates, & Early, 2008). Information gleaned from these interviews was used to define the domain boundaries for the present empirical and modeling approaches. Each interview lasted roughly one hour. The SMEs provided input on various aspects of planning and executing plans in search and rescue operations. Some sample interview questions include,

1. How are plans constructed for search operations?

2. How detailed are search plans?
3. During search operations, to what extent are operators allowed to exercise initiative to deviate from the plan?
4. During search operations, how do operators adhere to the plan?
5. During search operations, in what ways do the actual search trajectories differ from the planned trajectories?
6. During search operations, how do operators remember where they have searched, and what they have found?
7. Do you currently use any tools for constructing a plan, logging your routes, and logging what you have found?
8. *Simulation Interview*: Without naming any individuals, locations, or organizations, please describe a typical search and rescue operation from start to finish.

The basic structure of the interviews followed a structured interview for knowledge pertaining to search, followed by an incident-based simulation interview based upon the Knowledge Audit technique (Militello & Hutton, 1998). However, the present line of questioning treats plan construction, rather than the domain-specific job (e.g., bail bonds enforcement), as the expert task. These interviews departed from applied cognitive task analysis methods, such as the Knowledge Audit and Critical Decision Method (Crandall, Klein, & Hoffman, 2006), because their goal is not to elicit job-specific knowledge, but rather to elicit domain-independent knowledge about planning and provide ecological context to the human subjects studies and modeling. The subject matter experts

interviewed include an experienced bail bonds enforcer, two WiSAR operators, and a geospatial intelligence analyst, and two military specialists (see Table 2).

Table 2. Subject matter experts, their relevant experiences, and critical observations

SME Interview Number and Date	Domain & Experience	Critical Observations
SME 1	MI Upper Peninsula Wilderness Search and Rescue, ~2 years	Optimal exploration vs. exploitation strategies may include executing both simultaneously (i.e., conducting a hasty search while constructing the probability map for a more detailed search).
SME 2	Bail Bonds Enforcer	Information constraining search offloaded to databases wherever possible. Exploration vs. Exploitation: Optimal time spent in each phase determined by domain-specific characteristics. In bail bonds enforcement, most time is devoted to establishing the target's location and confirming that location via multiple information sources.
SME 3	Geospatial Intelligence Analysis in US Army, Navy, Air Force, and Marine Corps. Enlisted 21 years, officer 12 years.	Probability map construction and search execution are independent. Probability map construction depends upon training (geospatial analysts vs. imagery analysts). Probability maps are solved (i.e., directing intelligence-gathering assets) informally by the flight commander using heuristics. Optimization criteria includes probability of detection, not just target presence in a given location.
SME 4	Newfoundland, Canada Wilderness Search and Rescue, 2.5 years	Optimal exploration vs. exploitation strategies may

		<p>include executing both simultaneously.</p> <p>Probability maps constructed using combination of human input and software, such as GIS programs.</p> <p>Human operators test-fit lost person model parameters. Trust in automation is relatively low.</p> <p>Searchers employ technology to offload information about where they have searched (using GPS trackers) and what they found in those locations (using digital cameras, sometimes with geotagging)</p>
SME 5	<p>US Army, 13 years, 11B (infantry), Sergeant First Class.</p> <ul style="list-style-type: none"> • 4 years Airborne Infantry (2 deployments to Iraq) • 9 years National Guard (NJ; 1 deployment to Iraq, 2 to Afghanistan) • Currently an Operations Advisor to Republic of Albania's Defense Forces training Afghan National Army in Afghanistan • Experience as Platoon Regimental Tactical Officer, Team Leader, Squad Leader in a Recon Platoon, Weapons Squad Leader, Iraq Army Transition Team NCO (2008-2009) Assistant Battalion Operations NCO, Heavy Weapons Platoon Sergeant, and Advisor Team Engineer NCO (Afghanistan, 2011-2012). • Experience in nearly all aspects of light infantry warfare and urban mobility / breaching operations. 	<p>Probability map construction expertise depends heavily on the unit's command level, with the highest quality intelligence training found at the battalion level and above, and the greatest area of operations familiarity at the platoon level and below.</p> <p>Plans are often constrained by doctrine, usually in service of preventing fratricide.</p> <p>Patrols relay information on targets and location in real-time at the appropriate unit level, so there is no need to retain that information in memory / it is not subject to decay or error upon retrieval.</p> <p>Training deficiencies present a serious boundary to appropriate probability map construction.</p>
SME 6	<p>US Army, 9 years, Cav Scout to Brigade and Theatre-level Intelligence via Branch Detail Program. 15 Years contractor to US Army.</p>	<p>Plans provided to assets are kept as simple as possible to maximize agent initiative in local decision making.</p>

	<ul style="list-style-type: none"> • Experience as an intelligence officer in operations from theatre to brigade level, with experience in policing actions in South America. 	<p>Problem space is mathematically more complex than time-discounted reward. Searching for mobile targets involves temporally and spatially pairing named areas of interest (NAIs) with assets in a collection matrix.</p> <p>Optimization incorporates probability of detection based upon the expected signatures at the NAI (i.e. heat, electronic, etc.) and the asset's capabilities (e.g., thermal or electro-optical)</p> <p>Assets are not always deployed to high probability locations, but often to routes of travel.</p>
--	--	--

These SMEs were selected for a number of reasons. First, they span the civilian, law enforcement, and military domains in which search occurs. Second, the search problems facing operators in all of these domains are similarly time-discounted in some way, unlike many routing problems in which all locations must be visited, and the order of visitation does not matter. Finally, while they operate at various organizational and command levels (i.e., bounty hunters often act as individuals, WiSAR operators in small teams of roughly squad size, or 6-12 men, and military operations can occur at all unit sizes up to brigade-level), they are all executing roughly the same task – find a missing person.

3.2.1.1 Wilderness Search and Rescue

WiSAR operations begin with the report of a missing person, often by a loved one, last known to be in a remote environment. WiSAR operators are often volunteers, but they employ (wherever possible) other assets such as rotary and fixed-

wing aerial assets, and search and rescue dogs. To aid their search operations, WiSAR operators construct a probability map of where they would expect to find the target, centered around the point where the target was last seen, and weighted by environmental factors, such as terrain and weather, and characteristics of the missing person, such as physical fitness, mental state, and wilderness experience (Ferguson, 2008; Lin & Goodrich, 2010; Perkins, Roberts & Feeney, 2003). Typically these structured search operations are conducted in parallel with a heuristic hasty search, which checks the highest probability locations independent of other factors.

3.2.1.2 Geospatial Intelligence Analysis

Military intelligence operations are often aided by manned and unmanned aerial reconnaissance assets. These assets are capable of capturing real-time visual data, from the air. Geospatial intelligence is collected as a two-step process. In the first step, analysts (either geospatial analysts or imagery analysts) construct probability maps of the search space. In the second step, the flight commander executes the search with or without referencing the analysts' specific recommendations as a guide. Importantly, the specific route of the intelligence asset (e.g., UAV) is planned and executed heuristically as a time-discounted TSP. The main difference in probability map construction lies in training – geospatial analysts use databases that cross-reference layers of evidence to build target location probabilities, whereas imagery analysts construct probability maps in a heuristic fashion, relying mainly on visual cues from satellite images. For example, the SME in the present study, trained in geospatial analysis, reported participating in the high profile

search for Osama Bin Laden. Based upon Bin Laden's medical requirements as a dialysis patient, the geospatial analysts discerned that he should be located in an area with direct access to shelter, fresh water, and power. As these locations are relatively sparse in remote regions of Afghanistan and Pakistan, the analysts succeeded in cutting down two nations totaling over 550,000 square miles to a number of discrete regions totaling roughly 3 square miles.

3.2.1.3 Bail Bonds Enforcement

Bail bonds enforcers, also known as bounty hunters, are charged with recovering fugitives from justice. Similarly to WiSAR operators, they are charged with locating individual persons, however the problem spaces in which they operate are somewhat different. While WiSAR operators often search for missing persons who attempt to increase their visibility to searchers, bail bonds enforcers often search for fugitives who not only will attempt to evade capture, but will actively deceive the searcher by providing false addresses and contacts who can alert them to the bondsman's attempts at investigation.

Fugitive recovery operations begin with the report of a "skip," or a suspect on bail, who has violated the terms of their release or failed to appear for court. At a minimum, the bail bonds enforcer begins his investigation using the information contained in the bail application that the suspect submitted to the court. However, the information contained within the application is rarely verified, and is often deliberately deceptive. Due to the lack of available up-front information, and the high risk associated

with physically searching the wrong locations, bail bonds enforcers spend far more time in exploration (i.e., research) than exploiting that information (i.e., searching addresses). Research substrates include, but are not limited to, telephone numbers, physical addresses, social media accounts, billing records, and internet activity. While researching suspects, bail bonds enforcers will also research the suspect's friends and family. To help them deal with this large quantity of data, verify pieces of information using multiple sources, and juxtapose individual pieces of data to build evidence for specific theories, bail bonds enforcers often use proprietary databases to offload and manipulate the information.

3.2.1.4 Search in Military Operations

Deployed infantry units plan search missions for high value targets, weapons caches, and other enemy assets within their area of operations with the help of maps and databases of significant activity (SIGACTS), such as the Tactical Ground Reporting System (TIGR) and the Combined Information Data Network Exchange (CIDNE). Intelligence gained during patrols is logged into these databases, and intelligence staff, non-commissioned officers (NCOs), and officers at various levels of command (battalion, company, and platoon) use these databases to plan future operations. As with WiSAR operations, mission planning based on intelligence consists of first constructing the probability space during a phase called Intelligence Preparation of the Battlefield (IPB), and then executing that probability space and updating the intelligence resources with information gleaned from that mission.

During operations, soldiers relay intelligence in real time to the appropriate unit-level command, and also discuss SIGACTS occurring during the operation in debriefings. In this way, the intelligence gathering process can be viewed as a cycle. During the search process, at higher levels of command (i.e., battalion level and up, though soldiers with intelligence training operating at the company level as part of a company intelligence support team, or COIST, will also do this), search is planned by assigning assets (i.e., units) to named areas of interest (NAIs), analogous to high probability regions on a probability map. This task is accomplished by creating a collection matrix, and plotting available assets against NAIs. Assets are assigned to NAIs (i.e., data collection plans are optimized) based on a combination of factors (i.e., optimization criteria), such as congruency between the NAI's signature and the asset's capabilities (for example, an aerial asset cannot collect intelligence on search areas that are underground, though they can collect intelligence on ingress and egress routes from underground locations), factors of convenience such as the distance between the asset's staging area and the NAI, and temporal aspects such as overlap between unit availability and capabilities, and expected signature at those times from the NAIs. While it is rare that one unit will be assigned to multiple NAIs at a time, an aerial asset such as a UAV may be tasked with flying a route or visiting multiple open air locations, and in these cases the flight plan is constructed heuristically by the UAV commander.

At lower levels of command, such as Company and below, the assets will often be infantry units operating with little real-time information. Company-level missions sometimes receive help from battalion-level reconnaissance assets, but at the platoon

level, units must conduct their own reconnaissance, or engage in search missions without prior reconnaissance on the ground. In dismounted search operations, such as through small neighborhoods or villages, one of the strongest criterion for optimization is to minimize the probability of fratricide, and search plans are often heavily constrained by doctrine.

3.2.2 Study 1 - Results

These interviews have revealed a number of interesting findings. First, in applied settings, search is sometimes planned heavily before it is executed, and the extent of that planning is domain-specific. Second, due to the complex nature of the search environments and human memory constraints, operators often offload information to tools such as databases or mapping interfaces. Third, in many of the domains described above, the agents executing the search are not the same agents creating the probability map. Fourth, many of these domains incorporate software that cannot be modified on-the-fly to incorporate the most contemporary techniques, or cannot easily incorporate information that operators have learned or been trained to detect. Therefore, operators' trust in automation is somewhat low, and the use of technology is tempered in these domains by adjusting probability maps post-hoc.

3.2.2.1 Identifying Target Locations versus Search Implementation

Operators attempting to solve search problems must maintain a balance between time spent identifying better target locations (i.e., constraining the search space) and time

spent in planning and execution. This tradeoff bears dynamics similarities to classic exploration versus exploitation problem in reinforcement learning, in which the agent must balance strategy evaluation and implementation. In applied search tasks, the situation is such that the operator must balance time spent researching the target (i.e., constraining the search space) and searching those available areas for the target.

Appropriate exploration-exploitation balance appears to be domain specific. Due to the nature of their targets, bail bonds enforcers want to be certain of the target's location before searching the possible locations. Because they do not want to be found, fugitives will often lie about their addresses, or list addresses of friends and family who will warn them if a bail bondsman inquires regarding their whereabouts. These conditions necessitate a great deal of time spent constraining the search space (i.e., exploration) before the operator can plan a sequence in which to visit those locations (i.e., exploit that information).

Wilderness search and rescue operators address this problem by conducting both early and late searches. Immediately upon arriving on scene, WiSAR operators will conduct a hasty search in which they will quickly search the most obvious locations. While the hasty search is occurring, the incident commander will research the target individual and construct a detailed probability map that can be used to guide a more carefully planned search, should the hasty search fail. Finally, it is important to note that WiSAR operators are sometimes looking for targets who attempt to evade detection as well (M. Hoekstra, personal communication, July 25, 2012), so it may be more

appropriate to describe search location constraint versus search tradeoffs as mission-specific rather than domain-specific.

In military operations, the great wealth of personnel available for a particular mission often exceeds that seen in bail bonds enforcement and WiSAR, permitting command staff to allocate operators to a small number of locations, perhaps even a single location, as circumstances permit. The extent to which the problem space is explored prior to search (i.e., IPB construction) is largely a function of the unit-level of the operation, with larger-scale operations (i.e., battalion and above) having access to the best trained intelligence analysts and assets for collecting intelligence. Conversely, in small unit operations, the unit's leaders often make intelligence products themselves. These small unit leaders have first-hand information of the problem space, similar to experienced, WiSAR operators, and so may have a better understanding of the area of operations than higher level command and intelligence staff despite less sophisticated training.

These data indicate that a dichotomy lies in location selection versus prioritization, which may be temporally separated. Operators select possible subject locations based upon a search which is largely driven by semantic information in a process akin to location selection in foraging (Hills et al., 2008), and represents the operator exploring the problem space. This early effort is contrasted with planning for the actual search operations, in which the operator selects an order in which to visit the possible locations. This later planning and search execution effort may be accomplished using sequence memory cognitive machinery, and indicates that the operator is exploiting

knowledge learned during exploration. Study 2 (see below) is designed to test this hypothesis.

3.2.2.2 Offloading Search Information and Constructing the Probability Space

During the course of a search, operators benefit from access to a wealth of information ranging from aspects of their search through the environment (i.e., where they have searched and what they found) to characteristics of the target, which may be dynamic (e.g., preparedness for weather conditions, or quantity of perishable foodstuffs remaining). To alleviate the memory burden, operators use tools designed to offload and organize this information. This information then factors into the construction of the probability space.

In bail bonds enforcement, operators must constrain the search space through heavy research. In order to store and organize the necessary information, operators create proprietary databases or purchase commercial software. The key benefit to these databases is that they allow bail bondsmen to crosscheck information from multiple sources, which is critical when some sources of information are either intentionally deceptive or merely inaccurate.

WiSAR operators use tools for both constructing probability maps, and for planning and logging search information. Probability map construction is based on analyses of the lost person, and operator prior knowledge of the target area (M. Hoekstra, personal communication, July 25, 2012). The latter point relates directly to past search plans and implementations. WiSAR operators often use mapping tools, called geographic

information systems (GIS), to log past and current searches, and use digital cameras (sometimes equipped with geotagging capability, that links images to specific GPS coordinates) to document what they found in specific locations. Offloading precise information is advantageous in this domain, and even before the use of computers, many WiSAR organizations mandated the use of notebooks to log search information. Offloading information permits operators to overcome human memory limitations and create collective target area-specific histories. In addition to improving information sharing, this means that, should an operator retire or move, their knowledge of the target area is not lost. These histories allow WiSAR operators to analyze the spatial data to detect target area-specific patterns, driving future predictions and planning.

In military operations, patrol intelligence is collected at the Company level. The Company Command Post radio telephone operator (RTO) tracks elements moving through that Company's area of operations, logging the patrol's path. During the patrol, NCOs often write written notes or relays SIGACTS to the Company RTO. These actions permit the patrol to offload information from memory during the search. After the search, patrols conduct internal debriefings, and if the findings are minor they are passed up the chain of command via email, or if they are significant, discussed immediately between the patrol's platoon leader and sergeant, and the Battalion S-2. Collecting information in this way prevents the patrol from having to store important information in memory.

Prior to search missions in military operations, intelligence and command staff use intelligence collected during patrols and uploaded to TIGR to complete the IPB task. The IPB consists of geographical information (often maps derived from Geoquest), and

intelligence comprising a Modified Combined Obstacle Overlay, containing multiple overlays such as key terrain, mobility corridors (including egress points from buildings), and engagement areas (including weapons ranges and effects). This task is analogous to the probability map construction tasks in WiSAR and military geospatial intelligence analysis, however wherever possible the prior intelligence is augmented by real-time reconnaissance assets providing up-to-date information on the area of interest.

Importantly, the quality of assets available to the unit depend upon the command level of the mission assignment. For example, highly trained mapping and geospatial intelligence analysts are assigned only at the brigade level. At the battalion level, IPB is conducted by the battalion S-2, an intelligence branch captain, and senior NCOs and officers who are not necessarily trained. At the company level, COIST conducts the IPB in conjunction with the company Executive Officer (XO). COISTs are often comprised of untrained infantrymen who do not understand the flow of operations as well as the NCOs and officers doing IPB at the battalion level, and COISTs are not necessarily trained in using the required databases and mapping software. The result is that most company-level intelligence products consist of nothing more than a Geoquest map screenshot with graphics overlaid in MS Powerpoint. Finally, at the platoon level, IPB is done by the platoon leader (an officer) and platoon sergeant. Typically requests for intelligence pushed up the chain of command are not fruitful, so the IPB product is again typically a map with Powerpoint graphics overlays. Importantly, however, the relatively small size of the operations permits the platoon leader and sergeant to create a highly detailed product compared with those produced at the company level.

3.2.2.3 Solving the Probability Space

In WiSAR, probability maps are often constructed using a combination of software (to model lost person behavior given personal and environmental variables) and heuristic processes that rely upon familiarity with the operational environment. In military intelligence analysis, probability maps are constructed differently depending upon the analysts' training (i.e., as imagery analysts or geospatial analysts). In WiSAR, the problem space is often created and solved by command staff, whereas in aerial military intelligence, analysts create the probability spaces, but the problem spaces are solved by pilots or commanders operating the aircraft. Sometimes the flight commanders discard the analysts' recommendations altogether. The significance of this dissociation or integration is beyond the scope of this dissertation, but if plan formation is tied to plan execution and the memory for targets and the traveled route, then it stands to reason that dissociating these roles will produce memory effects that differ from those observed in domains where they are integrated.

In ground military operations, after the unit's intelligence team has completed the IPB task, problem space solutions are far more constrained than those generated in WiSAR due to the number of assets available to military units and domain-specific requirements. Physical plans in military search operations are constructed differently at various levels of command. As discussed above, trained intelligence officers, typically found at the battalion level or higher, will construct a collection matrix plotting NAIs (high probability locations) against assets, and pairing assets with locations based on optimization criteria such as probability of detection, temporal factors such as availability

and expected time of arrival if the target is mobile, and the asset's exposure during the task. That is to say that the task difficulty generally lies in assigning assets to individual locations rather than routing elements through multiple high probability locations. To minimize the probability of fratricide, solutions are necessarily constrained by doctrine, and therefore often consist of searching areas such as small neighborhoods or villages in rows or columns so that unsearched structures will remain in front of the unit, with friendly forces remaining behind or parallel, and therefore permitting commanders to engage the enemy without accidentally engaging friendly forces.

3.2.2.4 Trust in Automation versus Training

In WiSAR, operators benefit greatly from the use of software, but do not trust model recommendations blindly. For example, when constructing probability maps, command staff will often use software to build evidence for or against their own mental models, developed via experience and training in lost person behavior. This process consists of adjusting model parameters to account for a variety of possible conditions, and modifying model predictions (i.e., probability maps) post-hoc in a heuristic fashion.

In military intelligence, geospatial intelligence analysts use GIS software to model probable target locations. Again, model parameters are often test fit to compare to the analysts' mental models that incorporate additional information. SME 3 reported that GIS software he had used in searching for targets did not incorporate psychogeography (Debord, 1955), therefore he would refine the model predictions post-hoc to incorporate this type of information. The issue presented here is that software may not allow for the

most contemporary techniques to be used in modeling target behavior. To overcome this shortcoming, operators adjust model predictions to incorporate this information.

3.2.3 Study 1 – Conclusions

The results of Study 1 indicate that memory plays a key role in generating problem spaces used for naturalistic tasks, and wherever possible, the information required for generating that problem space is offloaded into a database of some sort to unburden the operator. For this reason, the problem spaces used in the aforementioned naturalistic tasks do not need to be “chunked,” or hierarchically subdivided, and therefore are protected from bias at the problem space generation level, much the same way that instruments such as GPS systems and compasses protect direction and distance estimates from bias as they remove “error baggage” between actual and believed spatial relationships (Colledge et al., 2000).

While problem space generation in military intelligence and WiSAR benefits greatly from technology, solving those problem spaces is largely accomplished by treating the problem as a combinatorial optimization problem, and applying traditional problem solving heuristics. Though operators will generally define a search plan in terms of segments drawn among “hot spots” or NAIs on a probability map, the actual trajectories traversed by the searching parties tends to be malleable, and subject to “recon-pull” or information available to the searcher once it reaches the destination. Problem space generation is incredibly complex, and requires domain-specific knowledge, and the greatest opportunities for advancement most likely lie in

technological, rather than human, advances. Therefore, this dissertation will focus on coarse-grain high-level plan formation and fine-grain low level plan execution for solving the combinatorial optimization problems presented by domain-independent problem spaces of discrete points in Euclidean space.

The results of Study 1 hold modeling implications for a general model of naturalistic planning and spatial problem solving. First, since planning and execution operations are separate in many operational domains (e.g., WiSAR, geospatial intelligence collection and analysis, and military operations), the model should be capable of planning and execution in isolation. Many existing models of problem solving (e.g., Pizlo et al., 2006; MacGregor et al., 2000) provide an integrated solution that is at odds with this dissociation. Second, models of planning must generate plans that are sufficiently flexible to allow for changes in operational considerations. For example, in military operations, plans are ideally concrete enough to produce the desired outcome, but sufficiently malleable to permit changes. Therefore, the modeling effort in this dissertation will focus on producing a model that can develop plans strategically based upon optimization criteria, with functionality for multiple approaches to local decision making.

3.3 Study 2 – Laboratory Investigation of Planning and Memory Representations

One serious barrier to modeling planning in search and navigation tasks is understanding the mechanism participants use to construct high level plans for search. Even for a single task, such as TSP for example, authors propose multiple dynamics all of

which are biologically plausible. No study to date has attempted to compare these high level planning dynamics directly. The second study of this dissertation will provide direction for my modeling efforts, distinguishing between cluster and linear planning dynamics, by characterizing participants' plan representations through error identification. Error distribution parallels between spatial and other types of memory, for example verbal memory (e.g., Hirtle & Jonides, 1985) and episodic memory (e.g., Huttenlocher, et al., 1991), have been used in prior studies to infer characteristics of spatial memory representations. Therefore, wherever possible I will draw parallels between spatial memory, and characteristics of other types of memory such as episodic and semantic memory. The experimental stimuli for this study will be implemented using a navigation task coded in the Psychology Experiment Building Language (PEBL).

Prior research (Perelman & Mueller, 2013a) has identified two types of spatial errors committed while participants attempt to retrieve target locations in a naturalistic TSP. Each of these error types maps onto one of two distinct types of memory representations in the literature: general spatial confusions, which suggest that target memory shares characteristics of semantic memory, and predict Gaussian error distributions around the actual target location, linear spatial confusions, transpositions (decreasing in recall probability with serial position), and primacy and recency effects, which suggest a link between target memory and episodic memory, and predict linear or ellipsoidal (skewed) error distributions centered on the actual target location, and drawn along the organism's route of advance through the problem space.

Identifying which of these distributions, or in what proportion, best describe human errors in this task is critical not because this knowledge would reduce operator error, but because it provides insight into the characteristics of target memories and plan representations. For the purpose of this dissertation, the term “cluster” will be used to describe Gaussian distributions whereas the term “linear” will be used to describe the aforementioned ellipsoidal skewed distributions. These two distinct possible representations map closely onto different approaches to modeling planning in navigation, with cluster models (e.g., Pizlo et al., 2006) using dynamics very similar to semantic memory (c.f., Hills et al., 2006), and linear and convex hull models (e.g., Kong & Schunn, 2007; MacGregor et al., 2000) using dynamics similar to episodic memory.

The aforementioned errors have been identified in spatial tasks, such as the Corsi block task and a dynamic aerial search task, and non-spatial tasks such as list memory. Mueller (*results not yet published*) noted spatial transposition errors in a Corsi block task that were similarly identified by Perelman and Mueller (2013a) in a simulated aerial search task through multiple possible target locations. Through qualitative analysis of the data, the authors identified two types of errors: transposition errors (i.e., when a subject reported that a target was found at a possible target location previously or subsequently visited to the actual target location) and proximity errors (i.e., when a subject reported that a target was found in a possible target location adjacent to the actual target location). Transposition errors are commonly seen in the verbal memory literature (Hurlstone et al., 2013), and proximity errors may be viewed as related to semantic proximity errors (i.e., phonological similarity in the verbal domain and item similarity in the visual domain,

though Hurlstone et al. [2013] note that these types of spatial errors have yet to be identified in spatial problems involving sequential presentation). We suspect that a third type of error, plan-based errors, may account for some error variance that occurs as a result of the participant resorting to the simpler “plan” representation during recall (c.f., Category-Adjustment Model; Huttenlocher et al., 1991). This error involves participants recalling target positions as erroneously closer in proximity to the higher level simplified plan trajectory. By quantitatively and systematically investigating these errors, this study will drive the modeling efforts to either cluster or linear dynamics, or perhaps a hybrid of both.

3.3.1 Study 2 –Method

This study is methodologically informed by prior tests of human performance on problems designed to force participants to strategically adapt search strategies based on instruction and confuse existing biologically inspired computational models (Perelman, 2014; Perelman & Mueller, 2014). For these experiments, we created spatial problems (see Figure 3, panel A) designed to disentangle the aforementioned memory errors (i.e., proximity, transposition, and plan-based errors). In two experiments, 22 (Experiment 1) and 46 (Experiment 2) Michigan Technological University undergraduate students were recruited via SONA systems, incentivized to satisfy course credit for introductory psychology courses. In Experiment 1, participants used a blue “plan line” to provide global guidance through the problem space. In Experiment 2, participants were allowed

to solve each problem without any guidance. For more on the differences between these two experiments, see Table 3 (below).

Table 3. Experimental Protocol Differences – Experiments 1 and 2

Experiment 1	Experiment 2
<i>n</i> = 22	<i>n</i> = 46
Participants guided via blue “plan line”	Participants received no global guidance
5 Conditions, 15 trials each <ul style="list-style-type: none"> • Condition 1: Cued recall, proximity and plan foils only • Condition 2: Cued recall, targets only • Condition 3: Free recall • Condition 4: Cued recall, proximity foils and targets only • Condition 5: Cued recall, plan foils and targets only 	4 Conditions, 20 trials each <ul style="list-style-type: none"> • Condition 1: Cued recall, proximity and plan foils only • Condition 2: Cued recall, targets only • Condition 3: Free recall • Condition 4: Cued recall, proximity foils and targets only
Targets generated a fixed distance from “plan line”	Targets generated randomly within each “leg”

Participants completed the PEBL-coded experimental task under two main conditions differing in terms of the required memory task: cued recall with and without foils, and free recall. These conditions were further subdivided (see Table 3, above) to include varying types of foil and target combinations.

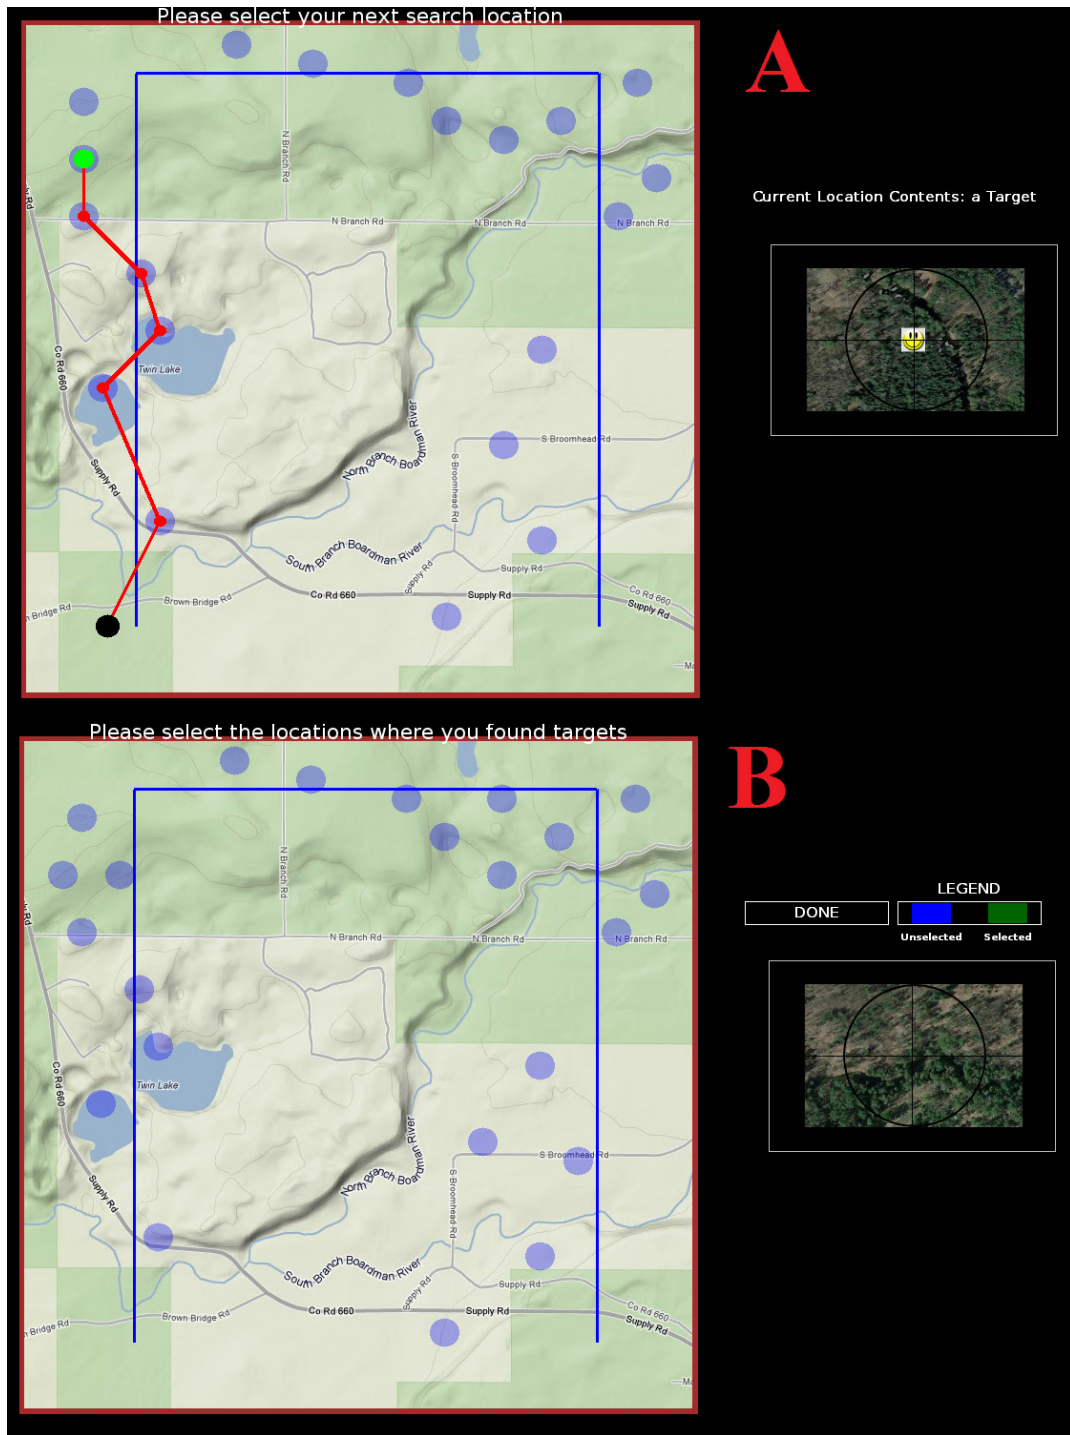


Figure 3. Screenshots of PEBL experimental task used in Experiment 1. Panel A shows the search task, while Panel B shows the memory task in the multiple decision with no correct choice condition (condition 1). In the second condition, participants manually place the target locations. Errors are coded as proximity-based, plan-based, or

transposition errors based upon their distance to the aforementioned foil locations. Note that Experiment 2 was visually identical, however the blue plan line was omitted.

In the search task, participants search (Figure 3, panel A) for targets (simulated camera view right of the topographical map) by clicking on the possible target locations (blue circles). When the search task is complete, participants see a mask over the display (100 ms) followed by the memory test for that condition. In the memory test, some target locations are added and others removed systematically to disentangle error types (see below for an in-depth explanation). In the first condition, participants complete a memory task in which they indicate which of the possible locations contained targets (Figure 3, panel B). In cued recall conditions, participants select which of the possible target locations presented contained targets. The cues presented depend upon condition (see below for an explanation). In the free recall condition, participants must manually indicate the target locations.

The foils used in this study (see Figure 4) are designed to disentangle three types of errors, transpositions (corresponding to linear errors), proximity errors (corresponding to Gaussian errors), and plan-based errors (corresponding to problems with hierarchical processing). These errors hold functional relevance for modeling human planning. Transposition errors are indicated when participants recall the target as located in a possible target location preceding or following the actual target location in terms of serial order of visitation (i.e., of locations that the participant actually visited). These errors indicate that the memory representation is based on sequential mental simulation (i.e., a rough position along the actual path encodes the target's location), and suggest that

planning would be better modeled using trajectory-based models (e.g., Kong & Schunn, 2007; Perelman, 2014). Note that these errors are derived from the serial position of visitation and therefore are not added to the map during the recall task. Spatial proximity errors are indicated by participants recalling the target in a possible target location near the actual target location that they did not ever actually visit (i.e., the location is added for the memory test), and does not lie close to the plan line used to constrain the search trajectory. These errors are expected if the target memory representation is merely low fidelity (i.e., the memory representation is coarse and confined only to a rough location on the topographical map). Such representations are ideally modeled using a cluster model (e.g., Perelman & Mueller, 2014). Finally, in Experiment 1 (where participants are shown the plan line), plan-based errors are indicated when the participant recalls the target in a location that is close to the plan line rather than the target location. These errors are expected if the memory representation used to store target information is bound to the higher order memory representation of the path. Reasons for this include the possibilities that the plan may provide proactive interference for encoding and retrieving the actual target's location, or that it may provide a cue to which the participant can scaffold the target's location. In the second experiment, plan-based errors indicate generalization of the target's position as they are positioned at the center of the target cluster in each leg.

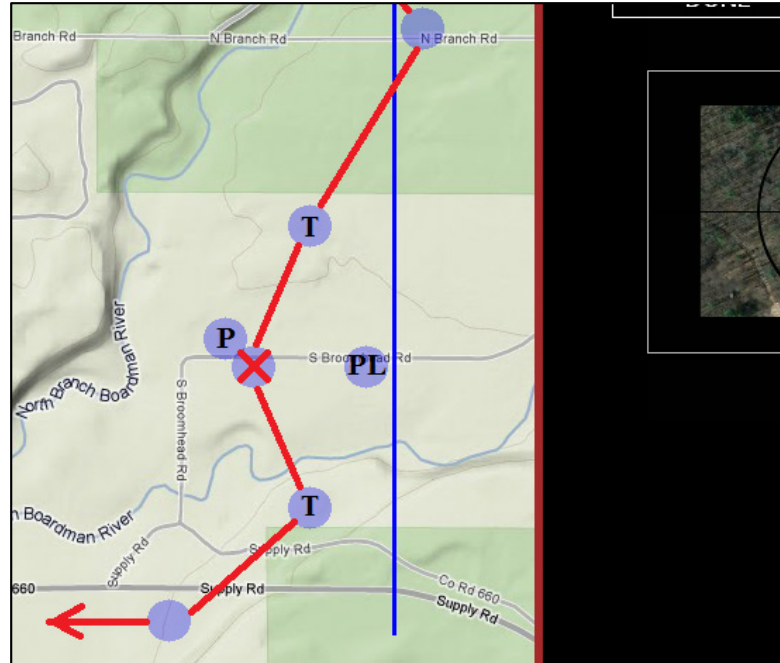


Figure 4. A magnified section of the search space with the actual trajectory and “plan” indicated by the red and blue lines, respectively. The target location is indicated with the red X. Note that the actual target location’s presence during the memory test depends on the condition, while across both forced choice conditions (i.e., 1 and 2) the spatial proximity error location (P) and plan-based error location (PL) are added during the memory test. Incorrect recall of the target location indicates either a spatial proximity error (P), a transposition error (T), or a plan-based error (PL).

After completing the above task, participants completed the Berlin Numeracy Test (BNT; Cokely, Galesic, Schulz, Ghazal, & Garcia-Retamero, 2012). Numeracy is relevant to decision making in search in a number of ways. First, numeracy is highly correlated with working memory capacity, which we expect plays a critical role in planning. Therefore, BNT scores should positively correlate with performance on the memory task. Second, judgments such as those made during the Decision Test paradigm (Perelman & Mueller, 2013a) require weighing distance-discounted rewards. In that study, in aggregate, participants probability-matched the cost ratio associated with the

two possible routes. While the aforementioned problems should constrain the participants' search paths, it is unlikely that we will find individual differences in search performance. However, it is worth mentioning that since numeracy correlates with decision making in risk evaluations (Cokely & Kelley, 2009), we should also expect that BNT scores will positively correlate with performance in general search tasks.

3.3.2 Study 2 – Data Collection and Analysis

The computer task script continually logs participant trajectories and response times for each location clicked during the search task, and participant responses during the memory task. In the cued recall conditions, participants' responses matched discrete points on the map, making error type identification simple (i.e., if a participant erroneously suggests a location indicating a transposition error, then the response is logged as a transposition error). In the free recall condition, participants' responses were characterized in terms of the raw response information, and the responses were also discretized to the closest actual target location presented to them in the search task. This transformation was trivial as the PEBL stimulus script logs all pertinent data, including target and foil locations, participant trajectories, and participant clicks during the memory task.

3.3.3 Study 2 - Hypotheses

The objective of these two experiments was to (1) plot serial position and Euclidean distance curves for the errors, and map them in relation to the target locations,

the traveled trajectory, and the higher order plan, in order to (2) determine the extent to which human memory representations in this task share characteristics with sequences of events, as in episodic memory, or noisy encoding and retrieval, as in semantic memory.

The primary hypothesis pertains to the error distributions relative to the target locations, and states that if recalled locations are normally or uniformly distributed around the actual target location (H_0), then human memory representations for target locations in this task are not tied to the route through the environment (i.e., sequence of events). However, if these recalled locations are not normally or uniformly distributed around the actual target location (H_a), but are rather skewed linearly, it seems reasonable to conclude that participants' memory representations are tied to the path they took through the environment. For the cued recall conditions, this means a greater proportion of errors along the route of travel than perpendicular to it. For the free recall conditions, this means a skewed distribution with less variance along than route of travel than perpendicular to it. While this admittedly a dual processes approach (c.f., Cokely, 2009) to what is most likely a far more complicated problem, the utility of this approach is to provide evidence toward one of two possible modeling directions – one which represents higher level plans linearly, and the other which represents them as clusters.

Additionally, we should expect that the error distributions are consistent with existing literature on serial memory. Specifically, if the cognitive machinery employed in this task is similar to that used for list memory, we should expect typical serial order effects (i.e., decreasing probability of recall with increasing serial position from the actual location), and primacy and recency effects. Finally, since performance in these

tasks requires risk calculation, we expect a significant correlation between task and BNT performance.

7.3.4 Study 2 - Experiment 1 Results

Participant performance in a cued recall condition without foils (Condition 2) provides an estimate of normative performance in this task. Generally, participants correctly remembered the locations of ~75% of the targets (see Table 4, below). The errors followed a Gaussian distribution around the actual target location, with probability of erroneously recording the target location decreasing in terms of both serial position (Figure 5) and Euclidean distance (Figure 6) from the actual target location, and no strong primacy effect, but a possible recency effect, were observed in comparing the probability of successful recall with serial position of the target in Condition 2 (Figure 7). Due to the random nature of target generation in the present study, however, set sizes of targets at each serial position were grossly inconsistent so this data should be interpreted with caution (see Figure 7, left panel). Similar distributions were also seen in the free recall condition (Condition 3). Two-sample Kolmogorov-Smirnov tests indicated significant differences between the recall probability distributions for cued recall and the raw free recall points ($D = 0.83, p < .001$), but this difference disappeared when the freely recalled points were discretized to real locations in the map space ($D = 0.05, p = .15$). Similarly, the two conditions showed no significant differences regarding the probability of recall given serial position ($D = 0.05, p = .19$). Levene's test revealed differences in homoscedasticity between the cued and discrete, $F(1, 2028) = 5.60, p = .018$, but not the cued and raw, $F(1,$

2028) = 2.35, $p = .13$, target recall probability distributions. Levene's test also revealed significant differences between the cued and free recall probability distributions in terms of serial position, $F(1, 1915) = 10.14, p = .001$. Participants effectively used the plan line to encode their responses, as in the cued condition (Condition 2) they never provided a response that lay on the side of the line opposite the target, and in free recall (Condition 3) they committed this type of error in only 7.17% of trials.

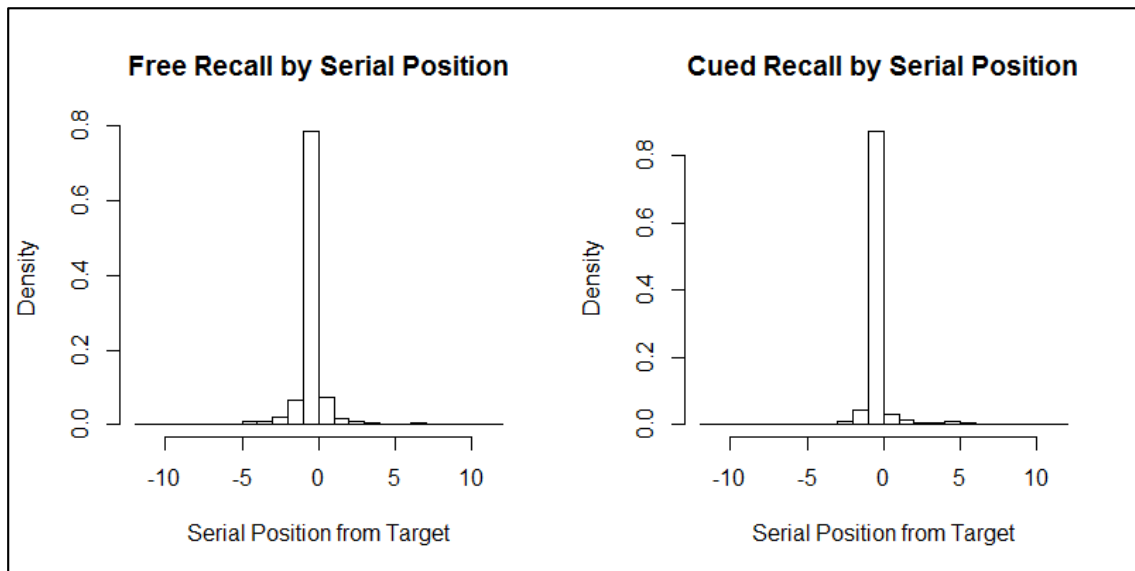


Figure 5. Free and Cued Recall Serial Position Effects.




Figure 6. Free and Cued Recall Euclidean Distance Effects.




Figure 7. Probability of recall given serial position of the target. The panel on the left indicates the number of trials in which a target appeared in that serial position (whole bar), and the proportion of those trials on which participants responded correctly (dark portion of bar). The light portion of the bar represents error. The panel on the right shows the probability of recall given serial position.

Conditions 1 (Plan and proximity foils; no targets), 4 (Proximity foils and targets), and 5 (Plan foils and targets) each provide different clues to the memory representations participants created and used during recall (see Table 4). The results of Condition 4 indicate that participants were able to discriminate between the actual target location, and a proximity foil placed touching the actual target location. In fact, participants were more likely to commit a transposition error than to confuse the proximity foil with the actual target location, despite the close proximity. In Table 4, in Conditions 1 and 4, participants were more likely to erroneously select the proximity foil when it lay along, rather than perpendicular to, the line of travel. However, a Chi Square test of independence showed

that these differences for Conditions 1, $X^2 = .03$, $p = .84$, and 4, $X^2 = .11$, $p = .74$, were not significant.

Participants responded most accurately in Condition 5. This may be because the plan foils, placed at the same distance along the plan line as the target, provided a retrieval cue to the actual target locations. These results should be contrasted with those of Condition 1. Note that participants were instructed to recall the *closest point to the actual target location*. According to these instructions, recalling the target location as the proximity foil would actually be correct. In Condition 1, by this criterion, participants were able to achieve equal accuracy using proximity foils and the actual target locations; again, presumably this effect was caused by the retrieval cue provided by the plan foils.

The proximity foils were generated by deviating the actual target location by 30 pixels across the X and/or Y axes. Therefore, proximity foils could be subdivided into three groups – foils perpendicular to, parallel to, and diagonal to the plan line relative to the target location. The free recall results indicate that participants were able to encode the target distance from the plan line more accurately than the target distance *along* the plan line, though Mardia's test did not detect significant multivariate skew (*Mardia's skewness coefficient* = .004, $p = .27$). This null finding is unsurprising, as Mardia's test often produces high rates of Type II error when applied to short-tailed non-normal elliptical distributions (Gutjahr, Henze, & Folkers, 1999), such as the one produced by participants in the present study.

If participants' memory for distance relative to the plan line was the sole cue used for retrieval, we might expect more errors in proximity foils falling parallel with the plan

line, as participants' free recall responses indicated more parallel ($S.D. = 60.92$ pixels) than perpendicular variance ($S.D. = 47.48$ pixels). This effect was not seen (see Table 4) in Condition 4, where participants were presented with the actual target locations, or Condition 1, where the proximity foil was the most correct answer as per the instructions. The relatively low number of target-proximity foil confusions in Condition 1 indicates that participants draw on more than just line-relative spatial cues to discriminate between locations they actually visited and those that they did not.

Table 4. Percentage of responses, by condition, for each of the locations. Note that responses for Condition 3 (free recall) have been recoded as discrete by assigning the free responses to the closest real point in the map space. The light gray cells refer to the proportion of trials, featuring the proximity foil in a particular orientation with respect to the plan line, on which the participants selected that foil.

Condition	Correct Answers	Other Errors	Plan Foils	Proximity Foil: Parallel to Plan	Proximity Foil: Perpendicular to Plan	Proximity Foil: Diagonal	All Proximity Foils
Foils Only (1)		14.11	10.13	80.97	79.83	70.30	75.77
Cued Recall (2)	75.34	24.66					
Free Recall (3)	78.61	21.39					
Targets and Prox Foils (4)	76.00	12.05		14.90	13.47	9.26	11.95
Targets and Plan Foils (5)	84.60	11.92	3.48				

3.3.5 Study 2 - Experiment 2 Results

The purpose of Experiment 2 was to replicate the findings of Experiment 1 without providing participants with global guidance, to determine whether or not participant-generated global plans would impact memory performance. In addition, the task environment was slightly more challenging; as targets were generated at random

distances from an invisible center line, participants could neither anchor on that center line (the blue plan line in Experiment 1), nor on the target's location on the screen relative to static terrain features.

Participant memory performance in this task was better than that in the first experiment, with participants correctly remembering the locations of ~88% of the targets in the cued recall task, and ~83% of the targets in the free recall task (see Table 5, below). The errors were distributed similarly to those in Experiment 1, approximating a somewhat Gaussian distribution according to serial position (Figure 8) and Euclidean distance (Figure 9) from the actual target location. In Table 5, in Conditions 1 and 4, note that as in Experiment 1 participants were more likely to erroneously select the proximity foil when it lay along, rather than perpendicular to, the line of travel. However, a Chi Square test of independence showed that these differences for Conditions 1, $X^2 = .01, p = .92$, and 4, $X^2 = .79, p = .37$, were not significant.

In the second experiment, two-sample Kolmogorov-Smirnov tests indicated significant differences between recall probability distributions between cued and free recall using both raw ($D = 0.88, p < .001$) and discretized ($D = 0.05, p < .001$) free recall points (see Figure 8). However, this test showed no significant differences between cued and free recall regarding the probability of recall given serial position ($D = 0.05, p = .19$). Levene's test was used to identify differences in homoscedasticity between the above experimental distributions, and found significant differences between the cued and discretized free recall probability distributions, $F(1, 5218) = 11.37, p < .001$, but no significant differences between the cued and raw free recall probability distributions, $F(1, 5218) = 0.95, p = .33$,

and a near-significant difference between the free and cued recall serial position distributions, $F(1, 5200) = 3.69, p = .055$.




Figure 8. Free and Cued Recall Serial Position Effects, Experiment 2.

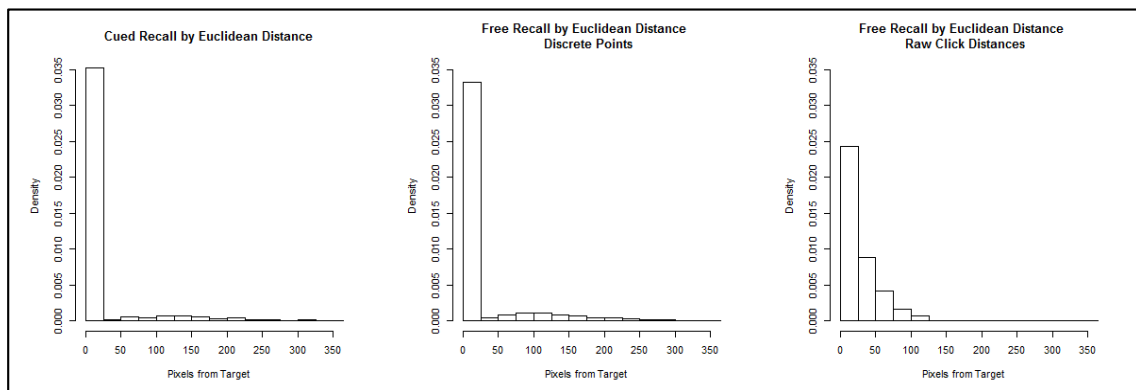


Figure 9. Free and Cued Recall Euclidean Distance Effects, Experiment 2.




Figure 10. Probability of recall given serial position of the target. The panel on the left indicates the number of trials in which a target appeared in that serial position (whole bar), and the proportion of those trials on which participants responded correctly (dark portion of bar). The light portion of the bar represents error. The panel on the right shows the probability of recall given serial position.

During free recall in the first experiment, participants' responses showed a higher variance along their route than away from it. This indicates two possibilities – either this effect is dependent upon the presence of the plan line (i.e., they used the plan line as a retrieval cue, and it was easier for them to determine distance *from* that line than the target's distance along it), or it is independent of that line. In the second experiment, conducted without the presence of the plan line during search and recall, variance along the route of travel was higher (S.D. = 55.44 pixels) than variance perpendicular to the route of travel (S.D. = 35.95), and Mardia's test detected a significant multivariate skew in the distribution (*Mardia's skewness coefficient* = 0.12, $p < .001$). This demonstrates that this effect is not an artifact of the line used to guide participants' paths, and that participants are able to encode the target location's relative distance to the plan line better than that target's

location along the plan line, lending credence to a linear quality of their memory representations used in this task.

Table 5. Percentage of responses, by condition, for each of the locations. Note that responses for Condition 3 (free recall) have been recoded as discrete by assigning the free responses to the closest real point in the map space. The light gray cells refer to the proportion of trials, featuring the proximity foil in a particular orientation with respect to the plan line, on which the participants selected that foil. Condition 5 has been eliminated from Experiment 1. Most interestingly, the absence of a plan line caused the number of plan-based errors to increase dramatically in Condition 1, indicating that participants were using the plan line to detect plan foils.

Condition	Correct Answers	Other Errors	Plan Foils	Proximity Foil: Parallel to Plan	Proximity Foil: Perpendicular to Plan	Proximity Foil: Diagonal	All Proximity Foils
Foils Only (1)		16.47	41.63	47.47	47.07	35.17	41.90
Cued Recall (2)	88.14	11.86					
Free Recall (3)	82.77	17.23					
Targets and Prox Foils (4)	74.81	12.40		16.08	14.29	9.51	12.79

Finally, BNT performance on this task, measured in terms of percent correct, was slightly lower than the theoretically ideal score of .50 ($M = 44.19$, $SD = 23.68$), and roughly close to the expected normative score in empirical research (see Cokely et al., 2012). Pearson's correlations were used to investigate a relationship between task and BNT performance. No significant correlations were detected between BNT performance and Condition 2 (cued recall), $r(41) = .08$, $p = .63$, BNT performance and Condition 3 (free recall, responses discretized to closest actual point), $r(41) = -.07$, $p = .64$, or aggregate Condition 2 and 3 performance scores, $r(41) = -.01$, $p = .93$. However, participant performance between Conditions 2 and 3 was found to be highly

intercorrelated, $r(41) = .84, p < .001$, indicating that strong performance in one task often corresponded to strong performance in the other.

3.3.6 Study 2 - Discussion

The purpose of Study 2 was to determine chiefly if the traveled route will influence a person's memory for target locations in a search task, to inform computational models of human performance in these search tasks. Specifically, if participants' free (Hypothesis 1) and cued (Hypothesis 2) recall error distributions cluster normally around actual target locations, it would suggest that their memory for targets and their locations is not tied to the traveled route. However, if there is a skew in the error distributions such that participants are more likely to place targets at incorrect distances *along* the traveled route, rather than perpendicular to it, it would suggest that participants' memory for targets is tied to the traveled route. It is important to note here that each possible outcome has an analogue in the memory literature; a route-dependent memory should share characteristics, such as context-dependency, with episodic memory (Godden & Baddeley, 1975). Contrarily, an absence of route-dependency in memory for this task, leaving only distance effects (such as normal serial position and Euclidean distance error distributions), would indicate the recruitment of more general visuospatial memory structures for this task.

The results of Study 2 indicate that serial order and distance do influence the probability of recall (see Figures 4-9), and there is a detectable effect of the participants' traveled route whereby targets are more likely to be erroneously recalled along the route

of travel than perpendicular to it. The Condition 3 results strongly support this conclusion – free recall of the targets is clearly skewed along participants’ higher level representations of the problem space.

The statistically significant skew in freely recalled (i.e., Condition 3) target locations along the route of travel in both Experiments 1 and 2 supports the hypothesis that participants’ traveled routes will influence their memory for targets. It could be argued that participants’ error distributions were driven largely by the presence of the blue plan line – perhaps participants were able to judge distance *from* the path better than the target’s location *along* the path, and that contributed to the skew error distribution along the route of travel. Therefore, Experiment 2 was conducted without the presence of the blue plan line. Still, effects detected in Experiment 1 persisted through Experiment 2. The clear effect of coarse-grained plan indicates that these tasks require, and use, comparatively low fidelity high-level problem space representations to minimize cognitive burden.

In the cued recall conditions (i.e., Conditions 1 and 4) of both Experiments 1 and 2, when presented with proximity foils, participants were somewhat more likely to erroneously select foils lying along the traveled route than those perpendicular to it (see Tables 4 and 5, respectively), though these differences were not statistically significant. One possible explanation for the lack of significant results here is that in Condition 1, participants were not provided with the actual target location as a possible choice, so would select the “best fit” point whether or not they recognized that it was incorrect. Both proximity foils lying along, and perpendicular to, the route of travel would be considered

the “best fit” in this situation. In Condition 4, the error rate was relatively low (see Tables 4 and 5), so while the data were trending in the expected direction, a small number of observations (38 vs. 33 in Experiment 1, 123 and 103 in Experiment 2). Note that the Experiment 2 analysis presented above approached statistical significance, owing largely to the greater number of observations.

One interpretation of the difference in statistical significance between the cued (1 and 4) and free (3) recall conditions is the number of observations. In the free recall condition, every recalled point provides a data point whereas only the errors provide clues to route dependent memory in the cued recall conditions. However, the results presented above all trended in the same direction, lending credence to the hypothesis that encoding target locations depends on the traveled route.

Finally, in Experiment 2, far more participants selected the plan foil than in Experiment 1 (see Table 5). One potential reason for this is that, because the target locations were generated randomly with respect to their distance from the plan line, the plan foil was often as good a fit for the target location as the proximity foil. Another interpretation of this data is that, in the absence of a plan line to provide a visual landmark against which to test candidate target locations, participants resorted to the coarse-grain, higher level encoding of the problem space. This interpretation is consistent with the point estimation and wayfinding literature (see Chapter 1, Section 1.5).

3.4 Study 3 – Naturalistic Investigation of Planning and Memory Representations

The third and final human subjects study of this dissertation involves a naturalistic implementation of the above experiment that is functionally identical to the task of a hasty search (i.e., following trails close to the missing person's point last seen) in wilderness search and rescue (SME 1, May 15, 2014; SME 4, December 14, 2014), and shares commonalities with other physical search tasks such as trail following in hiking. This implementation is necessary to demonstrate (1) the extent to which plan representations in the first person overlap with those generated in the third person, and (2) the generalizability of the experimental design used in Study 2 to ecologically relevant tasks.

A naturalistic implementation is a critical departure from the PEBL implementation in a number of ways. First, a naturalistic implementation allows us to test whether plan representations differ in terms of first person and third person spatial navigation. In third person tasks, such as the experimental task in Study 2, the entire environment is available to the participant at any given time, and the path is constructed using purely visual information. In first person navigation, only a subset of the environment, determined by local view, is available to the agent. However, the agent can draw upon stimuli across many sensory modalities (e.g., proprioceptive and olfactory feedback) to bind targets with locations. Second, there are temporal differences between first and third person navigation. In the task described in Study 2, the agent crosses instantly among points, whereas in naturalistic setting navigating the environment takes time, which can be a critical parameter in tasks such as orienteering (Eccles et al., 2002),

foraging (Hills et al., 2006), and WiSAR (Lin & Goodrich, 2009). Testing participants in a naturalistic setting will indicate the similarities between actual navigation and a route planning task, given the same experimental paradigm, and demonstrate the extent to which the findings of Study 2 can generalize to naturalistic, ecologically valid task environments.

While it is possible that participants will perform similarly in both tasks, performance discrepancies will highlight task-specific differences in plan representations. Regarding the modeling effort, these differences will constrain the extent to which the model is generalizable across tasks relevant to applied science communities. Empirically, these differences will highlight task-specific differences in spatial navigation, informing pathfinding research in the basic science community.

3.4.1 Study 3 - Method

18 Michigan Technological University undergraduate students were recruited via SONA systems, incentivized to satisfy course credit for introductory psychology courses, and from the Michigan Technological University Applied Cognitive Science and Human Factors graduate program, incentivized with food. The Portage Loop section of the Michigan Technological University's Nordic Trails (1.1 km) was selected as a test venue for similar problems to those in Study 2. During the course of this study, participants walked the Portage Loop in the Nordic Trails searching for targets (small colored rubber ducks) placed in red plastic cups on the side of the trail. The Portage Loop was divided into three legs, each containing seven targets (see Figure 11), then further subdivided into

problems containing two, three, four, five, or seven targets (see Figure 12). At the end of each problem, participants located a station at which they answered questions about the completed problem. To prevent subvocal rehearsal, participants executed a counting task similar to that used in the Brown-Peterson task (Brown, 1958; Peterson & Peterson, 1959) in which they counted backwards by three from a given number for each step they took while searching for targets on each leg.




Figure 11. The Portage Loop, divided into three legs. The green route is the GPS route traveled by each participant, and each target's location is indicated by a green flag.

After completing each problem and arriving at that problem's station, participants (Task 1) attempted to recall the serial position of each target on the leg (i.e., sort them into the correct order given the target's color), then (Task 2) were shown an overhead abstracted image of each leg (see Figure 12) and provided forced-choice responses to which of two possible locations contained the target. The first task was designed to test for serial position effects, while the second task is a pen-and-paper implementation of the cued recall task in Study 2.

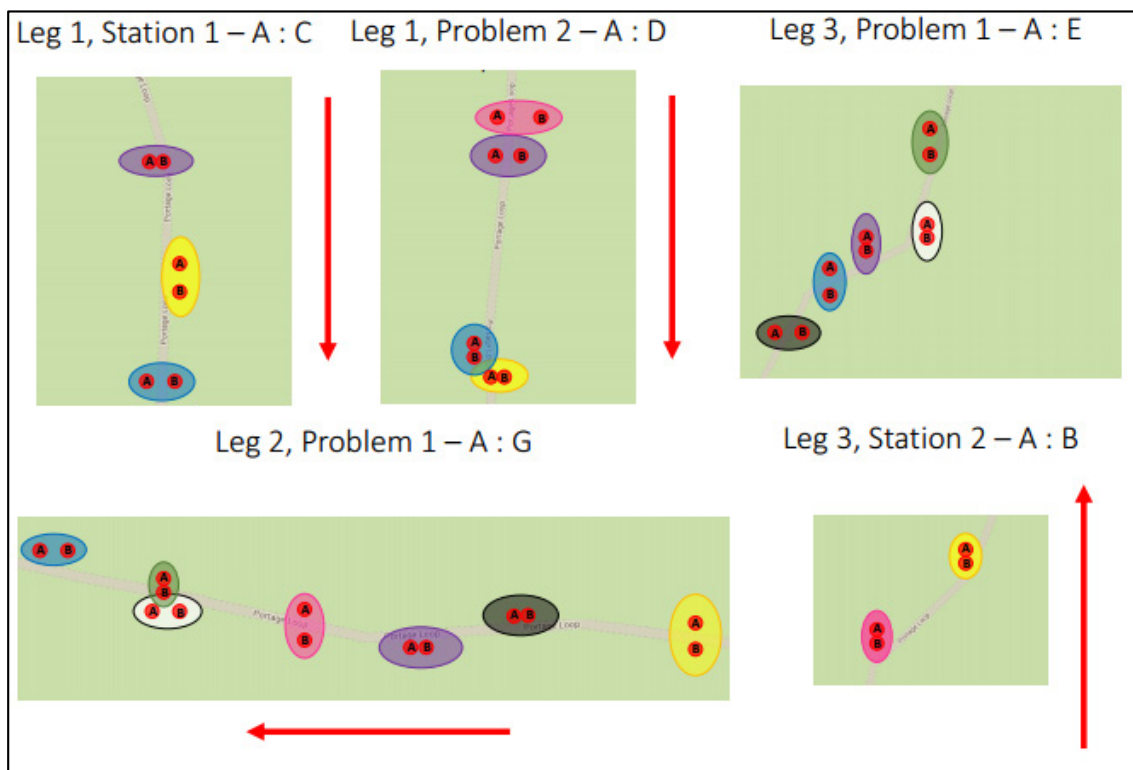


Figure 12. Forced choice stimuli used in Study 3, Task 2.

3.4.2 Study 3 – Hypotheses

The objective of the third study is to determine the extent to which the results of Study 2, a third person task executed via computer, will generalize to a naturalistic search task. To the extent that the two processes are similar, I expect that error distributions, by both serial position and Euclidean distance, will be similar to those found in Study 2. Specifically, Task 1 suggests an effect of serial position on target recall such that the probability of participants recalling a target in an incorrect location will decrease with increasing serial position from its actual location. A second hypothesis regarding Task 1 performance concerns primacy and recency; if the present task uses the same cognitive machinery and serial order tasks, we should expect clear primacy and recency effects. Finally, in Task 2, if the participant's route influences his memory for targets, participant errors are more likely in forced choice options when the foil lies along, rather than perpendicular to, the participant's direction of travel.

3.4.3 Study 3 – Task 1 Results

Task 1 tested participants' ability to recall found targets in the correct order, and the results are intended to compare with the serial position results of Study 2. Generally, overall performance in Task 1 was similar to participant performance in Study 2 (percent correct $M = 83.73$, $SD = 27.01$). One-way ANOVA was used to investigate Task 1 performance, and found an effect of problem length, $F(1, 87) = 25.64$, $p < .001$, whereby participant performance decreased with increasing numbers of targets when using either elements or problems correct as a metric (see Figure 13).

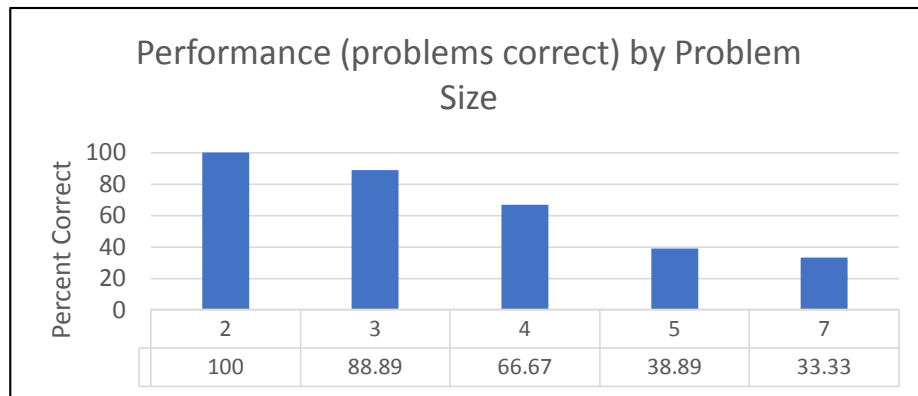
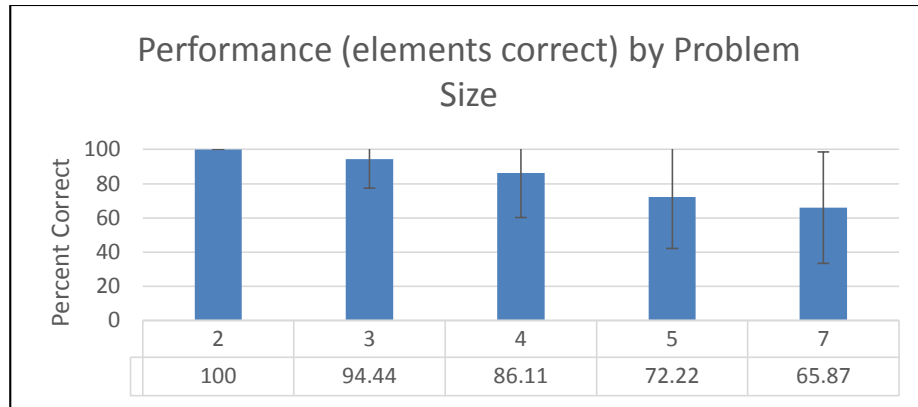


Figure 13. Performance, as percent correct, by problem size (number of targets). Actual percent correct values and problem sizes are shown in the table. Error bars reflect standard deviation.

In addition, participant serial position curves in this naturalistic study matched the serial position results of Study 2 (see Figure 14). Kolmogorov-Smirnov tests failed to find any significant differences between the present Study's (naturalistic) serial position distribution and those generated in Experiments 1, $D = 0.03$, $p = .995$, and 2, $D = 0.02$, $p = .992$, of Study 2. Similarly, Levene's test revealed no differences in homoscedasticity between the present Study's serial position distribution and Experiments 1, $F(1, 1304) = 0.99$, $p = .321$, and 2, $F(1, 2939) = 0.31$, $p = .576$, of Study 2.




Figure 14. Probability of recall given serial position from the actual location for Study 3, Task 1.

In order to investigate the probability of target recall by distance, GPS coordinates of the test site were transformed into Euclidean coordinates (100 pixels measures 13 meters) to allow a more direct comparison to the results of Study 2. Two-sample Kolmogorov-Smirnov tests investigated the similarities between these transformed probability distributions (see Figure 15), and the free and cued probability distributions generated in Study 2, finding no significant differences between it and cued, $D = 0.82$, $p = .055$, and free, $D = 0.65$, $p = .219$, distributions from Experiment 1.

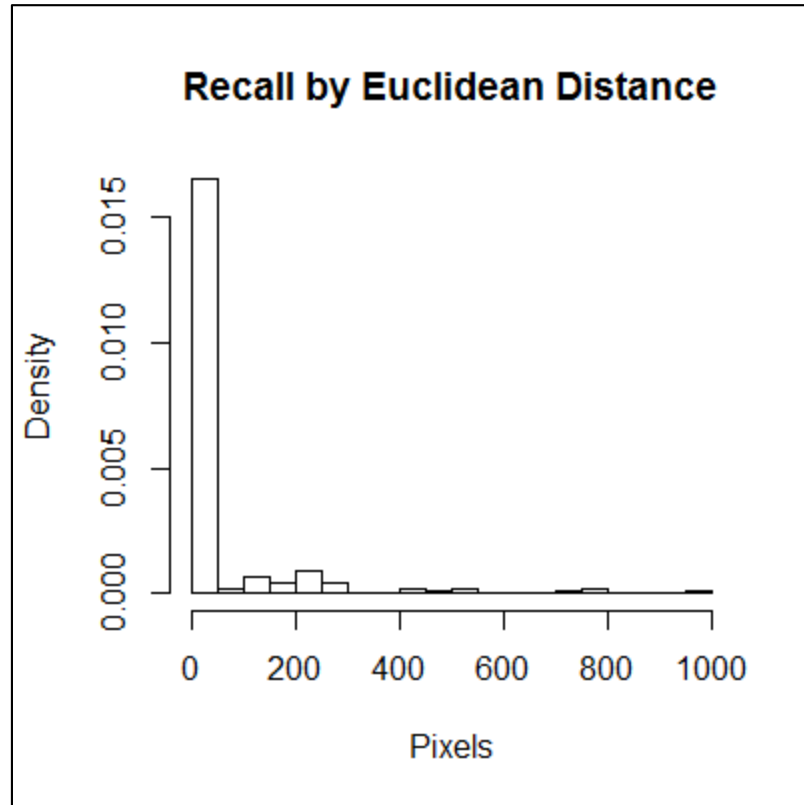


Figure 15. Probability of recall for Study 3 locations, transformed to Euclidean coordinates. At this scale, 100 pixels corresponds to 13 meters of actual space.

However, the tests detected significant differences between the Study 3 distance distribution and cued, $D = 0.10$, $p = .004$, and free, $D = 0.09$, $p = .011$, distributions in Experiment 2. Visual inspection of the aforementioned distributions suggests that the differences lay mainly in the tail, with the (naturalistic) Study 3 distribution having a slightly elevated (but still very low) probability of recall for targets up to 1000 pixels away. Note that both Experiment comparisons were trending in the same direction, but Study 2 Experiment 2 had greater statistical power due to its greater n .

Serial position curves for each station are presented in Figure 16. Participants were able to recall the serial order of targets in the two- and three-target problem sets.

Problems consisting of four or five problems produced inconsistent serial position curves. However, participants exhibited both primacy and recency effects on the seven-target problem. We may be seeing a ceiling effect in the smaller trials that, when removed, reveals that memory for targets in a naturalistic search task has similar characteristics to that of list (i.e., serial) memory.




Figure 16. Serial position curves by problem set size. Title indicates the number of targets in the problem. Data series indicate serial position of that target, while the Y axis indicates the probability of recall for the target at that serial position. Clear primacy and recency effects are visible in the largest problems.

3.4.4 Study 3 – Task 2 Results

Task 2 tested participants' performance in a forced-choice localization test. This Task provides an analogue to Study 2, testing the dependency of participants' memory for target locations on the traveled route. Across all three legs, 10 foils lay parallel to the traveled route, while 11 lay perpendicular to it. Of those 11, seven crossed over the traveled path (i.e., the actual target location and the foil location lay on opposite sides of the trail). This task proved exceptionally difficult, with performance just above chance as measured by percent correct ($M = 51.00$, $SD = 49.62$). In partial support of the aforementioned hypothesis, participants made slightly fewer correct answers when the foils were presented along the route of travel ($M = 50.56$, $SD = 49.62$) than against it (M

= 54.55, $SD = 49.47$), however this difference is not statistically significant, $X^2 = .45$, $p = .50$. Similarly, of the cases where the foils lay perpendicular to the participants' direction of travel, participants tended to perform better when the foil was located on the opposite side of the trail from the target ($M = 56.35$, $SD = 49.79$) rather than on the same side ($M = 51.39$, $SD = 50.33$), though this difference was also not statistically significant, $X^2 = .28$, $p = .60$. Some foils were presented directly adjacent to the actual target locations, whereas others were separated by a space equal to the diameter of the "hot zone" of each of the possible target locations. Whether or not the foil was touching the actual target location did not influence localization accuracy, $X^2 = .77$, $p = .38$. While it is difficult to draw any statistically significant conclusions from participant performance in Task 2, likely due to reasons of statistical power, statistically speaking, the data trended in the directions suggested by the aforementioned hypotheses, and the morphology indicated characteristics consistent with the study hypotheses.

3.4.5 Study 3 – Discussion

The purpose of this study was to replicate the findings of Study 2 in a different, but similar, naturalistic search task. It is often difficult to collect sufficient data in naturalistic settings to generate statistically significant results for a number of reasons, including time required for data collection, subjects' willingness to participate in a physically rigorous study, and environmental constraints such as weather (note that the present study's n of 18 is roughly half that of each of the experiments in Study 2).

Therefore, the data above should be evaluated insofar as they trend in the appropriate directions given the experiment hypotheses.

Task 1 was a serial memory task in which participants attempted to recall the targets in the correct order. This task maps onto the serial memory aspects of Study 2. Performance in this task was found to be the same as in Study 2, with respect to serial position, but slightly different in terms of the error distributions by Euclidean distance; specifically, the error distribution had a longer tail, indicating that participants erroneously recalled targets at greater distances than they had in Study 2. There are two possible explanations for this. First, it is possible that this effect is an artifact of the problem space; the maximum distance in Study 2 (from corner to corner) is 848.53 pixels, whereas the maximum distance between targets in Study 3 is 1067.95 pixels. The second explanation is that participants' memory for target locations in this naturalistic task was simply not as good as their memory in Study 2. This explanation is supported by task factors, such as task duration (each three-legged trial of Study 2 required less than one minute, whereas each participant took about an hour to complete the Study 3 problem), the presence of a distractor task in Study 3, and the use of a greater number of targets (7 vs. 1 per leg) in Study 3. This result should be interpreted cautiously, as while it is statistically significant, the tail consists of only a handful of observations. Nevertheless, gross performance in both Studies 2 and 3 was largely similar, producing statistically similar serial position error curves and morphologically similar error curves by Euclidean distance. Finally, the serial position data given the larger problem sizes indicated clear primacy and recency effects in the serial position curves, providing

evidence for the theory that this task uses similar cognitive machinery to that employed for list (i.e., serial) memory.

Task 2 consisted of a forced-choice localization task, in which participants attempted to assign points to their proper locations given an overhead map of the previously completed problem. Translating from first to third person, in the search and recall tasks, proved exceptionally difficult, and in most cases participants performed barely above chance. Other studies featuring spatial forced-choice tasks indicated that participant performance in such tasks is typically poor (Allen, 1981). However, in support of the study hypotheses, participants were more likely to erroneously select foils lying along their route of travel than those lying perpendicular to it. It is worth noting that, juxtaposing these results with those of Study 2, this effect, seen in both Studies 2 and 3 tasks, is relatively small but consistent. These results, taken together, indicate that participants' traveled route influences their memory for targets they found along the way.

Chapter 4: Modeling Search Behavior

This chapter contains two main sections, the first (4.1) consisting of a discussion of prior modeling work, and the second (4.2) presenting the development of the computational model developed over the course of this dissertation. Section 4.1 contains first a description of lessons learned from other authors' attempts to model human performance in navigation and search, followed by a discussion of our prior modeling efforts. The purpose of this section is to elucidate the successes and shortcomings of

existing modeling approaches, and to frame the process by which we decided upon the modeling approach described in Section 4.2.

4.1 Prior Efforts Modeling Navigation and Search

Taken together, the modeling efforts described in Chapter 1 Section 5 reveal a number of interesting points regarding the suitability of existing models for approximating human behavior. First, approximating human behavior requires a hierarchical structure. The literature above demonstrates this necessity in route planning, but other authors (e.g., Takahashi, 2013) have proposed that other functions, such as encoding and retrieving contexts used in episodic memory, reflect hierarchical organization. In terms of neurobiology, hierarchical structures are ubiquitous in human information processing systems in both the peripheral and central nervous systems. For example, visual information is processed hierarchically as early as the human retina. Similarly, the hippocampus and its connections to the EC, perirhinal cortex, and parahippocampal cortices are hierarchically organized (Lavenex & Amaral, 2000).

Second, existing computational models (e.g., Pizlo et al., 2006) that use hierarchical structure to decompose the problem space into simpler representations typically resort to brute force solutions to the reduced problems. This approach may not reflect the heuristic nature of human problem solving. For computational reasons, it is unlikely that humans use brute force optimization at any stage of route planning in most situations. The wayfinding literature specifically avoid brute force computation, but many computational models of wayfinding focus on validating individual heuristics (e.g.,

Hochmair & Frank, 2000) despite a generally accepted view within the community that multiple heuristics are often used in service of satisfying optimization criteria (Golledge et al., 2000). The computational model presented in this dissertation involves multiple optimization criteria fitted using a single heuristic, weighted with free parameters. That is, the modeling approach presented below can be described as integrative with respect to the extent to which it permits the use of multiple heuristics simultaneously.

Third, biologically inspired models used for navigation and problem solving assume that agent behavior is driven by hard-wired automatic processes, whereas the modeling approach described herein makes no such assumption. Research by Perelman and Mueller (2014) suggests that humans are able to optimize for different requirements, such as estimated time to find versus shortest overall path length. Certain target (or city, in the case of E-TSP) layouts can confuse algorithms that rely solely on hard-wired, structure-based tactics, forcing them to produce solutions that fail to approach optimality in predictable ways. Furthermore, even models that attempt to use heuristic solutions to these problems focus on one heuristic at a time, which limits their applicability (e.g., the least-angle heuristic and its application within network-spaces but not Euclidean spaces; Hochmair & Frank, 2000). Analysis of data collected by Mueller, Perelman, Tan, and Thanasuan (*manuscript under review*) suggests that multiple heuristics are employed to satisfy optimization criteria, even if those criteria are irrelevant. Solving these problems optimally requires some degree of top-down control, permitting the agent to use multiple heuristics either in sequence or in various proportions as optimization criteria, which is lacking in existing models. I propose that (perhaps multiple) strategic optimization

considerations in a higher level planning stage stage allow humans to optimize for different types of spatial problem solving tasks – TSP and search tasks.

Fourth, models must be capable of prioritizing individual target locations. Functionally, in existing E-TSP models (e.g., Pizlo et al., 2006), all target locations are valued equally. In reality, an agent may have to prioritize these locations on abstract bases, such as appetitive considerations in the case of rodents, or mission requirements in the case of humans conducting large-scale searches. Furthermore, such a plan may allow for distance-reward tradeoffs (for example, a comparatively small reward may be visited first due to its proximity) or the mission requirements may be rigid (i.e., a far location must be visited before a closer location). Spreading activation models (e.g., Mueller et al., 2013) are capable of this, but E-TSP models (e.g., Pizlo et al., 2006) are not. The model presented in this dissertation optimizes plans based upon a cost function that could easily be adapted to account for locations of different values. Data from naturalistic studies further indicates strategic control that is not appropriately represented by providing the agent with discrete goals. For example, Tenbrink & Seifert (2011) studied vacation planning trajectories and noted that participants' strategic goals were often ill specified (i.e., deciding to visit the west coast generally, rather than a specific location in it).

Fifth, many existing models using mechanisms such as spreading activation, clustering, and nearest-neighbor are deterministic, whereas evidence from human-generated TSP solutions suggests that these processes should be modeled stochastically, or at least in a manner that produces a range of solutions. Among subjects, humans

produce a range of solutions to TSP problems. Furthermore, one human will rarely generate the same solution to a given TSP problem if it is presented multiple times in a mirrored and rotated form (see Mueller, Perelman, Tan, & Thanasuan, *manuscript under review*; Experiment 3). While solution replication did occur in that study, it was far more frequent in problems to which the individual created optimal solutions. While many of the navigation models are stochastic, many problem solving and neural network models are not. Uniquely, convex hull models are deterministic but capable of producing a range of solutions (MacGregor et al., 2000), and can achieve stochasticity by means of perceptual noise (Best, 2005). These approaches have influenced our modeling efforts (see below).

Sixth, there is a great degree of incongruity between agent capabilities and task requirements that prevents existing models from being used in applied settings. At one extreme, models with the greatest degree of complexity in terms of agent intelligence (e.g., Gorchetchnikov & Hasselmo, 2002) are still only intended to model rodent behavior. These models are used to explore agent decisions, and the neural activity that gives rise to those decisions, in great detail. However, they have little utility in describing human behavior, because they seek to answer questions such as, “what are the dynamics of theta rhythmicity that provide the rat with a direction in navigation?” Conversely, many models used in applied settings, such as emergency egress (e.g., Hajibabai et al., 2007; Klüpfel, 2003), use comparatively unintelligent agents, focusing instead on the flow of bodies through spaces. That is, they are physics models, not cognitive models. These models tend to ask questions such as, “how many people can fit through a doorway

at once?” Therefore, the current state of modeling is divided between two extremes, one which is too specific (and in the wrong way, i.e., to rodent navigation), and another which ignores important aspects of agent cognition. The present effort bridges that gap, creating a computational model that is capable of human-like spatial navigation based upon strategic considerations.

Finally, interviews with SMEs working in operational domains (see Study 1) hold implications that influence the present modeling effort. Specifically, in operational domains, different agents are responsible for creating the problem space, generating plans for solving that space, and executing local decisions. While this effort does not seek to address human performance in probability map construction, the model should produce plans independently of local decision-making, and have a means to evaluate those plans without executing point-by-point solutions.

Furthermore, the model should be capable of producing plans that can be solved according to multiple local decision making strategies, reflecting different mental models (Perelman & Mueller, 2015, and Chapter 1 Section 6) or strategic changes in local decision making necessitated by domain-specific considerations, such as changes in plans due to weather condition or recon-pull factors (such as the search team noticing something in the field that warrants further investigation). The following sections contain descriptions of prior modeling efforts, followed by descriptions of the lessons learned from each modeling effort, and the ways in which those models fell short of accomplishing the stated goal of this dissertation: a computational model capable of approximating human performance in the task used by Perelman and Mueller (2014).

4.1.1 Modeling Planning - Spreading Activation Model

Our first modeling effort was derived from a pathfinding model described by Samsonovich and Ascoli (2005). This model will be referred to as the *spreading activation* model. Mueller et al. (2013) adapted this model to approximate human search trajectories in a ground-based search task, and Perelman and Mueller (2013b) applied the same model to human performance in a synthetic task environment (STE). In that experiment, participants flew a simulated UAV, using a north-up topographical map, and searched for targets using a track-up image representing the view through the UAV's sensor package. The STE was built to approximate the task requirements of WiSAR in two separate tasks, (1) a multiple target search task, which bears a strong resemblance to the E-TSP and (2) a decision task in which participants preferentially searched one of two terrain features varied in terms of their distances to the starting location. Lessons learned in these tasks are assumed to generalize to other tasks in which the subject must plan routes through multiple target locations using a traditional map. Human behavior in these tasks was compared to performance of the aforementioned spreading activation model, and an optimal model which generated cost ratios associated with each parameter of the decision task. This comparison yielded a number of observations,

1. The spreading activation model did not generate solutions as close to optimal as human performance in the multiple target search task because, without walls to constrain the model's decisions, it behaves as a nearest neighbor model. Subsequent analysis using a nearest neighbor model confirmed this suspicion (*results not published*).

2. Human preference for one terrain feature over the other in the decision task, across all parametric settings, as measured by the proportion of participants choosing one terrain feature over the other, matched the cost ratio generated using an analytical model. That is, performance in aggregate matched the cost ratio.
3. The spreading activation model overweighed the importance of distance to first contact in the decision task. One reason for this is that the spreading activation model makes only local decisions, whereas humans have access to the entire problem space.

Experience with this modeling approach indicated that this model, capable of making decisions only at a local level, was insufficient to account for human performance. The spreading activation algorithm, which performed very similarly to nearest-neighbor, takes into account only local information while the results of the decision task described above indicate that humans base their spatial decisions on global information. Therefore, our next modeling effort focused on a hierarchical model, using the same spreading activation mechanism, which was capable of accounting for global information at a higher level, represented by clusters. In this way, the model could solve the problem at multiple levels using a single, relatively inexpensive, mechanism.

4.1.2 Modeling Planning - Hierarchical Spreading Activation Model

In our second modeling effort (Perelman & Mueller, 2014), we applied cluster dynamics to the aforementioned spreading activation model, resulting in a *hierarchical*

spreading activation model. These dynamics permitted the model to view the problem globally and break the problem space into more manageable clusters, then solve for all of the targets within each cluster sequentially (see Figure 17). This clustering is accomplished as follows: First, the agent places five evenly spaced clusters, referred to as *Gaussians*, across the problem space (i.e., one in each corner, and one in the center). The model then creates five values for each goal location, one associated with each Gaussian. This value is the Gaussian-to-goal location distance across a multivariate normal (i.e., Gaussian) distribution. Therefore, goal locations close to the center of the Gaussian receive a high value relative to goal locations farther away. Goal locations are clustered to each Gaussian by means of a winner-takes-all mechanism among all of the Gaussians (i.e., each goal location “belongs” to the Gaussian with the highest associated value).

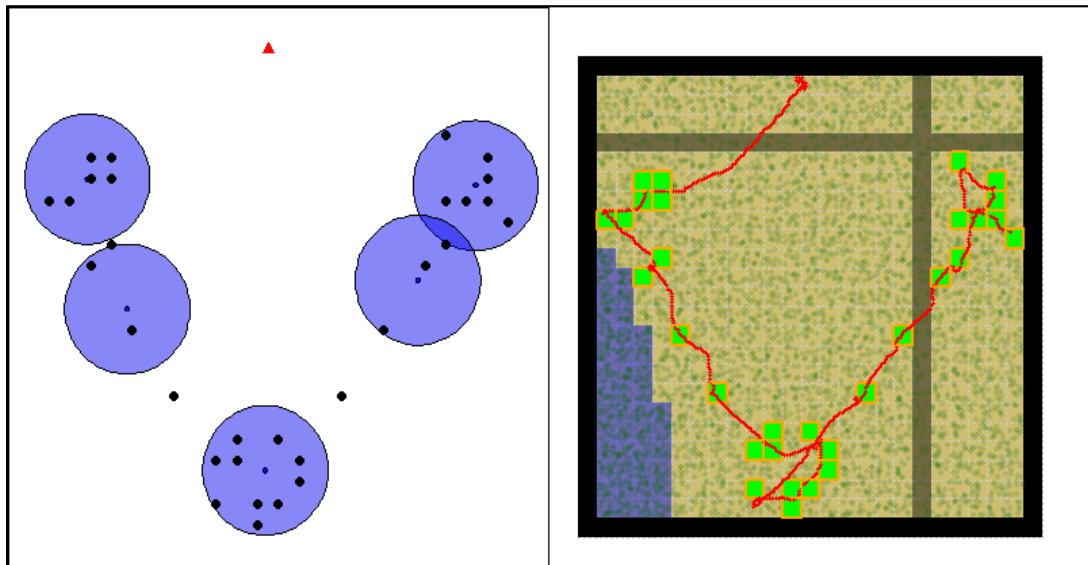


Figure 17. Hierarchical spreading activation model. The red triangle (left) indicates the agent’s starting position. The left panel shows the agent’s logic in clustering the problem space, and the right panel shows the agent’s actual trajectory used in solving the problem.

To optimize the clustering distribution (i.e., Gaussian locations), the Gaussians are iteratively perturbed to reduce each Gaussian's distance to the centroids of its associated goal locations. The way an agent selects the sequence in which to solve the higher level problem is determined by applying a nearest-neighbor algorithm, treating the Gaussians as goal locations, rather than applying brute force mechanics used in other cluster models (e.g., Pizlo et al., 2006). Spreading activation then allows the agent to solve for the goal locations in each cluster sequentially. One problem with this approach is that the model does not solve each cluster in a way that puts the agent in an optimal position for visiting the next cluster. Attempts at weighting the spreading activation mechanism to permit this behavior have proved unsuccessful largely because the agent is still driven primarily by distance to first contact (Figure 18). In order to account for these shortcomings, we have devised a linear plan-based model that is similar in many ways to the Visual Working Memory TSP model (Kong & Schunn, 2007) and convex hull models (MacGregor et al., 2000).




Figure 18. The hierarchical spreading activation model solving a 50-city TSP. Crossovers indicate suboptimal model performance.

The results of this modeling effort, juxtaposed with the results of Study 2 of this dissertation (Chapter 3), indicate that representing the higher level problem space in clusters may be inappropriate as it encourages crossovers. While this could be addressed using a simple rule, there are some cases where crossovers might be advantageous, specifically when creating time-discounted reward solutions (see Figure 2 – note crossovers in the human solutions to the Z-shaped problems). Attempts at using a hierarchical spreading activation model, coupled with the empirical results presented in Chapter 3, prompted us to explore a hierarchical model that represents the problem space at a higher level linearly, using line segments rather than clusters.

4.1.3 Modeling Planning - Linear Plan Model – Random Evolution

Since cluster models fail to produce adequate solutions to search problems when paired with spreading activation, we considered an alternative linear approach to planning that is biologically plausible, computationally feasible, follows the linear structure of sequence memories, does not promote crossovers (c.f. the hierarchical spreading activation model, which does not account for the next cluster's information while solving the current cluster, and therefore is prone to generate solutions with path crossovers), and relates to mechanisms used to solve similar visual problems (Kwon, Li, Scheessele, Michaux, & Pizlo, 2014). This approach is also justified by Study 2, Experiments 1 and 2 detailed in Chapter 3 of this dissertation.

Convex hull approaches to modeling solutions to TSP provide a superior account for the human data compared with cluster models (e.g., Pizlo et al., 2006) because they produce a range of solutions (e.g., MacGregor et al., 2000), and therefore can account for aggregate human data or test-retest data in an individual (Mueller, Perelman, Tan, & Thanasuan, *manuscript under review*). Convex hull models, and other linear planning models (e.g., VWM-TSP; Kong & Schunn, 2007) use a quasi-hierarchical approach in which a line comprises the global plan, freeing the agent to make local decisions. The difference between the two approaches lies in global to local decision-making - in convex hull models, the higher level plan is laid across points on the perimeter of the problem space, and interior points are added to the tour using various insertion methods, whereas the VWM-TSP model (Kong & Schunn, 2007) uses K-means clustering to derive clusters of targets in the environment, then plots a spline-curve through the centroids of those

clusters. Additionally, convex hull models are target-based in that, even at the seemingly higher level representation, individual targets are still accounted. Our linear plan model is a true hierarchical approach, holding many similarities to VWM-TSP, whereby a rough loop (i.e., the higher level plan) is laid across the problem space itself (Figure 19, panel B), irrespective of points, and iteratively deformed (Figure 19, Panel C) to provide the ideal path through the environment (Figure 19, panel D).

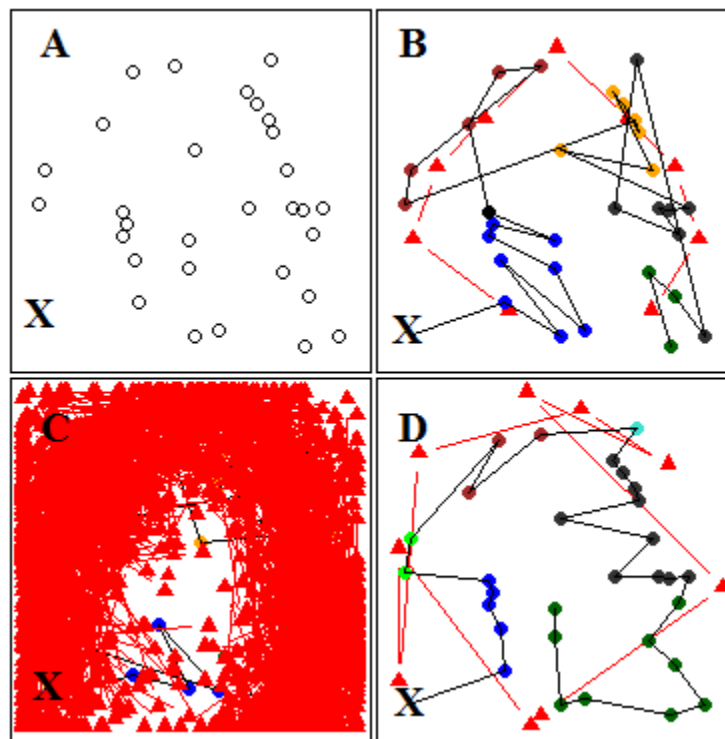


Figure 19. The linear plan model solving a 30 city TSP. X indicates the agent’s starting location. The red line punctuated with red triangles indicates the higher level plan. Panels A and B show the problem space without and with the plan overlay and city assignments, respectively. Panel C shows the results of 200 plans, and panel D shows the selected optimal plan after 200 iterations.

As indicated in Figure 19, the linear plan model is capable of providing heuristic, computationally feasible local solutions to TSPs. The linear plan model works by first

laying a higher level plan across the entire search space. Every target in the problem space is assigned to one of the segments of the plan by calculating the perpendicular of the target to each segment then choosing the minimum (i.e., closest segment). Within each segment, targets are visited in ascending order of a discounted distance value, defined as the distance from the segment origin to the perpendicular intersect of the target (subsegment A) to the segment plus the target-segment perpendicular distance (subsegment B), multiplied by the length of the hypotenuse created by those points (i.e., the segment origin to point distance). The basic result is that the agent travels generally along the higher level plan and visits the targets according to their distance from the origin. Another mechanism, such as spreading activation, would likely provide a similar, perhaps better account for human data.

Despite the merits of this approach discussed above, this model is still inadequate for achieving the stated goal of this dissertation – to account for human performance in naturalistic navigation and search tasks, specifically the results of Perelman and Mueller (2014). Solving those problems requires a model with cognitive control that permits strategically adapting solutions to task requirements, whereas the linear planning model described in this section merely attempts to minimize path length. The SME interviews presented in Chapter 3 (Study 2) of this dissertation indicate that strategic considerations influence spatial solutions during the planning stage, rather than during lower level execution. Furthermore, analysis of data collected by Mueller, Perelman, Tan, and Thanasuan (*manuscript under review*) indicated that multiple optimization criteria may influence TSP solutions, suggesting that the model should incorporate multiple criteria

simultaneously rather than switching between strategies (i.e., optimization criteria). There results, taken together, indicate that a model using a mixture of optimization criteria to constrain plan formation at the higher spatial representation level may be sufficient to account for human performance in the problems designed by Perelman and Mueller (2014).

4.2 Modeling Planning using Mixed Optimization Criteria

The modeling portion of this dissertation focuses on modeling human performance in path planning, and memory tasks related to search and planning, the end result being a model that uses multiple optimization criteria substantiated by research on human problem solving. This modeling effort addresses the following questions,

1. Are traditional memory encoding and retrieval modeling approaches sufficient to account for the human subjects data? Where are the deficiencies?
2. Do linear models provide a better account for the human data than cluster models? To what extent is the use of linearity in planning related to the linear nature of human memory for sequences?
3. Given a problem, can we predict where the agent can trust his “gut instincts” for path generation, and his memory for the trajectory?
4. Are there situations where an agent might want to switch between cluster dynamics and trajectory-based dynamics? Do these two different dynamics correspond to memory effects associated with each type of memory trace (i.e.,

transposition effects for trajectory plans versus grain size effects for cluster plans)?

The task used by Perelman & Mueller (2014) provides an empirical measure against which we can assess model functionality. The model will be judged as satisfactory in the extent to which it can successfully approximate human performance in that path planning and search task given two distinct strategies, one which attempts to produce the shortest overall path length (TSP), and another which minimizes the estimated time to find a target (search).

4.2.1 Creating a Plan and Plan Evaluation

The model presented in this dissertation, referred to as the *mixed criteria model* (MCM), is a two-layer hierarchical linear model (c.f., algorithms for generating UAV routes automatically, such as that presented by Sousa, Simsek, & Varaiya, 2004), mechanically similar to Kong and Schunn's (2007) VSM-TSP model, that follows a route generated by minimizing a cost function based upon multiple strategic criteria. This approach is somewhat similar to multi-criteria optimization in decision-making (Emmerich & Deutz, 2006), though the model presented below is a novel application of the approach to modeling *spatial* decision-making and problem solving.

At a high level, the goal of this modeling effort is to create a model, informed by both naturalistic and laboratory research, that bears characteristics of planning and execution in the real world. In operational environments, planning and execution are often

conducted in isolation by separate agents. To solve spatial problems, the agent analyzes the problem space and creates a general plan (represented linearly) that he, or another agent, can follow while solving the problem locally. This plan is created according to strategic goals, referred to in this dissertation as optimization criteria.

Perelman and Mueller (2014) showed that humans are capable of adapting optimization criteria based on written instructions. Analysis of interview data from operators working in search domains indicates that multiple optimization criteria influence human decision-making (see Chapter 3, Study 1), and prior research (i.e., the data collected by Mueller, Perelman, Tan, & Thanasuan, *manuscript under review*) shows this to be the case even in tasks where some of those optimization criteria are irrelevant. Therefore, we propose a model that is capable of weighting these optimization criteria to various degrees to empower heuristic decision making. These optimization criteria provide weighted inputs to a cost function that drives selection from random deformation of a higher level plan. At present, the MCM evaluates these higher level plans based on five parameters, mapping onto known heuristics used for spatial problems:

- 1. Log number of segments.** The higher this weight, the more the algorithm penalizes the number of segments in a given plan. Increasing this value will result in less plan complexity.
- 2. Goodness of fit.** GOF = mean of the distances of all points to their closest segments squared. The higher this weight, the more the algorithm penalizes the plan's ability to comprehensively account for points in the problem space. Increasing this value will result in more plan complexity.

3. **Path length.** Overall plan length + the sum of the distances of all points to their closest segments, measured in pixels. Increasing this value will encourage shorter plans.
4. **Time-discounted path length.** Mean of the cumulative sums of each plan segment's length * the proportion of the problem space accounted for by that segment. Increasing this value will reward plans that account for more of the problem space earlier in the search trajectory.
5. **Average Angle.** The scaled angle cost, computed as the mean of the absolute value of $\pi - \text{all plan angles}$, scaled by π , will fall between 0 and 1. Increasing this value will reward smoother plans with shallower angles.

4.2.1.1 Parameters - Log Number of Segments

Fewer segments corresponds to fewer waypoints or turns at compass headings. Minimizing segments is important given constraints on information capacity, and may represent a default heuristic across many domains. For example, in providing spoken directions, given a number of routes, it may be reasonable to provide a slightly longer but simpler route as the driver will experience less load in rehearsing the directions and is less likely to forget a turn.

4.2.1.2 Parameters - Goodness of Fit

The extent to which a plan strives to fit individual data points will likely be domain specific, and depend upon the agent's strategic goals. If a very high degree of

plan fidelity is necessary, then the agent will prefer plans with greater goodness of fit. However, this value should be constrained by known limitations of human memory (i.e., weighting the log number of segments parameter heavily to penalize complex plans). A simple heuristic in regards to this parameter is that a plan should fit the individual locations well enough to satisfy the mission requirements, and no better.

4.2.1.3 Parameters - Path Length

The rough measure of path length used by the model represents the agent's feel for the overall tour length given a plan without having to resort to point by point calculations. Agents may want to minimize path length in operational environments involving limited range before refueling, or over homogeneous ground if walking on foot in order to minimize energy expenditure.

4.2.1.4 Parameters - Time-Discounted Path Length

In military intelligence, targets are mobile and attempting to evade capture, so the probability that they will be found in a designated location decreases with time. In WiSAR, the target is often lost in rough terrain under hazardous weather conditions, and the probability of finding the target alive decreases with time. Time-discounted path length is analogous to estimated time to find; in these domains and others, the ideal plan is the one that locates the target in the minimum amount of time. In probability space, this translates to favoring routes that accumulate as much probability in the problem space as early as possible.

It is also worth noting that, in certain problem spaces, the agent may want the inverse of a time-discounted solution. For example, if the agent has to pick up items at every location and carry them to the end point, then the agent benefits from solutions that minimize the overall burden (i.e., the duration of time or the portion of the path for which the agent has to carry the most material). In order to produce these solutions, a negative value supplied for this weight will subtract it from the plan path and favor solutions that save the majority of the probability space for later in the path.

4.2.1.5 Parameters – Average Angle

Angles are important in domains where a plan consists of moving generally in a single direction, such as applying a “plow the field” strategy in WiSAR. Doubling back is generally considered deleterious in this type of problem. Heuristics such as the least-angle strategy (Hochmair & Frank, 2000) place a high cost on angles, and produce solutions that favor few angles. However, this heuristic was only applied by those authors in network spaces (i.e., city street networks); the present approach permits this heuristic to be applied in Euclidean space.

4.2.2 Fitting a Plan

The MCM begins the plan fitting sequence by creating a rough plan; K-Means clustering detects the centers of clusters of points in the problem space (c.f., VWM-TSP; Kong & Schunn, 2007) and the model plots segments through those centroids using a nearest-neighbor algorithm. This represents a heuristic approach to forming a plan “at a

glance” of the problem space, and is consistent with other models of problem solving. From this prototype, plans are iteratively modified and the evaluated, and the plan with the highest current fitness (i.e., lowest cost) is selected as the new prototype. Plan modification is probabilistic, and is accomplished by selecting among available strategies, which consist of,

- 1. Adding a point**
- 2. Deleting a point**
- 3. Moving a point**
- 4. Swapping points**
- 5. Discard & generate new plan**

Furthermore, at every iteration, the algorithm checks for intersections at the plan level and breaks them, if necessary. Eliminating path intersections is a key component of combinatorial optimization algorithms (e.g., 2-Opt TSP algorithm; Croes, 1958).

It is important to note that this optimization algorithm is not intended to exactly replicate the process by which humans create plans for search. Likewise, a goal of this optimization process is not to produce aggregate data consistent with human behavior, but rather to provide a range of solutions across which parallels between those solutions, and human solutions, can provide clues as to how humans produce those particular solutions, for example, whether certain solution patterns are due to differences in global or local decision making. While this optimization algorithm will occasionally settle on the known optimal solution of the problem space, given specific optimization criterion, it should not be expected to do so in proportions identical to the aggregate human data.

4.2.3 Following a Plan - Local Decision-Making

The MCM constructs the higher level plan without direct bottom-up feedback. That is, plan-to-point calculations are conducted in service of approximating path length and time-discounted path length associated with a given plan. However, point-to-point distance calculations are not incorporated into the optimization criteria as we expect the model to work under conditions of limited information access. The MCM, like humans, should not resort to brute force (i.e., exhaustive) optimization at any step of the planning process.

Presently, the MCM uses two distinct local decision-making strategies that map onto those seen in human data – nearest-neighbor and a plan-following algorithm referred to here as *leash* (see Figure 2, and Section 2.1). For example, the routes shown in Clusters 1 and 3 are similar to those produced by the present model using a nearest-neighbor algorithm, whereas the model produces the shoelace paths seen in Clusters 2 and 4 when using the leash algorithm for local decision making (see Figure 20). Simulations using these two different strategies indicate that differences in participants' mental models for TSP-like problems (Perelman & Mueller, 2015) can be explained by differences in local decision making strategies.

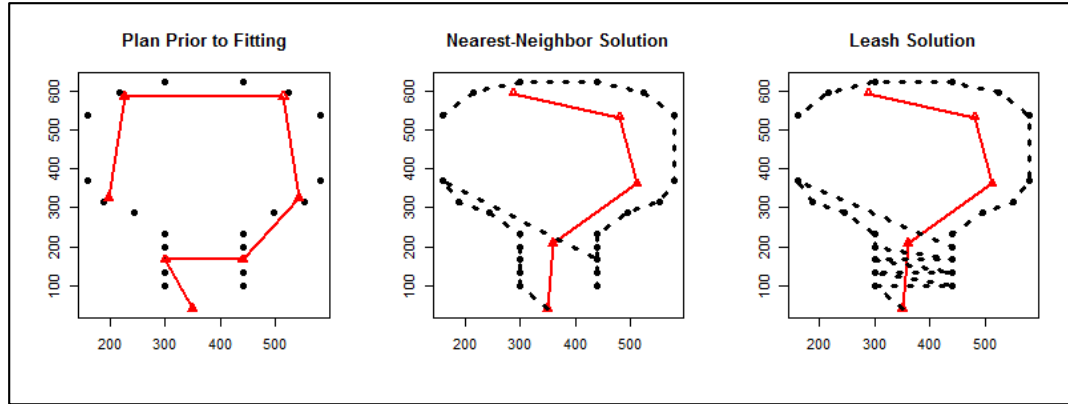


Figure 20. Sample MCM trajectory for Problem 2 (Perelman & Mueller, 2014) produced by the mixed optimization model given parameter values intended to optimize for estimated time to find (i.e., time-discounted reward; log number of segments = 25,000, GOF = 1, path length = 1, time-discounted path length = 250, average angle = 1000). Lines in red indicate the higher level plan, whereas the dotted black lines indicate the point-by-point solutions. The left panel shows the plan prior to fitting, generated via k-means clustering and nearest-neighbor ordering. The center and right panels show the fitted plan, given nearest-neighbor and leash local solutions, respectively. Refer to the text for the appropriate human subjects data mappings.

The MCM follows the plan using a nearest neighbor algorithm for each sequence. That is, from the starting location, the agent solves for all points belonging to each plan segment in sequence using a nearest-neighbor strategy. This provides a computationally inexpensive way to handle local decisions. Nearest-neighbor solutions do sometimes produce path crossovers, and so a heuristic for minimizing these might be incorporated into future versions of this model.

4.2.4 MCM Results

The purpose of this modeling effort is to construct a computational model of problem solving that can account for aspects of human performance in the naturalistic

TSP task used by Perelman and Mueller (2014). To this end, the following hypotheses will be used to gauge model fitness.

1. **Observed differences in planning and optimization:** The model should, through different parameter fits, be able to reasonably approximate human trajectories given instructions to weight different optimization criteria, such as path length or estimated time to find.
2. **Observed differences in local decision-making:** The model should account for fine-grain differences in human solutions by differences either in the higher level plans, or in local decision-making strategies. For the purpose of this dissertation, two local decision making algorithms, nearest-neighbor and a *leash* algorithm, are expected to account for the bulk of the local differences in human trajectories.

4.2.4.1 Observed Differences in Planning and Optimization

Perelman and Mueller (2014) showed that participants were able to adapt to different optimization criteria in problems designed to encourage separate solutions when optimizing for path length or estimated time to find. To test MCM performance in those problems, 250 model runs were executed for each problem and using parameter weights (see 5.4.1, above) designed to optimize for path length and estimated time to find, for a total of 500 model runs per trial. Local solutions, given these higher level plans, were generated using a nearest-neighbor algorithm, and the leash algorithm described in Section 4.2.3.

Regarding the high level plans, the MCM generally settled into a small number of solutions, differently slightly in terms of directionality. Of the 250 plans generated for each of the optimization criteria, 10 are shown overlaid on each problem in Figure 21 to provide a sample for the types of plans generated by the model for each problem, provided each optimization criteria.

When provided with targets as seen in the Problem 2 (Figure 21), the model chiefly provides one solution for path length (a triangle around the problem space beginning and ending at the starting location) and two separate solutions optimizing estimated time to find (question mark-shaped plans that direct the agent either right or left initially). Note that even in a relatively simple problem, the optimization algorithm failed to optimize successfully for path length in one of these samples trials. For Problem 4, the model settled upon two main plans when optimizing for estimated time to find, both driven largely by distance to first contact with the line at the top of the Z, then exploring across that line before solving the rest of the problem. Comparatively, solutions produced for path length optimization in Problem 4 were generally poor, with the modeling settling in rarely to the optimal solution. Path length optimization in Problem 5 proved similarly difficult. Estimated time to find solutions for Problem 5 interestingly followed one of two main plans – either traveling first to the cluster of points at the top right vertex of the Z, then solving the rest of the clusters and hitting the linear points along the way (the optimal solution), or heavily weighting distance to first contact and solving the problem as a normal Z. Finally, model optimization for both estimated time to find and path length given Problem 5 generally encouraged the appropriate optimal solution, with the model

first solving the clusters at the tips of the V when optimizing for estimated time to find, and solving it in the shape of the V when optimizing for path length. Note, however, the one example in path length optimization for this problem where the optimization algorithm failed strongly, placing the second point far into the lower left hand corner of the problem space.

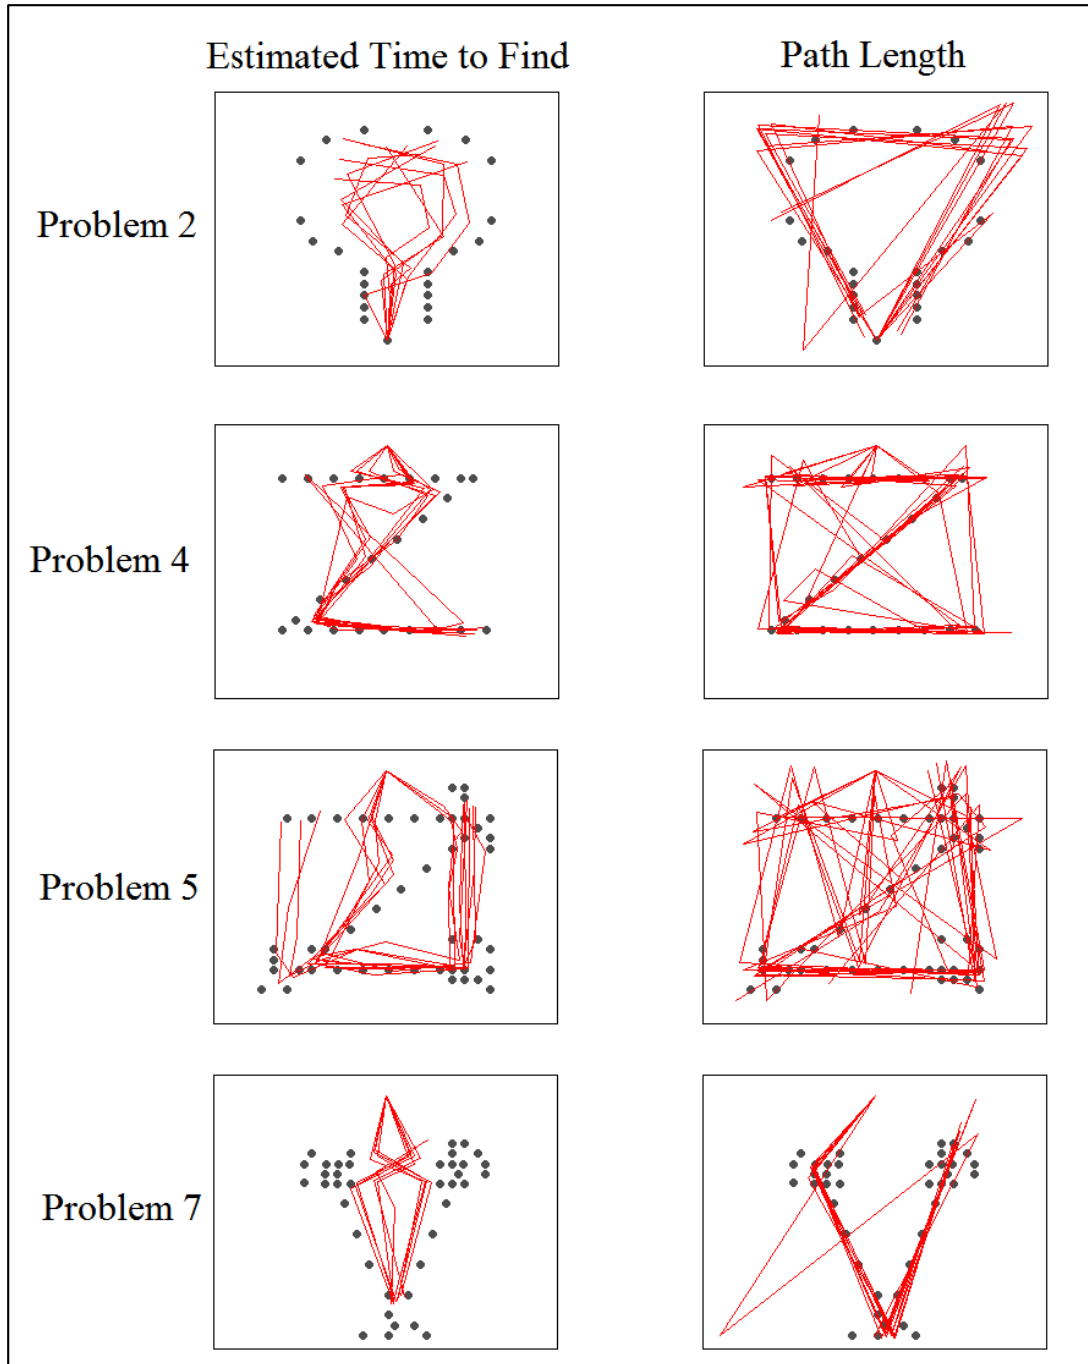


Figure 21. 10 sample plans generated for each problem, given parameter values (see Table 6, below) mapping to the two sets of optimization criteria provided to participants by Perelman and Mueller (2014) – estimated time to find / time-discounted reward, and path length.

These results, taken together, indicate that the MCM is capable of producing plans sensitive to different optimization constraints inherent to different types of navigation and search problems, and that those solutions produce better cost function values than the poor solutions. The limitation currently lies in the optimization algorithm used to generate these high level plans for these problem spaces, which presents an opportunity for future research.

In order to compare MCM and human local solutions for the problems used by Perelman and Mueller (2014), representative model solutions were generated by applying the two candidate local decision making algorithms (i.e., nearest-neighbor and leash) to representative high level plan solutions generated using the parameter settings used for the large-scale simulations and examples shown in 21 above (see Table 6, below). The resultant solutions are available in Figure 22 below.

Table 6. Parameter weights used in the simulations presented in this chapter, given optimization criteria (path length or estimated time to find / discounted reward).

Problem	Optimization Criteria	Log Number of Segments	Goodness of Fit	Path Length	Discounted Reward	Angle Cost
2 (Loop)	Path Length	50000	500	1	.000001	1000
	Disc. Reward	25000	500	1	250	1000
4 (Z)	Path Length	5000	500	5	.000001	1000
	Disc. Reward	5000	5000	.1	500	1000
5 (Z with clusters)	Path Length	5000	500	5	.000001	1000
	Disc. Reward	5000	5000	.1	500	1000
7 (V with clusters)	Path Length	50000	500	1	.000001	1000
	Disc. Reward	25000	500	.001	250	1000

Figure 22 shows representative point-by-point solutions generated by the model using the parameter weights in Table 6. The following sections contain comparisons of those solutions with the human solutions presented in Figure 2.

4.2.4.1.1 Problem 2 – Lightbulb

Referring to Figure 2, human solutions generally clustered as follows. When provided with instructions to optimize for path length, human solutions were identical to model solutions given that optimization criteria and either location decision making strategy (Figure 22, row 1, right two panels). The remainder of these, and roughly half of the solutions generated when optimizing for estimated time to find (Figure 2, Cluster 4), contain shoelace-style solutions and are very similar to the model's solution optimizing for estimated time to find and using the leash local decision algorithm. Finally, Clusters 1 and 3 (Figure 2) appear to be similar forms of the solution produced by the model when using the nearest neighbor algorithm and optimizing for estimated time to find.

4.2.4.1.2 Problem 4 – Z

Human solutions to the Z-shaped problem fell into three clusters, with nearly all of the path length optimization solutions, and roughly half of the estimated time to find solutions (Figure 2, Row 2, Cluster 3), were identical to the model's solution (Figure 22, row 2) when optimizing for path length and using the leash algorithm. While a few humans produced the optimal solution (Cluster 3), the majority of solutions prioritizing estimated time to find fell into Cluster 1, which contained solutions that appeared to heavily emphasize distance to first contact, but varied greatly in local decisions. This problem, and especially the similar Problem 5 (Z with clusters) which was very similar, produced a great deal of variance in the model as well (see Figure 21, noting the stark

differences between example plans generated by the model). Model solutions, when optimizing for estimated time to find, similarly weighted distance to first contact, and did not vary between the local decision making strategies due to the model's ability to account for the problem space almost entirely at a higher level (i.e., the linear nature of the problem allowed the model to obtain very high goodness of fit with little effort).

4.2.4.1.3 Problem 5 – Z with Clusters

Problem 5 is visually similar to Problem 4, only with clusters appended at the vertices and ends of line segments. Human solutions (see Figure 2, Row 3) to this problem varied wildly. When optimizing for path length, most participants solved the problem as a Z (Cluster 1) with variance in the local solutions. Optimizing for estimated time to find, participant solutions fell into Clusters 2 and 4, which were largely similar but varied in terms of where, on the middle segment of the Z, participants decided to move to the cluster of points at the tail of the Z. Cluster 3 contained roughly equal numbers of solutions from both instruction sets. Model solutions (Figure 22, Row 3), given parameter weights emphasizing estimated time to find, were similar to the human solutions with the exception that the model prioritized points in the top segment of the Z more so than people. That is, the model did not typically generate plans that involved hitting those points at the end of the solution. MCM solutions optimizing for path length followed similar solutions to those generated by humans with the exception that they overemphasized distance to first contact, hitting points on the top-right cluster en route to the left side of the top segment of the Z.

4.2.4.1.4 Problem 7 – V with Clusters

Participant solutions to Problem 7 (Figure 2, Row 4), provided with instructions to minimize path length, generally involved following the V around left or right with some local differences (Clusters 1 and 2, respectively), while Cluster 3 contained the majority of the estimated time to find optimization solutions, and Cluster 4 roughly equally represented both instruction styles. Model solutions to this problem (Figure 22, Row 4) were similar to the human solutions, either following the V when provided with parameter values to minimize path length, or prioritizing clusters when provided parameters optimizing for estimated time to find. Most notably, as in Problem 4 (the Z), the plans fit the problems well due to their linear nature, with the majority of the variance between runs accounted for by local decision making. Interestingly, the shoelace solutions in the human data have analogues produced by the model that can be accounted for by local decision making strategies (i.e., using the leash algorithm) applied to the same higher level plan (see Figure 22, Row 4, left panels).

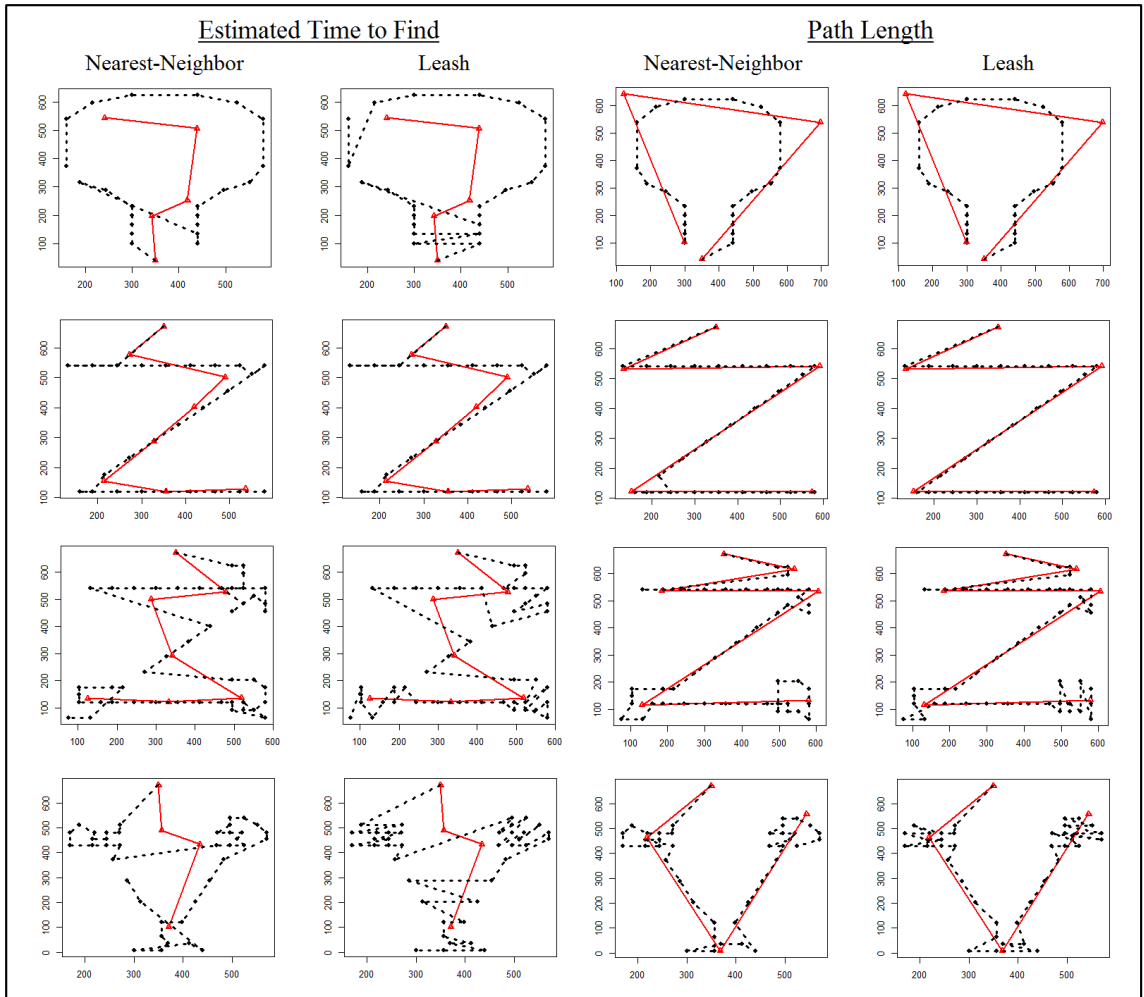


Figure 22. Representative MCM solutions for Problems 2, 4, 5, and 7 (see Figure 2 for human solutions) given parameter weights (see Table 6) designed to encourage solutions optimizing either estimated time to find (left two panels) or path length (right two panels). Within each panel, local decisions are made using either a nearest-neighbor (left) or leash (right) algorithm. Specific solutions, and their analogues in the human data, are presented in the text.

4.2.4.2 Simulations & Model Fitness

To permit quantitative comparisons of MCM and human performance, 250 model simulations were run on each of the four sample problems above, given each of the two optimization criteria described in Table 6 (above), yielding 500 plans to solving each

problem. To each of these 500 plans, local solutions were computed using both nearest-neighbor and leash algorithms, for a total of 1000 local solutions to each problem.

Efficiency, calculated as solution path length / optimal solution path length, and time-discounted efficiency, provide performance metrics against which to compare the human solutions to these problems. Path length efficiencies associated with each instruction set are available in Table 7, while estimated time to find efficiency is shown in Table 8.

Successful selection of optimization criteria can be defined by the agent producing results consistent with the optimization criteria or instructions with which it is provided.

Therefore, when provided with optimization criteria optimizing for path length, the path lengths of the agents' trajectories should be shorter than when the agent is optimizing for estimated time to find, and vice versa.

Table 7. Means and standard deviations of efficiency (Path Length) for humans and the model, given the two experimental optimization criteria (model) or instructions (human) associated with each problem described in Table 6, and the two local decision making strategies described in the text.

Problem	Path Length Optimization			Estimated Time to Find Optimization		
	Model		Humans	Model		Humans
	Nearest Neighbor	Leash		Nearest Neighbor	Leash	
2	1.10 (0.13)	1.21 (0.38)	1.05 (0.11)	1.21 (0.08)	1.56 (0.20)	1.18 (0.19)
4	1.37 (0.27)	1.23 (0.18)	1.01 (0.04)	1.19 (0.13)	1.23 (0.16)	1.12 (0.13)
5	1.45 (0.23)	1.29 (0.14)	1.03 (0.05)	1.26 (0.09)	1.28 (0.07)	1.12 (0.13)
7	1.24 (0.12)	1.36 (0.04)	1.11 (0.13)	1.34 (0.05)	1.96 (0.63)	1.21 (0.13)

Comparing agent performance in path length, given instructions to optimize either path length (left cells) or estimated time to find (right cells), humans were successful in optimizing for path length judging by the greater efficiency when provided with those

instructions. In Problems 2 and 7, the model was also successful in doing so. MCM solutions, when optimizing for estimated time to find, were superior in terms of path length when compared to those optimizing for path length in Problems 4 and 5, however it is important to note that the model had a great deal of difficulty with these problems compared with humans, who exhibited near perfect efficiency when provided with instructions to do so. Referring to the qualitative data (see Figures 20 and 21), the model was capable of settling into the optimal solutions given each instruction set, and those optimal solutions do produce superior cost functions, however the model's inability to achieve those solutions consistently over many simulations indicates that the current optimization process (see Section 4.2.2) is not ideal.

Table 8. Means and standard deviations of efficiency (Estimated Time to Find) for humans and the model, given the two experimental optimization criteria (model) or instructions (human) associated with each problem described in Table 6, and the two local decision making strategies described in the text

Problem	Path Length Optimization			Estimated Time to Find Optimization		
	Model		Humans	Model		Humans
	Nearest Neighbor	Leash		Nearest Neighbor	Leash	
2	1.14 (0.12)	1.29 (0.32)	1.27 (0.19)	1.14 (0.10)	1.74 (0.24)	1.31 (0.36)
4	2.31 (0.65)	2.01 (0.35)	1.00 (0.01)	2.91 (0.73)	13.15 (0.89)	1.15 (0.13)
5	1.42 (0.24)	1.36 (0.22)	1.02 (0.04)	1.26 (0.11)	1.30 (0.08)	1.07 (0.09)
7	1.55 (0.12)	1.70 (0.11)	1.29 (0.26)	1.45 (0.11)	2.39 (1.05)	1.25 (0.23)

The results shown in Table 8 indicate that, only on Problem 7, were humans, in aggregate, able to optimize for estimated time to find. The model was able to produce better optimization for the appropriate criteria in Problems 5 and 7, with no meaningful

difference seen in Problem 2, and the reverse effect seen in Problem 4 (which proved exceptionally difficult for the model). Qualitative analysis of the human solutions (see Figure 2) shows that their failures are often the product of misapplied strategies at the local decision making level. For example, in Problem 1, while some participants were able to obtain the optimal solution, many participants' solutions exhibited a highly suboptimal shoelace strategy, which is also produced by the present model using the leash local decision making algorithm. This same is true for the model, where the suboptimal strategy (i.e., the leash local decision making algorithm used by the model) is responsible for producing the poor solutions in Problem 7, for example (see Table 8, comparing poor performance on the leash algorithm with a sample solution available in Figure 22, Row 4, Column 2).

4.2.5 MCM Discussion

The results described above should be viewed with the following caveats. First, the optimization algorithm used in the present study is not intended to directly model any process theoretically used by humans during problem solving. That is, I do not believe that humans are generating a candidate plan based on statistical clustering, then deforming it iteratively. Second, the local decision making strategies employed by the model are computationally inexpensive and, while they do sometimes produce the types of trajectories seen in the human data, they are not intended to capture the full range of variance in human problem solving. Finally, the problems presented here were selected because they were relatively easy for humans to optimize, and comparatively

exceptionally difficult for models to optimize. Two of the presented problems (Problems 4 and 5) proved difficult for the model, while two others (Problems 2 and 7) proved comparatively simple.

Regarding the optimization algorithm, juxtaposing the results of the qualitative and quantitative simulations, it is clear that the optimization algorithm is sufficient for solving Problems 2 and 7, but was insufficient for solving Problem 4 and 5 (Z). That is not to say that the algorithm never found the correct solution in the more difficult problems, only that these problems proved particularly difficult for the model given its current optimization algorithm described in Section 4.2.2, and the optimization algorithm was unable to settle into the optimal solution in the majority of trials. This is unsurprising as the current optimization algorithm uses random rather than strategic evolution. Despite this, the cost functions associated with the superior solutions produce the appropriate values, indicating that given an improved optimization algorithm the model would more reliably achieve those solutions. Nevertheless, the purpose of this present modeling effort is not to design the best optimization algorithm, but rather to demonstrate that such a model is capable of producing solutions that err similarly to humans. To that end, the model achieved that goal as it, like humans, produces a wider range of solutions to the more difficult problems.

Based upon clustering of actual human solutions (see Figure 2), it appears that people approach these problems with different types of mental models that reflect, in some cases, differences in local decision making strategies. Nearest-neighbor and the leash algorithm are nowhere near exhaustive and are not intended to holistically model

human decision making. Rather, they represent two candidate strategies that humans might use. For this purpose, the model was capable of producing similar solutions to humans given each of these strategies (see Figure 22), even when the approach was globally optimal. For example, for Problem 7, when optimizing for estimated time to find, the model produced two solutions commonly seen in the human data (see Figure 2) despite the fact that the shoelace-style approach is suboptimal. It is, however, important to note that the model's leash algorithm should not be considered universally suboptimal and sometimes it produced better efficiency in select circumstances (e.g., Table 7, Problems 4 and 5, Path Length Optimization instructions).

The results of these simulations indicate a number of important conclusions about the mixed criteria approach. First, a two-layer model that incorporates strategic planning based upon optimization criteria, at a higher level, and comparatively simple algorithms for local decision making, is sufficient to produce both human-like successes, and failures (see Figure 22). Second, while the MCM is capable of optimizing for specific goals, perhaps more importantly, it is capable of optimizing for multiple goals simultaneously, and in varying proportions. Therefore, we might expect that, given traditional TSP problems, the model would be capable of producing the types of effects seen in Chapter 2 (i.e., bias toward irrelevant optimization criteria), whereas other TSP models would not. Third, while the local decision making algorithms used in the present model do not fully describe the range of human strategies, they can account for some of the variance in approaches to local problem solving, given the same global plan, observed in the human data. A model intending to capture the full range of human local decision making

strategies should draw upon heuristics such as avoiding overlaps, which are especially present in the solutions created using the nearest neighbor algorithm for local decision making, and are the basis of algorithms such as 2-Opt (Croes, 1958).

One potential criticism of the modeling effort as presented is that there is no comparison between this model and other TSP models. I chose not to compare the present model against TSP models for two reasons. First, existing TSP models would be incapable of generating solutions that optimized for more than one criteria at a time, and generally do not optimize for any criterion other than path length. For these reasons, I chose to compare model and human performance against global optimal performance in each of the problems. Second, I wanted to avoid using the traditional (random) TSP problems to which most TSP models are calibrated, and instead use problems that are known to be exceptionally difficult for models to solve (Perelman & Mueller, 2014). The problems used in the present modeling effort generally encouraged one of two types of global solutions, and so while the costs associated with each of those two solutions were very similar in some problems, the differences resided in global aspects of the problems, not local aspects, such as a long tail to the last point in the problem space as seen in the analysis in Chapter 2.

Broadly, the MCM is consistent with existing theories of global versus fine grained encoding and retrieval. However, specifically as it pertains to human search and navigation, the model represents a mathematical implementation of a mixed criteria theory of performance maximization. A single cognitive system, capable of incorporating multiple optimization criteria in service of similar goals (i.e., search, navigation, routing,

planning, and other related spatial behaviors), is computationally convenient and perhaps adaptively and evolutionarily plausible. Future research could focus on constraining the space to which this theory applies. This research would provide evidence for or against this theory by testing for the presence of the invariants described in Chapter 2 in human solutions to problems with validity in naturalistic tasks. For example, will irrelevant optimization criteria intrude on solutions when subjects are presented with non-visual versions of TSP (i.e., a laboratory task with validity in nonvisual routing tasks pertinent to operations research)? The search for these invariants would further define boundaries for applying the model and theory.

Chapter 5: Using the Model – Recommendations for Implementation

WiSAR provides a naturalistic task against which to test a naturalistic computational model's performance. One main reason for selecting this specific task, however, is that the search and navigation behaviors are common to human behavior in other fields and the optimization criterion (i.e., estimated time to find) is well specified. This chapter contains design recommendations for adapting the aforementioned naturalistic computational model to other domains, including emergency egress, military operations such as cordon and search, and urban design.

While these design recommendations are by no means intended to be exhaustive instructions, they provide guidance for future research and modeling efforts. In order to approximate human behavior in these domains, I believe that the first step is integrating the planning mechanism described in Section 4.2 with a hippocampus-based network

model such as that used by Mueller et al. (2013). This discussion is followed by a proposal for an empirical study in which the present planning model is used to test human factors effects in pilots and command personnel.

5.1 Integrating Planning and Navigation

As outlined in previous chapters, current models of navigation are insufficient to account for human behavior for a number of reasons. First, many attempt to describe rodent behavior in great detail, but behaviorally speaking are capable of making only local decisions. Second, most modeling approaches attempting to predict human behavior, such as those seen in emergency egress, are predominantly flow (i.e., physics) models that are acceptable for predicting the behavior of multiple agents who are all familiar with the environment, but fail to account for agent familiarity, or lack thereof, with the environment. In these models, agents are assumed to obediently follow signage. One consequence of this approach is that it assumes that all exits will be equally viable, which in emergency egress is not necessarily the case. For example, exits may be blocked by debris or fire. Third, in 3rd person routing tasks with a deep structure very similar to combinatorial optimization problems (i.e., TSP), such as plotting search routes through probability spaces in military intelligence gathering operations (SME 3, November 18, 2014) and WiSAR (SME 1, May 15, 2014; SME 4, December 14, 2014), existing modeling approaches optimize for only one criteria at a time, where the research discussed in Chapter 2 suggests that humans use multiple optimization criteria simultaneously.

To address these shortcomings, it is necessary to merge existing network models of navigation, such as those used by Mueller et al. (2013), with the MCM described in Section 4.2 of this dissertation. In their current form, the models used by Mueller et al. (2013) made navigation decisions based upon hard-coded target locations alone, analogous to cities in a TSP, and local decisions were handled using spreading activation. In the proposed integrated model, the MCM's linear planning mechanism would provide a higher-level plan for navigating through the problem space based upon the task-relevant optimization criteria (e.g., estimated time to find, path length, and other optimization criteria) while the network model provides both local decision making, through spreading activation, and experiential learning mechanisms to constrain the problem space based upon what the agent knows about the environment. This is especially important for modeling first person navigation, such as that seen in emergency egress, as individual agents may not be sufficiently familiar with the environment to follow the optimal route through the problem space.

Implementing this change would be rather simple, and consist of (1) running the linear planning algorithm over the problem space as “understood” by the network model, (2) mapping target locations to the resulting plan segments, and (3) allowing the network model to use spreading activation to solve for each segment in sequence. Note that spreading activation is functionally very similar to the nearest neighbor mechanic (Perelman & Mueller, 2013b) implemented in Section 4.2 of this dissertation.

5.2 Integrating Planning and Memory

The utility of neural network models lies not only in target location encoding and retrieval, but in context encoding and retrieval. While targets are used specifically to investigate participants' ability to bind items with contexts in the environment, targets are merely environmental features. The probability of an agent encoding or retrieving a target is functionally equivalent to the probability of an agent encoding or retrieving any salient feature in an environment that could define a context. Therefore, a cognitive agent capable of both spatial navigation using plans, and target memory encoding and retrieval which can be perturbed based upon characteristics of the environment, will reasonably predict human behavior in any search task in which the operator must explore an environment and retain a memory for where he has been, and what he found in those locations. The most straightforward approach to modeling memory in search is to integrate probabilistic encoding and retrieval within a neural network model.

In such a model, the agent would probabilistically encode target locations (i.e., associate a specific target with a specific location) with a probability that scales with the saliency of the location, serial position of the target (derived from Studies 2 and 3 of this dissertation), set size (i.e., the number of items to be remembered), and other factors. This same saliency weight could be used to compute the agent's encoding of location transitions. That is to say that one saliency weight would impact both target-location binding (i.e., CA1-CA3 connections) and location-location binding (CA3-CA3 connections), making the model useful for both environment modeling (e.g., egress

modeling) and target memory modeling (e.g., WiSAR). Finally, a decay parameter can be used to modify the saliency weight to produce recency effects during retrieval.

When asked to navigate among points in an environment, the agent would generate a plan using the mechanisms described in Section 4.2 of this dissertation applied to *known* locations in the environment. That is, the neural network model (used for local decision-making and target-location and location-location binding) provides the constrained problem space based upon what the agent knows, and the MCM generates a plan for navigating through that space given optimization criteria. Higher level plans can easily be constrained to known CA3-CA3 connections between which travel is possible, allowing for learning-based navigation and, from a programming perspective, easy integration of barriers.

When asked to recall a specific target's location (i.e., CA1 location), the agent will try to retrieve the memory representation for that target. The results of Studies 2 and 3 (Chapter 3) indicate appropriate probability functions for encoding and retrieval, which indicate that the agent's memory should (1) exhibit primacy and recency effects, at least in applications modeling first person navigation, (2) exhibit decreased probability of recall with serial position from the actual target location, (3) exhibit decreased probability of recall given increasing spatial distance from the actual target location, and (4) exhibit a linear quality whereby errors *along* the route of travel are more likely than errors perpendicular to the route of travel.

To achieve these effects, I make the following recommendations. First, weight the probability of encoding by the saliency rating of a target's location (i.e., target-context

binding), making targets found in highly salient locations more easily remembered. Second, addressing point 1 above, account for primacy and recency effects by applying, in some combination, working memory load and decay functions, respectively. Third, targets should be confused with other targets along the route with a frequency that decreases with increasing serial distance from the actual target location (point 2 above). This can be accomplished by representing the targets ordinally according to serial position, and introducing this error as possible transpositions during recall according to the probability distributions in Studies 2 and 3. Fourth, memory for the precise locations of targets should be perturbed during recall by applying *skewed* Gaussian noise to the target's location according to those same probability functions (points 3 and 4, above). That is, the probability of recall should be skewed elliptically along the agent's route of travel.

5.3 Potential Domains of Application

Though this dissertation has focused to some extent on WiSAR, there are other domains in which an integrated model, capable of learning and planning, would be necessary. One target domain is emergency egress. As previously mentioned, most emergency egress models are based upon the flow of bodies through a space, which may depend to some extent on environmental landmarks and signage. However, such models do not adequately describe effects of familiarity with a space. For example, structures like hospitals contain both staff, who may be intimately familiar with the building layout, and patients and their family members, who rarely frequent the space. Traditional egress

models would assume that both populations would egress from the building with equal ease or difficulty. A computational model capable of experientially learning an environment would provide a more realistic account of egress. Another domain is modeling human behavior in military and law enforcement environments, to permit predictive analysis. Finally, a model capable of planning based upon known features of an environment would have utility in urban design.

5.3.1 Emergency Egress

As previously mentioned, many emergency egress models account for group dynamics by focusing on the *flow* of typically unintelligent agents through an environment (e.g., Klüpfel, 2003). While these models are useful for identifying choke points in the movement of large crowds, they fail to account for the individual of behaviors based on their strategic goals and past experiences. Two examples highlight this problem. First, egress models typically assume that all agents move directly to the exits with the sole goal of escape, whereas actual accounts of human behavior in fire situations sometimes involve parents running back into burning buildings to find children (e.g., Kapp, Venugopal, & Kern-Jedrychowska, 2013), or searching within the building before evacuating. Second, the fire at The Station night club in Rhode Island, which took the lives of 100 people, highlights the importance of recency traces in modeling egress behavior. When the fire started, most people attempted egress through the front door through which they had entered the venue, despite an available exit on the West side of the building (Edward, 2007).

The proposed integrated (i.e., neural network and MCM) model is capable of addressing both of these considerations. First, agent behavior is controlled by strategic considerations that are incorporable into a plan. Agents with alternative goals (e.g., searching for children) can execute those search plans before attempting to egress. These types of goals can be accounted for by incorporating hard points into the plan optimization process. Modeling this behavior might identify new choke points, as such goals may be shared by a large proportion of the agents present in the simulation (e.g., an auditorium fire with a large number of parent-child dyads or triads in attendance). Incorporating this functionality into egress models involves translating the search plan generation and execution algorithms described above into the agent AI. Egress simulation software already containing AI functionality (e.g., Virtual Egress and Analysis System [VEgAS] and Legion) should provide a suitable base for implementation. The end result is agent AI that permits strategic plan generation and execution rather than simple “follow the exit signs” dynamics to dictate movement trajectories and determine choke points.

Second, a computational model capable of modeling environment encoding and retrieval, and by extension recency traces in movement trajectories, can account for the return-via-entrance behavior described above. The spreading activation model used by Mueller et al. (2013) and Perelman & Mueller (2013b) suppressed activation in previously visited locations. Therefore, return-via-entrance behavior can be replicated simply by using the inverse of the default hill-climbing mechanism. By exhibiting a preference for the areas of *least* activation (i.e., the recently suppressed locations), the

agent will prefer to return via its recency trace. Additional model parameters using the two-layer (i.e., a higher level plan driving lower level execution) nature of the proscribed computational model can also be used to induce this behavior, such as a preference for least-resistance linear movement (i.e., rather than following along its recency trace, the model draws an azimuth cue from its recency trace at the coarser grain size then backtracks in that direction). This behavior, and the behavior described above, should provide a more accurate account of human behavior during emergency egress.

5.3.2 Military and Law Enforcement Operations

Cordon and search often involves systematic search techniques shared among all branches of the military (FM 3-06-20) and perhaps extending to law enforcement communities. Therefore, the high level planning involved in these operations is the product of established tactics, techniques, and procedures (TTPs), and has already been modeled by Mueller et al. (2013) using experimenter-coded high level plans. However, the present model retains utility in modeling memory for areas searched and the items found in those locations during cordon and search operations. This information could be used to refine established TTPs.

Not all military operations are constrained by such well-defined TTPs, especially Platoon and Company-level actions. For example, after patrols, army units often provide debriefings and discuss their findings with their COIST (see Study 1). The recent emphasis on COISTs is part of a new “recon pull” approach to intelligence, which takes a grass roots approach to intelligence as compared with the typical “command push”

approach (Poole, 2003). In Platoon and Company-level operations, plans are often kept sufficiently coarse-grained so as to be flexible and easily followed, in order to allow all elements to remember their role in the plan, but also to permit individuals to exercise initiative where appropriate (SME 6, March 10, 2015). One potential application for the MCM, given its two-layered structure, is modeling individual differences in local decision making (i.e., plan execution) given a single higher level plan to allow commanders to account for a broader range of potential element behaviors in operational environments. That is, given a single higher level plan, what are the ways in which subordinates may execute the plan, and are all of these possible solutions acceptable? Or would some of these local solutions be unacceptable, for example, because they increase the probability of fratricide or increase exposure unacceptably?

In addition to promoting initiative, new doctrine declares that “every soldier is a sensor” (FM 3-21.75) tasked with sensing, remembering, and subsequently reporting information to the COIST during patrol debriefings. Therefore, it is paramount that soldiers remember what they saw, and where they saw it, during the patrol. By characterizing agent memory representations during spatial navigation, the present effort provides both an empirical account and model for understanding error during recall in post-patrol debriefings. Some information is offloaded from memory during the patrol via radio (see Study 1), but smaller pieces of intelligence noticed by individual members of the patrol may not be sufficiently significant to report to a commanding officer; these would be discussed in the post-patrol debriefing (SME 5, March 7, 2015). While accounts gathered during the debriefing process are usually taken at face value, understanding and

modeling the memory representations would provide a probabilistic characterization of event locations along the patrol route, to help a COIST better understand the accuracy of post-patrol intelligence, or in simulations to design better TTPs for gathering this information.

5.3.3 Urban Design

Certain characteristics of the integrated model offer advantages over existing approaches to modeling wayfinding through urban environments during the design process. Specifically, the proposed integrated model makes predictions about pedestrian flow through environments on the basis of cue saliency and prior exploration, and also uses cue saliency to determine the probability of the agent encoding particular sections of the environment.

Because the proposed integrated model is capable of weighting associations among locations in the cognitive map by the saliency of the locations (i.e., contextual cues in environmental locations aid encoding and retrieval of those locations), the model could be used to evaluate designs for their “learnability.” Designs containing more salient environmental cues should be easier for agents to navigate, as they provide more context cues for encoding and retrieval, and the proposed model should approximate these performance differences.

Egress models such as Legion are sometimes used to determine flow through urban spaces, as facilitating pedestrian flow is often an important design consideration. Since the proposed model makes specific predictions about agents’ preferences for highly salient

locations, it may provide a useful perspective for determining flow between pedestrians' vehicles in parking lots, and their destinations in large urban centers, such as shopping malls.

5.4 Trust in Automated Plan Generation – a Proposed Framework for Application

The WiSAR and geospatial intelligence SMEs interviewed in Study 1 of this dissertation revealed that they typically use probability maps generated using software, and that the software also provides a route that is mathematically optimized according to a number of criteria. Despite this, incident (in WiSAR) and UAV commanders (in military and intelligence operations) often manually generate solutions to the probability map even though they are provided with presumably mathematically superior solutions by the software (according to the software's criteria; SME 3, November 18, 2014). The literature on heuristic judgment and decision-making is rich with examples of expert human decision-makers succeeding in complex naturalistic environments (Keller, Cokely, Katsikopoulos, & Wegwarth, 2010). Therefore, we should not assume that ignoring the automatically-generated routes (or routes generated by other agents) is the *incorrect* course of action. Instead, this research problem warrants empirical investigation.

Are the commanders who plan routes in search operations creating superior routes to those generated by the software? Or are these commanders simply biased against using computer-generated solutions because, despite their mathematical superiority, those solutions bear no resemblance to the solutions that the commander would generate? The

evaluation algorithm presented in this dissertation provides a means of evaluating solutions based on criteria by which humans generate plans. And, the present MCM may therefore provide a superior means of generating plans, as these plans may be more likely to fit commanders' mental models for search than those generated via a software optimization algorithm. In this section, I propose two experiments designed to answer these empirical questions.

5.4.1 Automatic Route Generation

Automatic route generation has received a great deal of attention recently due to a push toward increasing UAV autonomy, and reducing operator burden in UAV control. Domain specific software focuses on optimizing flight paths for considerations like weather (Frew & Brown, 2008), radar and surface to air missile exposure (e.g., Kabamba, Meerkov, & Zeitz, 20006; Sousa et al., 2004), distance from home (to aid in reestablishing communications with the UAV in cases of link loss; Stansbury, Vyas, & Wilson, 2008), probability of detection given a specific sensor package (De Filippis, Guglieri, & Quagliotti, 2012), and path congruence with another agent (for UAVs tasked with providing aerial cover; Ding, Rahmani, & Egerstedt, 2010).

Approaches to generating these flight paths are typically not cognitive, and focus on global optimization. This is important because the pilot's input is considered important even in the one-to-many operator-to-UAV setup that serves as a goal for technology development (Frew & Brown, 2008). Human control is especially important for UAVs carrying weapons payloads, as policy makers are understandably reluctant to

entrust shoot-no shoot decisions to automation (Mills, 2007). Therefore, understanding the extent to which pilots trust automation in flight planning and execution is beneficial to the mission of the United States Air Force and other end users of UAV technology, such as WiSAR operators and the intelligence community.

The mechanisms by which flight plans are generated are specific to software used in particular domains. However, the algorithms used bear some similarity to the work presented in this dissertation. For example, in cases where time is critical, such as autonomous UAV target tracking and engagement, researchers in this domain employ heuristic methods such as genetic algorithms to provide acceptable solutions to the combinatorial optimization problems (e.g., Shima & Schumacher, 2009). Therefore, these proposed experiments are furthermore necessitated by the need for computationally feasible solutions to combinatorial optimization problems.

5.4.2 Experiment 1: Measuring Alternative Optimization Criteria in Search Plans

The plan evaluation algorithm presented in this dissertation was used, in conjunction with a stochastic optimization algorithm, to generate higher level plans based upon five optimization criteria that are relevant in human planning, (1) plan simplicity, as measured by log number of segments, (2) goodness of fit, (3) path length, (4) distance-discounted reward, and (5) average angle. If human solutions are generated to minimize the cost of these five criteria, then it is possible that discrepancies between computer- and pilot-generated flight plans are due to differences in optimization criteria. The first experiment in the proposed study will consist of analyzing computer- and pilot-generated

flight plans based upon these five criteria to explore the possibility that pilots' solutions are superior in terms of these criteria. The alternative hypothesis in this experiment is that human solutions, when compared to the software-produced solutions, more heavily weight optimization criteria that the software does not take into account, such as plan complexity. The null hypothesis is that human plans are suboptimal in terms of all of the aforementioned optimization criteria because they are generated by agents with more limited perceptual and cognitive resources. The implication of this outcome is that the computer-generated paths are in fact superior, and the problem therefore lies in convincing pilots to trust the automation.

5.4.3 Experiment 2: Trust in Automation – Path Congruency

One possible reason that pilots do not trust routes provided by software is that the software-produced routes are very dissimilar from what they, other pilots, or even other non-experts, would likely produce. This dissimilarity between a pilot and another agent is referred to herein as congruency. To test the effect of path congruency on pilots' trust in the route, the second proposed experiment involves eliciting pilots' and commanders' qualitative and quantitative feedback on plans generated by a number of different agents, and then asking them to plot their own routes given that problem space. These flight plans would come from three sources: pilot SMEs, traditional domain-specific software that these operators currently use, and the MCM described in this dissertation. Path congruency would be evaluated using a pathmapping approach (Mueller et al., 2015). Pilots' trust ratings would be analyzed by calculating statistics by path source (to which

the pilots would be blind), as well as via regression analysis using the trust ratings against divergence values calculated using pathmapping to compare each participant's path with the other agent-generated paths.

The alternative hypotheses in this proposed study are that pilots' trust and preferences for plans will be highest for plans generated by other pilots, then by the MCM, then by the domain-specific optimization software, and that pilots' trust and preferences will increase with correspondence between their own flight plans, and flight plans generated by other agents. Evidence for the null hypothesis (i.e., no effect of congruence) would indicate that path congruency does not influence pilots' willingness to accept plans provided to them, and perhaps that human pilots' mistrust in automation is not due to some difference in models with which the agents approach the problem space, but rather due to biases against the automated agent itself. These possible outcomes have implications for human factors research in operational domains.

5.4.4 Expected Implications and Discussion

The goals of these proposed experiments are to (1) evaluate human pilots' considerations (i.e., optimization criteria) in planning flight routes as compared with factors considered by domain-specific software, and (2) determine the extent to which correspondence between human- and domain-specific software-generated flight plans impacts pilots' trust in automation for that agent. The implications of the alternative hypotheses described above are that (1) humans may be integrating information into flight plans that domain-specific software is not, and that information may be useful in

driving the development of automation software in the future, and (2) if trust in automation is based upon path congruency, then domain-specific software should incorporate algorithms that produce flight paths similar to what pilots would produce in order to increase pilots' trust in the software and willingness to use the routes, and the model presented in Section 4.2 of this dissertation may be capable of producing that functionality.

References

- Aguirre, G. K., Detre, J. A., Alsop, D. C., & D'Esposito, M. (1996). The parahippocampus subserves topographical learning in man. *Cerebral Cortex*, *6*, 823-829.
- Akdogan, I., Unal, N., & Adiguzel, E. (2002). Estimation of the number of neurons in the hippocampus of rats with penicillin induced epilepsy. *Image Anal. Stereol.*, *21*, 117-120.
- Allen, G. L. (1981). A developmental perspective on the effects of "subdividing" macrospatial experience. *Journal of Experimental Psychology: Human Learning and Memory*, *7*, 120-132.
- Ascoli, G. A., & Samsonovich, A. V. (2013). A spiking-network cognitive architecture inspired by the hippocampus. *Biologically Inspired Cognitive Architectures*, *3*, 13-26.
- Baddeley, A. D., Eysenck, M. W., & Anderson, M. C. (2009). *Memory*. New York: Psychology Press.
- Barry, C., Lever, C., Hayman, R., Hartley, T., Burton, S., O'Keefe, J., Jeffrey, K., & Burgess, N. (2006). The boundary vector cell model of place cell firing and spatial memory. *Rev Neurosci*, *17*, 71-97.
- Basso, D., Bisiacchi, P. S., Cotelli, M., & Farinello, C. (2001). Planning times during traveling salesman's problem: Differences between closed head injury and normal subjects. *Brain and Cognition*, *46*, 38-42.

- Best, B. J. (2005). A model of fast human performance on a computationally hard problem. In *Proceedings of the 27th Annual Conference of the Cognitive Science Society*, 256-262.
- Blake, R. & Sekuler, R. (2006). *Perception (5th Edition)*. New York: McGraw-Hill.
- Blum, A., Chawla, S., Karger, D. R., Lane, T., Meyerson, A., & Minkoff, M. (2007). Approximation algorithms for orienteering and discounted-reward TSP. *SIAM Journal on Computing*, 37, 653-670.
- Brown, J. (1958). Some tests of the decay theory of immediate memory. *Quarterly Journal of Experimental Psychology*, 10, 12–21.
- Burgess, N. & O’Keefe, J. (1996). Neuronal computations underlying the firing of place cells and their role in navigation. *Hippocampus*, 7, 749-762.
- Burgess, N., Recce, M., & O’Keefe, J. (1994). A model of hippocampal function. *Neural Networks*, 7, 1065-1081.
- Caduff, D. & Timpf, S. (2008). On the assessment of landmark salience for human navigation. *Cognitive Processing*, 9, 249-267.
- Cannon, R., Hasselmo, M., & Koene, R. (2002). From biophysics to behavior: Catacomb2 and the design of biologically plausible models for spatial navigation. *Neuroinformatics*, 1, 3-42.
- Clark, R. E., Feldon, D., Van Merriënboer, J., Yates, K., & Early, S. (2008). Cognitive task analysis. *Handbook of research on educational communications and technology*, 3, 577-593.

- Cokely, E. T. (2009) Beyond generic dual processes: How should we evaluate scientific progress? *PsycCritiques*, 54.
- Cokely, E. T., Galesic, M., Schulz, E., Ghazal, S., & Garcia-Retamero, R. (2012). Measuring risk literacy: The Berlin Numeracy Test. *Judgment and Decision Making*, 7, 25-47.
- Cokely, E. T. & Kelley, C. M. (2009). Cognitive abilities and superior decision making under risk: A protocol analysis and process model evaluation. *Judgment and Decision Making*, 1, 20-33.
- Cokely, E. T., Schooler, L. J., & Gigerenzer, G. (2009). Information use for decision making. In M. N. Maack & M. J. Bates (Eds.), *Encyclopedia of Library and Information Sciences*.
- Crandall, B., Klein, G., & Hoffman, R. R. (2006). *Working Minds: A Practitioner's Guide to Cognitive Task Analysis*. Cambridge MA: Bradford Books.
- Croes, G. A. (1958). A method for solving traveling salesman problems. *Operations Research*, 6, 791-812.
- Cutini, S., Di Ferdinando, A., Basso, D., Bisiacchi, P. S., & Zorzi, M. (2008). Visuospatial planning in the traveling salesperson problem: A connectionist account of normal and impaired performance. *Cognitive Neuropsychology*, 25, 194-217.
- Dantzig, G. B. & Ramser, J. H. (1959). The truck dispatching problem. *Management Science*, 6, 80-91.

- Debord, G. (1955). Introduction to a critique of urban geography. (K. Knabb, Trans.). *Les Lèvres Nues*, 6.
- De Filippis, L., Guglieri, G., & Quagliotti, F. (2012). Path planning strategies for UAVS in 3D environments. *Journal of Intelligent and Robotic Systems*, 65, 247-264.
- De Jong, L. W., Gereke, B., Martin, G. M., & Fellous, J. (2011). The traveling salesrat: Insights into the dynamics of efficient spatial navigation in the rodent. *Journal of Neural Engineering*, 8. doi:10.1088/1741-2560/8/6/065010.
- Dijkstra, E. W. (1959). A Note on two problems in connexion with graphs. *Numerische Mathematik*, 1, 269–271.
- Ding, X. C., Rahmani, A. R., & Egerstedt, M. (2010). Multi-UAV convoy protection: An optimal approach to path planning and coordination. *IEEE Transactions on Robotics*, 26, 256-268.
- Eccles, D. W., Walsh, S. E., & Ingledew, D. K. (2002). The use of heuristics during route planning by expert and novice orienteers. *Journal of Sports Sciences*, 20, 327-337.
- Edward, P. P. (2007). *Tally of a tragedy: 462 were in The Station on night of fire*. The Providence Journal.
- Ekstrom, A. D., Kahana, M. J., Caplan, J. B., Fields, T. A., Isham, E. A., Newman, E. L., & Fried, I. (2003). Cellular networks underlying human spatial navigation. *Nature*, 425, 184-188.
- Emmerich, M., & Deutz, A. (2006). *Multicriteria Optimization and Decision Making*. LIACS. Leiden university, NL.

- Eriksen, C. W. & St. James, J. D. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception & Psychophysics*, 40, 225-240.
- Evers, L., Dollevoet, T., Barros, A. I., & Monsuur, H. (2012). Robust UAV mission planning. *Annals of Operations Research*, 222, 293-315.
- Frew, E. W. & Brown, T. X. (2008). Networking issues for small unmanned aircraft systems. *Journal of Intelligent and Robotic Systems*, 54, 21-37.
- Gaisbauer, C. & Frank, A. U. (2008). Wayfinding model for pedestrian navigation. In *Proceedings of the 11th AGILE International Conference on Geographic Information Science*, Spain: University of Girona.
- Godden, D. R. & Baddeley, A. D. (1975). Context-dependent memory in two natural environments: On land and underwater. *British Journal of Psychology*, 66, 325-331.
- Golledge, R. G., Jacobson, R. D., Kitchin, R., & Blades, M. (2000). Cognitive maps, spatial abilities, and human wayfinding. *Geographical Review of Japan*, 73, 93-104.
- Gorchetchnikov, A. & Hasselmo, M. E. (2002). A model of hippocampal circuitry mediating goal-driven navigation in a familiar environment. *Neurocomputing*, 423-427.
- Graham, S. M., Joshi, A., & Pizlo Z. (2000). The Traveling Salesman Problem: A hierarchical model. *Memory & Cognition*, 28, 1191-1204.
- Gutjahr, S., Henze, N., & Folkers, M. (1999). Shortcomings of generalized affine invariant skewness measures. *J. Multivar. Anal.*, 71, 1-23.

- Hafting, T., Fyhn, M., Molden, S., Moser, M.B., & Moser, E.I. (2005). Microstructure of a spatial map in the entorhinal cortex. *Nature*, *436*, 801-806.
- Hajibabai, L., Delavar, M. R., Malek, M. R., & Frank, A. U. (2007). Agent-based simulation of spatial cognition and wayfinding in building fire emergency evacuation. In *Geomatics Solutions for Disaster Management*, pp. 255-270, Berlin: Springer Verlag.
- Hargreaves, E. L., Rao, G., Lee, I., & Knierim, J. J. (2005). Major dissociation between medial and lateral entorhinal input to dorsal hippocampus. *Science*, *308*, 1792-1794.
- Hart, P. E., Nilsson, N. J., & Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, *4*, 100-107.
- Hill, A. J. & Best, P. J. (1980). Effects of deafness and blindness on the spatial correlates of hippocampal unit activity in the rat. *Exp. Neurol*, *74*, 204-217.
- Hills, T. T., Todd, P. M., & Goldstone, R. L. (2008). Search in external and internal spaces: Evidence for generalized cognitive search processes. *Psychological Science*, *19*, 802-808.
- Hirtle, S. C. & Jonides, J. (1985). Evidence of hierarchies in cognitive maps. *Memory & Cognition*, *13*, 208-217.
- Hirtle, S. C. & Mascolo, M. F. (1986). Effect of semantic clustering on the memory of spatial locations. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *12*, 182-189.

- Hochmair, H. & Frank, A. U. (2000). Influence of estimation errors on wayfinding-decisions in unknown street networks – analyzing the least-angle strategy. *Spatial Cognition and Computation*, 2, 283-313.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79, 2554-2558.
- Hurlstone, M. J., Hitch, G. J., & Baddeley, A. D. (2014). Memory for serial order across domains: An overview of the literature and directions for future research. *Psychological Bulletin*, 140, 339-373.
- Huttenlocher, J., Hedges, L. V., & Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review*, 98, 352-376.
- Huttenlocher, J., Newcombe, N., & Sandberg, E. H. (1994). The coding of spatial location in young children. *Cognitive Psychology*, 27, 115-147.
- Jensen, O., & Lisman, J. E. (1996). Hippocampal CA3 region predicts memory sequences: accounting for the phase precession of place cells. *Learning & Memory*, 3, 279-287.
- Kabamba, P. T., Meerkov, S. M., & Zeitz, F. H. (2006). Optimal path planning for unmanned combat aerial vehicles to defeat radar tracking. *Journal of Guidance, Control, and Dynamics*, 29, 279-288.
- Kapp, T., Venugopal, N., & Kern-Jedrychowska, E. (2013). Father rushes into burning home to save children in fresh meadows fire. *DNAinfo.com*. Retrieved August 10,

2014 from <http://www.dnainfo.com/new-york/20130401/fresh-meadows/four-critically-hurt-fresh-meadows-fire>.

- Keller, N., Cokely, E. T., Katsikopoulos, K., & Wegwartwh, O. (2010). Naturalistic heuristics for decision making. *Journal of Cognitive Engineering and Decision Making, 4*, 256–274.
- Klüpfel, H. (2003). *A cellular automaton model for crowd movement and egress simulation* (Doctoral dissertation, Universitätsbibliothek Duisburg).
- Knierim, J. J., Lee, I., & Hargreaves, E. L. (2006). Hippocampal place cells: Parallel input streams, subregional processing, and implications for episodic memory. *Hippocampus, 16*, 755-764.
- Kong, X. & Schunn, C. D. (2007). Information seeking in complex problem solving. In *Proceedings of the 8th International Conference on Cognitive Modeling*, Oxford: UK, 261-266.
- Kosslyn, S. M., Pick, H. L., & Fariello, G. R. (1974). Cognitive maps in children and men. *Child Development, 45*, 707-716.
- Kwon, T., Li, Y, Scheessele, M., Michaux, A., & Pizlo, Z. (2014). *Spatially-global interpolation of closed curves*. Presented at the 4th Annual Midwestern Cognitive Science Conference, OH.
- Lavenex, P., & Amaral, D. G. (2000). Hippocampal-neocortical interaction: A hierarchy of associativity. *Hippocampus, 10*, 420-430.

- Levy, W.B. (1989). A computational approach to hippocampal function. In R. D. Hawkins & G. H. Bower (Eds.), *Computational models of learning in simple neural systems* (pp. 243-305). New York, NY: Academic Press.
- Levy, W. B. (1996). A sequence predicting CA3 is a flexible associator that learns and uses context to solve hippocampal-like tasks. *Hippocampus*, 6, 579-590.
- Levy, W. B., Colbert, C. M., & Desmond, N. L. (1990). Elemental adaptive processes of neurons and synapses: a statistical/computational perspective. In M. A. Gluck & D. E. Rumelhart (Eds.), *Neuroscience and connectionist models* (pp. 187-235). Hillsdale, NJ: Lawrence Erlbaum Assoc., Inc.
- Lin, L. & Goodrich, M. A. (2009). A Bayesian approach to modeling lost person behaviors based on terrain features in wilderness search and rescue. In *Proceedings of the 18th Conference on Behavior Representation in Modeling and Simulation*. UT: Sundance.
- Loomis, J.M., Klatzky, R.L., Golledge, R.G., Cicinelli, J.G., Pellegrino, J.W. & Fry, P.A. (1993). Nonvisual navigation by blind and sighted: Assessment of path integration ability. *Journal of Experimental Psychology: General*, 122, 73–91.
- Lynch, K. (1960). *The image of the city* (Vol. 11). MIT Press.
- Lyon, D. R. & Gunzelmann, G. M. (2011). Functional equivalence and spatial path memory. *The Quarterly Journal of Experimental Psychology*, 64, 2081-2087.
- Lyon, D. R., Gunzelmann, G. M., & Gluck, K. A. (2008). A computational model of spatial visualization capacity. *Cognitive Psychology*, 57, 122-152.

- MacGregor, J. N. & Ormerod, T. C. (1996). Human performance on the Traveling Salesman Problem. *Perception & Psychophysics*, 58, 527-539.
- MacGregor, J. N., Ormerod, T. C., & Chronicle, E. (2000). A model of human performance on the traveling salesperson problem. *Memory & Cognition*, 7, 1183-1190.
- Marr, D. (1971). Simple memory: a theory for archicortex. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 23-81.
- Militello, L. G. & Hutton, R. J. B. (1998). Applied cognitive task analysis (ACTA): A practitioner's toolkit for understanding cognitive task demands. *Ergonomics*, 41, 1618-1641.
- Miller, G. A. (1956). The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological review*, 63, 81.
- Mills, M. E. (2007). Challenges to the acceptance and proliferation of tactical UGVs. *RUSA Defence Systems*, 10, 28-30.
- Montello, D.R. (1993): Scale and Multiple Psychologies of Space. In A.U. Frank and I. Campari (Eds.), *Spatial Information Theory: A Theoretical Basis for GIS*, pp. 312-321. Berlin: Springer-Verlag.
- Mottet, M., & Saury, J. (2013). Accurately locating one's spatial position in one's environment during a navigation task: Adaptive activity for finding or setting control flags in orienteering. *Psychology of Sport and Exercise*, 14, 189-199.

- Mueller, S. T., Perelman, B. S., & Simpkins, B. G. (2013). Pathfinding in the cognitive map: Network models of mechanisms for search and planning. *Biologically Inspired Cognitive Architectures*, 5, 94-111.
- Mueller, S. T., Perelman, B. S., & Veinott, E. S. (2015). An optimization approach for mapping and measuring the divergence and correspondence between paths. *Behavioral Research Methods*, [Epub ahead of print].
- Multi-Service Tactics, Techniques, and Procedures for Cordon and Search Operations (2006). FM 3-06-20. April http://armypubs.army.mil/doctrine/Active_FM.html.
- Newell, A. & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hill.
- O'Keefe, J. & Burgess, N. (1996). Geometric determinants of the place fields of hippocampal neurons. *Nature*, 381, 425-428.
- O'Keefe, J. & Dostrovsky, J. (1971). The hippocampus as a spatial map: preliminary evidence from unit activity in the freely moving rat. *Brain Research*, 34, 171-175.
- Pan, X., Han, C. S., Dauber, K., & Law, K. H. (2006). Human and social behavior in computational modeling and analysis of egress. *Automation in Construction*, 15, 448-461.
- Perelman, B. S. (2014, May). *Modeling problem solving in search tasks and the traveling salesman problem*. Presented at the 2014 Midwestern Psychological Association Annual Meeting, IL.

- Perelman, B. S. & Mueller, S. T. (2013a). Examining memory for search using a simulated aerial search and rescue task. In *Proceedings of the 17th International Symposium on Aviation Psychology*, OH.
- Perelman, B. S. & Mueller, S. T. (2013b). Modeling human performance in simulated unmanned aerial search. In *Proceedings of the 12th International Conference on Cognitive Modeling*, Ottawa, CA.
- Perelman, B. S. & Mueller, S. T. (2014). *Hierarchical spatial representations in search and planning*. Presented at the Michigan Technological University Graduate Research Colloquium, MI.
- Perelman, B. S. & Mueller, S. T. (2015). Identifying mental models of search in a simulated flight task using a pathmapping approach. In *Proceedings of the 18th International Symposium on Aviation Psychology*, OH.
- Peterson, L. R. & Peterson, M. J. (1959). Short-term retention of individual verbal items. *Journal of Experimental Psychology*, 58, 193–198.
- Pizlo, Z., Rosenfeld, A., & Epelboim, J. (1995). An exponential pyramid model of the timecourse of size processing. *Vision Research*, 35, 1089–1107.
- Pizlo, Z. & Stefanov, E. (2013). Solving large problems with small working memory. *The Journal of Problem Solving*, 6, 5. DOI: <http://dx.doi.org/10.7771/1932-6246.1155>
- 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, 8.

- Plumert, J. M. & Hund, A. M. (2001). The development of memory for location: What role do spatial stereotypes play?
- Poole, H. J. (2003). *The Tiger's Way: A U.S. Private's Best Chance for Survival*. MD: Posterity Press.
- Ragni, M. & Wiener, J. M. (2012). Constraints, Inferences, and the Shortest Path: Which paths do we prefer? In *Proceedings of the Annual Meeting of the Cognitive Science Society*, Sapporo, Japan.
- Rolls, E. T., & O'Mara, S. M. (1995). View-responsive neurons in the primate hippocampal complex. *Hippocampus*, 5, 409-424.
- Rolls, E. T., Robertson, R. G., & Georges-François, P. (1997). Spatial view cells in the primate hippocampus. *European Journal of Neuroscience*, 9, 1789-1794.
- Rüetschi, U. (2007). Wayfinding in scene space: Modeling transfers in public transport. (Doctoral dissertation).
- Saab, B. J., Georgiou, J., Nath, A., Lee, F. J., Wang, M., Michalon, A., Liu, F., Mansuy, I. M., & Roder, J. C. (2009). NCS-1 in the dentate gyrus promotes exploration, synaptic plasticity, and rapid acquisition of spatial memory. *Neuron*, 63, 643-656.
- Sadalla, E. K. & Montello, D. R. (1989). Remembering changes in direction. *Environment and Behavior*, 21, 346-363.
- Samsonovich, A., & McNaughton, B. L. (1997). Path integration and cognitive mapping in a continuous attractor neural network model. *The Journal of Neuroscience*, 17, 5900-5920.

- Schrijver, A. (2005). On the history of combinatorial optimization (till 1960). *Handbooks in Operations Research and Management Science: Discrete Optimization*, 12, 1.
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions. *Journal of neurology, neurosurgery, and psychiatry*, 20, 11.
- Shallice, T. (1982). Specific impairments of planning. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 298, 199-209.
- Shima, T. S. & Schumacher, C. J. (2005). Assignment of cooperating UAVs to simultaneous tasks using genetic algorithms. In *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, CA: San Francisco.
- Skaggs, W. E., McNaughton, B. L., Wilson, M. A., & Barnes, C. A. (1996). Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus*, 6, 149-172.
- Sousa, J., Simsex, T., & Varaiya, P. (2004). Task planning and execution for UAV teams. In *Proceedings of the 43rd IEEE Conference on Decision and Control*, 3804-3810.
- Sperling, G. (1960). The information available in brief visual presentations. *Psychological Monographs*, 74, 1-29.
- Stansbury, R. S., Vyas, M. A., & Wilson, T. A. (2008). A survey of UAS technologies for command, control, and communication (C3). *Journal of Intelligent and Robotic Systems*, 54, 61-78.
- Stevens, A. & Coupe, P. (1978). Distortions in judged spatial relations. *Cognitive Psychology*, 10, 422-437.

- Takahashi, S. (2013). Hierarchical organization of context in the hippocampal episodic code. *eLife*, 2. DOI: 10.7554/eLife.00321
- Taube, J. S., Muller, R. U., & Ranck, J. B. (1990). Head-direction cells recorded from the postsubiculum in freely moving rats. I. Description and quantitative analysis. *Journal of Neuroscience*, 10, 420-435.
- Tenbrink, T. & Seifert, I. (2011). Conceptual layers and strategies in tour planning. *Cognitive Processes*, 12, 109-125.
- Tenbrink, T. & Wiener, J. (2009). The verbalization of multiple strategies in a variant of the traveling salesman problem. *Cognitive Processes*, 10, 143-161.
- Timpf, S., Rüetschi, U., & Caduff, D. (2005). Cognitive wayfinding agents in public transportation networks. Geographic Information Science Center, University of Zurich, Switzerland.
- Tolman, E. C. (1948). Cognitive maps in rats and men. *Psychological review*, 55, 189.
- The Warrior Ethos and Soldier Combat Skills (FM 3-21.75)
- Viswanathan, V., Lees, M., & Ayt, H. (2012). An information-based model of pre-evacuation behavior for agent based egress simulation. *Proceedings of Pedestrian and Evacuation Dynamics*.
- Wang, Q., Gao, E., & Burkhalter, A. (2011). Gateways of ventral and dorsal streams into mouse visual cortex. *The Journal of Neuroscience*, 31, 1905-1918.
- Weng, M., Jiang, S., & Qu, R. (2008). Hierarchical spatial reasoning and case of wayfinding. *Geo-spatial Information Science*, 11, 269-272.

- Werner, S. & Diedrichsen, J. (2002). The time course of spatial memory distortions. *Memory & Cognition*, 30, 718-730.
- Whitlock, J. R., Sutherland, R. J., Witter, M. P., Moser, M., & Moser, E. I. (2008). Navigating from hippocampus to parietal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 105, 14755-14762.
- Witter, M.P., Naber, P.A., van Haeften, T., Machielsen, W.C., Rombouts, S.A., Barkhof, F., Scheltens, P., Lopes da Silva, F.H. (2000). Cortico-hippocampal communication by way of parallel parahippocampal-subicular pathways. *Hippocampus*, 10, 398-410.
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, 453, 233-235.