

The Importance of Being Expert: A Multi-Method Approach
to Modeling Expert Cognition in Naturalistic Environments

by

Warren Korey MacDougall, B. Sc. (Hons)

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs

in partial fulfilment of the requirements for the degree of

Doctor of Philosophy

Institute of Cognitive Science

Carleton University

Ottawa, Ontario, Canada

April 08, 2016

© Warren Korey MacDougall, 2016

ABSTRACT

The following work aims to contribute to the study of experts and expertise. It is concerned with three principle questions. First, what is expertise? Second, which methodologies can be usefully applied to the study of expertise as it operates in naturalistic contexts? Third, how can the scientific community evaluate and refine these methodologies, as well as develop new ones? These questions are addressed using three frameworks: the Universal Architecture of Expertise, Methodological Pipelining, and Lakatosian Analysis.

Three studies of expert cognition are presented, each of which uses a different methodology. The first study presents a novel method for building macro cognitive models of experts, and applies this method to an analysis of individuals playing a fast-paced video game (Gears of War 3). The second study is a communication analysis of teams playing Gears of War 3 and Counter Strike: Global Offensive, with a focus on communication relating to coordination and interruption handling. The third study is a Python ACT-R model of team coordination using a simplified simulation environment based on Gears of War 3 gameplay. Taken together, these three studies are intended as a demonstration of a research programme built upon the aforementioned three frameworks.

Acknowledgements

A great many people have helped me along the way through graduate school. I am deeply grateful for the assistance I have received, and would like to thank the people who have made this journey possible. I apologize in advance to those whom I have forgotten to include here; I know there are many.

First, I would like to thank Robert West. Your supervision and guidance through the years has been instrumental to my intellectual development and my progress through the PhD program. I would have never finished without your help.

I would also like to thank the entire faculty of the Institute of Cognitive Science. All of the professors with whom I studied and worked provided me with a wonderful education which I am both proud and grateful to have received. You have also provided me with models of brilliance, humility, generosity, and grace to which I will aspire for the rest of my career.

Many thanks to my dissertation committee members: Kasia Muldner, Ida Toivonen, and Brian Greenspan. Your feedback and insights vastly improved the quality of my dissertation, and the conversations leading to its writing were invaluable to me. Thank you for taking time out of your mind-bogglingly busy schedules to help me with this process. A special thanks to my external examiner Frank Ritter for offering your time and providing such extensive feedback.

I owe a great deal to the administrative staff of ICS: Liane Dubreuil, Colleen Fulton, May Hyde, and John Tracey. Without your help, the byzantine mazes of

university protocol would have destroyed me. Thank you so much for helping me with the countless forms, submissions, deadlines, meetings, and milestones that make up a degree.

Thank you also to my lab mates and student colleagues: David Leibowitz, Sterling Somers, Kathy Van Benthem, Francis Jeanson, Janine Fitzpatrick, Nathan Nagy, Matthew Martin, and the rest of the CogSci graduate students. The conversations in the lab and at the pub, and the joint projects running long into the night, have been centrally important to my experience of graduate school.

I would like to thank Corrie Vendetti for her generous support and patience, particularly in the final stages of writing the document. You brought a great deal of light to an otherwise absurd and difficult time.

Finally, I would like to thank my parents and my family. Without their ongoing support, none of this would have been possible. I owe you more than I could ever say.

TABLE OF CONTENTS

Abstract.....	ii
Acknowledgements.....	iii
Table of Contents.....	v
List of Tables	ix
List of Figures	x
List of Appendices.....	xi
Chapter 1 – Introduction and Overview of the Thesis.....	1
Three Frameworks	2
Universal Architecture of Expertise	2
Methodological Pipelining	3
Lakatosian Analysis	4
Three Studies	7
Macro cognitive model – SGOMS & video annotation of game playing	7
Communication Analysis - Coordination and the division of labour	8
Computational implementation of interruption and communication mechanisms	10
Outline of thesis.....	12
Chapter 2 – The Study of Expertise.....	13
Overview	13
Introduction	13
Approaches to Classification.....	15
Recursive Nature of Expert Studies	16
Review of Methods	18
Psychological Methods	18
Psychology Laboratory Methods.....	20
Familiar versus Contrived Tasks	21
Social Psychology Methods.....	26
Neuroscientific Methods.....	27
Sociological, Anthropological, and Economic Approaches	31
Observation of work practices	32
Task analysis	34
Distributed cognition	36

Game theory	37
Sociology of professions –	39
Historiometric Methods.....	41
Artificial Intelligence & Computational Approaches.....	42
Expert systems –	43
Ethology and Computational biology –	44
Computational Cognitive Modeling -	46
Conclusion.....	47
Chapter 3 – Modeling Expertise in Naturalistic Settings	49
Overview	49
Introduction	49
Naturalistic Cognitive Modeling.....	51
Macro Cognition	53
SGOMS and the Expert Architecture	54
SGOMS Framework.....	55
Method	58
Participants and Performance Criteria	60
Procedure.....	61
Data Collection and Preliminary Models	62
Model Overview.....	66
Model Elements	66
Situational Awareness Model –	66
Chunks, (Context) Buffers, Modules, Production Rules –	68
Model Evaluation	70
Filmed gameplay	70
Filmed interview	71
Video annotation	73
Results.....	79
Discussion	90
Chapter 4 – Expert Cooperation	95
Overview	95
Introduction	95
Division of Labour – Importance to Expertise.....	97

Division of Labour - Genesis.....	102
Communication.....	108
Caveat: Trust and Narratives of Professionalism	114
Organization.....	120
Distributed Cognition and the Division of Cognitive Labour	124
SGOMS Ontology and Methodological Pipelining	127
Cognition, Architectures, and Levels as Ontological Entities.....	129
Modularity	130
Communication Analysis – Gears of War.....	134
Overview	134
Participants	135
Data	135
Results.....	135
Discussion.....	140
Communication Analysis – Counter Strike: Global Offensive	146
Overview	146
Data	148
Discussion.....	148
Conclusions	151
Chapter 5 – An SGOMS ACT-R Model of Interruptions and Communication	152
Overview	152
Introduction	153
Interruptions	153
Communication.....	155
Practical difficulties of modeling interruptions	157
Methodology.....	160
Architecture Changes.....	160
Bottom-up Buffering -	164
Bottom-up Production Systems and Interruption -	167
Integration with SGOMS	169
Description of Model Functioning	171
Functioning of the Model – Interruptions:	172
Functioning of the Model – Communication	172

Evaluation of the Model	173
Evaluation of Interruption Mechanism.....	173
Evaluation of the Communication Mechanism.....	174
Quantitative Analysis of Single Agent Pilot.....	175
Quantitative Analysis of Wander Model.....	179
Quantitative Analysis of Communication Model.....	182
Conclusions	185
Chapter 6 – Summary of Contributions, Conclusions, and Future Directions	186
Summary of Contributions.....	186
Conclusions and Future Directions	188
References	194

List of Tables

Table 1. Example of context buffer contents written in Python ACT-R pseudocode	72
Table 2. Example planning unit - Defend Position	80
Table 3. Measurements of occurrence for each planning unit.....	81
Table 4. Measurements of occurrence for each unit task	82
Table 5. Measurements of occurrence for each failure type.....	84

List of Figures

Figure 1. Screenshot of Gears of War 3.	59
Figure 2. Post-It note model of Gears of War.	63
Figure 3. Screenshot of SGOMS GUI	65
Figure 4. Screenshot of the Anvil annotation software.	78
Figure 5. Comparison of planning unit use	87
Figure 6. Comparison of planning unit use in terms of map.....	88
Figure 7. Comparison of planning unit use by wave	89
Figure 8. An ACT-R model with two arbitrary modules	163
Figure 9. Schematic of the ACT-R expert cooperation model.....	163
Figure 10. Grid world with maximum initial delta	176
Figure 11. Average time to completion for a single agent in Wander Model, based on size of world.....	178
Figure 12. Average time to completion for a single agent in Wander Model, based on initial delta.	178
Figure 13. Average time to completion in multi-agent wander model, based on world size. ...	180
Figure 14. Average time to completion in multi-agent wander model, based on agent count.	180
Figure 15. Average time to completion in multi-agent wander model, based on world size and agent count.	181
Figure 16. Average time to completion in multi-agent communication model, based on world size	183
Figure 17. Average time to completion in multi-agent communication model, based on number of agents, averaged across world size.	183
Figure 18. Average time to completion in multi-agent communication model, based on world size and number of agents.....	184

List of Appendices

Appendix A. Example ANVIL annotation XML.....	251
Appendix B. Sample Gears of War Team Communication Transcriptions.....	256
Appendix C. Sample Counter Strike: Global Offensive Transcript	262
Appendix D. ACT-R Model Code.....	271
Appendix E. Sample Output of the Model with Debugging Print Statements.....	288

Chapter 1 – Introduction and Overview of the Thesis

The following work aims to contribute to the study of experts and expertise. It is concerned with three principle questions. First, what is expertise? Second, which methodologies can be usefully applied to the study of expertise as it operates in naturalistic contexts? Third, how can the scientific community evaluate and refine these methodologies, as well as develop new ones?

The solutions offered herein to these questions are given in the form of three frameworks which can be thought of as the three layers of abstraction that structure this document. The three frameworks are the Universal Architecture of Expertise, Methodological Pipelining, and Lakatosian Analysis, constituting the layers of theory, method, and philosophy of scientific practice, respectively.

This thesis presents three studies of expert cognition, each of which uses a different methodology. The first study presents a novel methodology for building macro cognitive models of experts, and applies this method to an analysis of individuals playing a fast-paced video game (Gears of War 3). The second study is a communication analysis of teams playing Gears of War 3 and Counter Strike: Global Offensive, with a focus on communication relating to coordination and interruption handling. The third study is a Python ACT-R model of team coordination using a simplified simulation environment based on Gears of War 3 gameplay.

This chapter is structured as follows. The next section examines the three frameworks that are used to structure this thesis, and which drove the three studies

that are herein reported. The following section examines these three studies briefly. The final section outlines the structure of the rest of the thesis. We turn now to an examination of the three frameworks.

Three Frameworks

Universal Architecture of Expertise

The first framework used to structure this work is based upon the idea that human beings share a common expert system, a Universal Architecture of Expertise (UAE). Simply put, this means that different individuals and different forms of expertise rely on the same cognitive structures and abilities. More precisely, this means that at a particular level of analysis, important commonalities underlie superficially distinct manifestations of expertise, such as performing surgery, playing basketball, and conducting scientific research.

It is important to be clear that this hypothesis is closely tied to the concept of levels of analysis, and to the notion of macro cognition. The macro cognitive level sits above the biological and micro cognitive levels of analysis, and below the knowledge level. Without this context, the idea that all forms of expertise rely on the same cognitive structures is either, A) obviously true and too broad to be interesting (all human experts are composed of neurons, muscles, tongues, etc.), or B) patently false (one's legal expertise is wholly unrelated to one's facility with tensor calculus or knowledge of baseball statistics). Three examples of relevant macro cognitive functions are the abilities of: dealing effectively with interruptions; creating complex, multi-part plans; and dynamically balancing global and local demands and opportunities when

working in teams. These are the sort of domain-general abilities that are important in very different forms of expertise. More will be said about the notion of macro cognition in the following chapters, particularly Chapters 2 and 3.

This notion of a UAE is an unorthodox position in the expertise literature. If there is truth in it, this idea would have significant implications for learning and collaboration in real-world environments (it might suggest, for example, why we are so well suited to social learning and team work) and could have substantial impacts on the field of expert studies. To properly address this question, I believe it is necessary to integrate insights from the many different fields that have turned their attention to the study of expertise, and to employ methods that can be applied to a range of different types of expertise. This brings us to the next framework: Methodological Pipelining.

Methodological Pipelining

Many different fields have turned their attention to the study of experts: psychology, artificial intelligence, economics, sociology, history, education, sport and performance studies, and neuroscience, among others. Each of these offers unique methods and insights to tackle the complex set of phenomena related to expertise, from developmental factors to social determinants affecting expertise, from the economic conditions underlying skill development to the genetic and neurological basis of high performance.

In an attempt to integrate the range of ideas produced by these various fields, I am proposing in this work the idea of a methodological pipeline. In outline, this involves

conducting a study using one set of theories and methods to generate results and novel ideas, and then feeding these results forward into another paradigm, each time with the intent to further refine and expand the scope of the working hypotheses. This thesis is a test of a methodological orientation that is explicitly multi-disciplinary. The intent is not to conduct multiple studies using different tools and ideas, simply for the sake of variety. The plan, instead, is to begin systematizing a methodology that facilitates the integration of multiple approaches, demonstrating one possible path by which insights from different fields may be leveraged to generate novel ideas and techniques. This methodology is aimed at the end goal of developing a unified theory of expertise.

Lakatosian Analysis

The final framework concerns the philosophy of advancement that we apply to our research programme. It is a view onto what constitutes an appropriate model of scientific progress. The framework was developed by Imre Lakatos (Lakatos, 1970) in response to the philosophies of Popper (Popper, 1963) and Kuhn (Kuhn, 1962), and it has been largely adopted by the cognitive modeling community, beginning with the work Alan Newell (Cooper, 2007; Ernst & Newell, 1969; Newell, 1973). The principle point of relevance for the purposes of the present work is that a simple model of empirical falsification is not an adequate strategy for developing unified theories of cognition, particularly when drawing on the collection of concepts and techniques used in research with cognitive architectures. The basic reason for this is that work of this nature relies on a high number of “free variables” and interacting parts, and it is frequently not possible to isolate a single variable and test it through conditional

manipulation, as is the intent in much psychological laboratory experimentation. This general idea is captured in a quote from Douglas Adams, who wrote: “If you try and take a cat apart to see how it works, the first thing you have on your hands is a non-working cat” (Adams, 2002, p. 135).

Instead of direct tests for falsification, this framework suggests an approach of iterative development, using rolling hypotheses and refinement over time. It is similar in spirit to the approach taken in software development, which follows a pattern of architect-build-test-refine-repeat. In other words, the intent of a particular experiment or study is to find a way to improve the existing theory if possible. If a theory repeatedly runs up against unsolvable issues, it eventually ought to be discarded, but reaching this consensus requires the analytical work of a community. Due to the nature of these theories, such a conclusion cannot be reached through a single failed hypothesis test.

This is not to say that falsification is not relevant for work with cognitive architectures. When such tests are possible, they are applied to peripheral commitments of a theory, rather than to the “hard core” of the theory (Cooper, 2007). The idea here is that we ought to employ an approach that Lakatos calls “sophisticated falsificationism”, as distinct from naive or methodological falsificationism. The naive form of falsificationism, according to Lakatos, rests upon the idea that any statement or idea that can be constructed in such a way that it may be falsified qualifies as a scientific statement (as opposed to an unscientific one), and by subjecting these falsifiable statements to experimental scrutiny we can marshal evidence in favor of or against a

particular hypothesis or theory. The progress of a field, under this view, depends upon a collection of such exercises in falsification.

Alternatively, under the view of sophisticated falsificationism, before an existing theory/hypothesis can be rejected there must first be a replacement position available that can explain the empirical content of the predecessor theory, as well as account for novel or excess empirical content that the predecessor theory cannot account for. In other words, falsification in practice requires that something better be available before we jettison an existing paradigm.

This position is partly descriptive and partly prescriptive. On the one hand, it is Lakatos' account of how scientific work is actually conducted. Scientists rarely, if ever, immediately abandon a position upon discovery of counter evidence: more studies may be ordered, the evidence may be reinterpreted, or the theoretical framework may be expanded or modified to encompass the newly discovered facts. On the other hand, it is a view on how we ought to conduct research on complex, multi part systems, particularly when this work is highly distributed throughout a community. For problems of this nature, it is important to be systematic in our efforts to develop unified theories and to communicate the results of research, and this requires more than conducting batteries of hypothesis-testing experiments. This is the point made by Alan Newell in his paper "You can't play 20 questions with nature and win" (Newell, 1973). This approach is particularly germane to the study of expertise, as it is a complex phenomenon that has received attention from a wide range of disciplines. This issue is addressed more thoroughly in Chapter 2.

Three Studies

Three studies were conducted for this work. At the theoretical or conceptual level, they were all aimed at further developing the idea of a Universal Architecture of Expertise, and taken together, they are intended as a demonstration in practice of the Methodological Pipeline. In other words, the results from Study 1 serve as the starting point for Study 2, the results of which serve as the starting point for the final study. The 3 studies are introduced below.

Macro cognitive model – SGOMS & video annotation of game playing

The first study, detailed in Chapter 3, presents a novel methodology for generating and evaluating macro cognitive models of naturalistic tasks. The method was tested with individuals playing a fast-paced video game, Gears of War 3 (Epic Games, 2011). This task was chosen because of the cognitive demands placed on the players: the game is complex and chaotic, requiring the players to respond to regular interruptions, to track numerous environmental variables, and to generate and execute complex plans. These are the sorts of environmental pressures to which experts in the “real world” (i.e., beyond the confines of the laboratory) must routinely respond. In other words, the method is aimed at capturing and analyzing real-world task performance by experts.

The study uses recorded audio-video data of gameplay. Analysis begins with the theoretical base of SGOMS, Sociotechnical Goals, Operators, Methods, Selection rules(Gray, John, & Atwood, 1993; West & MacDougall, 2014; West & Nagy, 2007), to generate a cognitive task model of gameplay. The model is repeatedly improved using a process of iterative model tracing and video annotation. With each iteration of the

model, data is reanalyzed to examine changes in coverage of the model over the data. The final version of the model presented in Chapter 3 was the seventh such iteration.

The results of this study were a model that predicted approximately 85% of the actions recorded in the gameplay video, and a hypothesis about the importance of “graceful interruptions” in expert behavior. These two outcomes were carried forward into the second study.

Communication Analysis - Coordination and the division of labour

The second study, presented in Chapter 4, combines the model and ideas generated from Study 1 with approaches drawn from anthropology, sociology, and economics, to examine the importance and functioning of coordination, communication, and division of labour in expert communities. The data for the study comprised transcripts of communication from teams playing Gears of War 3 and Counter Strike: Global Offensive. These data were analyzed to further refine the SGOMS model developed in the first study. The focus was on testing ideas about the SGOMS macro cognitive architecture in a multi-agent scenario, and to develop a method suited to scaling up the individual SGOMS model to a team model.

The result of Study 2 was an analysis of two forms of communication that mapped onto theoretical ideas from the SGOMS framework, specifically planning units and graceful interruptions. Planning units are theoretical cognitive control structures that are meant to model complex, goal directed behavior and are thought to underlie effective collaboration between teams. Planning units are, further, hypothesized to be

the mechanism by which experts manage “graceful interruptions”; that is, mitigating the cost of disrupted action.

Concerning the first pattern of communication that was analyzed, it was observed that communication patterns differed according to whether the individual agents were modeled as currently within a planning unit or between planning units. It appeared that, while executing a given planning unit, communication occurred in short bursts, and could be usefully modeled as supporting a form of “pooled perception”, an idea that is inspired by Klein et al.'s notion of common ground (Klein, Feltoich, Bradshaw, & Woods, 2005). In contrast to this, when team members were between planning units, communication was more strategic and involved higher level negotiation of team plans. The importance of identifying this pattern in the communication is that it lends support to the notion that the idea of the SGOMS macro cognitive architecture can be usefully applied to the study of multi-agent scenarios.

The second pattern of communication that was interesting concerned responses to failures or interruptions. The plans carried out by the teams were often highly complex, and required for their successful execution precise spatial and temporal coordination on the part of each team member; this precision was particularly evident in the analysis of Counter Strike gameplay, where the teams analyzed were professionals playing in international competitions. When an individual failed to execute their role within the larger plan, the individual and the team had to make rapid decisions about whether to abandon or rethink the plan, and this required quick and precise communication. The tension between local and global reactivity is a central

point of interest in the SGOMS research programme, and this is thought to be mediated through the use of shared planning units. This communication concerning failure and interruptions was thus an interesting test-bed for further developing ideas about communication and coordination in expert communities.

These two patterns identified in the analyses, concerning the relation of communication to planning units and response to interruptions, were carried forward into the third study, which centered around the construction of a computational model of (a simplified version of) this sort of team-based game play. This model is described next.

Computational implementation of interruption and communication mechanisms

The final study was a computational implementation of the primary points of interest derived from the first two studies. It was built using Python ACT-R (Stewart & West, 2007) and ideas from the SGOMS modeling framework.

The model served as a test of two primary ideas: a hypothetical mechanism of coordination based on shared planning units within a team, and a graceful interruption mechanism that allowed computational agents to respond locally to unexpected events without crashing, both of which, it is argued, are important macro cognitive functions underlying expert performance in the real world.

The route by which these ideas were tested was to minimally modify the ACT-R architecture by implementing a dual, bottom-up/top-down vision system that could gracefully handle interruptions, and a communication mechanism which allowed agents

to communicate planning units to one another. One of the points of interest that arose from the theoretical treatment of division of labour in Chapter 4 was that proper coordination of teams (i.e., effective division of labour) can lead to greater than additive gains in productivity. In other words, teams can be greater than the sum of their parts. The communication mechanism implemented in the computational model served as a preliminary test of a cognitively plausible mechanism by which these gains may be realized in real-world division of labour; namely, the communication of shared planning units. The model demonstrated substantial (that is, greater than additive) gains in performance when this communication mechanism was implemented.

At the theoretical level, the successful computational implementation of these mechanisms may provide some early, supporting evidence for the idea of the Universal Architecture of Expertise. If experts of various types do in fact share a common macro cognitive architecture, the packaging, communication, and unpacking of shared planning units may be a principle mechanism by which the value of this shared system is realized, and the research programme presented herein offers a route by which these ideas can be tested and refined.

At the methodological level, the successful functioning of the model is evidence in favour of the value of the methodological pipelining approach. It is a case-study demonstration of the viability of the feed-forward approach to iterating and generating useful insights across disciplinary lines. More work will, of course, need to be conducted in order to make any claims about the general usefulness of this approach and these ideas, but the present work is a promising start.

Outline of thesis

The remainder of this document is structured as follows. Chapter 2 reviews the literature concerning experts and expertise. It examines the wide range of methods and theories that have been applied to the study of expertise in a variety of scientific disciplines. The chapter concludes with a look at how several of these methods are combined in the present work.

Chapter 3 details a methodology that has been developed to create macro-cognitive models of individuals carrying out a complex task. This method consists of recording video data and iteratively constructing and testing a cognitive model against that data. A study of individuals playing a fast paced video game is presented as a test case of the methodology.

Chapter 4 examines the centrality of social structures in expertise, and the problem of moving from modeling individual experts to modeling teams of experts. The chapter presents a study of expert cooperation built upon the work presented in Chapter 3.

Chapter 5 presents a computational model of expert cooperation. It examines the role of interruptions and communication in real-world expertise, and presents a computational solution to modeling these in an SGOMS ACT-R model.

Chapter 6 is an overview and summary of the work, and a discussion of potential future directions for this research programme.

Chapter 2 – The Study of Expertise

Overview

This chapter presents an overview of the theories and methods that have been applied to the study of experts and expertise. The chapter is structured in 3 sections. The first section presents an introduction to the issue of classifying approaches to studying expertise. The second section lays out a taxonomy of theories and methods for studying experts. Each approach is described in its general goals and intents, and a brief caveat about its limitations are offered. The chapter concludes with a note about the variety of methods.

Introduction

Studying expertise presents a number of challenges of scope. First, expertise spans many areas, domains, or disciplines, and this makes demarcating or defining expertise in a general sense difficult. Second, many different fields have studied the processes and products of expertise, bringing their own theories, tools, and methodologies to the problem, often in ways that are difficult or impossible to compare, reconcile, or integrate. In this sense, expertise is like the proverbial elephant being studied by a group of blind men: each man touches a different part of the animal, and each comes to different conclusions about what he has encountered. Third, expertise occurs over varied scales of time and space. Consider the expertise of building and operating the International Space Station versus finding checkmate in a particular chess position, or the development of the medical profession over centuries versus the ability to ascertain a correct diagnosis by examining an MRI scan.

The products of expertise are so varied and multifaceted, and the fields that have turned their attention to the problem so diverse, that it is sensible to ask whether expertise is a unitary phenomenon at all, whether there is any value or validity in discussing it as such. In the Cambridge Handbook of Expertise (K A Ericsson, Charness, Feltovitch, & Hoffman, 2006), the authors state that:

The premise for a field studying expertise and expert performance is that there are sufficient similarities in the theoretical principles mediating the phenomena and the methods for studying them in different domains that it would be possible to propose a general theory of expertise and expert performance (p. 9).

Whether this premise is true or false has yet to be shown conclusively. Of the many approaches that have been applied to the study of expertise, some have taken the positive stance on this question: expertise is a concept that can be usefully generalized beyond domain-specific competence (Eccles & Feltovitch, 2014; K A Ericsson & Smith, 1991; Henrich, 2015). Others have taken the negative stance: all forms of expertise are more or less unique, and commonalities are superficial or illusory (Baer, 2012; Owen et al., 2010; Silvia, Kaufman, & Pretz, 2009). The present work takes the position that important commonalities exist between different forms of expertise, and that a combination of macro-cognitive modeling, task and communication analyses, and ideas from sociology and distributed cognition is a useful approach for examining these commonalities. The next section examines the various ways by which approaches may be classified, and the remainder of the chapter reviews the range of theories and methods that have been applied to the study of experts and expertise.

Approaches to Classification

Classifying the approaches that have been used to study expertise can be accomplished in different ways. I will be classifying the various approaches with two tiers: at the top level, according to academic disciplines or sciences, and at the second level, the set of methods used. Alternative classification schemes might focus on the phenomena of interest (e.g., pattern perception or reasoning or expert communities), theoretical frameworks (behaviorism vs. cognitivism) or detail the evolution of expert studies historically. Another approach is to classify the various approaches according to the view taken on what constitutes expertise. The Cambridge Handbook of Expertise (K A Ericsson et al., 2006) reviews 5 different perspectives on this question: expertise can be conceived of as resulting from differences in innate mental capacity (Galton, 1869; Howe, Davidson, & Sloboda, 1998; Sternberg, & Grigorenko, 2003), an extrapolation of everyday skill to extended experience (Chase & Simon, 1973b; de Groot, 1965; Ernst & Newell, 1969), a qualitatively different representation and organization of knowledge in some domain (Chi, Feltovich, & Glaser, 1981; Chi, 2006; Starkes & Allard, 1993), elite achievement resulting from superior learning environments (Bloom, 1985; K A Ericsson, 2009), or reliably superior performance on representative tasks (Bolger & Wright, 1992; de Groot, 1965; Van Der Maas & Wagenmakers, 2005). These views are not all mutually exclusive, and there are yet others beside.

As a final example, one influential sociological perspective has examined expertise in light of the power that knowledge communities can wield by maintaining a monopoly on useful information, such as tradecraft secrets or medical technique

(Abbott, 1988; Brante, 1988; Groenewegen, 2006; Klegon, 1978). The view taken on this question largely informs the methods and theories one applies to research into expertise.

Recursive Nature of Expert Studies

Before examining the various methodologies that have been applied to the study of experts, a note about how the study of expertise may serve as an illuminative microcosm for the more general issue of group cohesion in scientific practice. Imposing some order on the diverse range of work ostensibly about “experts”, “expertise”, “expert performance”, or “expert cognition” makes tractable an otherwise unwieldy body of literature. That is the intent, at least. In using such a classificatory scheme, we must address, implicitly or explicitly, the following question: how do knowledge communities hang together? What is the core around which such groups form? The answer one arrives at in the context of organizing this literature reflects interestingly back upon the nature of expertise. Academic groups studying expertise are, after all, communities of experts.

To illustrate the significance of this issue in the context of the present work, consider the organization of a university. Universities, by and large, share the same disciplinary structure: departments for psychology, mathematics, physics, biology, etc. Within those departments there are sub-specialities: a psychology department might employ specialists in memory, mathematical cognition, neuroimaging, and child development. And, of course, those specialities are further subdivided into more targeted sub-specialities or topical communities. So the question is: why are universities

organized this way? A simple and intuitive answer, though it is perhaps inaccurate or inadequate, is that members of a department share theoretical commitments, methodologies, and specialized languages, all of which facilitate cooperation and the sharing of resources. From a slightly different angle, Böhme (1975) identifies four types of “scientific solidarity”: solidarity in common methods, goals, theories, or objects. Perhaps the most influential view on this matter is that expounded by Thomas Kuhn in his landmark book, *The Structure of Scientific Revolutions* (1962), in which he examines the centrality of “paradigms” in binding scientific communities together (see also Bizzell, 1992; Horgan, 2012), where paradigms are, roughly, exemplar theories, methods, and standards shared by the community.

The point to be made with these example perspectives is two-fold. First, the expert communities that we call academic disciplines are connected according to a complex set of interacting shared interests, and any lines drawn in the sand are bound to be, to some degree, artificial or incomplete. This is particularly true in the study of experts, because the objects of interest (i.e., the experts and their abilities) are enormously varied, and the variety of tools that can be brought to bear upon their study arguably includes all methods that can be applied to the study of human beings, including methods of psychology, neuroscience, anthropology, linguistics, sociology, ecology, history, mathematics, genetics, kinesiology, economics, education, literary studies, and artificial intelligence. There is, therefore, a great deal of overlap between fields that might in other cases be considered distinct, such as artificial intelligence and ecology, or economics and literary theory. Differences in approach, method, and jargon

can obscure important shared interests, and thus synthesis of this literature presents some unique challenges.

Second, and consequently, the attempt to construct a unified theory of expertise requires simultaneously addressing two parallel forms of differentiation: between different types of expertise, and between different methods/theories that have been used to studies these various types of expertise. The form of the study reflects its content.

With that said, let us turn to a review of the literature about experts.

Review of Methods

Psychological Methods

The most extensive work on expertise has been done within psychological paradigms. This seems to be due to two factors: the nature of expertise, and the toolkit of psychologists. Regarding the first factor, expertise tends to involve a great number of different psychological capacities, including working memory (Guida, Gobet, Tardieu, & Nicolas, 2012; Liu, Schallert, & Carroll, 2004; Moore, Cohen, & Ranganath, 2006; Sohn & Doane, 2004), long-term memory (Beilock, Wierenga, & Carr, 2003; K A Ericsson & Kintsch, 1995; Kolodner, 1983; Vicente & Wang, 1998) , pattern perception (Bialic, Turella, Campitelli, Erb, & Grodd, 2012; Cheung & Bar, 2012; Curby & Gauthier, 2010; Herzmann & Curran, 2011; Kellman & Garrigan, 2009; Loveday, Wiggins, Festa, Schell, & Twigg, 2013), self-monitoring (Cleary & Zimmerman, 2005; Glaser, 1996; Webster & Schempp, 2007; Zimmerman, 2006), sustained attention or heightened focus

(Farrington-Darby & Wilson, 2006; Memmert, Simons, & Grimme, 2009; Memmert, 2009, 2011; Tsang & Shaner, 1998; N. Weber & Brewer, 2003), routinized cognition and behavior (Kimball & Holyoak, 2000; Lazaric, 2008; Paletz, Kim, Schunn, Tollinger, & Vera, 2013), gross and fine motor coordination (Roberts, Bain, Day, & Husain, 2013; Schack & Mechsner, 2006; Starkes & Allard, 1993), motivation (Laverick, 2007; Van de Wiel & Van den Bossche, 2013), social cognition (Byrne & Whiten, 1989; Henrich, 2015), reasoning and problem solving (Carbonell, 1986; Chi, Glaser, & Rees, 1982; Nokes, Schunn, & Chi, 2010), abstraction and categorization (Johnson & Mervis, 1997; A. Mason & Singh, 2011; Noss, Hoyles, & Pozzi, 2002; Tanaka & Taylor, 1991), analogical or metaphorical thinking (Carbonell, 1986; Marchant, Robinson, Anderson, & Schadewald, 1991; Novick, 1988), planning (Gilmore & Green, 1988; Koedinger & Anderson, 1990; Mutton, Hagger, & Burn, 2011), and improvisation (A. L. Berkowitz & Ansari, 2010; Kleinmintz, Goldstein, Mayseless, Abecasis, & Shamay-Tsoory, 2014; Pinho, de Manzano, Fransson, Eriksson, & Ullén, 2014; Vera & Crossan, 2005).

Different forms of expertise require these capacities in different degrees, but many domains of expertise, such as medicine, science, chess, and martial arts, require nearly all of them (and perhaps others beside). This implies that the study of any particular psychological capacity may lead quite naturally into the study of experts more broadly, because each of these capacities seem to be co-present, often in highly

developed form, in the operations of experts¹. Which brings us to the second factor accounting for psychology's extensive work on expertise: psychologists have developed a rich array of methodologies specifically tuned to investigating the behavior and mental functioning of human beings (and, of course, non-human animals), and thus they are well-equipped to investigate the various phenomena underlying expert cognition and behavior. While expertise may be largely dependent upon economic, social, and historical factors, it is psychology that has been most naturally positioned and well-equipped to study individual experts.

The rest of this section examines the classes of methods that psychologists have used to study experts and some of the representative problems they have addressed.

Psychology Laboratory Methods - Psychology laboratories have produced a great deal of work examining expertise. The principal advantage of doing work in the lab is that environmental conditions can be controlled. This allows the investigator to manipulate variables of interest while holding others constant. It minimizes interference from factors deemed irrelevant, and allows the repetition of conditions across many trials and participants (generally with conditional variation), thereby making data amenable to statistical testing. This is, of course, what drives experimentation across many

¹This property of expertise nicely illustrates the message of Indra's Net, a fable of Hinduism about the interconnectedness of all things. The fable depicts an infinitely expansive net, with jewels at each vertex, each jewel being infinitely-many sided. Each face of each jewel reflects light from the other jewels, a sort of nested recursion. Each jewel face ultimately contains within it a reflection of everything else, and thus a close study of any part will reveal the whole. This notion has recurred in human thought for ages. William Blake famously began his *Auguries of Innocence* with the lines "To see a world in a grain of sand; And a heaven in a wild flower", which anthropologist Clifford Geertz (1973) references to ground his anthropological method when he says "seeing heaven in a grain of sand is not a trick only poets can accomplish", (p. 44). This idea is also closely related to the idea that "everything is connected to everything else", (Hutchins, 2010, p. 705), and to Plato's notion, presented in *Phaedrus*, that we must carve nature at its joints (Plato, 2003).

disciplines; it is not unique to psychology or expert studies. It is made explicit here because the strength of this control becomes a weakness in some contexts, and addressing this trade-off is a principal interest in the present work. This trade-off between control and ecological validity is discussed in more depth in Chapter 3. We will now look at specific psychology methodologies.

Familiar versus Contrived Tasks - One way of differentiating experimental methods is according to whether the expert participant is completing a “familiar task” or a “contrived task” (Hoffman, Shadbolt, Burton, & Klein, 1995; Hoffman, 1987; Vicente & Wang, 1998). Familiar tasks are those that are representative of what an expert would normally do as part of their expertise: a chess player selecting a move, a pilot flying an airplane or operating a flight simulator, or a professional counselor leading a discussion. These tasks have a degree of ecological validity, but they are difficult to control and extract meaningful data from, largely because of the complexity involved in expert cognition and behavior, and thus these are less commonly used in experimental scenarios. Contrived tasks, on the other hand, are simplified, artificial tasks that give the experimenter more control over testing conditions. These are generally simplified tasks that are intended to allow the investigator to focus on a subset of psychological abilities that are of interest. Problem solving tasks are a common method by which to assess participants’ domain knowledge (Chi et al., 1981; A. Mason & Singh, 2011), and may allow for evaluation of speed or accuracy in solving the problems and subsequent comparison between experts and novices (Brand-Gruwe, Wopereis, & Vermetten, 2005; Jacobson, 2001; Kirlik, Plamondon, & Jagacinski, 1993; Spence & Brucks, 1997). The raw

solution data acquired from such methods do not always provide meaningful insight into the knowledge upon which experts may draw, nor on the processes used in solving the questions, however, and thus more controlled paradigms have been developed. Chi (2006) examines 4 sub-classes of contrived tasks to elicit the structure of expert knowledge: these are tasks testing recall, perception, categorizing, and verbal reporting, detailed below.

Recall tasks have been among the most important in developing generalized models of how expert knowledge differs from that of novices. These tasks involve briefly presenting experts and novices with some domain-specific content, and asking them to recall as much as they can. This allows for tabulation of response accuracy, while speeded conditions allow for measurements of response time. These tasks are often done with visual stimuli. A landmark study by de Groot (1965) presented chess masters with chess positions and asked them to recall as many pieces as they could, and found that experts were able to recall many more pieces and configurations than were novices. This work served as inspiration for Chase and Simon (1973a) to study the memory of chess experts, which led to the development of the chunking theory of memory which persists today and informs the ACT-R cognitive architecture used in Chapter 5 (J. R. Anderson, 1996; Gobet et al., 2001). Other recall tasks use verbal stimuli. As an example, Morrow et al. (2001) tested recall of pilots and non-pilots who had listened to air traffic control messages and found that the pilots were able to recall significantly more information than the

non-pilot participants. This superior memory effect for domain-specific information has been demonstrated repeatedly, across many different types of expertise; Vicente and Wang (1998) catalogue 51 different studies across 19 domains of expertise in which this effect has been demonstrated, ranging from computer programming to medical diagnosis, basketball to musical notation. See (Vicente, 1988) for a review of this literature.

Perception Tasks are used to test differences between experts and novices in the initial stages of processing stimuli, often to determine whether certain features or patterns are differentially salient, and to examine the process by which experts internalize information presented to them (e.g., what is attended to first, where in the process they may pause to reflect, etc.). This approach has been used in a variety of domains of expertise, including X-ray analysis (Lesgold et al., 1988), weather forecasting with satellite infrared images (Hoffman, Trafton, & Roebber, 2006), and teachers watching video footage of classroom instruction (Sabers, Cushing, & Berliner, 1991). Studies consistently find that experts are quicker than novices to fixate on task-relevant details.

Categorization Tasks have participants sort “objects” into various categories, according to a particular feature. An influential application of this method was the work done by Chi et al. (1981) which had physics graduate and undergraduate students sort physics problems into categories, and then describe the dimensions according to which they did the sorting. The graduate students, i.e., the relative experts in the sample, sorted the problems according to deep

domain principles, such as whether the problems required invocation of Newton's Second Law or the law of the conservation of energy, whereas undergraduates sorted the problems according to more superficial similarities, such as whether the problem description involved inclined planes or frictions. This has been interpreted as evidence that experts have richer knowledge representations that are structured differently than that of novices. This paradigm has also been applied to eliciting the knowledge structure of expert computer programmers (Weiser & Shertz, 1983), counsellors (Mayfield, Kardash, & Kivlighan, 1999), and commercial fisherman (Shafto & Coley, 2003). Among the most well-known varieties of categorization tasks is the Wisconsin Card Sorting Task (Laws, 1999; Nyhus & Barceló, 2009; Vayalakkara, Backhaus, Bradley, Simco, & Golden, 2000), which has participants sort cards into piles according to colour, number, or shape. While the WCST is not used in direct assessment of expert knowledge, it has been used to study other mediating mechanisms of expertise, such as cognitive flexibility (Yudes, Macizo, & Bajo, 2011) and working memory (Hartman, Bolton, & Fehnel, 2001). Another variant of categorization tasks is metaphorical thinking or analogical reasoning tasks, which assess whether participants can generalize information from one domain or problem type to another. A notable example of this method is the research by Duncker (1945) and Gick and Holyoak (1980), that presents participants with a story about a general who captured a fortress by spreading his men out to approach from multiple angles simultaneously (a radius of attack). The task then requires them

to solve a problem about treating a patient's tumor. The tool to be used is a special kind of ray that will destroy the tumor at high intensity, but that will also destroy surrounding tissue at the necessary intensity. The solution is to spread the rays out such that they approach the tumor from different directions: a solution nearly identical in structure to the General's solution for capturing the fortress. Gick and Holyoak (1980) found that without the analogical story as a prime, roughly 10% of participants came up with the correct response, but when provided with the General's story before the radiation story, roughly 30% of participants found the answer.

Verbal Reporting tasks attempt to elicit the structure of participants' knowledge through speaking. This can take various forms. Protocol analysis (Smagorinsky, 1998; Trickett & Trafton, 2009; Walker, 2005) is a speak-aloud method that has subjects report as thoroughly as they can their thinking while accomplishing some task, such as selecting a chess move. This can provide the investigator with information about what is being attended to, in what sequence, categories being formed, proposed plans, and solutions being entertained by the expert. Another commonly used verbal task is conducting retrospective interviews (Hauw & Durand, 2007; Sosniak, 2006), which ask participants to reflect on and communicate the reasoning behind their actions. Both of these methods, concurrent think-aloud protocol analysis and retrospective interviews, form part of the methodology used in Study 1 of this thesis, which will be presented in Chapter 3.

Various issues with verbalization methods have been identified and debated. The principal three are: people do not have perfect access to everything that goes on in their brain/mind, and thus will be unable to report certain elements or processes that may be important (Nisbett & Wilson, 1977) ; that by asking participants to verbalize what they are doing, they may change the processing underlying the task (H. L. Dreyfus & Dreyfus, 1986); and the “differential access hypothesis” (Cheatham & Lane, 2002), which asserts that different modes of eliciting information (e.g., retrospective interviews vs. think-aloud protocols) may lead to differences in that nature of what is recalled, which implies that the analyst must be careful not to confuse the artificially elicited knowledge structures with a transparent or direct representation of the participants’ cognition.

Issues: Doubts have been raised about whether the results of psychology experiments can be scaled up out of the lab to help us understand cognition as it occurs in the real world, often discussed in terms of “external validity” (Anderson & Bushman, 1997; Berkowitz & Donnerstein, 1982; Calder, Phillips, & Tybout, 1982). Macro-cognition is a sub-field that is largely concerned with addressing this concern (Klein & Hoffman, 2008), and is one of the approaches that informs the current work; it will be discussed in greater detail in Chapter 3.

Social Psychology Methods - Another body of psychological work has concentrated on the social variables and personality factors that predispose individuals to high achievement and elite performance. Much of this work consists of survey or

questionnaire based research, and often aims to connect the development, maintenance, and exercise of expertise to factors such as general intelligence (Grabner, 2014; Horn & Masunaga, 2006), introversion/extroversion (Balthazard, Potter, & Warren, 2002; Kayaoglu, 2013), motivation (Laverick, 2007; Van de Wiel & Van den Bossche, 2013), access to educational resources (Litzinger, Lattuca, Hadgraft, & Newstetter, 2011), socioeconomic status (McCoy, 2005; Noble, Wolmetz, Ochs, Farah, & McCandliss, 2006), and so on. Other work in social psychology has examined communication and collaboration in the context of expertise, and the role of experts in communities (Brown, Coman, & Hirst, 2009; Mieg, 2001; Smith & Ellsworth, 1987; Smith-Jentsch, Johnston, & Payne, 1998). The social components of expertise are discussed at length in Chapter 4.

Issues - Social psychology methods provide evidence of which personality, social, or environmental factors are important to the development or exercise of expertise, but tend to be silent about the processes underlying expert cognition and behavior.

Neuroscientific Methods

The neuroscientific study of expertise is a natural extension of the psychological approaches. If it is accepted that the human cognitive apparatus is fundamental in enabling expertise, and further, that the physical basis of cognition is the structure and function of the brain, it is a small step to use psychological research on expertise as a launching pad for research into the neuroscientific basis of these capacities. This work is primarily done under the rubric of cognitive neuroscience (Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1999; Maurer, Zevin, & McCandliss, 2008; Rossion, Kung, & Tarr,

2004; Vuust et al., 2005) and neuropsychology (Rezlescu, Barton, Pitcher, & Duchaine, 2014; Turk, Handy, & Gazzaniga, 2005), but it can also be connected to the neuroeconomics of decision making (Glimcher, Dorris, & Bayer, 2005; Sanfey, Loewenstein, McClure, & Cohen, 2006).

Neuroimaging techniques such as functional magnetic resonance imaging (fMRI) and electroencephalography (EEG) allow recording of neural activity in real-time, allowing researchers to examine the effects of training or skilled performance on the structure and function of the brain. These non-invasive imaging techniques can be combined with behavioral data to model differences in neural functioning between experts and novices (Doppelmayr, Finkenzeller, & Sauseng, 2008; Lutz, Greischar, Perlman, & Davidson, 2009; Milton, Solodkin, Hlustik, & Small, 2007), or to observe changes in brain structure or functional patterns over time as a response to training (Erickson et al., 2007; Kassubek, Schmidtke, Kimmig, Lücking, & Greenlee, 2001; Lee et al., 2012; Xiong et al., 2009).

Neuroscientific experimental methods for studying expertise must be complimented by psychological perspectives, as there is currently no satisfactory explanatory bridge from neural activity to complex cognition and ultimately to expert behavior. In other words, we cannot yet extract from EEG or fMRI data the fact that during those 3 seconds, the chess expert decided to move her Bishop to g4 because it will pin the opponent's knight to his king, or that a stock trader decided to short the mortgage market because he foresaw widespread failure of the American housing and credit markets (which is one explanation of what happened in the United States housing

bubble in 2007/08; see Lewis [2010], for more). We do not yet fully understand what sorts of patterns, regularities, structures, or processes exist in expert cognition, and thus psychological and cognitive models are crucial in making sense of the neuroimaging data.

On the other hand, neuroscientific research can inform psychological theorizing, as it has at its disposal a number of theoretical constructs and empirical techniques which quite powerfully complement cognitive models of expertise. Two of the most important of these are neuroplasticity and neural modularity. Neural plasticity is the mechanism by which the brain is reorganized according to some internal or external stimulus (Meyer, Elmer, & Jäncke, 2012; Münte, Altenmüller, & Jäncke, 2002), and is a natural complement to psychological studies of learning. Understanding how the brain changes in response to instruction or deliberate practice (Ericsson, Krampe, & Tesch-Römer, 1993) may be useful in creating models of expertise development, retention, and loss.

One relatively recent finding from neuroscience which may prove to be quite important in our psychological models of expertise is that the capacity for neuroplasticity is retained more robustly into old age than was previously thought (Kheirbek & Hen, 2015; Spalding et al., 2013; Valkanova, Eguia Rodriguez, & Ebmeier, 2014). This suggests that training in later years of life may be more effective than previously believed. This idea is supported by some early evidence that training elderly participants on new skills can lead to increases in gray matter (Boyke, Driemeyer, Gaser, Büchel, & May, 2008).

Concerning modularity, mainstream views in both psychology and neuroscience treat the mind/brain as a modular organ, with specialized regions or modules for accomplishing specific tasks, but also possessing certain domain-general properties or abilities (Bullinaria, 2007; Redies & Puelles, 2001). An example illustrating the interaction of modularity and generality from the view of psychology is a deep knowledge of visual patterns of chess positions: pattern recognition and storage is domain general, knowledge of chess patterns and their significances is domain specific. An example illustrating the difference from the perspective of neuroscience is the combination of specialized areas of the brain devoted to language processing, such as Wernicke's area and Broca's area (DeWitt & Rauschecker, 2013; Keller, Crow, Foundas, Amunts, & Roberts, 2009), and domain-general mechanisms like attention and memory, which underlie language learning and use, but seem to be much more structurally diffuse (Bressler & Menon, 2010; Shipp, 2004) . This view of the brain/mind as a set of interconnected modules has been influential in the theories underlying modern cognitive architectures such as ACT-R (Anderson, 2007; Anderson, 2010), a topic which we will return to in Chapter 5.

Issues: As noted above, neuroscience's understanding of how we learn, remember, and think is still in its infancy. Much more work will need to be done before we fully understand the complexity involved in rewiring a brain to pilot an airplane or lead a team of engineers.

Sociological, Anthropological, and Economic Approaches

A great deal of expert behavior in the real world takes place in complex environments characterised by interruptions, unexpected events, failures, and imperfect information. Furthermore, it often involves multiple individuals coordinating to accomplish complex tasks that are beyond the scope of any individual expert, and often beyond the abilities of even a single type of expert (De Langhe, 2010; Fiore et al., 2010; Hutchins, 2010; Keil, Stein, Webb, Billings, & Rozenblit, 2008; Weisberg & Muldoon, 2009). As a simple example, consider the operation of a modern university, which requires professors, administrative staff, maintenance personnel, network technicians, cooking staff, transportation infrastructure, fire alarm systems, inter-departmental mail protocols, security staff, and so on. A good deal of the expertise possessed by these individuals is knowledge about how to operate in this environment.

To understand how expertise is developed and deployed in the real world, disciplines such as sociology, anthropology, and economics have constructed models of experts operating in all the complexity and chaos that characterize their “natural habitats”. This is in contrast to the psychological method, which derives a large part of its value from its techniques for controlling and isolating variables in the environment. Because so much of expert cognition and behavior is dependent upon interacting effectively with specialised environments (e.g., hospitals or flight cockpits), important aspects of expertise can be obscured in studies using implied and/or artificial tasks and environments. On the other hand, simply observing a surgical team operating on a patient or a military unit executing a plan may not yield much useful information,

possibly because there is too much to take in at one time, there is too much variability across “trials”, or the experimenter is unsure about what to attend to. Some degree of methodological rigor, however slight, is therefore necessary to get a handle on observation and data collection. Once the data is collected, it must also be analyzed, and thus the development of qualitative and quantitative analytical techniques tailored to the analysis of the collected data is required. The remainder of this section will examine the approaches that have been used to study “cognition in the wild” (Hutchins, 1995) that have come out of anthropology, sociology, and economics.

Observation of work practices – Many types of expertise form the basis for careers: medicine, law, science, engineering, athletics, navigation, aviation, and musical performance are notable examples. While some parts of these types of expertise may extend beyond a professional context (a professional musician may play a song for friends, and a doctor may offer free advice to a family member), a great deal of the skill possessed by these experts is only possible in the specialized context of work. One reason for this is social, in that collaboration with others is what enables the expertise. Consider, for example, a medical doctor working with other specialists, nurses, and technicians to conduct a surgery, or a conductor leading a symphony. Many skill sets cannot be exercised by an expert acting alone. In other cases, it may be that the workplace provides tools and technologies needed to exercise the skills in question, such as an airplane for a pilot or a nuclear reactor for an engineer. Other properties of the environment may also play a role in the possibility of deploying expertise, such as

physical characteristics of the environment (a speed skater requires ice), or legal conditions (a police officer is unable to conduct an arrest outside their jurisdiction).

The importance of the environment in which expertise is exercised has led researchers to develop techniques for gathering data from observations of experts operating in their natural work environments. These methods are often referred to as ethnomethodology, workplace studies, or field research (Atkinson, 1988; Garfinkel, 1986; Maynard, 1991; Reeves, Brown, & Laurier, 2009; Rouncefield & Tolmie, 2011), and derive from research traditions in sociology. The principle

The roots of applying these techniques to the study of expertise stretch back to cognitive anthropologists and sociotechnical analysts operating in the 1950's (Clancey, 2006; Emery, 1959). It has since found a place in analyses of automation practices in business (Gregory & Nussbaum, 1982; Olson, Lucas Jr., & Kling, 1982), organizational dynamics and climate in schools (Hoy, 1990; Weiss, 1999), computer science studies of knowledge acquisition (Buchanan & Shortliffe, 1984), human-factors psychology and cognitive task analysis (Vicente, 1999; see next section), and in the discipline of human-computer interaction (Blum, 1996; Kling & Star, 1998). These techniques have largely been adopted by communities with a focus on applied research, notably the human factors engineering community (Schraagen, Chipman, & Shalin, 2000; Seamster, Redding, & Kaempf, 1997).

Workplace studies exhibit a great diversity of data-gathering and analysis techniques. Data of interest may include video footage, audio recordings, photographs,

interview transcripts, field notes, schedules, documents, and tools (Button & Sharrock, 2009; Clancey, 2006; Luff, Hindmarsh, & Heath, 2000). Analysis of these data may involve video annotation and interaction analysis (Derry et al., 2010; Greenbaum & Kyng, 1991; Ruhleder, 1997; Luff & Heath, 2012), schedule and process modeling (Donald, 1998; Ley et al., 2008; Van Stijn & Wensley, 2001), data mining on document sets (Rivas et al., 2011), conversational analysis of gaze and gesture (Heritage, 1984) and discourse modeling (Koester, 2006; Stubbe et al., 2003), among others.

Issues: While workplace studies can provide insight into the ways that experts behave in the real world, perhaps in a manner that laboratory experimentation is unable to provide, the method set is quite fragmented. There are few standards of practice that guide data collection and analysis, and this may impede communal evaluation of a paradigm's value.

Task analysis – Task analysis, probably not surprisingly, involves the deconstruction of key tasks to be accomplished by an individual or team. It is an approach that has much in common with workplace observation techniques; in some instances the two approaches may blend together or co-occur. The principal difference is that task analysis has tended to focus more on developing micro-cognitive models of the capabilities and actions of the agents, and somewhat less on how they are coupled with the environment and other people (Diaper, 2002; Kirwan & Ainsworth, 1992). This is in part because of the community's interest in developing generalized models of cognition. This work also tends to be more focused on theoretical unification than on applications, and

the associated methods are primarily used in cognitive psychology research paradigms (Schraagen et al., 2000).

Cognitive Task Analysis (CTA) is a sub-species of task analysis with particular relevance to research into expert cognition and behavior. CTA is defined as “the extension of traditional task analysis techniques to yield information about the knowledge, thought processes and goal structures that underlie observable task performance” (Chipman, Schraagen, & Shalin, 2000). These techniques are often used in conjunction with micro-cognitive modeling techniques, which allows for context and task generalization and models specified to millisecond level predictions (Clark, Feldon, Van Merriënboer, Yates & Early, 2007). In some computational cognitive modeling, such as that informed by GOMS (Amant, Freed, & Ritter, 2005; Gray et al., 1993; John & Kieras, 1996), the modeller will conduct an initial cognitive task analysis, from which they derive hypotheses about what to model in the cognitive architecture. This informs the approach taken in Chapter 3 of this work.

Issues: Micro-cognitive task analysis has proven a useful approach to modeling cognition in many diverse contexts, but the methods are often not well-suited to modeling cognition and behavior of experts operating in the real world. The set of techniques for evaluating the accuracy of a micro-cognitive model by comparing it against human data does not work well when evaluating macro-cognitive models. Micro-cognitive models tend to be of simplified tasks that can be repeated many times. Macro-cognitive models are of complex tasks, and the “repetitions” of complex tasks by humans or computational models are rarely (if ever) identical. Addressing the difficulty

of evaluating macro-cognitive models is one of the aims of the present work; this is addressed further in Chapter 3.

Distributed cognition – The Distributed Cognition (DC) framework studies cognition as it occurs among a group of actors and artifacts (Hollan, Hutchins, & Kirsh, 2000; Hutchins, 1995). Rather than treat the individual mind as the locus of cognition and object of study, DC takes as its unit of analysis teams of agents with their tools and environments (Rogers & Ellis, 1994). This line of thinking was heavily influenced by methodologies and ideas from anthropology (Bender, Hutchins, & Medin, 2010; R. Williams, 2013) and ecological cognition (Cooke, Gorman, & Rowe, 2009; Hutchins, 2010; Yoerg, 1991). Work in this area focuses principally on the relationship of an individual or team to the greater context in which they are embedded, including other social actors, environmental factors, and technological aids. Topics of interest include: how memory is enhanced by various devices and social structures (Michaelian & Sutton, 2013; Sutton, Harris, Keil, & Barnier, 2010); how cognitive load is managed through interface design (Dror & Harnad, 2008; Jiajie Zhang & Patel, 2006); how cooperation is facilitated (Artman & Garbis, 1998; Rogers, 2006); and how information is communicated between specialists (Giere, 2002). Chapter 4, in its discussion of the social dimension of expertise, draws heavily upon this literature.

Issues: Because so many forms of expertise involve teams acting in complex environments (science, military, healthcare, athletics), distributed cognition is a powerful framework for the study of expert cognition and behavior. The focus of this research is largely on high-level descriptions of teamwork, cooperation, and

communication², and less so on theories and methods aimed at developing a comprehensive model of individual cognitive abilities. It remains unclear where the capacity for distributed cognition comes from, though work in this area has made great strides in the past two decades. The data collection and analysis pipeline presented in this dissertation is, in part, an attempt to extend some of the insights from this area of research to the domain of computational cognitive modeling.

Game theory – Game theory offers an economic approach to studying the motivations and mechanisms underlying strategy selection, competition, and cooperation in human behavior (Brandenburger & Nalebuff, 1995; Gintis, 2000; Von Stengel, 1996). Expert behavior frequently involves all three elements. Sports and other games are a straightforward example of this, but the same elements are evident in corporate planning, scientific communities, political action, and educational design. Game theoretic models tend to be mathematical formalisms or computational, agent-based simulations (Osborne, 2000; Shoham, 2008).

Game theory differs from psychological approaches to the study expertise largely in the background assumptions that underlie model creation. Early instantiations of the game theoretic approach were almost exclusively derived from the “rational man” hypothesis in economics (also known as *homo economicus*), which posits that people will attempt to act rationally to maximize their own reward (Argyris, 1973; Harsanyi, 1966; Simon, 1955; Zizzo, 2005). Models of this principle have taken the form of many different modes of interaction, perhaps the most notable of which is the iterated

²This is likely because of its roots in anthropological research.

prisoner's dilemma (IPD) game³, (Newth, 2009; Press & Dyson, 2012). Modern psychology, on the other hand, tends not to treat human beings as perfectly rational, but as swayed by heuristics, emotions, imperfect information, and so on (Chant, 1963; Miljkovic, 2005; Stich, 1985). Behavioral economics is a modern approach that (in part) attempts to reconcile these two views on human cognition and psychology (Camerer, 1999; Kahneman, 2003; Vernon L. Smith, 2005), replacing the image of *homo economicus* with the notion of "bounded rationality".

The conception of human beings as rational agents had been around long before game theory formalisms, but the concurrent treatment of rationality and action in a competitive marketplace is a useful lens onto how expertise operates in a world with limited resources. Specifically, it is an interesting approach to modeling competition between industries and scientific sub-fields, both of which rely heavily on effective development and utilization of expertise.

Expertise is, arguably, largely a product of market forces. The character of expertise in a given region at a particular time is affected by the type of college and university programmes offered, the success stories propagated through a culture, the financial incentives for pursuing one form of expertise as compared to another, the power wielded by particular professions, and so on. A full understanding of how expertise is developed and deployed through a culture requires an understanding of the

³A Google Scholar search for "iterated prisoner's dilemma" currently returns 18.6k results, which gives some indication of how widely used the paradigm is.

role of various economic forces in the process. The macro-economics of expertise is a topic which has received little or no attention in the academic community, but it has been a topic of concern for policy makers and politicians in recent decades. One of the most far-reaching examples of this can be seen in the nationalization of science education and research programmes after the Second World War. Many nations in North America and Western Europe developed substantially more centralized scientific and engineering infrastructures following the war, and the rapid proliferation of new technologies and scientific advancements in the 3 decades that followed were, arguably, largely the result of this centralization, (Drori, 2003; Franzmann, Jansen, & Münte, 2015; Shinn, Sörlin, & Crawford, 1993). Such wide scale development and mobilization of expertise and expert-potential can be enormously powerful.

Issues – While game-theoretic models of behavior can provide useful approximations for how experts cooperate and compete, they often do not attempt to provide a window onto the cognitive mechanisms underlying these capacities. Furthermore, these models tend not to capture the close coupling of the expert with the environment; the agents are typically distillations of rationality (and perhaps some set of heuristics or simplified cognitive mechanisms) in a simplified environment.

Sociology of professions – Many forms of expertise are socially solidified through the establishment of professional groups: sciences, engineering, medicine, law, theology, athletics, education, transportation, financial services, etc. These professional bodies facilitate the transfer of knowledge and skills through education and training programs, the granting of accreditations by which the public can easily recognize who belongs to

the professional group, the control of financial and other resources with which to influence broader markets, and (potentially) a good deal of control in shaping how the profession is perceived by the larger public. Studying how professions develop and operate can therefore provide a useful window onto how expertise operates in the real world. Sociological methods of relevance include philosophical inquiry (H. L. Dreyfus & Dreyfus, 1991; Nowotny, 2003; Selinger & Crease, 2007), interviews (Bogner, Littig, & Menz, 2009; Dorussen, Lenz, & Blavoukos, 2005), social network models (Nunn, 2008; Jun Zhang, Ackerman, & Adamic, 2007), and historical research (see historiometric methods, below).

Among the most interesting questions tackled by the sociological approach, and most relevant to the present work, concerns how the division of labour (DoL) functions in society. This has been a prominent concern in sociology for quite some time (Durkheim, 1893; Marx, 1978; M. Weber, 1978). It seems that human societies leverage DoL and specialization (expertise) to greater effect than any other species on the planet, yet it is not clear what enables this capacity, and enables it to such a heightened pitch.

Issues – Sociology provides an important lens onto the social phenomena surrounding expertise, including how people are trained and knowledge is disseminated, how groups form and work together, and how competition affects professional landscapes. It does not, however, provide much guidance in the construction of cognitive process models by which to study the cognitive abilities in individual agents that underlie communication, cooperation, and competition. The challenge of moving

from sociological ideas to computational cognitive models will be a principal concern in Chapter 4.

Historiometric Methods

Less prominent than other methodologies applied to the study of expertise, historiometric methods attempt to derive general principles about the acquisition and exercise of expertise by studying historical experts. According to Simonton (1990), “historiometrics is a scientific discipline in which nomothetic hypotheses about human behavior are tested by applying quantitative analyses to data concerning historical individuals” (p. 3), where nomothetic hypotheses concern “general laws or regularities of human behavior” (Simonton, 2006, p. 320) The motivation to apply this method set to expertise is that many individuals who reach high levels of achievement, in various fields, become notable public figures (e.g., chess world champions, Nobel laureates), and thus there is often informative, readily available data to work with. This approach is in stark contrast to the methods of the psychology laboratory, where participants are conceptualized as fungible representatives of some idealized population. Another differentiating factor is that historiometric methods will often focus exclusively on populations of deceased individuals, which allows examination of the population across the entire life-span (Raskin, 1935; White, 1931).

One of the most valuable contributions of historiometric methods is their focus on the various indicators of expertise developed by communities of experts. According to Simonton (2006; see p. 323 for citations), individual competence or achievement has been studied through the proxies of:

- (a) eminence as recorded by space allotted in reference works
- (b) the receipt of major honors such as the Nobel Prize or Olympic medals
- (c) total lifetime productivity or the output of highly influential works
- (d) objective scoring systems such as those used to rate chess players and athletes
- (e) the attainment of high offices and positions, such as president, prime minister, pontiff, patriarch, or company CEO
- (f) subjective assessments based on surveys of scholars and other experts

The point to be made here is that communities evolve a variety of mechanisms by which to signal and recognize expertise, and these play an important role in connecting the citizenry with specialists. We will return to this topic in Chapter 4.

Issues – Historiometric methods are useful in historical data exploration and hypothesis formation, but they are less valuable as hypothesis testing methods. They take a long-view of history, and thus compliment sociological perspectives, but do not permit detailed, low-level modeling and testing.

Artificial Intelligence & Computational Approaches

Computational approaches to expertise take as their starting point one of two ideas: the idea that the human mind is an information processing machine and thus is best modeled using computer programs (Card, Moran, & Newell, 1986; Heyck, 2008), or the belief that computational modeling is a useful way to simulate the dynamics of expert cognition and behavior, but not necessarily claiming the strong architectural similarities posited in the first approach (Pitt, Myung, & Zhang, 2002; Sun, 2009). Three camps of this methodological approach are presented here: expert systems in artificial intelligence, computational biology models, and computational cognitive models.

Expert systems – Expert systems are decision making systems, implemented as a computer program, that are meant to mimic the decision making processes of a human expert (Anjaneyulu, 1998; Gallant, 1988). These systems are built by first eliciting, as thoroughly as possible, the knowledge that a human expert would use in solving a particular problem and the rules for using or manipulating that knowledge. Such knowledge and rules are then implemented this information in a computer program, organized primarily as a series of if-then statements about how to reason over the available information. Expert systems have been built to model the expert knowledge of physicians (Díez, Mira, Iturralde, & Zubillaga, 1997; Osuagwu & Okafor, 2010; Schectman, Schorling, Nadkarni, & Voss, 2005; Silva et al., 2014) and lawyers (Cammelli & Socci, 1992; Leith, 2010; Tan & Thoen, 2000), nutritionists (Chen, Hsu, Liu, & Yang, 2012), smoking cessation specialists (Velicer & Prochaska, 1999), machine vibration analysts (Ebersbach & Peng, 2008), and fish disease diagnosticians (Li, Fu, & Duan, 2002), among many others.

Expert systems were an important development in early AI work, and have since played an important role in the study of how knowledge is organized and encoded (Aikins, 1983; Wagner, Otto, & Chung, 2002), and how it is retrieved from memory or informational tools (Brajnik, Guida, & Tasso, 1990; Yao, Zeng, Zhong, & Huang, 2007). Expert systems are still widely used in various domains, but they now enjoy less attention than they once did as models of human expert cognition (Rowe & Watkins, 1992; Sahin, Tolun, & Hassanpour, 2012).

Issues - Perhaps the most serious limitation in constructing expert systems is the bottleneck of getting information into the system: the so-called “knowledge acquisition problem” (Wagner et al., 2002). The most problematic feature of these systems, once they are built, is that they tend to be “brittle”: they are typically unable to handle novel scenarios or intelligently incorporate new information (Holland, 1986). While experts in the real world are continually updating their knowledge base and improvising in new circumstances, expert systems are unable to do either of these very convincingly. These limitations cast doubt on the validity of the expert system model of expert cognition.

Ethology and Computational biology – A number of ethological and computational biology models have examined the division of labour (DoL) in social insects such as ants, wasps, and honeybees (Beshers & Fewell, 2001; Bonabeau, Sobkowski, Theraulaz, & Deneubourg, 1997; Ferguson-Gow, Sumner, Bourke, & Jones, 2014; Gordon, 1996; Jeanson, Kukuk, & Fewell, 2005; Waibel, Floreano, Magnenat, & Keller, 2006). These animals frequently exhibit behavioral specialization, with some individuals foraging for food, others defending the nest, and others cleaning the living space.

As stated in the Sociology of Professions section above, the present work asserts that the division of labour is a key component in the development and exercise of expertise in human societies. Biological studies of this phenomenon in “simpler” organisms provide a view onto how these capacities may have evolved in the natural world and eventually have come to exist in human societies. Whereas sociological or anthropological models of the division of labour might privilege the cultural or uniquely human aspects of collective living in attempting to understand DoL, biological models

look to the natural world for cues. Using the data gleaned from ethological surveys, computational modellers try to account for how incentive structures or communication strategies could have evolved to support this differentiation. For example: concerning incentives, one influential study, by Bonabeau et al. (1997) suggests that a form of DoL can arise out of assigning each individual a “response threshold”, which is the point at which that individual will respond to some stimuli, such as debris around the nest, and take action to reduce the stimulus intensity. They take this computational mechanism to reflect a potential reward mechanism in insects with a strongly genetic basis, thus placing the roots of DoL in the genes. In another study examining communication strategies in relation to the division of labour, Manfredini et al. (2014) examined fire ants and found evidence that they use pheromonal signaling to communicate agent density around the nest, and that this can account for shifts in task allocation across the colony. This relatively simple, spatialized mechanism of communication could serve as a base upon which to build more complex network models of communication-mediated division of labour in human societies.

Issues – Biological models of the division of labour allow careful examination of possible mechanisms by which cooperation could have evolved. These models have trouble in accounting for the many complexities of human cooperation, however, such as our use of mass media, inter-generational instruction, and economic and legal structures to help professional groups maintain power. Combining the insights about division of labour from computational biology and sociology is one of the goals of the present work.

Computational Cognitive Modeling - Cognitive modeling sits between the disciplines of cognitive psychology and artificial intelligence (Sun, 2008). It uses concepts derived from the study of human and animal cognition in conjunction with computational methods developed in AI research to construct models of cognition that can be run as computer simulations. Often the intent of building such models is to assess whether the model as constructed matches human data that has been collected, and to thereby determine whether the model is a reasonable representation of the cognition underlying the task or domain in question (Ritter & Larkin, 1994). This is the principle route by which computational cognitive models are evaluated (Pitt et al., 2002).

A central feature of research using computational models is the overarching attempt to build unified cognitive models that can capture a wide range of cognitive activities and abilities, and to work towards a complete model of a cognitive system. This is the motivation behind the construction of cognitive architectures, which are meant to serve as explicit representations of the general cognitive capabilities of the agents in question (usually humans). Many different architectures have been constructed for this purpose, such as ACT (Anderson, 1996), EPIC (Kieras & Meyer, 1994), and SOAR (Laird, 2012), each intended to test a particular theory of the mind. In Chapter 5, the ACT-R architecture, which stands for Adaptive Control of Thought – Rational (Anderson et al., 2004) will be used to construct cognitive process models of video game playing. To reiterate: the goal of these architectures is to offer a unified model of the mind. My use of such an architecture in the creation of the expert model presented in Chapter 5 represents my endorsement of the value of unification. This

dissertation is, in part, an attempt to articulate a mode of methodological combination that aims at the same goal of unification.

Conclusion

Expertise is a topic that has been investigated from many different angles. A diverse array of research methodologies has been directed at uncovering the mechanisms underlying human beings' incredible abilities, and this has led to an exceptionally rich literature on the topic. This diversity of work can be met in one of two ways: one can attempt to determine the most appropriate line of inquiry for the question of interest (as broadly or narrowly as one wishes), or instead, attempt to determine where dissimilar lines of research are complementary, and attempt to combine their insights or methods. I have chosen the latter approach. I believe that scientific research communities can coalesce either around a topic of interest or according to a shared set of techniques. Contrast, for example, the community of memory researchers with that of fMRI researchers. These two dimensions co-occur in groups and individuals, of course, as methods and topic are often tightly bound; one would be hard pressed to study cellular motility without a working knowledge of microscopy, or pharmacodynamics without knowledge of statistical and experimental techniques. Furthermore, individual researchers may simultaneously belong to many different research communities. The distinction between method- and theory-centered research communities is therefore merely one of convenience.

In the case of expert studies, there remain so many open questions, spread across so many domains, that it seems premature to limit one's self to a single

methodology for reasons of convenience or dogma. Approaching the topic from multiple angles promises (at least the possibility of) a more varied set of intellectual challenges and delights, but such an approach is not self-evidently more valuable, fruitful, or interesting than a mono-method approach. It is my hope that the present work can provide at least modest evidence of the value in such an enterprise.

Chapter 3 – Modeling Expertise in Naturalistic Settings

Overview

The following chapter addresses the following problem: how can we create and evaluate cognitive models of experts operating in the “real world”? The solution offered is an iterative model-tracing method that uses a macro-cognitive modeling framework, audio-video recordings, retrospective interviews, and preliminary model metrics. The method is a combination of several approaches examined in Chapter 2. The chapter also addressed the practical hurdles of communicating the findings of cognitive modeling research.

Introduction

This chapter addresses the difficulty of constructing models of expert cognition and behavior in naturalistic contexts, where interruptions, failures, unpredictability, and other complications are regularly present. Laboratory experimentation with experts effectively controls the complexity of the task environment, gaining an analytical foothold in a difficult problem, but this typically incurs a cost in the form of reduced ecological or external validity of the experimental findings (Anderson & Bushman, 1997; Berkowitz & Donnerstein, 1982; Calder et al., 1982). On the opposite end of the methodological spectrum, anthropology- and sociology-inspired methods for studying expertise tend to take a hands-off approach in order to study maximally “natural” behavior patterns, but observation that is ontologically agnostic can less readily be used to generalize from observation or make predictions about future behavior. To put this another way, pure or unbiased description must eventually give way to systematization,

and *truly* pure description is probably impossible anyway: the analyst is always burdened with a suite of preconceptions (Elgin, 2008; Lapadat & Lindsay, 1998; Schindler, 2013).

The present work thus aims to walk a middle ground, drawing inspiration from these two approaches in the following way. Initial construction of cognitive models is to be constrained by preliminary hypotheses and theoretical constructs drawn from the cognitive process-modeling literature (J. R. Anderson et al., 2004; Kieras & Meyer, 1994; Laird, 2012), but the models so-constructed will be iteratively evaluated and modified, allowing naturalistic observation to feed back into the model. The intent is to test whether a “methodological pipeline” can be developed to leverage insights from dissimilar methodologies.

One of the interesting features of modern cognitive science is the fragmentation of the field. Departments and journal papers routinely refer to “cognitive science”, but arguably there are cases in which “the cognitive sciences” more accurately captures the state of affairs. Michael Dawson (2013), suggests that this fragmentation may be due to incommensurable theories in different camps of cognitive science, leading to communication breakdowns between different “language communities”, such as cognitive modellers and anthropologists. Alternatively, he suggests, it may be that the field in toto is pre-paradigmatic, in the Kuhnian sense, and that we lack “a unifying body of belief” (p. 14). Dawson quotes Kuhn, who writes: “In the early stages of the development of any science different men confronting the same range of phenomena, but not usually all the same particular phenomena, describe and interpret them in

different ways”, (Kuhn, 1962, p. 17). Setting aside the question of the full extent of these traits in cognitive science more broadly, they are surely evident in the study of expertise. In the Cambridge Handbook of Expertise (Ericsson et al., 2006), the first editor claims that the overcoming of such fragmentation in the study of experts has enabled a handbook to be published on the topic. He writes: “A substantial body of empirical findings, distinctive theoretical concepts and frameworks, and a set of new or adapted methods justify a unifying volume.” (p. 3). While I believe there is great merit in compiling the diverse studies presented in the handbook, I think it is a step too far to present the handbook as a “unifying volume”. It remains unclear whether meaningful unification is possible in the study of so complex a topic as the heights of human ability. I wish to explore in the following three chapters the possibility of meaningful methodological combination, rather than simple juxtaposition.

The rest of this chapter presents the first study in a set of 3, each of which aims to build on the insights derived from the previous method. This chapter focuses on an iterative model building process whereby cognitive models are built from observation of experts in minimally constrained (i.e., naturalistic) task environments and the model is continually improved through rolling observation and testing.

Naturalistic Cognitive Modeling

One of the principal difficulties in modeling experts is that many forms of expertise are fundamentally situational, embedded in particular environments that are often quite complex. Consider the expertise of a doctor, lawyer, professional athlete, or pilot. A great deal of their competence involves the ability to navigate the social and

technical components of their workplaces and broader environments. A full understanding of expert capabilities, therefore, requires a model of their actions as enabled by and enacted within these highly specialized environments (Cooke et al., 2009; Vicente & Wang, 1998; Yoerg, 1991). The model constructed in this chapter attempts to connect the internal cognitive operations of the expert with the elements of the environment relevant to this task. The model of the agent's use of information in the environment is referred to as the "situational awareness" (SA) model. The SA model is constructed using the approach of "low fidelity" environment representations (McMahan, Bowman, Zielinski, & Brady, 2012; Walker, Takayama, & Landay, 2002). Low-fidelity refers to the characteristic of only representing those elements in the environment which are relevant to the task. This is in contrast, of course, to high fidelity environment representations, which would aim at capturing as closely as possible everything contained within a task environment.

It is worth pointing out here the distinction between high-fidelity environment models and naturalistic environments. It is perhaps natural to conflate the two because they tend to co-occur, i.e., the complexity of an environment is represented in a complex model. Although real-world environments *tend* to be much more complex than environments presented to participants in laboratory experiments, this need not be the case. Naturalistic environments may be rather simple. An expert chess player engaged in a match, for example, is dealing with a rather simplified environment. In fact, the purpose of the chess board and pieces is precisely to function as a simplified interface with a complex functional layer (i.e., the legal movements of the pieces and

their relationships). While there may be fluorescent lights overhead, spectators in stands, a muttering opponent, and the ticking of a clock in the peripheries of awareness, these environmental factors are arguably ignorable when modeling the task environment. This is the approach taken here. The next two sections detail the notion of a macro cognitive architecture and the SGOMS modeling framework, which underlies the model presented later in this chapter.

Macro Cognition

The study of cognition can be divided into micro cognition and macro cognition (Klein et al., 2003; Schraagen, Militello, Ormerod, & Lipshitz, 2008; West, Hancock, Somers, MacDougall, & Jeanson, 2013), where micro cognition refers to cognition as it is studied in cognitive psychology experiments and macro cognition refers to cognition in complex, real world tasks. This division was motivated by skepticism over whether or not the results of experimental psychology experiments, which use highly simplified and artificial tasks, can be usefully applied to complex real world behaviours. This general concern is not unique to macro cognition. It has been voiced by numerous groups concerned with scaling up from lab based experiments to real world activities (Gelder & Port, 1995; Risko, Laidlaw, Freeth, Foulsham, & Kingstone, 2012; Turvey & Carello, 2012)

Research addressing this concern has looked at the aspects of cognition in the real world that are eliminated in laboratory experimentation, and attempted to construct accounts for how people deal with these conditions. Two such aspects are particularly important in later chapters: dealing effectively with interruptions, and

engaging in cooperative action with other cognitive agents. The next section examines the framework in which these later investigations will be presented.

SGOMS and the Expert Architecture

SGOMS (West & Nagy, 2007) is a macro-architecture designed to model expertise in dynamic, multi-agent tasks. SGOMS stands for sociotechnical GOMS, and is an extension of the GOMS modeling framework (Gray et al., 1993) to address complex environments characterized by multiple actors, imperfect information, and frequent interruptions. It is offered as an alternative to the currently dominant mode of thinking in both expert studies and expert cognition, which treats each domain of expertise separately. A great deal of research has examined whether and to what degree components or types of expertise are domain-general, and the current consensus is that expertise is tightly domain dependent (see Chapter 2). The two principle lines of inquiry here are those examining transference, whereby training and improvement in one skill leads to gains in additional areas, and those examining analogical reasoning, whereby an individual's representation of a particular problem structure can be generalized to another task domain, thereby improving performance. The best known example of the latter is Duncker's (1926) fortress/tumor study. He presented participants with a story about a general who was able to capture a fortress by having his army approach it from many directions at once. If he were to send all of his troops along a single path, the buried landmines would detonate and kill the poor soldiers. After reading this story, participants are presented with a story that is structurally similar: a doctor needs to use x-rays to kill a tumor, but if all of the rays are passed through along the same path, the

intervening tissue will also be killed. The solution is radial, distributed approach, as in the fortress scenario.

These paradigms have largely shown that such high level, cross-domain expertise is rare (Baer, 2012; Silvia et al., 2009), though some work has found evidence that runs in the other direction (Eccles & Feltovitch, 2014; McGraw & Pinney, 1990). SGOMS is an attempt to further investigate whether such domain generality can be captured by reference to a macro cognitive architecture: a common set of cognitive control structures, combined in a particular way, to prop up human performance in a wide range of skills.

The broader question of interest is whether human beings are predisposed to become experts. In other words, how do we account for human beings' ability to become so skillful across so many different domains? Some interesting recent work looking at the intersection of biological evolution and culture has suggested that "Natural selection has equipped our species with a wide range of mental abilities that allow us to effectively and efficiently acquire information from the minds and behavior of other people" (Henrich, 2015, p. 35). SGOMS is offered as a model of the cognitive architecture that may underlie those abilities, and Chapter 4 addresses the importance of learning from others in the development of expertise. We turn next to an overview of SGOMS.

SGOMS Framework

SGOMS was created by extending the GOMS modeling system (which is good for modeling uninterrupted, solitary expert tasks) to model expert tasks in dynamic, multi-

agent environments. GOMS has difficulty handling the frequent interruptions, task switching, and re-planning that happen in real world tasks (West & Nagy, 2007). To address this, SGOMS modifies the definition of the unit task by adding the criterion that a unit task should be small enough so that it will most likely not be interrupted. The unit task is therefore defined as a control structure that functions to avoid overload, downtime, and interruptions. This modification allows the unit task to continue to serve its original function: to define islands of work that can be executed in an uninterrupted, well defined way (MacDougall, Martin, Nagy, & West, 2015; MacDougall, West, & Hancock, 2014; West et al., 2013; West & MacDougall, 2014).

SGOMS also adds a second control structure: the planning unit. Planning units are chains of unit tasks directed towards some larger goal, and are envisioned as the cognitive control mechanism to enable individuals to handle interruptions and task switching. Planning units also allow efficient communication and coordination between agents by functioning as the building blocks for creating plans and modifying them. For example, planning units are theorized to have names collectively used by the members of a group, and thus to be identifiable by examining communication. This idea motivated the communication analyses presented in Chapter 4.

The simplest and most common form of planning unit is thought to be an ordered list of unit tasks. If a planning unit is interrupted, the current unit task is either finished or abandoned and the situation is assessed. The task can be resumed or a new planning unit can be chosen based on the current constraints. When a planning unit is interrupted the progress on the planning unit is stored in memory so that it can be

resumed. This notion of interruption handling is examined further in Chapter 5, and is tested through an implemented SGOMS ACT-R computational model.

In the SGOMS theory, the highest level of decision-making is the constraint satisfaction process used to choose planning units based on the agent's current representation of context, which is continuously updated⁴. In addition, each planning unit is associated with a set of constraints. The goal of this system is to allow agents to react locally and independently as much as possible without disrupting coordination between agents.

SGOMS has the following hierarchical structure of representations. Each is associated with a different set of mechanisms. Although the SGOMS macro-architecture extends only down to the unit tasks, because SGOMS is also a theory of how the micro and macro levels interact, we have also included the micro cognitive GOMS elements in this list:

- Context - constraint based decision-making rules
- Planning units - list or plan, can be modified
- Unit tasks - production systems, smart but brittle
- Methods - fixed set of actions, executed sequentially
- Operators - basic units of perceptual and motor actions
- Bottom up monitoring - when not busy the system checks the environment

⁴ The appropriate frequency at which to model context updating is currently under investigation. For the study presented here, the position taken is informed by both the standard ACT-R production cycle timing of 50 ms per production (Anderson, 2005) and the constraints of the video footage, which was recorded at 30 frames per second (or 33.33 ms per frame). The footage was analyzed frame by frame, though changes in context were typically annotated at a time scale of about 1 second.

The level above controls the level below but the resources are shared; for example, different planning units can call on the same unit task.

In summary, SGOMS is proposed as a way of conceptualizing the macro cognitive regularities that underlie a diverse range of human skills. It represents an effort to systematize the study of expert cognition in complex environments.

Method

The current study used an iterative macro-cognitive modeling procedure to model the cognition of individuals playing a fast-paced video game: Gears of War 3 (Epic Games, 2011), played on the Microsoft XBOX 360 video game console, on “Horde Mode”, which requires players to defeat a continuous stream of enemies. The game is fast-paced and chaotic, and the virtual environments are information-dense and unpredictable. The challenge for the player is to make correct strategic decisions and execute them, which requires situational awareness, an understanding of the task structure, accurate timing, and precise hand-eye coordination. The challenges in modeling this are to determine which factors in the environment are important to the player, how that information is used to make and execute plans, and how to represent this information in a cognitive model. A screenshot of the gameplay is presented below, in Figure 1.



Figure 1. Screenshot of Gears of War 3.

Audio-video data of gameplay was recorded and players were subsequently interviewed to glean some preliminary insight into their decision making processes. These data were used in conjunction with various software tools to construct and refine a cognitive model of their play.

This study served as a testbed for the first step in a methodological pipeline for generating data from a complex task from which to build a macro-cognitive model of situated expertise. The tools and techniques used are examined in more detail below.

Participants and Performance Criteria

Three individuals participated in the study on a voluntary basis, one of whom was the author of this report, and another who was affiliated with the lab. All participants were male, between the ages of 20 and 28.

All participants met the skill threshold of being able to complete the 5th stage of Horde Mode in GoW3. This threshold for expertise was chosen because the stage was difficult enough to require the players' concentration and sustained effort, and they were able to complete the stage successfully approximately 80% of the time. Novices involved in pilot testing were rarely able to complete this stage, and thus were not included in analysis. The reason for the focus on experts and the exclusion of novices is detailed in the discussion section.

It should be noted here that the differentiation between expert and non-expert is largely situational, and often somewhat arbitrary. Various models detail the progression from novice to expert as a series of stages (Alexander, 2003; S. E. Dreyfus &

Dreyfus, 1980; Germain, 2006), which may be of use in particular scenarios or types of expertise. Some forms of expertise, such as chess and tennis for example, have solidly established rating systems which allow for accurate ranking of individuals according to skill level (Neumann et al., 2011). There exists no such differentiating criterion or scale for expertise in Gears of War 3, so we chose what seemed to be a reasonable cut-off to differentiate competent players from beginners. We will return to this difference between competent players and professionals in Chapter 4.

Procedure

Data collection and analysis was conducted with the iterative model tracing methodology referred to above. The data collection process began with observation of experts and novices playing the video game Gears of War 3. After recording, we conducted interviews with the players to glean some insight into their decision making processes while playing.

We then began developing the cognitive model, conducting a cognitive task analysis (Kieras & Meyer, 2000; Schraagen et al., 2000) to determine the relevant unit tasks (UT) and planning units (PU), as well as the environmental conditions associated with each. This model was first represented with paper and pencil, and later using software developed in-house. From here we evaluated the model to determine where it could be improved upon, then implemented these changes. These latter stages (evaluate, implement) were cycled through several times; this was the process of iterative model tracing. Seven versions of the model were created in total. Analyses presented below were all conducted using Version 7 of the model.

Data Collection and Preliminary Models - In the early stages of the modeling process, we experimented with multiple representation formats to encode and communicate the structure of the working model. We began the modeling process with a Post-It note representation, moved to a spreadsheet structure, and finally constructed the model using a graphical user interface (GUI) developed in-house. Note that one of the challenges addressed by the current research programme is that of communicating the structure of a cognitive model in order that it may be evaluated by the broader community; this is addressed further in Chapter 5. The stages of model representation were as follows.

Post-It Notes - From our initial observations, we constructed a preliminary model of unit tasks and planning units to describe the behavior of the players. The first few iterations of the model were visualized as a nested set of Post-It notes attached to the wall. Each proposed unit task and planning unit was represented as a single Post-It. These were arranged hierarchically, with the unit-tasks below their respective planning units. Figure 2 represents an early version of this Post-It note model.

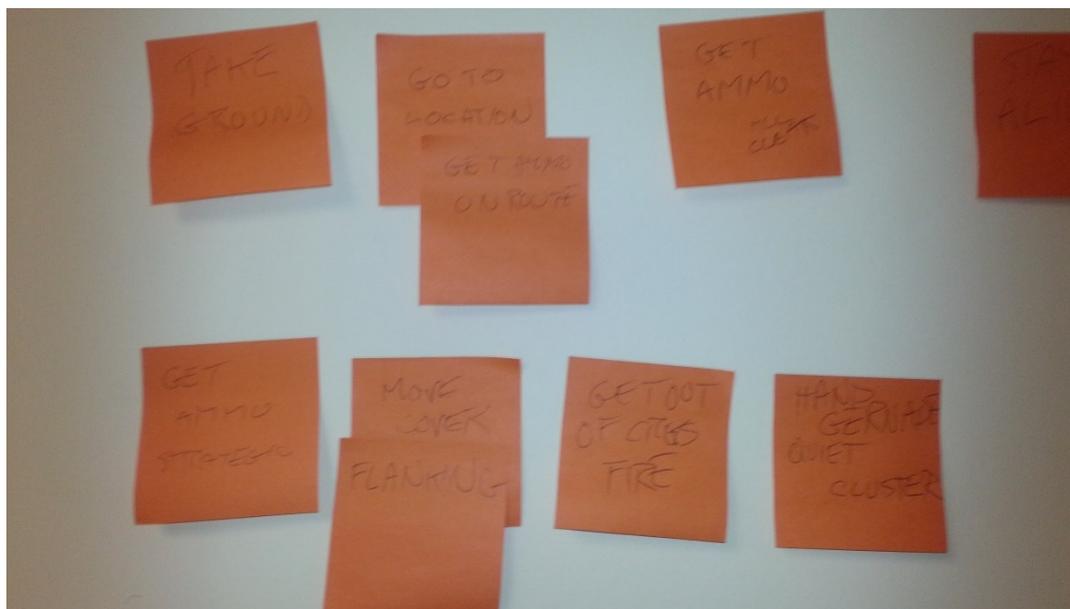


Figure 2. Post-It note model of Gears of War. Top – Overview of model. Bottom – Close-up of one section of the model (the top left corner).

The principal shortcoming of the Post-It method was that of unit-task reuse. One of the primary claims of SGOMS is that unit tasks can occur in multiple planning units (West & MacDougall, 2014). To represent this re-use, a distinct Post-It was used to represent a particular unit task as it occurred in each separate planning unit. For example: if the “reload” unit task occurred in the “hold ground”, “retreat”, and “hunt” planning units, there would be three separate items representing the same theoretical construct in the model. To address this issue, the model was re-factored into a spreadsheet representation.

Spreadsheet – As the cognitive model came to include more planning units and unit tasks, a spreadsheet representation became a more efficient way to store and manipulate the model. A simple representation placed planning units in one column, and the component unit tasks associated with each planning unit in the column adjacent to it. A separate sheet was used to list all unit tasks, independent of the planning units.

While the spreadsheet representation was useful for inventorying the elements of the initial mode, it did not facilitate modeling the internal structure of the planning units and unit tasks, such as the firing conditions or goals. To accomplish this, the model was constructed in a graphical user interface (GUI) that was developed in-house for this purpose. This GUI is described below.

SGOMS GUI – A graphical user interface was developed in house to facilitate more detailed depiction of SGOMS macro-cognitive models (M. Martin, 2014). The interface is shown in Figure 3, with some elements of the model visible on screen.

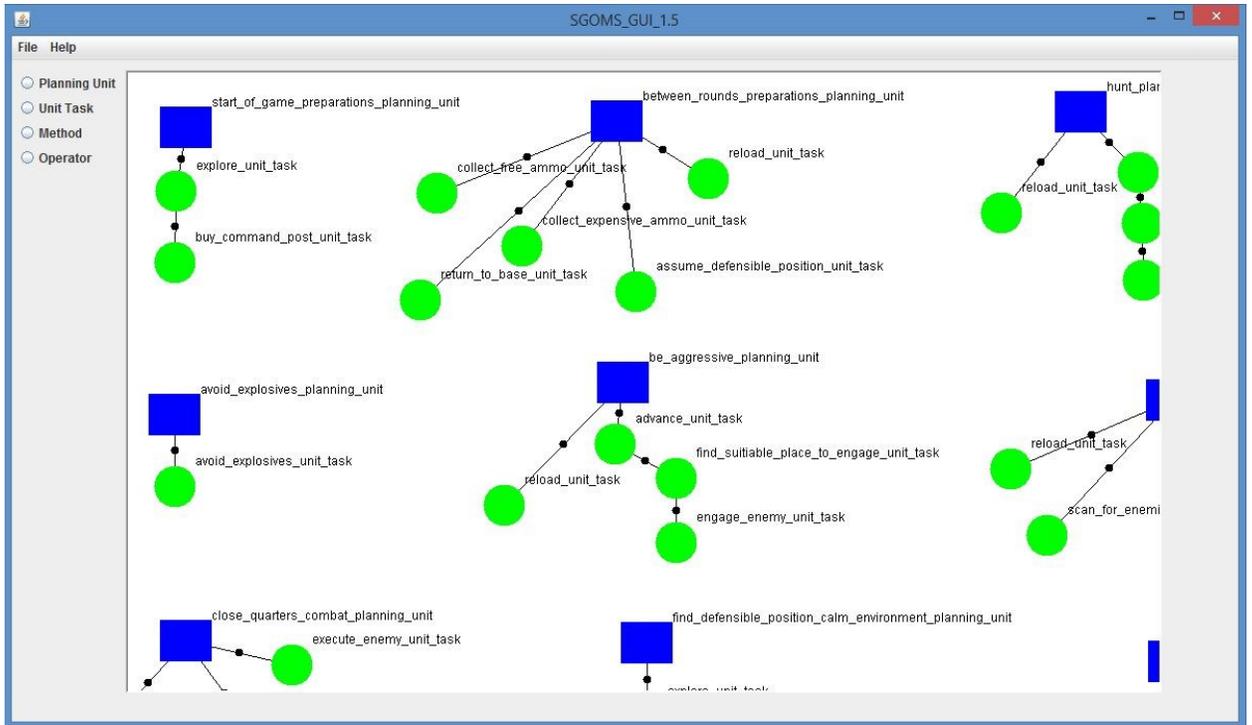


Figure 3. Screenshot of SGOMS GUI displaying a model. The blue squares represent planning units, and the green circles represent unit tasks.

In the GUI, planning units, unit tasks, methods, and operators are represented as connected nodes with internal structures that roughly reflect Python ACT-R syntax (Stewart & West, 2007). Each element is represented as a production rule, with a left-hand “*if*” rule, or firing condition, and a right-hand “*then*”, or behavior/output.

Schematically: Reload unit task: *If*: ammo is low → *Then*: reload The intent of the GUI is to allow the graphical construction of SGOMS cognitive models that can be automatically output to Python ACT-R code, thereby simplifying the modeling pipeline, and helping to bridge the divide between micro- and macro-cognitive modeling. This interface is still under development. The next section examines the principal constructs of the cognitive model. **Model Overview**

The model presented below is of an individual playing Gears of War 3 on Horde Mode, based on the SGOMS modeling framework. The cognitive model includes a Situation Awareness Model (representing the context of gameplay) and a Cognitive Task Model (composed of planning units and units tasks). Together, these make up the cognitive model, which is described in more detail below.

Model Elements

Situational Awareness Model – In order to link the agent models to the environment, it was necessary to construct a model of the environmental factors involved in shaping a player’s behavior and cognition. This was among the most complex aspects of this modeling process, and also among the most important, given the intent to model cognition and behavior in naturalistic environments.

To accomplish this, it was necessary to link behavioral patterns to observable features of the environment (i.e., elements on screen). In some cases, the features of importance were transparently relevant. For example: when the player runs out of ammunition in their weapon, the player-character (i.e., the avatar being controlled by the player) will automatically enter a reloading animation, and ammunition will be replenished. This animation occurred every time the state of “no ammo” was reached, and was thus a simple indicator that could be used to describe the current state of the game environment. The cognitive model can incorporate such simple environmental variables with a discrete, binary state representation: “has-ammo/no-ammo”.

In other cases, the information presented on screen could not be mapped transparently to some state variable. For example: when the player was injured by an enemy, the screen would gradually become more tinted with red, with greater damage levels being represented by a deeper red. While the implementation of this feature in computer code would include a discrete state value, the representation presented to players (and investigators) appears continuous. Mapping the information on screen to hypothetical cognitive constructs is more difficult in such cases than in the ammunition example above, because while it is evident that players use the displayed damage information to make decisions (e.g., if their health becomes dangerously low, they will take shelter somewhere safe to recover), their response thresholds for various elements are invisible and differentially variable. In other words, it is not always clear at which point a player realizes “my health is running low”. The modeller is therefore forced to make a judgement about the *player’s* judgement, to which no direct access is possible.

A binary state of “high-damage/low-damage” may be used (and was in preliminary versions of the model), but this introduces an important possible source of error. It is worth repeating here that this is a complication because these cognitive constructs (“high-health”, “no-ammo”) are meant to represent the investigator’s best guess about the player’s internal representation of the environment, not the environment itself nor the investigator’s post-hoc awareness of relevant features. For the purposes of modeling, this and other continuously-represented features of the environment were discretized, and represented as variables with a fixed number of states, typically two or three. The situational awareness model was defined as the aggregate of all possible states of the context buffer, described in the next section.

Chunks, (Context) Buffers, Modules, Production Rules – The structure of the cognitive model and the situational awareness model that ties the agent to the environment draw on several conventions in cognitive modeling, detailed here. In ACT-R and other cognitive architectures (e.g., SOAR; see Laird, 2012) symbolic information available to a cognitive agent is often represented as “chunks”. These are represented as slot-value pairs. For example: knowledge about one’s pet dog might be represented by the following chunk:

Spot – Isa:Dog Color:Black Breed:Spaniel Weight:70lbs

The chunk construct is also used in a great deal of psychological memory research (Gobet et al., 2001). It derives from Chase and Simon’s work (Chase & Simon, 1973a) with chess experts.

Declarative memory is often modeled as a largely passive store of information (i.e., a list of chunks), and in order to “use” the chunks stored therein, cognitive models posit various processes that act on these chunks, either in storage, retrieval, modification, or some other operation. In ACT-R, the on-line or real-time use of this information is accomplished with the use of “buffers”, theoretical constructs that roughly parallel notions of working memory, the visuo-spatial sketchpad, or phonetic loop (Baddeley, 2003; Bruyer & Scailquin, 1998; Pastells & Roca, 2003). In other words, buffers are neural-psychological functional structures that can temporarily hold a limited amount of information of some particular format (symbolic, auditory), and pass that information to some other structure in the cognitive apparatus, such as the motor system or central executive (depending on the cognitive ontology in use). The buffers therefore serve as an intermediary between the various components or modules of the cognitive system. The “declarative memory buffer”, for example, represents a structure that can hold a chunk from declarative memory, making that information available to some other module.

In the present study, the agents were imputed with a “context buffer”, which was modeled as a structure holding the chunks which represented the agent’s perception of the context⁵. An example context buffer state is shown below, in Table 1.

⁵ The complexities of how sensory-perceptual information is transduced and translated into symbolic information are skirted in the present study. The functional assumption is that cognitive agents in the real world form some representation of that world, and that we can usefully model this representation using symbolic structures such as chunks.

In ACT-R, the actions of modules are often modeled by production rules which take the form of “if-then” pairings of environmental conditions (internal or external) with an action or response, as mentioned in the previous section. In the present study, the function of the chunks in the various buffers serve the same purpose: when a buffer holds a chunk that matches the firing condition of a production rule, the appropriate action will be taken. For example, the *reload* unit task can be represented as:

if buffer_context - ammunition_level:low // *then* reload

In the above statement, the chunk is a single slot:value pair in the context buffer. The slot name is *ammunition_level*, representing a salient feature of the task environment, which it was assumed the agent would be tracking. The slot value is *low*. When this agent processes this condition, the result is to fire the *reload* unit task.

Model Evaluation

Creation and evaluation of the model involved filmed gameplay, retrospective filmed interviews with one of the players, and model-based annotations of the filmed game play. This section details each of these data sources.

Filmed gameplay – Six videos of gameplay of *Gears of War 3 Horde Mode* were used for the model. Footage was captured using a video camera positioned on a stand, pointed at the television screen. A number of other videos were recorded and discarded due to video quality, conversion issues, and compatibility problems with other software in the analysis pipeline, most notably issues with ANVIL (Kipp, 2014) not being compatible with particular video codecs.

Several factors of gameplay were fixed across videos to ensure comparability of runs. These factors were the map (or “level” or “scenario”), the difficulty level, and the “wave number”, which is a finer-grained setting of the difficulty level.

Participants played through a warm-up stage prior to the test run. The warm-up was an earlier stage of the game, i.e., a less difficult version of the task to be completed. Some minimal limitations were placed on the gameplay mechanics to simplify analysis.

Participants were instructed to use a sub-set of the weaponry available in the game, and to avoid purchasing “upgrades”. Neither of these mechanics change the fundamental character of gameplay, and thus should be straightforwardly incorporable into future versions of the model. They were excluded from preliminary models for the sake of expediency.

Filmed interview – One of the participants took part in an interview on the day following their recorded gameplay. He was asked to explain his decision making process during the gameplay phase. The participant was seated, viewing a playback of the previously recorded gameplay, and was asked a series of questions by the examiner. The examiner would play a snippet of the video, pause the video, then ask questions such as: “What are you doing here?”, “What is your assessment of the situation”, or “Do you have a plan in this spot?” The gameplay footage was approximately 3 minutes long. The interview lasted approximately 30 minutes.

Table 1. Example of context buffer contents written in Python ACT-R pseudocode

Categories	Variables
buffer_context = environment:	between_rounds / not_between_rounds
buffer_context = environment:	calm / manageable / chaotic
buffer_context = agent:	health_low / health_high
buffer_context = enemy:	alive / injured / dead

The other two participants were not interviewed in this portion of the study, as they were involved with the laboratory conducting the study, and thus provided similar information to that provided by the interviewee, but in a rolling fashion during analysis and model construction.

Information gained during this interview process was used in later annotation of the video. It was used principally a starting point, to orient the modeling process. It has been addressed at length elsewhere that there are various complications and shortcomings of protocol analysis (e.g., Trickett & Trafton, 2007, and see Chapter 2 of this work) and other verbal elicitation techniques for deriving information about the cognitive processes occurring at a given point during performance of some task. It is generally accepted that individuals do not have complete access to the machinations of their own cognitive apparatus, and thus utterances elicited in these contexts ought to be treated as data rather than truth. With that caveat in mind, these techniques can be useful in helping an investigator to see elements of the environment or task performance that are relevant but which were previously invisible, and thus serve as an important preliminary step in building a cognitive model of expertise.

Video annotation – Formal evaluation of the model began with video annotation. This was accomplished by three individuals (the author of the current report and two others) using the Anvil Video Annotation Tool (Kipp, 2014). This package is built to aid researchers using video data, and offers functionality for video markup and data export. For the present study, it was used to overlay multiple simultaneous annotation tracks over the videos, each of which captured a particular aspect or level of description of the

cognitive model. The tracks used were the Planning Unit (PU), Unit Task (UT), Context (C), and Failure (F) tracks. Annotation in Anvil requires the annotator to demarcate blocks of time with discrete start and end points, for blocks on all tracks. Blocks can be arbitrarily named, typed, and colored, and all are associated with a text field which could be filled with arbitrary text content. Markup data is stored as XML, or extended markup language, files. This markup process is described in more detail below. The initial stage of annotation was to describe the context or the state of the environment. To do this, a simple *primary* track (a markup feature of ANVIL) was used to annotate and describe the context at each point. The text file associated with each marked-up block contained a complete list of all *slot:value* pairs judged to be relevant during construction of the situational awareness model (i.e., those elements that would be make up the context buffer). Annotators judged when the context had changed enough to be considered a new state, and indicated this with a transition to a new block.

Once the context had been annotated, a similar procedure was followed to annotate the planning units on the PU track. Using the information encoded in the context track and the observed behaviors in the gameplay video, annotators indicated the planning units thought to be active at each point during the video, and marked transitions from one PU to another. This second step of annotating PUs involved the most iteration and refinement, with annotators returning to the C-track to make adjustments to initial annotations according to observed shifts in goal-directed behavior, represented in the model (and annotated on PU track) as a change in active planning unit. If the behavior being displayed on screen did not seem to match any of

the planning units currently in the model, the immediate action was to label the time chunk in the PU track as *none*, and the long-term action was to attempt to account for the planning unit in the next iteration of the model (though some instances of *none* remained in Version 7 of the model). Aside from these *none* cases, the entire gameplay video was marked with one of 13 possible planning units specified in the process model. A *primary track* was used again, although in this case they were typed according to the planning unit, which allowed for drop down menu selection of the 13 PUs, and more importantly. This also enabled data export for later analysis. The associated text field was not used for the PU track.

The unit task track was annotated similarly to the PU track. The principle difference is that the UT track was a child of the PU track: unit tasks occur as sub-components of planning units, and thus all UTs are bounded by the planning unit in which they occur. If, for example, a particular planning consisted of three component unit tasks, these three unit tasks would be laid on the UT track below the associated PU, and would extend exactly to the temporal boundaries of the parent track. Again the text field was unused for this track.

Finally, failures of the model to describe the action on screen were annotated on the Failure (F) track. Six types of model failure were identified:

- 1 – Planning Unit does not exist
- 2 – Planning Unit implausible
- 3 – Unit Task does not exist
- 4 – Unit Task implausible

5 – Unit Task out of order

6 – Miscellaneous

Failure types 1 and 3 indicated that the model lacked the appropriate structure to describe the behavior being observed in the video. Types 2 and 4 indicated that, while the observed behavior could be described using existing PUs or UTs, it would require a re-description of the context or environment (on the Context track) that did not seem reasonable. For example, one planning unit, *Switch guns because of ammo*, was intended to capture the behavior of changing weapons when the player had low or no ammunition in a particular weapon. If the player was observed changing weapons after using only 3 bullets from a 60 round clip, marking the context as “low ammo” would be unreasonable. When this (and similar cases) was first encountered, annotators marked the time slice as *Planning Unit implausible* (Failure Type 2), and in the next iteration of the model added a planning unit that fired off different environmental conditions, namely a *Switch gun strategic* planning unit. The same logic applies to instances of Failure Type 4.

Failure Type 5 marked cases where sequential planning units, i.e., those structured specifically as a series of unit tasks meant to be carried out in a particular order, were executed out of the “proper” order, as indicated by the model. Failure Type 6 indicated all other forms of incongruity between the model and the observed behavior on screen. For example, if a section of video was annotated with 3 planning units, and further investigation indicated that section could be more parsimoniously described by a

single repeated planning unit, this might be marked as a “Misc.” failure, to be improved in the next round of modeling and annotation.

A screenshot of these annotation tracks overlaid on a frame of gameplay video footage is presented below, in Figure 4. Appendix A presents an example of the XML markup which underlies the track visualization; the contents of the SA model can be seen more readily in the XML format.

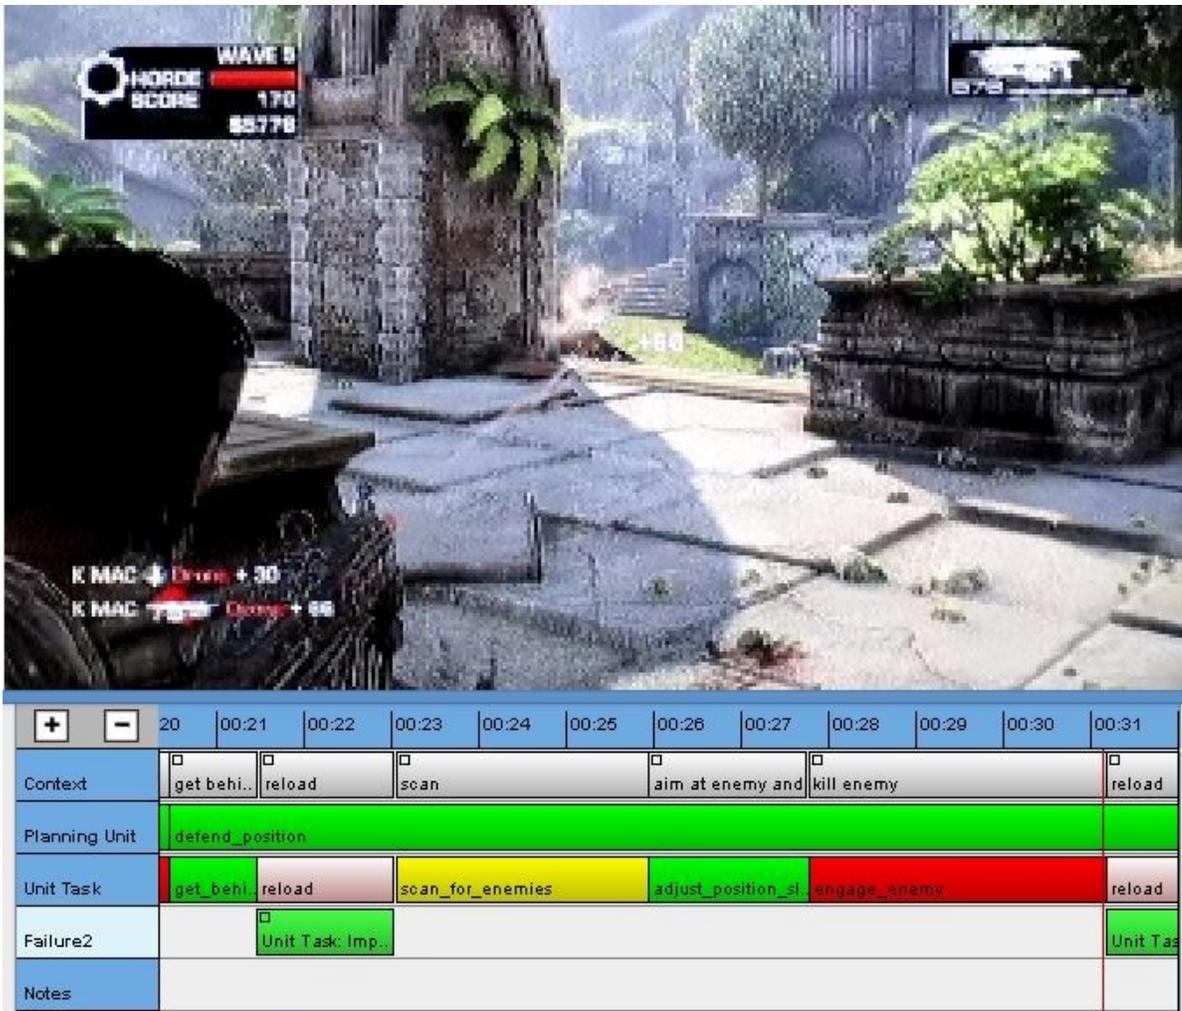


Figure 4. Screenshot of the Anvil annotation software. Frame annotated with context, planning unit, unit task, and failure tracks.

Results

Model Version 7 was used for analysis, and comprised 13 planning units and 25 unit tasks. Planning units were classified as either ordered or unordered, according to whether the constituent unit tasks were consistently executed in sequence, or were combined in different orders. An example planning unit, “Defend Position”, with its associated unit tasks and firing conditions is presented below, in Table 2

Six (6) videos were annotated in ANVIL, constituting a total 1604 seconds of audio-video footage (mean length = 267.3s; range 238-341s). Annotation data are currently hosted at https://github.com/mattmartin256/SGOMS_Annotations.

Twelve of the 13 planning units in the model were observed in the annotated video. The missing PU was “*Start of game preparations*”, which only occurred in the initial round of a game. This planning unit was used by each player in their warm-up round, but analysis of these rounds is not included here. Similarly, of the 25 unit tasks, only two were not observed in the present data set: “*Buy command post*” and “*Collect ammo opportunistic*”. The first of these UTs is also only applicable in the initial round of a game.

The video annotation procedure allow for comparison of total and relative amounts of time spent executing each of the planning units and unit tasks in the model. Table 3 presents a quantitative summary of the planning unit occurrences in the data, and Table 4 presents the same information regarding the unit task occurrences.

Table 2. Example planning unit - Defend Position

Firing conditions:

- Agent is in defensible position.
- The environment is manageable.
- Enemies are approaching.

<u>Unit task:</u>	<u>Firing conditions:</u>	<u>Behaviours:</u>
Adjust position slightly	<ul style="list-style-type: none"> • Target’s presence is known. • Target is not visible or not vulnerable. 	<ul style="list-style-type: none"> • Agent shifts position so that the target can be engaged.
Collect ammo opportunistic	<ul style="list-style-type: none"> • Agent’s stored ammo is low or medium. • Ammo is accessible, can be gathered without much risk. 	<ul style="list-style-type: none"> • Agent moves and grabs ammo.
Engage enemy	<ul style="list-style-type: none"> • Target is vulnerable. • Agent is in a suitable engagement position. 	<ul style="list-style-type: none"> • Agent shoots at target (aiming is optional).
Get behind cover	<ul style="list-style-type: none"> • Agent is not behind cover. 	<ul style="list-style-type: none"> • Agent is behind cover. • Agent is in suitable engagement position (optional outcome).
Reload	<ul style="list-style-type: none"> • Agent’s gun ammo is low. • Agent has stored ammo. 	<ul style="list-style-type: none"> • Agent reloads, resulting in full gun ammo.
Scan for enemies	<ul style="list-style-type: none"> • Agent has not located a target. 	<ul style="list-style-type: none"> • Camera movement • Target located (optional outcome).

Note. Cell elements here (e.g., “Target’s presence is known”), are simplified representations. In the model, these are described in ACT-R pseudo-code.

Table 3. Measurements of occurrence for each planning unit

Planning unit	Frequency	Percent of occurrences	Average duration (s)	Percent of video duration
Avoid explosives	9	5.3	2.6	1.4
Be aggressive	21	12.3	7.5	9.8
Between rounds preparations	8	4.7	18.6	9.3
Close quarters combat	28	16.4	5.7	10
Defend Position	53	31.0	14.5	47.9
Find defensible position calm environment	6	3.5	5.0	1.9
Find defensible position chaotic environment	17	9.9	4.5	4.8
Hunt	14	8.2	13.8	12.0
None	1	0.6	5.0	0.3
Recover health	2	1.2	3.0	0.4
Replenish supplies chaotic	3	1.8	5.0	0.9
Start of game preparations	0	0.0	0.0	0.0
Switch guns out of ammo	1	0.6	3.0	0.2
Switch guns strategic	8	4.7	1.6	0.8
Total	171	100	6.9	99.8

Note 1. The total average duration represents the average duration of any given planning unit; planning units that did not occur at all were not counted towards the average.

Note 2. The percent of video duration does not add up to 100 due to small errors in the annotated files; very small portions of the videos were not covered by the annotation.

Table 4. Measurements of occurrence for each unit task

Unit task	Frequency	Percent of occurrences	Average duration (s)	Percent of video duration
Adjust position slightly	32	6.1	2.1	4.2
Advance	17	3.2	2.2	2.3
Assume defensible position	23	4.4	3.0	4.4
Avoid explosives	9	1.7	2.7	1.5
Buy command post	0	0.0	0.0	0.0
Create distance close quarter	19	3.6	2.6	3.1
Collect ammo opportunistic	0	0.0	0.0	0.0
Collect free ammo	7	1.3	5.1	2.2
Collect expensive ammo	2	0.4	4.0	0.5
Engage enemy	165	31.4	3.7	38.2
Engage enemy close quarters	13	2.5	3.0	2.4
Execute enemy	10	1.9	4.4	2.7
Explore	4	0.8	4.5	1.1
Find suitable place to engage	22	4.2	1.1	1.4
Get behind cover	38	7.2	1.7	4.0
Locate enemy	16	3.0	7.1	7.1
None	6	1.1	5.0	1.9
Reload	77	14.6	1.7	8.0
Return to base	6	1.1	8.3	3.1
Run and scan	16	3.0	3.9	3.9
Scan for enemies	33	6.3	3.0	6.2
Switch to rifle	3	0.6	1.7	0.3
Switch to shotgun	1	0.2	2.0	0.1
Switch to rifle strategic	3	0.6	1.3	0.3
Switch to shotgun strategic	3	0.6	1.7	0.3
Wait until healthy	1	0.2	4.0	0.3
Total	526	100.00	3.3	99.5

In the analyzed footage, the frequency and duration of the individual planning units and unit tasks varied considerably. For example, the “*Defend Position*” planning unit accounted for 48.0% of the total recorded action and 31.0% of the total number of planning unit occurrences (53/171), whereas the planning unit “*Recover health*” accounted for only 0.4% of the time and 1.2% of occurrences (2/171). The same inequalities were observed between the set of unit tasks. The significance of this will be addressed in the discussion.

On 3 of the 6 videos, a failure track was implemented. The intent of this was to gauge the usefulness of explicitly annotating errors of the model, in order to indicate where future iterations of the model must be refined. In the videos which used this failure track, “failure blocks” accounted for 14.3% (range 11.2-17.2%) of the total video duration. Frequency and duration data for failure tracks is presented in Table 5.

As mentioned above, failures of the model were classified into 6 types. Of these, the most commonly observed was “*Unit Task: Implausible*”, accounting for 38.6% of failures, and of these, the most commonly observed implausible unit task was “*Reload*”. These instances were marked when the player reloaded their weapon while still having many bullets in the gun. The “*Reload*” unit task initially had a firing condition of “ammo:low”. Observation of several instances of players reloading without low ammunition indicated the necessity of accounting for these behaviors with a different unit task. Unit tasks and planning units are sensitive to specific conditions in the environment. In this case, the condition of interest was not that the player had low ammunition, but that there was a lull in the action that was judged long enough to allow

Table 5. Measurements of occurrence for each failure type

Failure type	Frequency	Percent of occurrences	Average duration (s)	Percent of video duration
Miscellaneous	4	9.1	1.0	0.5
Planning Unit: Does Not Exist	2	4.6	3.5	1.0
Planning Unit: Implausible	4	9.1	2.0	1.1
Unit Task: Does Not Exist	9	20.5	4.0	4.9
Unit Task: Implausible	17	38.6	1.7	3.9
Unit Task: Not In Order	8	18.2	2.8	3.0
Total	44	100	2.5	14.3

the player to safely reload, thereby avoiding the risk of running out of bullets during a firefight. Therefore, a future iteration of the model ought to differentiate between *“Reload – low ammo”* and *“Reload – strategic”*.

Individual differences observed during video review were reflected in data comparing relative frequencies and durations of planning units and unit tasks across the three players. The most notable difference observed concerned the amount of time spent in the *“Defend Position”* vs *“Be Aggressive”* planning units. Player 1 demonstrated a more aggressive style of play, reflected by 20.6% of their time being spent the *“Be Aggressive”* planning unit, and 37% of the time spent in the *“Defend Position”* PU. By contrast, players 2 and 3 spend 53.5% and 53.6% of their video data, respectively, in the *“Defend Position”*, PU, and only 1.2% and 8.9% of their time in the *“Be Aggressive”* planning unit. Other differences in play can be observed in Figure 5.

Annotation data can also be used to examine differences in play according to the map or environment selected, each of which has different characteristics. For example, Player 3 was recorded playing on two different maps, Checkout and Sandbar. During rounds on the Sandbar map, more time was spent in the *“Hunting”* planning unit than on the Checkout rounds. The reason for this was that the Sandbar level was laid out such that the player had to change locations more frequently to find new targets, whereas on the Checkout level, it was possible to stay roughly in the same location for the entire round. This and other differences in planning unit durations according to map selection can be seen in Figure 6.

Differences in play according to the difficulty level can also be seen in the dataset. Figure 7 depicts these differences on waves 5 and 6, with the latter being more difficult. Again, only Player 1 is represented in this data.

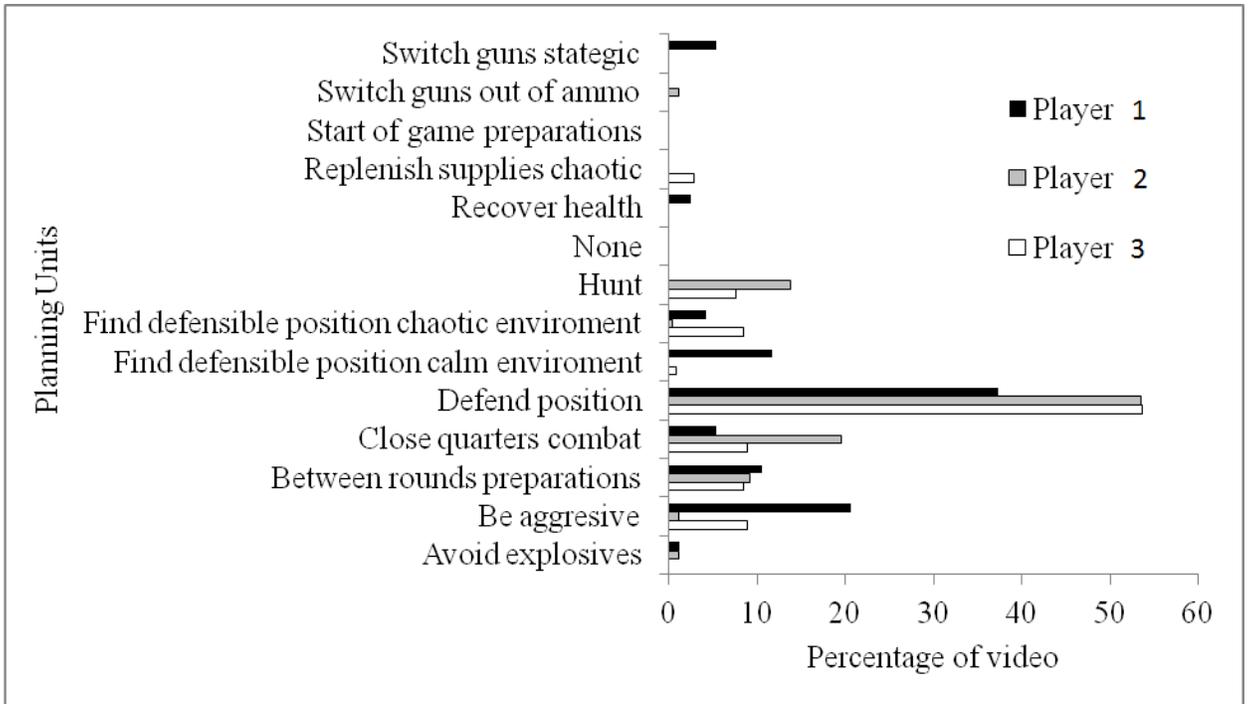


Figure 5. Comparison of planning unit use in terms of percentage of video duration between the three players for the map Checkout on Wave 5 difficulty.

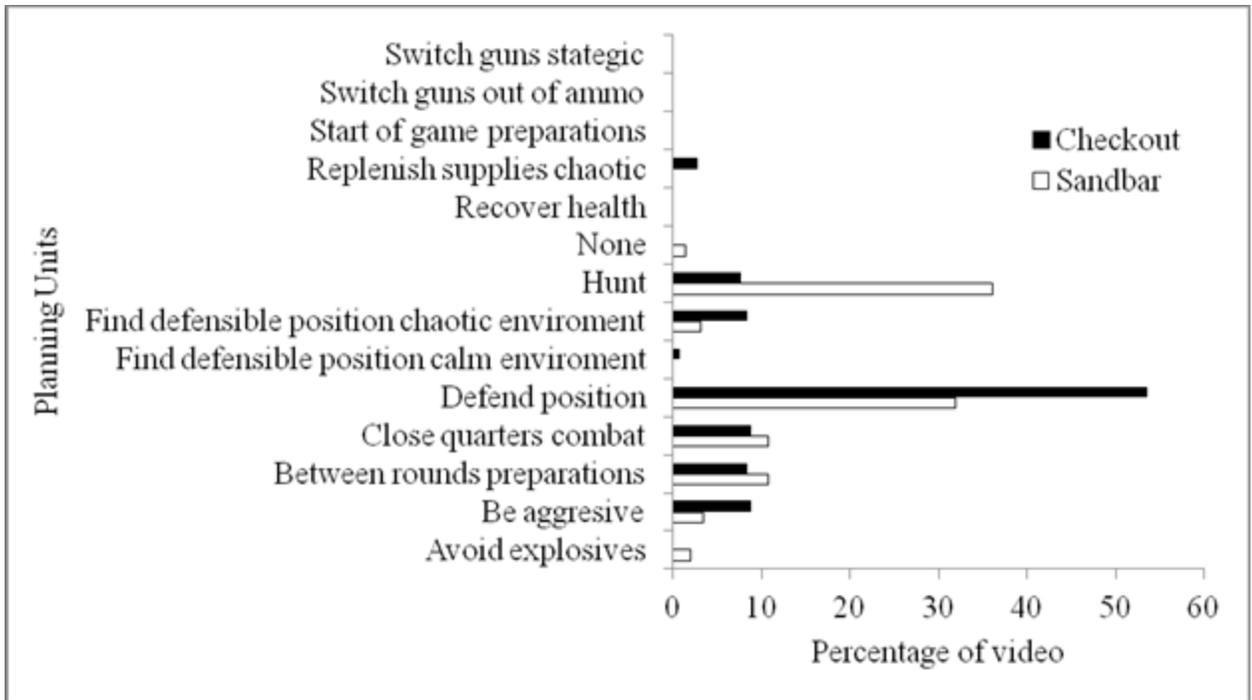


Figure 6. Comparison of planning unit use in terms of map for Player 1 at Wave 5 difficulty.

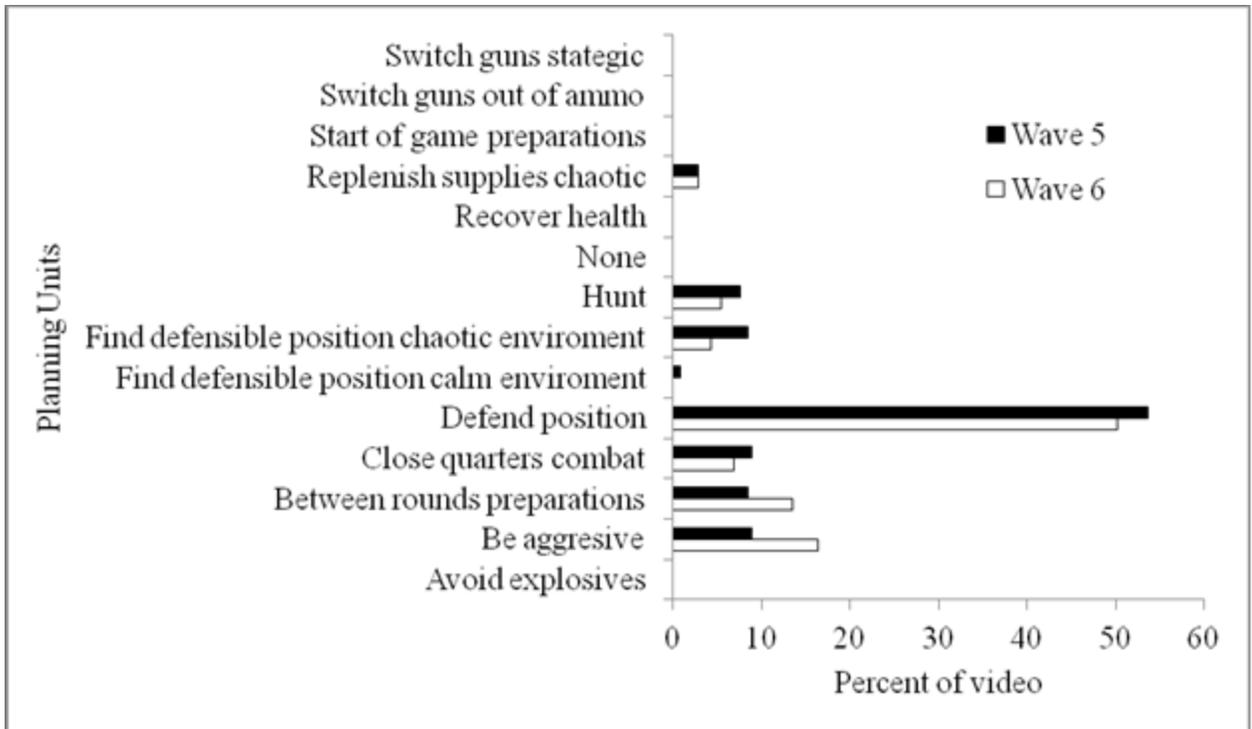


Figure 7. Comparison of planning unit use by wave (i.e., difficulty) for Player 1.

Discussion

The highly unequal distribution of frequency and duration of the unit tasks and planning units in the model highlights some interesting issues, relating both to methodology and to theory. Regarding theory, it suggests that those elements which cover the greatest portions of the data (such as “*Defend Position*”), can be considered the core element of the task. Indeed, the game structure of *Gears of War 3* is such that defending the chosen location from the advance of enemies forms a primary component of what the players are trying to achieve. On the other hand, a single planning unit covering such a wide swatch of the observed data might suggest that the modeling element used to describe and annotate that time lacks specificity. In other words, it may not just be that players spend a lot of time doing the dominant PUs and UTs, but that these elements can be further subdivided into more precise components, thereby affording the model greater specificity and predictive power. It is here that iterative analysis becomes important in model creation of this sort. If a future model (version 8) decomposes a particular planning unit into sub-components, this new model iteration can be applied to the same dataset, and differences in model failures can be quantified.

From a practical standpoint, one of the issues that became evident in the process of modeling this data is that some planning units are significantly rarer than others. This means that there may be very few data points for a given construct, such as *planning_unit:restore_health*. In these cases, it can be difficult to decide whether the planning unit is seen so infrequently because it is a poor descriptor of what agents are doing in the task environment, or instead that the planning unit is tuned to an event

that is, in fact, quite rare. One way of addressing this issue is simply to collect more data. If the event is consistently observed, even infrequently, it may suggest that retaining the relevant planning unit or unit tasks is wise. Regardless of the volume of data collected, however, this will likely not resolve the problem of whether the planning unit or unit task is appropriate to modeling the particular behavior in question. There is always the possibility that an alternatively imagined model, with more or different constructs, would better described the data. This is the *raison d'être* of the iterative modeling approach: a lone model compared against a single dataset can never be sufficient for determination of completeness or accuracy of the model. It is only through comparison of alternatives that this can be gauged.

Another general practical issue revealed by the present study is the difficulty of determining which cognitive actions/operations may be co-occurring. As an example, players would frequently flee from enemies when their ammunition was depleted. While running, they would simultaneously be searching for suitable cover, maneuvering the avatar, and initializing the reload procedure, perhaps among other things. This sort of multi-layered activity is common in this type of video game, and many other forms of expertise. Consider, as a simple example a group of musicians improvising a piece of music. Minimally, they must execute the motor actions to play or sing the appropriate notes, and at the same time process the sound of the others to maintain timing and key. This sort of “multi-tasking” has received a significant attention in psychology (Sanbonmatsu, Strayer, Medeiros-Ward, & Watson, 2013; Stoet, O’Connor, Conner, & Laws, 2013; Wood et al., 2012), and is part of the problem that the SGOMS hierarchical

structure of planning units and unit tasks is attempting to address: responsivity to environmental pressures of different forms or at different levels. It is difficult in practice, however, to account for multiple, concurrent cognitive processes. Most everyone can walk and chew gum at the same time, but it is not clear how this concurrency in general is best represented in a process model. This difficulty is particularly evident when using top down-driven cognitive architectures, such as ACT-R which *by design* must be executing a single action at any given instant (roughly speaking). This issue will be revisited in Chapter 5.

Finally, one of the more broadly interesting topics that this study has thrown up is the communication of scientific models in a larger community of researchers. The focus of this dissertation is on expertise in the real world, a significant component of which is the sharing of information. The case has been made that effectively sharing information is in fact *the key* to human expertise, that human beings individually are not particularly impressive animals, and it is only through the complex mechanisms of social learning and collaboration that we have become the most dominant and advanced species on this planet (Henrich, 2015). The importance of the social dimension of human expertise is the central theme of Chapter 4.

Scientific expertise is communal through and through, and it is thus imperative that scientific work is conducted with an aim towards sharing the products of that work with a wider community. In the case of computational cognitive models, the formula for doing this effectively is still being worked out. As computer and internet technologies mature, more options are becoming available for the creation and dissemination of such

models, and the cost of communicating data is being radically reduced, slashed to essentially zero. A portion of the work presented in this chapter, particularly in the construction of the SGOMS graphical user interface, was aimed at exploring the barriers to communication when dealing with cognitive models of this sort. While it is no longer expensive to transmit models or associated data, it is not always entirely clear how cognitive models can be used to enhance the collective discourse about cognition. How can the publication of a cognitive model be made maximally useful to the community?

To put this another way, I see a dual purpose in scientific communication. First, to pass on information in a transparent way. If a journal article reports a quantitative analysis of a particular experiment, those numbers and the results of the hypothesis test (depending on field) can be straightforwardly digested by the reader trained in the discipline. That result or those ideas can then be simply rolled into one's own work. This is, in effect, why we reference other writers. Second, and more importantly here, that same article will function in a more nebulous way, interacting with existing information and ideas in the head of a receptive reader. The value of the article is not simply to inform another researcher that some statistically significant effect was established through a particular set of stimuli presentations and recorded responses. The point is to extend the thinking and understanding of the community by (hopefully) triggering novel thoughts in the readers, which may lead to new investigations being conducted by those readers, investigations that were not envisioned by the writer. If the writer can predict what each of his/her readers will bring to the reception of the paper, everything can be anticipated and spelled out clearly ahead of time, as in the first

function. The work need be carried no further, and probably had no business being published in the first place. In the case of cognitive modeling, as well as other fields within cognitive science, it remains an open question as to how best to communicate the ideas and findings generated by a particular paradigm to maximally stimulate fruitful discussion and inquiry. Future work will explore the properties of cognitive process models that make them effective tools for thinking and communicating, and possible routes of improvement.

Chapter 4 – Expert Cooperation

Overview

This chapter addresses the following question: How do experts work effectively in groups? The question can be decomposed into three parts. First, how do individual cognitive agents operating in the real world coordinate their efforts? Second, when is coordinated, multi-agent activity more effective than individual activity or the activity of multiple independent agents, and when is it less so? And third, what methods are useful in studying this topic?

This chapter is structured as follows. An extended introduction will examine the importance of cooperative, coordinated action in the context of expertise. The latter half of the chapter will detail a study examining two data sets of cooperating video game experts. The first data set comprises video recordings of pairs of individuals playing Gears of War 3 (Epic Games, 2011), and the second set examines professional teams playing Counter Strike: Global Offensive (Hidden Path Entertainment & Valve Corporation, 2012). This study used an extension of the method presented in Chapter 3, moving from studying a solo player to studying teams of players. The chapter concludes with an examination of how these analyses can be used within a computational modeling paradigm, which will be the focus of Chapter 5.

Introduction

This chapter is based on the premise that human expertise is fundamentally social in both its genesis and exercise. Concerning the genesis of expertise: novices typically become masters through access to existing experts, or by learning from materials

created by experts, such as training manuals and university courses. An important difference between *Homo sapiens sapiens* and all other animals on the planet is our ability to systematize and institutionalize mechanisms for the transfer of knowledge and skills. We have built, among other things, nationalized grade school systems, colleges and universities with standardized curricula and accreditation procedures, craftsman guilds, encyclopaedias, and web platforms for sharing technical expertise. Our ability to store and transmit knowledge has allowed us to stand on the shoulders of giants in a way that bacteria, gnats, and leopards cannot, and this is the basis for the evolution of expertise across succeeding generations.

Concerning the exercise of expertise: once mastery is developed, experts tend to ply their trade in a greater social context. Consider the relationship of skills to professional activity. There are many forms of expertise that form the basis of a profession: e.g., medical, legal, athletic, management, computing, military, and scientific expertise. In many professional domains an expert both works with others who have similar or complimentary skill sets, on the one hand, and provides services to those who lack their specific skill set, on the other. As an example: medical doctors work day to day with nurses, technicians, specialists, etc., (complimentary skill sets), and deal with patients (lacking their skill set).

The broader context of the exercise of expertise is not limited to professional activity, however. It may also concern one's general social status (as certain forms of expertise, such as law and medicine, are afforded greater-than-average prestige in many cultures), one's role in conversation, and one's role in civic life. Concerning the last

point, in Plato's *Crito*, the dialogue set immediately preceding the death of Socrates, the discussion largely revolves around who the populace ought to consult and trust on various matters (Rosano, 2000). The party reaches the consensus, roughly, that we ought to trust the experts, and that it is the duty of the skilled to step forward when they are needed.

So both the inward-looking and outward-looking dimensions of expertise typically involve a network of other people. To misquote John Donne: no expert is an island.

A complex network of interlinked specialties characterizes the modern world, and this is largely due to the human capacity for the division of labour. The next section will examine this social phenomenon in the context of expertise. The subsequent sections addresses two supplementary aspects of expert activity that allow the division of labour to lead to complex multi-agent structures, rather than solitary, hyper-differentiated individuals: communication and organization.

Division of Labour – Importance to Expertise

Division of labour (DoL) is a key concept in sociology and economics (Durkheim, 1893; M. Weber, 1978), and is crucial to the role of expertise as it functions in the wider world outside the laboratory. It consists most simply in splitting up a task into pieces, and usually in assigning the different pieces to different individuals. This is enormously powerful in a wide range of activities: production of goods, sporting activities, scientific research, software engineering, political activity, and a host of other activities that require the efforts of multiple people.

The significance of the division of labour for a theory of society has been recognized at least since the appearance of Plato's *Republic* (Plato, 1966). The classic analysis of the principle is Adam Smith's examination of a pin manufacturing plant in his *An Inquiry into the Nature and Causes of the Wealth of Nations* (Smith, 1776), the title of which is typically shortened to *The Wealth of Nations*. This book may have marked the first occurrence of the term "the division of labour" in print. I will draw extensively from this work to set the stage for the following discussion of division of labour in the context of expert studies. Smith's illustration of DoL in the context of the pin plant is as follows:

One man draws out the wire, another straightens it, a third cuts it, a fourth points it, a fifth grinds it at the top for receiving the head; to make the head requires two or three distinct operations; to put it on is a peculiar business, to whiten the pins is another; it is even a trade by itself to put them into the paper; and the important business of making a pin is, in this manner, divided into about eighteen distinct operations, which, in some manufactories, are all performed by distinct hands, though in others the same man will sometimes perform two or three of them. (p. 8)

This is a relatively straightforward description of subdividing a complex task into smaller pieces. What is key in the context of expert studies is the notion that such division of labour leads to the development of greater expertise. The first line of Chapter 1 of *Wealth of Nations* states that "The greatest improvement in the productive powers of labour, and the greater part of the skill, dexterity, and judgment with which it is anywhere directed, or applied, seem to have been the effects of the division of labour."

This suggests that expertise is greatly facilitated by our ability to divide labour in intelligent or appropriate ways. He expands this claim in his analysis of the pin factory:

I have seen a small manufactory of this kind where ten men only were employed, and where some of them consequently performed two or three distinct operations. But though they were very poor, and therefore but indifferently accommodated with the necessary machinery, they could, when they exerted themselves, make among them about twelve pounds of pins in a day. There are in a pound upwards of four thousand pins of a middling size. Those ten persons, therefore, could make among them upwards of forty-eight thousand pins in a day. Each person, therefore, making a tenth part of forty-eight thousand pins, might be considered as making four thousand eight hundred pins in a day. *But if they had all wrought separately and independently, and without any of them having been educated to this peculiar business, they certainly could not each of them have made twenty, perhaps not one pin in a day; that is, certainly, not the two hundred and fortieth, perhaps not the four thousand eight hundredth part of what they are at present capable of performing, in consequence of a proper division and combination of their different operations.* (p. 9, emphasis added)

Importantly, this suggests that the division of labour leads to *more than additive* gains in the productive capacity of the group. In other words, the whole is greater than the sum of its parts. The three reasons that Smith suggests for these gains in productivity are that DoL allows individuals to focus on one or a few tasks while ignoring others and thereby become proficient more quickly; that they save time by not needing to

frequently switch tasks; and that workmen deeply engaged with a particular problem are more likely to invent new methods and machines to solve it. Plato offered different reasons, namely that there are natural differences in ability among individuals, that daily practice yields greater skill (Smith's first reason), and that because some tasks are highly time-sensitive, dividing labour in the proper way allows the right person for the job to always be available when necessary (i.e., availability is part of the division of labour). Modern, cognitivist hypotheses about the reasons for these gains in productivity might involve reduced mental interference or cognitive load as a result of circumscribed domains of responsibility, or might frame the problem as a function of learning and forgetting curves. A variety of approaches can be imagined. This question will be returned to in Chapter 5, in the discussion of the communication mechanism to be implemented.

One of the most important aspects of Smith's analysis, for present purposes, concerns the claim that the division of labour leads to a social order in which individuals rely on an enormously complex, distributed network of experts and professionals to meet their needs, and furthermore, that even the simplest artifacts may embody this complexity and interconnectedness. Smith describes the situation this way:

Observe the accommodation of the most common artificer or day-labourer in a civilised and thriving country, and you will perceive that *the number of people of whose industry a part, though but a small part, has been employed in procuring him this accommodation, exceeds all computation*. The woollen coat, for example, which covers the day-labourer, as coarse and rough as it may appear, is

the produce of the joint labour of a great multitude of workmen. The shepherd, the sorter of the wool, the wool-comber or carder, the dyer, the scribbler, the spinner, the weaver, the fuller, the dresser, with many others, must all join their different arts in order to complete even this homely production. How many merchants and carriers, besides, must have been employed in transporting the materials from some of those workmen to others who often live in a very distant part of the country! How much commerce and navigation in particular, how many ship-builders, sailors, sail-makers, rope-makers, must have been employed in order to bring together the different drugs made use of by the dyer, which often come from the remotest corners of the world! What a variety of labour, too, is necessary in order to produce the tools of the meanest of those workmen! To say nothing of such complicated machines as the ship of the sailor, the mill of the fuller, or even the loom of the weaver ... if we examine, I say, all these things, and consider what a variety of labour is employed about each of them, we shall be sensible that, without the assistance and co-operation of many thousands, the very meanest person in a civilised country could not be provided, even according to what we very falsely imagine the easy and simple manner in which he is commonly accommodated. (p. 12, emphasis added)

This degree of interconnectedness among the members of a society is a double-edged sword: it enables enormous gains in expertise and wealth for the group, but renders individuals increasingly dependent upon the social world because of the requisite forgoing of developing other abilities. This interdependence engenders a need for

communication about the respective roles and expectations of experts and the laity.

This is addressed further below.

Division of Labour - Genesis

How does division of labour arise? Smith claims that “the principle which gives occasion to the division of labour” is the “propensity in human nature ... to truck, barter, and exchange one thing for another” (p. 13). The notion that exchange can lead to differentiation in skill sets is the basis of the modern macro-economic concept of comparative advantage (Bernard, Redding, & Schott, 2007; Costinot, 2009; Hunt & Morgan, 1995).

Drawing heavily on Smith’s work, Emile Durkheim, the influential French sociologist, argued that the division of labour reflects a deep underlying principle of the natural world, rather than emerging from the human inclination toward trade. In his seminal work, *On the Division of Labour in Society* (Durkheim, 1893), he writes that the division of labour is “no longer a mere social institution whose roots lie in the intelligence and will of men, but a general biological phenomenon, the conditions for which must be sought in the essential properties of organized matter. The division of labour in society appears no more than a special form of this general development”, (p. 34). He also asserts that “the law of the division of labour applies to organisms as well as societies. It may even be stated that an organism occupies the more exalted a place in the animal hierarchy the more specialized its functions are” (p. 34). Durkheim is not arguing simply that the differentiation of functions within an organism, say between the functions of the heart and the brain, or between the functions of the nucleus and

mitochondria within a cell, are analogous to the division of labour at the social level, as exemplified by the variety of professions. He is instead making the stronger claim that some unseen principle of the natural world is expressed through the differentiation of biological functions within individual organisms, and that the same principle is expressed at a higher level of complexity or structural organization when we examine the division of labour in society. Various scientific approaches have been employed in trying to better understand what this property of the natural world might be, some of which are described below.

The ethological study of eusocial insects, such as honeybees (Siegel, Kaftanoglu, Fondrk, Smith, & Page, 2012; Toth & Robinson, 2005), wasps (Chandrashekara & Gadagkar, 1991; Judd, 2000), and leaf-cutter ants (Julian & Fewell, 2004; Orr, 1992; Wilson, 1980), has revealed some surprisingly sophisticated forms of differentiated labour in insect populations. In bee hives, for example, there are some bees who go out to scout for appropriate flowers from which to collect pollen and then communicate that information to other bees, some who focus on trucking the pollen back to the hive, some who stay within the hive to clear away debris and deceased bees, and others who work to place larva in the appropriate combs within the hive (Jeanson et al., 2005; Toth & Robinson, 2005). Many other species of animal also exhibit division of labour. It is seen in hunting behavior by hawks (Bednarz, 1988), lions (Stander, 1992), dolphins (Gazda, Connor, Edgar, & Cox, 2005), and others. Perhaps the most commonplace division of labour is the differentiated sexual behaviors among sexually reproducing animals.

While ethological studies of DoL have collected substantial data about what these divisions look like in the wild, they have not yielded an answer as to whether such a general mechanism exists as that postulated by Durkheim. To address this, computational biology has created models of differentiated labour that test various hypotheses about what characteristics within individuals might lead to specialization. One influential model has used the notion of “response thresholds” to account for DoL as a function of individual variability in the preference for completing certain task (Duarte, Pen, Keller, & Weissing, 2012). Response thresholds refer to the level of some stimulus which will cause an organism to act to modify the level of that stimulus. The hypothesis is that each individual has a profile of stimulus-response thresholds, at which it will act to remove/reduce/increase a particular stimulus if its threshold is exceeded. For example, in the case of honeybees, there would be some individuals who will engage in hive cleaning behaviors when they spot n deceased bees, whereas other individuals may have a response threshold of $2n$ dead bees; the first group will therefore be more likely to engage in cleaning behavior than the second group, and will seem to show a relative preference for cleaning. The interaction of various environmental factors and response thresholds will thus lead (by design in these models) to a differentiated labour landscape. Though the characterization of “response thresholds” remains fairly primitive and abstract compared to what would be required to account for the division of labour as seen in human societies, it does provide some conceptual traction into the generalized problem. This approach is complemented by research in neuroscience demonstrating that neurons exhibit variability in their likelihood of firing in response to

different stimuli (Ringach, Bredfeldt, Shapley, & Hawken, 2002). This a well-established feature in the visual cortex, for example, where some neurons respond selectively to vertical lines, others to horizontal lines, and others to lines at various orientations (Ferster & Miller, 2000). Neural selectivity of this sort may provide the mechanism by which such response thresholds are realized, but accounting for more complex forms of preference, such as the desire to pursue engineering or medicine, would obviously require a substantially expanded account.

Another approach to modeling this phenomenon is epitomized by Herbert Simon's work on complex, hierarchical systems. In his book *The Architecture of Complexity: Hierarchic Systems* (Simon, 1996), Herbert Simon offers a computational-thermodynamic view of how such differentiation might arise in the natural world through processes of biological evolution. Specifically, he examines the problem of intermediate, stable structures in evolution, which he argues must be necessary for complex entities to arise in the face of a chaotic world characterized by the Second Law of Thermodynamics, whereby closed systems ultimately move toward disorder. He offers this allegorical description of the process, in his now-famous two-watchmaker parable:

There once were two watchmakers, named Hora and Tempus, who manufactured very fine watches. Both of them were highly regarded, and the phones in their workshops rang frequently -new customers were constantly calling them. However, Hora prospered, while Tempus became poorer and poorer and finally lost his shop. What was the reason?

The watches the men made consisted of about 1,000 parts each. Tempus had so constructed his that if he had one partly assembled and had to put it down – to answer the phone say – it immediately fell to pieces and had to be reassembled from the elements. The better the customers liked his watches, the more they phoned him, the more difficult it became for him to find enough uninterrupted time to finish a watch.

The watches that Hora made were no less complex than those of Tempus. But he had designed them so that he could put together subassemblies of about ten elements each. Ten of these subassemblies, again, could be put together into a larger subassembly; and a system of ten of the latter subassemblies constituted the whole watch. Hence, when Hora had to put down a partly assembled watch in order to answer the phone, he lost only a small part of his work, and he assembled his watches in only a fraction of the manhours it took Tempus”

(Simon, 1996, p. 188)⁶

Simon’s approach comes closer than computational biology models have in providing a fully generalizable treatment of how differentiation might occur in the biological world, though it is unclear how to apply the mathematical techniques he draws from thermodynamics and evolution to the study of human social systems.

⁶ Interestingly, in a footnote in Adam Smith’s *Wealth of Nations*, editor Edwin Cannan also uses watchmaking to illustrate the notion of divided labour. Cannan reproduces, from Mandeville’s *Fable of the Bees* (1729), the following quotation: “watchmaking ... is come to a higher degree of perfection than it would have been arrived at yet, if the whole had always remained the employment of one person; and I am persuaded that even the plenty we have of clocks and watches, as well as the exactness and beauty they may be made of, are chiefly owing to the division that has been made of that art into many branches”, (p. 8).

A third approach is presented by John Holland in his *Signals and Boundaries* (2012), which is in part an extension of Simon's work. Holland's focus is on the mechanisms by which sub-assemblies combine and coordinate. In his system, boundaries are those structures that limit interaction with the external world (such as cell membranes and legal certification) in order to stabilize the internal structure or limit entropy-increasing forces from the outside; and signals are the mechanisms that allow information to pass through those boundaries, between assemblies or sub-components of those assemblies (such as protein channels in a cell membrane or protocol for legal communications between counsel and judge). What Holland's analysis makes clear is that boundaries and signals are powerfully complementary. Boundaries allow for a system to maintain internal stability, while signals across those boundaries allow the system to be integrated into a larger, more complex system. Crucially, the nature of the signals are shaped by the characteristics of the boundary. This is important in understanding how expert communities both differentiate themselves from the broader world, and then set up channels to interface with that world. This notion of signals and boundaries is highly relevant for discussions of "indicators of expertise", as discussed later in the chapter.

To wrap up this section: the "unseen principle" of differentiation that Durkheim examines points at a deeply mysterious property of the universe: why is there difference instead of sameness? It is the extension of the ancient question: why is there something rather than nothing? These questions are in the realm of metaphysics, and for practical purposes beyond the scope of this work. I mention them only because I believe that,

whether Durkheim was correct or not in asserting that the division of labour represents a general principle of organized matter, attempting to link the targeted questions of one's research to broader questions is valuable in a field as fragmented as the cognitive sciences. Whether such links turn out to be valid or not is perhaps less important than the act of setting up the localized research program to act as a center of gravity that can attract elements from other domains, revealing connections between fields that would be invisible under the rule of myopic pragmatism.

Communication

In order for advanced states of specialization to be realized in a particular group, there must exist mechanisms by which goods and services produced through various forms of expertise and division of labour can be circulated through the community and combined into larger units of activity. This section and the next address two such mechanisms: expert communication and organization.

A great deal of what experts do involves communication. This is the principal role of teachers of various sorts (as the most direct and obvious example), but communication is also centrally important in a variety of other forms of expertise that involve the provision of advice: lawyers, doctors, interior designers, career coaches, marital counselors, bankers, mechanics, political advisers, etc.

The internally-facing function of expert communication (i.e., between experts within the same expert group) is structured in part as a highly precise language, often called jargon (Hirst, 2003). Having unique or low-frequency words to refer to important concepts in a field allows experts to quickly and precisely communicate complex

concepts that might otherwise require circuitous and unwieldy constructions. A word like “transference” in the context of a discussion about Freudian analysis conveys an entire suite of processes and emotional responses (Levy & Scala, 2012; Zetzel, 1956). Similar complexity is conveyed by constructions such as anarcho-syndicalism (Henderson, 1997), Odonata (Corbet, 1980), saltatory conduction (Neishabouri & Faisal, 2014), and collateralized debt obligation (J. R. Mason & Rosner, 2007). The character of such language is familiar to scholars of all sorts, and is one of the principal tools at the disposal of knowledge workers and experts of many kinds.

Jargon, or field-specific language, is not the only communicative tool that experts use. One important supplement to expert language is the vehicle that will carry the expert language. Scientific journal articles are a prime example. The structure of a psychology article, for example, is relatively standardized, which allows a trained reader to more quickly find what they are interested in, whether that be methodological details, statistical results, or discussions about the wider implications of the paradigm. This regularized structure expedites the process of information extraction. An index in a reference book serves a similar function in organizing a complex body of knowledge. Another interesting example, one that quite explicitly links experts to the public, is the Yellow Pages service (Jackson & Parasuraman, 1986), which provides a list of the experts in a particular region, along with their contact information, sorted by trade or skill set. Today, internet platforms and forums offer highly precise and multi-dimensional systems for classifying questions and answers: on www.stackexchange.com, an internet help forum for computer programmers, problems and solutions can be classified according to

the computer language being used (e.g., Python, C++, Java), the conceptual area (e.g., machine learning, natural language processing, graphics), the poster of the question or solution, the date, and so on. Such meta-information or organizational strategies are a necessary part of providing access to complicated bodies of information and skill sets, and these are tightly linked to the medium in question. As Marshall McLuhan famously wrote, “the medium is the message” (McLuhan, 1964).

The externally-facing function of expert communication involves something of a reversal of jargonized communication. When experts provide advice or services to non-experts (that is, non-experts in the field in question), they often must make intelligible complicated issues for which their patient, customer, or client lacks the training to fully understand. If a doctor offers a patient a choice between two possible drug treatments, it is typically of little value for the doctor to relay to the patient the differences in pharmacological mechanism of action, whereas the success rate, cost, and side-effect profile must be communicated in such a way that the patient can be said to have granted informed consent. This ability to simplify and communicate with lay people allows experts act as an extended mind for the average person, offering a simplified interface with a complex pocket of the world. This is key in allowing the individual to access the collective set of abilities possessed by the community as a whole, and it is largely accomplished through the various media we have at our disposal (Knapp, 1999).

Note that the terms “internally-facing” and “externally-facing” are terms of convenience that are sensitive to scale or specificity of domain. In the above paragraphs, the distinction between internally and externally facing was made with the

example of medical doctors communicating with other doctors versus communicating with patients. This is, of course, a simplified description of a complex communicative web. The “in-group” could be all medical doctors and nurses, internal medicine physicians, virologists, or the Ebola virus disease specialists at a particular clinic. It is a testament to the sophistication of the human mind that we are sensitive to these gradations in shared expertise and are able to fluidly adapt our discourse to the perceived capacities of our interlocutors (Van Dijk, 2006). There is evidence that even children as young as 5 years old are sensitive to cognitive divisions of labour, recognizing that different types of experts are more likely to have certain types of information (Keil et al., 2008).

This aspect of expertise also raises some interesting questions about the possibility of a universal expert architecture (UAE), that is, something underlying expertise which is akin to Chomsky’s universal grammar (Cook & Newson, 1996). Perhaps humans are able to collaborate so effectively because we share a common “expert system”, and we understand at some (probably unconscious) level how to tap into that system in other people when we are communicating. That system may be responsible for organizing new bodies of information and skills in particular ways as we learn about the world. To push the analogy a bit further: perhaps the notion of the poverty of the stimulus can be extended to reflect on our ability to extract expertise from relatively limited samples of the world.

A great deal of work in social cognition has attempted to account for our ability as humans to extract information from other individuals, whether through conversation,

observation, or instruction (James R. Anderson, 1998; Bandura, 1971; J. S. Brown, Collins, & Duguid, 1989; Frith & Frith, 2012; Frith, 2008; Higgins, 2000). While other animals, such as chimps, demonstrate a modicum of this ability (Vogel, 2009), humans display incredible facility in this domain. If it is true that we are particularly well-tuned to learning from others (Henrich, 2015), this process would be greatly facilitated by some degree of standardization in the way that complex skills are encoded in the brain. A shared architecture of expertise, a common expert system, might offer part of the solution to why we are so good at teaching one another: we know intuitively how to package information such that it will be “un-packable” by and useful to a recipient. If such an architecture exists, it is certainly not perfect, as anyone who has attempted to teach can attest to. But it might be a powerful bootstrapping mechanism to get things off the ground. The SGOMS macro-architecture hypothesis is built upon this idea; we will return to it in Chapter 5.

In addition to the language used to communicate with the laity, experts also rely on a variety of non-linguistic indicators of expertise, such as titles and badges. The most rigorous and widely acknowledged of these indicators become official accreditations, including licenses to practice medicine and law, doctoral degrees, military rankings, and FIDE designations in chess. These are an important interface in communicating competence or reliability to consumers, clients, or patients, thereby helping to circulate expertise through a social group and reducing the cognitive demand on individuals attempting to access the appropriate experts.

An interesting phenomenon concerning these indicators in contemporary,

networked society is that there has been an explosion in the variety of such signals, due in large part to the reduced price of content production and distribution afforded by the internet and other digital technologies. Today, if an individual is interested in learning about diet or exercise, the current state of string theory, or the platform of a political candidate, the internet provides a dizzying array of sources that an inquisitive person may use to educate themselves. The various websites, forums, and blogs, will contain content spanning the entire gamut of quality and veridicality, including everything from op-eds written by professional newsprint journalists, to peer-reviewed journal articles written by Nobel laureates, to the ignorant rantings of a 10 year old boy in a basement. Frequently the source of the information will be unclear. Many online platforms have therefore implemented systems that indicate the quality of a post or reliability of the user. The computer programming forum, www.stackexchange.com, for example, uses up-votes and down-votes to help users quickly find the questions and answers voted most relevant and helpful by the rest of the community, while many other sites, such as www.digg.com, display a “shares” count, indicating some measure of positive valuation by the community.

This proliferation of sources has certainly made information more readily available, but it has also made more difficult the task of sorting the good from the bad, the useful from the useless, and this may be placing new cognitive demands on individuals. Traditional indicators of expertise function, I would argue, to simplify for individuals the process of accessing expertise in a complex social environment, characterized by extensive specialization. Those indicators serve to mark who is skilled

and trustworthy. But today's information landscape has engendered a host of novel indicators that require new strategies for adjudicating between valuable and worthless information, placing greater demand on users to be savvy to the source of the content. What does it mean if an individual has posted 363 messages to a fitness web forum and has a "rep power" of 2200 (at www.bodybuilding.com), or a comment on a physics lecture posted on www.youtube.com has 250 "thumbs-up"? It is no longer so clear who the "experts" are.

All of these communicative functions are nested within a larger narrative structure, positioning the experts and public with respect to one another according to expectations and abilities. This is the topic of the next subsection.

Caveat: Trust and Narratives of Professionalism

In a social world marked by extensive division of labour and interdependence, trust is a crucial currency. When lacking the ability to solve a particular problem, people frequently turn to others, as this is often the most expedient and effective way of resolving a problem in such societies. These 'others' are often strangers, or individuals who are known to the customer/client only in a professional or functional capacity. Because the relationships between experts and the public are typically impersonal, transiently established to satisfy some need on the part of the customer/consumer, and because the function being performed is often a highly delicate one, trust on the part of the individual seeking help is paramount. People must believe that the doctor has the knowledge and ability to heal them, but also trust that the doctor has their best interests at heart, rather than, say, intending to profit from prescribing an extended and

unnecessarily expensive course of treatment. People must trust that the priest will save their soul, that the therapist will respect their privacy, and that the mechanic will not rip them off too badly.

These beliefs and expectations are delivered largely in the form of narratives of expertise circulated by a social group (Knapp, 1999). I use the term narrative to convey two key aspects of the relationship we have to experts. The first is captured by the literary theorist Tzvetan Todorov's (Todorov, 1971) minimalist model of narrative structure: a narrative describes the transition from one equilibrium state to another, where the transition is occasioned by some event disturbing the initial equilibrium. In the case of interfacing with experts, this notion of narrative captures the idea that a person or group will often seek out expert assistance when they desire that some current state of affairs shift to some more desirable state, and believe that the expert is able to aid in that transition (e.g., you are sick and wish to be well again). The second, complimentary reason that the notion of narrative is used to frame this relationship to experts is that a great deal of narrative analysis has concerned prototypical or archetypal character roles, and the impersonal, functional "expert professional" plays precisely the role of this character-in-outline. Enumerating the range and roles of character types is an activity that figures prominently in the work of many scholars of narrative. Much of this work has been informed and inspired by analysts such as Vladimir Propp (Propp, 2012), whose analysis of the Russian folktale has had an enormous impact on modern literary studies; Joseph Campbell, whose model of the hero's journey drawn from mythology has been influential in both anthropology and psychology (Campbell, 1949);

and Carl Jung, whose work on archetypes has been an important contribution to the psychological theory of narrative (Jung, 1980).

To recast the point about narratives of expertise in simpler terms: we tell stories about who to approach in particular circumstances. For example: if your car breaks down, you should go see a mechanic, not a professional scuba diver or art restoration specialist, because the mechanic knows how to get it running again. This is a trivial example, of course, but not all such decisions are quite so obvious. Consider the problem of attempting to address debilitating emotional issues. Should one visit a priest, a shaman, a psychotherapist, a friend, a marriage counselor, or some other professional? Perhaps the problem is best solved without consulting anyone. The question is already enormously complicated, and without some guidance from the culture by which to narrow down the options for consideration, it would be even more difficult. The narratives provide a starting point for contextualizing experts in our lives, for informing individuals about who they should trust to help them with particular problems in their lives.

One troubling consequence of the role trust plays in mediating the relationship between experts and the laity is a growing feeling of dependence upon experts, and an associated feeling of helplessness. A number of recent books have addressed this tension as an important factor shaping public consciousness and decision making. The titles of these books are revealing: *“Wrong: why experts keep failing us -- And how to know when not to trust them”* (Freedman, 2010); *“Trust us, we're experts: How industry manipulates science and gambles with Your future”* (Rampton & Stauber, 2001); *Future*

babble: Why expert predictions fail - and why we believe them Anyway (Gardner, 2011).

If people feel, legitimately or not, that some important part of their life must be left in the hands of strangers or functionaries, a degree of discomfort is understandable, particularly if the expertise being relied on is so incomprehensible as to appear mystical, thereby raising the concern that the “expert” may be a con-artist.

One way of accounting for this growing unease is by reference to a recent shift in the macro-level narrative of expertise shaping the interface between experts and publics. Brint (Brint, 1994) observes a shift in the professional landscape in America during the mid-20th away from traditional “social trustee professionalism” toward “expert professionalism”. The former denotes a professional status largely constituted by the role in which one is expected to serve for the benefit of broader society, whereas the latter indicates that the key component of one’s professional status is the set of abilities that one can offer. If Brint’s analysis is correct, this would indicate the source of the unease: social trust is no longer constitutive of professional identity to the same degree that it formerly was, and experts have become faceless service providers.

Arguably, the issue runs much deeper, and is symptomatic of an economic, political, and social philosophy that centralizes and reifies selfishness and competition: survival of the fittest writ large. In such a world, trust becomes a de-personalized social imperative, to be instituted through regulatory and oversight committees, constantly on guard against the experts who will exploit their arcane knowledge for personal gain at

the expense of the community, ignorant and hapless⁷. A staggeringly complex infrastructure exists in modern society to combat this, to “keep us honest”, including court systems, traffic laws, accreditation boards, auditing bodies like the Canadian Revenue Agency and the Internal Revenue Service (IRS) in America, restaurant inspection bodies, and so on. Loosely speaking, these are like the auto-immune system of the social organism, protecting us from ourselves and from organizations that attempt to seize excessive power and exert monopoly control.

One particularly germane instantiation of this function is the scientific peer-review process, striving to maintain high standards of rigor and integrity in the publication of academic work, and serving as an indicator that the published material can be trusted. As with all such mechanisms, however, the process of peer review is subject to occasional failures and manipulations, which was colourfully demonstrated by the so-called “Sokal Affair” or “Sokal Hoax” (Sokal, 1998).

In 1996, Alan Sokal, a professor of physics at New York University and University College London, submitted an article to *Social Text*, an academic journal of postmodernist cultural studies. The article parodied the style of writing that Sokal saw as becoming more prevalent and problematic in humanities scholarship at the time, a style informed by deconstructionist or post-modernist writers such as Jacques Derrida and Jacques Lacan (Sokal, 1996b). On the day of the article’s publication, Sokal revealed it to be a parody, and stated that he published the article because he was concerned

⁷Incidentally, some analysts argue that it was this sort of nefarious obfuscation that led to the global financial collapse of 2007-2008. They argue that investment banks exploited the power of obscurity by building horrifically complex financial products that were designed to be incomprehensible ((Lewis, 2010; Posner, 2011)).

with “the proliferation, not just of nonsense and sloppy thinking per se, but of a particular kind of nonsense and sloppy thinking: one that denies the existence of objective realities, or (when challenged) admits their existence but downplays their practical relevance.” (Sokal, 1996, p. 4). More importantly for our purposes here, the experiment was an attempt to see whether this journal would “publish an article liberally salted with nonsense if (a) it sounded good and (b) it flattered the editors' ideological preconceptions”, (ibid, p. 2). In other words, Sokal was attempting to manipulate the review and publication process by mimicking a style of writing that he found harmful to the communication of ideas. The superficial features of the expert discourse in question were, according to Sokal, more important in the decision to publish the article than were the thoughts that were ostensibly being communicated by the article, which, by his own admission, contains within it nothing “...resembling a logical sequence of thought; one finds only citations of authority, play on words, strained analogies, and bald assertions.” (ibid, p.3). The article was accepted for publication, triggering a series of publications addressing the implications for the scientific and academic communities (Bouilloud, 2003; Feldman, 1999; Hilgartner, 1997; Slack & Semati, 1997).

There is a great deal of work being done on the issues engendered by highly specialized languages and their use by particular communities, much of it sharing with the Sokal Affair a focus on academic work. Much of this conversation is conducted under the banners of the “science wars” (Berube, 2011; Fuller, 2001; Latour & Noor, 2002; Rubin, 1998), the sociology of scientific knowledge (Bloor, 2004; Collins, 2001;

Shapin, 1995), the strong programme in the sociology of knowledge (Bloor, 1991; Haddock, 2004; Kochan, 2010), and discourse communities (Duff, 2010; Howorth, 2004; Imm & Stylianou, 2012; Little, Jordens, & Sayers, 2003). While expert communication may not be unilaterally positive, it is undeniably powerful, and I argue that it is important to try to understand its role in the broader contexts of expert activity.

Organization

Complimentary to communicative functions are the patterns of organization which facilitate group productivity beyond what a group of individuals working independently could achieve, as described in Smith's quotation above in the discussion of the division of labour. The theoretical question of interest is: what are the costs and benefits of experts acting cooperatively?

Consider the following 'dueling parables': 'many hands make light work', but 'too many cooks spoil the broth'. These are familiar idioms in the English language that seem to express two contradictory views of cooperation. Is one true and the other false? Are both false? Or is the truth more complicated?

The most straightforward way to resolve this tension is to assert that the former parable is true in some instances, while the latter is true in others. There are trade-offs in co-operative action: two heads are sometimes better than one, but not always. For example: it is often asserted that the comparative advantage of technology start-ups over existing, large companies, is that small organizations tend to be fast and agile, fluidly adapting their plans whenever necessary, while their lumbering counterparts in the corporate world struggle through mazes of red-tape and focus groups, share-holder

optics, and brand coherence (Ries, 2011; Thiel & Masters, 2014). On the other hand, large companies often have at their disposal large reserves of cash, deep talent pools, and extensive networks of partners and clients that allow the undertaking of ambitious projects that could not be accomplished by small organizations. It is David-vs-Goliath, Velociraptor-vs-Tyrannosaurus.

There are costs associated with communication, with standard protocols, best practices, and performance reviews, but there are also enormous gains to be had from arranging a group correctly, dividing the labour in an efficient way, and providing top-down guidance. A simple formulation of this relationship is that if the cost to the group of coordinating is lower than the cost incurred to the group of not coordinating (i.e., of individuals acting locally), the coordinated path ought to be preferred. Take as an example the decision to call a group meeting for a team of software engineers, product designers, and managers. These meetings often function as a way to quickly share information between members of the group, to allow direct discussion and clarification of unclear points, to negotiate future plans, etc. If the meeting involves 6 people, and lasts one hour, it has a cost of roughly 6 work hours, and therefore can be considered worthwhile if it saves at least 6 work hours distributed through the group. This calculation is, of course, schematic, as such a calculation in a real-world setting would be complicated by the fact that most teams of this nature are nested in larger organizations, and thus the cost/benefit hours must consider effects far downstream (e.g., the “6 hours” that are saved might show up in the form of a simpler marketing campaign 3 months later). This notion of costs/benefits of communication are revisited in Chapter

5, in discussing the cognitive model of cooperating experts.

As with expert communication, it is important to note the potential problems with expert organizations. Among the most recognizable are organizational bloat and monopoly. In the first case, we see losses in productivity or efficiency when organizations become too large and processes become cumbersome, when the costs of coordinating become greater than the benefits derived, as in the example above. This is the familiar problem of “red-tape”, a common symptom in universities, governments, and large corporations (Bozeman, 1993; Brewer & Walker, 2010; Pandey & Scott, 2002). The Canadian federal government, for example, lists on its website (<https://www.canada.ca/en/government/dept.html>) 214 different departments and agencies, and in 2015 employed 257,034 people (“Government of Canada - Population of the federal public service,” 2016). Sharing resources and information in a network this complicated may allow for enormous leverage, but it is also liable to introduce some delays in decision making and inefficiencies in operations. This issue is also relevant in business. Former CEO of Hewlett-Packard Company (HP) Lew Platt has famously said that “if only HP knew what HP knows, we’d be three times more productive” (Sieloff, 1999, p. 1).

The problem of monopoly is nearly the inverse of the productivity loss resulting from bloated organizational infrastructure: instead of internally slowing a team down, monopoly stifles outsiders⁸. Anti-trust laws in the United States of America, for

⁸Or so goes the theory in classical economics (Leismuller & Schimpf, 2012). The case can be made that localized monopoly might spur innovation indirectly (Thiel & Masters, 2014).

example, reflect the belief that monopoly is unhealthy for society. Whether this is true or not has been debated in political and economic philosophy for centuries (Boldrin & Levine, 2009; Bresnahan & Reiss, 1990; Ellingsen, 1992; Jonna, 2015; Mussa & Rosen, 1978; Skogh, 1976; A. M. Spence, 1975), so I will not offer a perspective on the merits or dangers of monopoly. The purpose of mentioning it is to highlight the fact that experts working in groups often find ways to acquire social power by seizing control of some aspect of collective living, such as medicine, law, or finance. This is often accomplished with the help of expert narratives to shape public opinion about the expert group in question, but can also be effected through patents, legal strategies, and other tactics. In an influential book by Andrew Abbot entitled *The System of Professions* (1988), he discusses the process by which professional groups often acquire monopoly over provision of some service by controlling the “abstractions” (concepts) used to characterize or operate within a domain. This is important in shaping conversation between members of the expert community in question, such as the role of jargon in indicating in-group/out-group status and facilitating collaboration, and in the training (indoctrination) of new recruits. But most importantly, these abstractions form the crucial interface between the expert group and the public: because the lay person lacks command of the relevant abstractions and thus does not understand how to go about solving their problem (or even how to think clearly about it), they must go to the masters who possess that capacity, seeking legal advice, medical treatment, spiritual guidance, or automotive assistance, and in order to do so, they must be taught where to look. This is the application of Abbott’s framework to the point made above concerning

narratives: the relationship between the public and its experts is conditioned and contextualized by narratives of expertise which in turn are shaped by the abstractions at the center of story (Knapp, 1999) .

We turn next to an examination of how these three factors – division of labour, expert communication, and group organization – have been studied under the banner of distributed cognition.

Distributed Cognition and the Division of Cognitive Labour

As discussed above, the division of labour is a key theoretical construct in sociology and economics. It has received considerably less attention in the cognitive sciences, but some excellent work has been done on the topic, particularly in the field of distributed cognition. Hutchins (2000) identifies two lines of thought that bear upon the issue, one which treats a division at the social level, and one which treats a division at the neural/cognitive level. In the former, the phenomenon of interest is how multiple cognitive agents cooperate and coordinate to accomplish complex cognitive tasks. In the latter, the concern is with how a task is accomplished by a cognitive system with various specialized sub-components. In other words, division runs in two directions: starting with the individual cognitive agent, we can draw lines of division radiating outward or moving inward. Hutchins (2000) finds this symmetry reflected neatly in the title of two important books bearing upon the topic that were published contemporaneously: Lev Vygotsky's *Mind in Society* (published in 1978 in English) and Marvin Minsky's *Society of Mind* (Minsky, 1986). At the same time, parallel distributed processing was again becoming a fashionable mode of investigation within cognitive science (Rumelhart,

McClelland, & Williams, 1986).

The question of how to get a theoretical and methodological handle on these divisions between social units and neural units has been a central concern in work conducted under the banner of distributed cognition. This research has aimed, in part, at understanding how cognitive processes can be distributed across agents, tools, and time (Hutchins, 2000; Michaelian & Sutton, 2013; Rogers & Ellis, 1994; Rogers, 2006; Sutton et al., 2010; R. Williams, 2013). It takes as one of its starting points the notion that many cognitive functions are augmented by technologies and social organizations. Memory, for example, is enhanced by tools that allow people to store information, arithmetic abilities are enhanced by calculators, and the scientific enterprise is enhanced by the functioning of universities that train new scientists and provide space (physical and cultural) for research.

One of the key challenges of this work concerns demarcating units of analysis that are tractable and informative. For example, in the case of memory enhancement mentioned above, what exactly is being enhanced? If we view the cognitive agent and their tools as a couple system, perhaps we can speak of the agent's memory being improved (A. Clark & Chalmers, 1998). Alternatively, if we see only the agent, it may be more sensible to note that the agent is accomplishing a different task if they are using the internet to retrieve a piece of information than they are by recalling it from their memory. Another view of the debate, others have argued that memory is fundamentally a social phenomenon, existing to the extent that we are enmeshed in community, and that to draw the bounds of around the individual agent is to miss the

point of memory entirely (Echterhoff & Hirst, 2009; Olick & Robbins, 1998).

Herbert Simon famously offered his “ant on the beach” metaphor, to illustrate the importance of selecting a unit of analysis (Simon, 1996b). The parable states that if you watch an ant walking along the sand of a beach, it may move in complex patterns, making circles, turns, reversals, and so on. This might suggest to an observer that the ant is engaged in some complex process of navigation, with sophisticated heuristics for dealing with various aspects of the environment. In reality, the ant may be doing no more than avoiding hills, and thus the complexity seen in the movement of the ant from the conjunction of agent and environment. Watching the ant move may indicate more about the features of the environment than about the internal operations of the ant, and attempting to understand the ant without reference to the environment is consequently sub-optimal, if not impossible.

The communication analyses presented later in this chapter assume that understanding how an agent is coupled to a particular environment is crucial to understanding the cognitive activities underlying behavior. In particular, the analyses are focused on how agents working in concert are, in some cases, better able to deal with complex environments than they would be acting independently. This assertion can probably be taken as obviously true, and thus the intent of the analyses is not to provide support for the assertion. They are aimed, rather, at describing the advantages accrued to such cooperatives in a form amenable to inclusion within cognitive process models. The effort relies on the ontology of the SGOMS macro-cognitive architecture to specify the function of particular communicative acts. In particular, the concern is with

analyzing local versus global reactivity, and agents' capacity to respond to events in one form or another largely as a result of sharing planning units. More about this in the next section.

The method of analysis draws upon that employed in distributed cognition, which itself is largely inspired by the methods of anthropology (Bender et al., 2010; Hutchins, 2000). The philosophy of this method rests on two principles. First, the notion that observation of "natural behaviors", i.e., *not* those which are artificially induced by experimenters, is an important step in deriving the ontology being used by the agents in question. Second, that it is impossible to move from observation to analysis without some framework, and therefore that being explicit about one's assumptions is important. These two principles, taken together, shift the burden of hypothesis generation from pre-experimental design to data analysis. This is a common approach in various forms of fieldwork, such as that conducted by anthropologists, ecologists, and historians (Fontana & Wolcott, 1999; Robben & Sluka, 2012)(Restif et al., 2012).

The next section presents the relevant elements of the SGOMS ontology and describes how they are used in analyzing the data.

SGOMS Ontology and Methodological Pipelining

One of the principle problems, perhaps the central problem, to be overcome in conducting empirical scientific work is moving from abstractions to measurements, from theory to data. Words or concepts may be necessary for a researcher to work out an idea or develop a method, but eventually these constructs must give way to

measurement processes and the resultant numbers. A common form of this in cognitive science is the process of operationalizing a concept in psychology experiments. More generally, we may call this process ontologizing: laying out the entities that fall within a theoretical system, and (in some cases) describing how they can be quantified.

This is a difficult, and emphatically non-trivial problem in a domain such as cognitive science, where the entity of interest is a black box. The current state of cognitive science is a fragmented landscape characterized not only by macro-level disputes between neuroscience, linguistics, psychology, philosophy, artificial intelligence, and anthropology, but also by micro-level disputes within each of these fields⁹. To a significant degree, these disputes revolve around what constitute valid entities for a scientific research programme aimed at understanding cognition. Are constructs such as working memory, consciousness, and representation useful, or are they hindrances? The stronger version of this question is: are these entities real or illusory? This question is further complicated by the fact that it cannot be adequately resolved (or so I would argue) without discussing the validity of the methods that will be applied to measuring and analyzing these constructs.

In one sense, this is simply the problem of the old philosopher's joke: ontologies are like toothbrushes; everyone needs one, but no one wants to use anyone else's. In another sense, it is an intriguing meta-problem in the context of the current work's focus: how do abstractions and methods function as currency in scientific expertise? How are these complex constructs passed back and forth through journal articles and

⁹Another illustration of Abbott's (1988) notion of contention over official abstractions.

conferences, and how do they shape scientific activity in the lab and the classroom? The most influential work on this topic is probably Thomas Kuhn's *Structure of Scientific Revolutions* (1962), as mentioned above, but there has been a great deal of writing produced on this topic since (Bricker & Bell, 2008; Collins & Evans, 2002; Latour, 1999; B. R. Martin, Nightingale, & Yegros-Yegros, 2012; Parry, 1996), and the question can be tackled from a variety of angles. Distributed cognition, for example, has inspired interesting writing analyzing the division of cognitive labour in scientific work, given that the scientific enterprise is a prime exemplar of distributed cognition (Giere, 2002; Magnus, 2007), and science studies have examined the issue from a sociological perspective (Collins & Evans, 2002; Collins, 1983; Wynne, 2015)

This issue is an eminently practical problem in the context of the present work: what should we be looking at if we wish to understand how experts work together? That is the focus of the rest of this chapter. The following section examines how cognitive architectures can be used in addressing this question.

Cognition, Architectures, and Levels as Ontological Entities

Consider the following question: if we wish to test Durkheim's assertion that the division of labour in society reflects the same principle of biological organization that accounts for differentiation of functions within the higher organisms (see above), what scientific or empirical methods are available to us? More precisely: how can we determine which methodologies are appropriate for addressing this problem? The problem is that this question concerns the degree of domain generality of a theoretical construct, but experimental methods tend to be tightly linked to scales, domains, or levels of analysis,

and collectively we lack a rigorous framework for the reconciliation of conclusions drawn from dissimilar methods. This leads to complications in theorizing and communicating, and may account for a significant portion of the recent interest in inter-disciplinary and multi-disciplinary research (Jones, 2009; Naiman, 1999; Repko, 2008).

The question can be more simply stated as: how do we generalize ideas and test them? Possibly there is no *a priori* solution to this problem, and we are left to answer it for each effort of induction. The next section will examine a case-study of this difficulty that is relevant to the study of expertise: the debate over modularity in cognitive science.

Modularity

Modules can be considered to be, at a first pass, sub-units of a larger structure. They can be variably self-contained and specialized, and can combine in various ways to enable more complex structures or functions. Most complex systems, almost by definition, are made up of modules or sub-assemblies: airplanes, skyscrapers, university departments, multicellular organisms, computers, etc.

One of the great debates in the cognitive sciences is whether and to what degree the mind is modular (Barrett & Kurzban, 2006; Bennett, 1990; Coltheart, 1999; Evans & Over, 2008; Fodor, 1989; Geary & Huffman, 2002; Prinz, 2006). On one far-end of the debate are the evolutionary psychologists who argue that the mind is “massively modular” (Machery, 2007; Samuels, 1998; Sperber, 2001), housing an enormous number of functionally specialized components, each narrowly tuned to processing some sub-set of the world, such as mate acquisition, foraging, or visual pattern recognition. On the

other far-end are theorists who argue that cognition is constituted by neuronally implemented domain-general information processing functions, that minds are meat-based universal Turing machines (Van der Velde, 1993; Zylberberg, Dehaene, Roelfsema, & Sigman, 2011). It is likely that the truth lies somewhere in the middle, or is a combination of both views. A great deal of the writing on this topic has concerned working out the precise definition of modularity (Fodor, 1989; Robbins, 2013), the properties and variety of modules (Coltheart, 1999; Geary & Huffman, 2002), and debating the degree of modularity in real minds (Barrett & Kurzban, 2006; Evans & Over, 2008). I will not attempt a deeper treatment of the debate here.

Whatever the details turn out to be, it seems that mental modularity is an excellent candidate for an instantiation of the general principle Durkheim discusses, sitting between the social-level division of labour and the organismal-level of functional differentiation. And now we return to the problem posed above: how could we possibly test this hypothesis? This chapter proposes the following answer: by chaining together methodologies targeted at multiple levels of the hierarchy, guided all the way through by a connecting theory. The theory used must provide bridging conceptual structures and methodological guidance concerning how to connect the various levels.

One possible alternative approach to this problem is to defer to some form of emergence, whereby dynamics at a lower level give rise to those at a higher level, obviating the need for distinct levels of analysis, and rendering bridging concepts and methods useless or impossible (Baas & Emmenche, 1997; Bullinaria, 2007; Chalmers, 2011). This is a complex and active area of investigation. I am not taking a strong stance

on this issue, as I think both ways of thinking about the issue are probably valuable, though perhaps only until the general problem is resolved one way or the other. Note that the general problem of differentiation, part of the conceptual background of this chapter, may be precisely the sort of phenomenon that is best analyzed using scale-independent techniques. But in the present work, the principal aim is at specifying methods for studying the cognition of experts in the real world, and this particular problem seems more tractable under the bridging approach. This reasoning is discussed further below.

The rest of this chapter presents an analysis of communication behaviors between 2 types of expert groups acting cooperatively. To recap: the larger goal of this work is to develop a methodological pipeline that can be used to study expert cognition in complex, real world environments, and to do so in a way that facilitates computational modeling. The shape of this approach in the current work is to use the SGOMS ontology to analyze individual behavior and cognition (Chapter 3), to then bridge up to group cooperation using the same constructs as well as ideas drawn from distributed cognition and sociology (Chapter 4), and finally to implement the ideas in an ACT-R cognitive model, imposing computational specification in the process (Chapter 5). The cycle is meant to be iterative, such that insights derived at each level can be leveraged at other levels in subsequent iterations.

The justification for using the bridging approach as a candidate solution for addressing the general problem of differentiation in the context of expert cognition is as follows. First, the methods used in Chapter 3 (iterative model tracing with SGOMS) and

Chapter 5 (computational cognitive modeling) are both shaped by an assumption of some degree of differentiation. As discussed in Chapter 2, the constructs of planning units (PU) and unit tasks (UT) in the theory of SGOMS are related to each other as in a hierarchy: a planning unit is constituted by one or more unit tasks. Each unit task and planning unit is constructed/defined to map onto a particular set of environmental circumstances and to lead to particular behaviors. These elements are centrally constituted by their differences in input and output. Specifying these differences is a key aspect of the method of building SGOMS cognitive models (and other forms of cognitive models; (Anderson, 1996; Laird, 2012), as choice among alternatives is an important form of differentiation in perception and cognition. In Chapter 5, much the same logic applies (in the ACT architecture the notion of modularity as discussed above is also present). Different functional structures, such as the focus buffer and the declarative memory, are operationally and structurally distinct, and a great deal of the work in building these models involves specifying how these modules interact (e.g., passing a chunk from memory to the proper buffer). Note that modeling these integrative functions is also a key component of the work in macro-cognition: in order to scale-up from cognitive psychology experiments, which tend to focus on isolated cognitive capacities or behaviors, macro-cognition attempts to build models of how multiple cognitive components and capacities operate in conjunction (Klein et al., 2003). Chapter 5 presents some ideas about how this might be accomplished in the construction of cognitive process models.

The final reason for the decision to use a bridging approach is that there exists an

enormous wealth of data concerning the cognition and behavior of experts, and I am interested in the problem of how to combine or integrate insights derived from disparate approaches in previous work. In other words, using the multi-method approach is a way of testing the viability of interdisciplinary research on the cognition of expertise.

We turn now to an examination of the 2 data sets of expert communication.

Communication Analysis – Gears of War

Overview

This section presents the results of a study that used an extended version of the methodology presented in Chapter 3. In this study, participants played the same video game discussed in the previous chapter, Gears of War 3 (GoW3) for the Microsoft XBOX 360 video game console. Participants played in teams of two. The point of interest for was the communication behavior of the two teammates. We were interested in the possibility of deriving insight into the strategic coordination of the two players by analyzing their speech while playing. The plan was three-fold. First, attempt to link the communications between players to the cognitive model created in Chapter 3 by identifying when (if ever) utterances mapped onto planning units or unit tasks derived from modeling the play of individuals. Second, use the recorded speech behavior to identify gaps in the individual cognitive model, by observing in-game actions that were either missed in the previous model but highlighted by the speech of the players, or which represented behaviors that were possible only in the context of team play (e.g.,

revive a fallen teammate, set-up a crossfire). Third, use this iteration of the cognitive model to formulate more general questions about the nature of cooperative expertise. This third and final step then served as a launching point for the discussion presented later in this chapter.

Participants

Data was collected from 6 participants playing in 4 different pair configurations. Each participant completed 3 to 6 rounds of play, each round lasting roughly 3 minutes. Individuals participated on a voluntary basis, receiving no remuneration or course credit. Participants included 2 graduate students, 2 undergraduate students, one post-doctoral fellow, and one professor.

Data

Collected data consisted of 15 audio-video recordings of pairs of participants playing GoW3, including the verbal communications between the two players. This footage was shot using a hand-held video camera, secured on a stand, and facing the television screen. Total duration of footage was 45 minutes, 13 seconds. Mean video duration was 3 minutes, 1 second (range = 1 minute, 3 seconds to 4 minutes 49 seconds). The audio content of these play sessions were transcribed, which generated 1683 words of dialogue, constituting roughly 321 utterances¹⁰. Samples of the transcribed dialogue are presented in Appendix B.

Results

Four principal observations were derived from analysis of the video data and

¹⁰ An utterance was taken to be roughly a sentence, or a sentence fragment expressing some idea. Many communications were sentence fragments, or were interrupted due to action on screen. A more precise delimiting of utterances was not attempted, due to the fact that players rarely spoke in complete, well-formed sentences.

associated communication transcripts. First, it was possible to accurately capture a great deal of the team-based gameplay activity by applying the model developed for single player mode as presented in Chapter 3. This was expected, as the game in multiplayer mode retains most of the same mechanics as those present in single player mode, and players are required to do many of the same things (locating enemies, defending a position, tracking ammunition and health levels, etc.). The fact that the single player model held up well in analysis of team play was taken as a promising start.

Second, there was one important mechanic present in cooperative mode that is not present in single player mode, and this had to be accounted for by modifying the SGOMS model. The mechanic is that when a player takes too much damage in cooperative mode, instead of immediately dying, as in single player mode, they are knocked down, and can be revived by their teammate for a limited amount of time if the teammate can reach them quickly enough. During the process of reviving their teammate, the standing player is unable to attack or defend, and is thus left vulnerable. If both players die, the round ends, and thus performing this revival put the team in jeopardy of losing the game. This introduced an interesting strategic tension for the player who remained standing: while it was generally advantageous for the team if the downed player could be revived, it was not always possible for the surviving player to immediately abandon what they were doing and revive their teammate. The surviving player is thus forced to manage their local conditions, by either killing nearby enemies or running away from them, while balancing this with consideration for the global condition of the team, which involved making a decision to revive their partner. This

decision to revive or not revive their fallen teammate turned out to be one of the most interesting and significant choices that players had to make, and was a key data point that led to a deeper consideration of interruptions in complex, real-time activity of this sort. How do players weigh the cost/benefit of maintaining their current course of action against switching to another plan that may be more advantageous or time-sensitive? And how should this ability to interrupt and resume be accounted for in a model of expert cognition, particular when serialized, routinized behaviors constitute such a large part of what is typically considered expertise (Paletz et al., 2013)? The usefulness of the planning unit concept in analyzing interruptions is examined in depth in the section below on Counter Strike team communications.

Third, the communications between teammates included many statements of the general form: [x] is at [y]. Examples include: “there is a monster behind you”; “Boomer on your left”; “I’m in the back room”; “ammo behind the couch”. This type of utterance accounted for most of the communication between teammates.

Approximately 75% of total utterances were of this form. Considering the highly spatial nature of the task, information concerning the location and movement of elements in the game world were of particular importance for constructing and executing plans. This type of information sharing is particularly amenable to a distributed cognition analysis. Direct information sharing of this sort can be considered a form of pooled perception, whereby each player functions like a second set of eyes for their teammate, and passes information deemed important on to their partner. The interesting bit about this type of communication is what happens before the communication occurs: namely, the player

extracts the strategically relevant information before expressing it and condenses it into a short utterance of 4 or 5 words. Players did this naturally, with no direction or coaching. This lends support to the idea that part of what constitutes expertise is the capacity to judge what information will be useful to collaborators and knowing how to communicate this effectively (see the section on expert communication, above). This idea is also revisited in the Counter Strike analysis below.

Fourth, higher level strategic communications between teammates revealed gameplay elements that could not exist in single player mode, and which were of a different type than the revive mechanic presented above. These utterances revealed a type of cooperation that was not forced by the design of the game (as with the revive mechanic), but which the players had to invent or discover themselves. These included cooperative actions like setting up a crossfire, having one player distract enemies while the other flanked them, and standing back to back to cover all angles of approach. Most of the player pairs (5/6) observed in this study had not spent time practicing together in this gameplay environment before being recorded, yet they exhibited the sort of environmental exploitation and intelligent coordination that is the foundation of highly specialized teamwork. Such advanced coordination is routinely demonstrated by professional sports teams (see the next data set) and emergency operations centers which are jointly staffed by police, fire, medical, and other professionals (see the third data set, below).

Analyzing player communications was particularly valuable for detecting that this kind of high-level coordination was taking place, whereas without the recorded speech,

these plans may have been invisible. There is no direct mapping from the behavior of the player avatars on screen to what the players were thinking, and what may have looked to be identical play behavior in both the solo and the teamwork conditions may have been underlain by substantially different cognitive operations, particularly with regard to the coordination that was occurring. If the same team play footage was observed without audio, the execution of crossfires and flanks may have been missed by the analyst; certainly a naive observer (i.e., naive to the video game or similar tasks) would have been unlikely to see them. While an analyst who is an expert in the activity being studied may be able to extrapolate from the data that more sophisticated coordination is occurring, there is a significant limitation of requiring such expertise on the part of the analyst: it imposes strict limits on who is able (or “qualified) to study a particular form of expertise. The reasoning is roughly: if one is not an expert in music, how could they possibly make headway into understanding the mind of a musician? If not a grandmaster in chess, how could they appreciate the subtleties of a theoretical novelty in a particular system? The perspective from the other side is that if the skills required to play a game of chess or compose a symphony are merely variations on universal cognitive capacities, such as memory retrieval, pattern recognition, or problem space search, then it matters not whether the analyst knows the minutest details of the domain so long as they understand the mental constructs in play. The present work makes the argument for a middle way: structure initial data collection and analysis according to a guiding framework, but then use feedback from true experts (who could be the analysts in some cases) to iterate upon and improve the model.

Discussion

Communication is a crucial component of coordinated expert action (Artman & Garbis, 1998; Hutchins, 1995; Rogers, 2006). The goal of this communication study was to explore how communication data could be incorporated into the SGOMS macro-cognitive modeling framework. The principal finding was that this could be achieved by analyzing utterances with particular attention to the planning units being communicated. The following section examines this idea in more detail.

One of the primary points of interest in the SGOMS framework is the tension between local and global reactivity. In team situations, there is often a macro-level (or global) goal shared by the team, and micro-level (or local) goals which individual agents or sub-teams must accomplish in service of the macro goal. The tension between these two levels plays out in countless scenarios: soldiers in battle are often expected to sacrifice their lives in service of a greater mission; professionals must balance personal career advancement with service to the group; and athletes must at times sacrifice the chance to be the hero in a big game by passing the ball to an open teammate. While negotiating these two demands is a source of potential conflict for experts in groups, it is also among the most powerful features of distributed expertise. The ability to react flexibly to local circumstances without the need for constant re-planning at the top makes groups much speedier and more robust in the face of complex and unpredictable environments, but doing so in service of a broader goal makes the group much more powerful than agents acting alone. The SGOMS framework, along with the communication analysis presented here, suggests that planning unit negotiation may be

the principle vector by which this tension is addressed in-situ, supporting distributed decision making that is sensitive to both global and local conditions. Interruptions are among the most common occasions for needing to balance local and global reactivity: if something goes awry locally and it becomes necessary to try a different plan, agents must often decide between re-planning on their own versus consulting management or the rest of the group. Different expert environments have radically different protocols for negotiating this choice. Consider the relative autonomy of a soldier, a professor, a professional athlete, or a lawyer. The next section, analyzing Counter Strike communication data, examines expert responses to interruptions in more detail, and Chapter 5 presents a method of modeling local replanning in response to interruptions.

Another insight derived from these analyses concerns the “greater than additive” gains in productivity described by Adam Smith in his analysis of the pin factory (Smith, 1776). Recall that much of the communication observed in this study involved statements that passed information from one player to another: e.g., “there is an enemy behind you”. This sort of simple-seeming communication potentially serves two important functions. First, it allows players to “pool” their perception, improving each player’s awareness of the environment by calling out relevant features. This draws on Klein et al.’s (2003) notion of common ground. Secondly, communication may reduce the cognitive load (Kalyuga, 2011; Schnotz & Kürschner, 2007; Sweller, Ayres, & Kalyuga, 2011) on the player receiving the information, because the data is in some sense “preprocessed”. Consider the following two variants on a scenario. In situation A, there is no communication between players. There is an enemy, a “Boomer”, on the left of the

first player. In order for that player to respond, let us assume that three schematic steps are involved: search for the enemy, determine what sort of enemy it is, and decide how to respond. The first two of these are largely perceptual tasks, and may occur quite quickly, though they take longer or be more difficult if the environment is full of distractors or the enemy is camouflaged, for example. The third task, largely cognitive (though see the model in Chapter 5 for a further discussion of this) may also be accomplished rapidly if the player has a strong sense of what to do when encountering that enemy, a ready-made response. The total response thus involves 3 steps. In situation B, Player 1 is told by Player 2 that “there is a Boomer on your left”. Assuming that the receiving player know what a “Boomer” is, he/she need now only decide how to deal with it. Steps 1 and 2 of the process are largely accomplished by player 2 before player 2 communicated the information (though perhaps neither step is completely eliminated, and they are just made quicker). If players and other experts are able to process this sort of information and package it into useful communicative chunks for their teammates, particularly if this can be done using “spare cycles” of cognition while performing other tasks, this may provide one route by which expert cooperation leads to greater than additive improvements in performance. This example is, of course, only a sketch, and a full assessment of this idea would require calculating the relative costs and benefits of processing such information and communicating it to one’s teammate, to determine when it is effective to communicate and when it is detrimental to the group.

Another interesting finding from the communication data is that the notion of narrative served as a useful framing concept for making sense of how teammates

created and communicated plans. They frequently followed the same simplified narrative model as that presented above:

- 1 – We are currently in state A
- 2 – We wish to transition to state B
- 3 – We must do X to move from A to B

Here are some examples of this from the recordings:

Ex 1:

Player 1 – “Heading to the center...here...whatever you call this place”

Player 2 – “I’ll watch the other side.” (i.e., provide covering fire)

Ex 2:

Player 1 – Can we destroy that middle segment?

Ex 3:

Player 1 – I totally lost track of you ... Can you get back to base?

The congruence between the narrative model and the form of the plans being shared by players was interesting not because it was surprising, but because the overlap was so clearly evident. This suggests that applying techniques of narrative analysis in studies such as this can be quite natural and effective. This is in line with recent thinking in the area of cognitive narratology, which attempts to understand both how we make sense of stories and how we use stories to make sense of the world (Herman, 2000, 2013).

One of the surprising features of applying this idea of narrative to analyzing planning units was that some of the actions revolved around maintaining an equilibrium state, and seemed to run counter to the notion of narratives being a transition to a new

equilibrium. These plans presented as a variation on the Todorov model:

- 1 – We are currently in state A
- 2 – We wish to remain in state A
- 3 – If X happens, we will be in state B
- 4 – We must therefore do Y to prevent X

This was the basic structure for the planning unit “hold ground” (see Chapter 3), which amounted to preventing enemies from encroaching upon a stronghold. Here is an example of this from the recordings:

Ex 4:

Player 1 – Force them to come to us

Player 2 – I’ll take the right side then?

The players here are coordinating to maintain their position, or more abstractly, to maintain the current equilibrium. While the simplified narrative model was not evident in these plans, the narrative approach still offered an interesting way of framing the problem. Furthermore, there are many different models of narrative structure (Mishler, 1995), and I am not attempting to argue in this work that Todorov’s narrative arc is the only way of making sense of narrative structures.

Finally, on-the-fly strategizing observed in the videogame play is an important part of what experts do in many fields, and these are often idiosyncratic or situational, perhaps never to be repeated again. This makes the general case of modeling complex expert behavior difficult, and recalls Simon’s ant-on-the-beach metaphor (Simon, 1996b). The hypothesis of the SGOMS framework, which allows for some traction into

the matter, is that expert strategizing can be modeled using the construct of planning units (see Chapter 3). SGOMS also claims that experts in many domains tend to share planning units, which allows them to coordinate so effectively (think of a football quarterback calling a play). At the highest levels of expertise in many domains, whether they be artistic, athletic, or scientific, experts of a particular domain will share many of the same skills (classical musicians, for example, can all read sheet music, recall and anticipate melodies, count time, have memorized multiple scales and pieces of music, etc.), but will often combine these skills in unique and characteristic ways, and this is what gives rise to individual variation at the top levels of expertise. The choice of planning units is thus at the top level of the theory of SGOMS. This study suggests that using the SGOMS ontology, particularly the notion of planning units, in analyzing expert communications can provide some traction into the problem of analyzing into how experts strategize about problems at a high level when operating in complex environments, as well as how they coordinate their efforts. In other words, the “unrepeatability” in how these scenarios play out introduces difficulty into analysis (e.g., statistical methods do not work particularly well in some cases), but perhaps looking for planning units in communication can help lend some shape to behavior. Chapter 5 will use this idea to model communication between cooperative agents as the sharing of planning units.

We turn now to a discussion of the second data set: communication from professional teams of *Counter Strike: Global Offensive* (CS:GO) players.

Communication Analysis – Counter Strike: Global Offensive

Overview

In the above study, we examined the communication behavior of video game players. These individuals were skilled and could reasonably be designated “expert video game players”. They were not, however, professionals. Their abilities were the result of recreational video game playing over the course of several years. It therefore remained an open question whether the same methods of analysis used in the previous study could be usefully applied to the study of professional players, who presumably would demonstrate significantly more developed expertise in this area. To examine this question, communication data was collected from professional Counter Strike teams.

CS:GO is a 5v5 team-based, tactical first person shooter video game. A match consists of a best-of-30 rounds structure, split into two halves of 15 rounds, each round lasting a maximum of 105 seconds (1:45). In each half, one team plays as the offensive team, and the other plays as the defensive team. After 15 rounds, teams switch roles. When in the offensive role (“terrorists”), the team’s objective is to plant a bomb at a particular spot on the map, and defend it from the defensive team until it detonates. The role of the defensive team (“counter-terrorists”), is to prevent the offensive team from gaining access to the “bomb site” for the duration of the round; if it has been planted, to defuse the explosive before it detonates; or to eliminate the entire enemy team. At the end of each round, the teams’ positions are reset, and the players are awarded money according to which team won the round, which is then used to buy equipment for the next round. This two-sided structure of offense/defense is very close to that used in baseball, American football, and cricket.

Counter Strike is currently a highly popular game in the electronic sports or “eSports” world (PC Games News, 2015). Live matches of CS:GO routinely draw 500,000 online spectators and teams compete in tournaments for million-dollar prize purses (Lahti, 2016). Electronic sports have quickly become big business. Revenue from eSports worldwide was estimated at 325 MM USD, and are projected to reach 465 million USD in 2017 (Newzoo.com, 2015). While video games have been a popular form of entertainment since Pong was released to the public in 1972, it is only in the past 10 years that competitive video gaming has begun to adopt the tropes of other forms of professional sporting, such as live casting with commentary, highly publicized matches, seasons, team standings, large prize purses, and hero narratives.

Verbal communication is key to success in online games such as CS:GO for two principal reasons. First, because players are typically not in the same physical space, but instead inhabit the same virtual space over a network, alternative channels for communication are limited. Players cannot rely on body language or other visual cues from teammates other than those expressible by their avatars, and must therefore use words to convey important information. Second, these games are designed to be played in teams, and their task structure is heavily dependent on coordinated action. This coordination often requires highly precise timing and positioning, and the only way to establish this coordination is through language. For this reason, we expected to find highly precise language used to convey strategic information.

To examine whether the techniques presented in the last section could be fruitfully applied to expert teams exhibiting a professional level of training and

teamwork, the methods were extended to a dataset of professional CS:GO players communicating in real time. This analysis is presented in the next section.

Data

The dataset comprised 9,104 utterances, totaling 41,959 words, extracted from audio-video recordings of 7 professional matches of CS:GO. The total duration of the video data was 333.5 minutes; average match duration was 47 minutes, 36 seconds. Mean number of utterances per match was 1300.5.

Data were examined qualitatively, with the intent of extracting insights that could later be used in constructing cognitive models.

Discussion

Examination of the data revealed two interesting aspects of coordination at this high level of expertise. The first concerned specialized terminology, and the second involved communication to deal with interruptions and unexpected events. These will be discussed in turn.

The first finding, supporting initial hypotheses, was that the teams used domain-specific terms to refer to precise locations and complex plans. These utterances tended to be short and direct, but conveyed enormous amount of information to the rest of the team. To take an example, one of the teams routinely used a strategy that they called “Lemon Dogs”, which is named after another Counter Strike team. At the beginning of a round, the team leader would say something like “Let’s run Lemon Dogs”. This statement was enough to coordinate, at a high level, the actions of all 5 team members for the round. Each individual had a particular role to play in the strategy, and could

immediately begin preparing for it. Other examples of these strategy names included “B-split”, “Default”, and “Envy”.

This sort of practiced and named strategy is common in many (possibly all) forms of team-based sports. One stereotypical example of the associated communication is the football quarterback shouting a play before the ball is hiked: “blue 42, hut!” The use of this sort of communication for coordination was therefore not particularly novel in itself. It was interesting in the current context, however, because it seemed to support the notion that the SGOMS notion of planning units was useful for describing such coordination.

Recall that planning units are modeled as cognitive control structures with two functions: group coordination and robustness against interruptions. Concerning the first, individual experts acting in groups must often balance the need for local and global responsiveness, and planning units are conceived of as a level of task organization within experts that permits establishment of common ground (Klein et al., 2003) through communication. In the case of the Lemon Dogs example above, the SGOMS account of the team behavior would state that the term “Lemon Dogs” refers to an asymmetric plan, where each member of the group may have different responsibilities, but built into that referent is knowledge of how individual action fits into a larger goal structure. Each player is continually balancing his/her own local actions and needs (e.g., staying alive, defeating an enemy) with the needs of the team (e.g., plant the bomb, defend the bomb site). Planning units provide a way of conceptualizing the line of communication between team mates that allow for this balancing. In the model presented in the next

chapter, it is through the communication of planning units that the cognitive agents coordinate and function more effectively as a team than independently.

The second function of planning units as a theoretical structure is to account for expert resilience in the face of chaotic circumstances or interruptions, which often characterize the real world environments that experts operate in. In the collected data, we observed numerous instances of teams quickly and collectively responding to such circumstances. The most obvious cases of this occurred when one component of a plan went awry, and the individual who made the mistake announced “cancel” or “abort”. This was an instruction for the other team members to pause the plan and negotiate an alternative strategy. Often this was a temporary delay, and the plan was resumed afterwards; other times the team chose a different strategy. The importance of the planning unit structure in these cases, from the perspective of the analyst, is that it helped make sense of the actions and communications of the players when they were interrupted. By viewing the player communication after unexpected events as a way of interrupting the current planning unit and negotiating a new one, the superficially “messy” behavior appeared more structured and directed.

Examination of the Counter Strike data suggests that the method of modeling expert communication using the SGOMS framework is potentially viable. In the next chapter, these two insights concerning coordination through planning units and interruption management through planning units are explored in a cognitive model of a simplified version of Gears of War.

Conclusions

Human beings are exceptionally capable animals not because of our individual intelligence or strength, but because of our ability to engage in sophisticated cooperation and cumulative learning (Harari, 2014; Henrich, 2015; W. M. Williams & Sternberg, 1988; Woolley, Chabris, Pentland, Hashmi, & Malone, 2010). While the expert individual is often held up as the image of human excellence, such great ones can exist only in the context of supportive communities, in human groups that have passed on knowledge, abilities, heuristics, techniques of survival and socialization. We do not yet understand all of the factors that contribute to an individual becoming an expert, but it is clear that they cannot do it alone.

This chapter has examined the role of communication and organization in the development and exercise of expertise, using the framework of macro cognition, the idea of a shared macro cognitive architecture, and SGOMS. The goal has been to support the analysis of expert social behavior and communication, such that data can be generated for use in the creation of multi-agent cognitive models set in complex environments. Several ideas and hypotheses were generated through the process detailed above, and thus the effort may be considered a success. These ideas, concerning planning unit communication and response to interruptions, will be picked up and implemented in the next chapter. This is the continuance of the methodological pipeline discussed above, representing the broader goal of the present work.

Chapter 5 – An SGOMS ACT-R Model of Interruptions and Communication

Overview

The following chapter presents a solution for modeling interruptions and communication during expert activity using Python ACT-R and the SGOMS framework. The model presented is a pared down version of the Gears of War 3 model detailed in Chapter 3. It is meant to serve as a proof of concept, and embodies a theoretical and a practical component, the former relating to a theory expert cognition (SGOMS), and the latter referring to implementation challenges in the ACT architecture.

The model is of an abstracted task environment with multiple cognitive agents. The environment is a simple grid world, and the goal is for agents to navigate to a particular square that is designated the goal square. The environment is also populated with enemies that agents must eliminate, and bunnies that can be safely ignored (i.e. non-threatening agents). Agents are capable of communicating the location of the goal square to the other agents once that location has been reached. These three behaviors – navigate, engage enemies, and communicate – define the behavioral repertoire of the agents.

Two key components of the model are presented and analyzed: interruptions and communication. The first demonstrates a mechanism by which a cognitive agent can interrupt a *navigation* planning unit to engage in an *eliminate enemy* planning unit. The solution is to use both a bottom-up and top-down visual buffer, which together constitute the visual module, and to allow bottom-up activity to interrupt top-down activity. The second component explores how communication between cognitive agents

can improve performance of the group. A quantitative analysis of the model is in support of the notion that communication leads to great gains in efficiency, as discussed in Chapter 4.

Introduction

ACT stands for Adaptive Control of Thought, and represents a theory of cognition as well as an implementation of this theory in a computational cognitive architecture (Amant et al., 2005; J. R. Anderson, 1996; Stewart & West, 2007). The architecture is used to create cognitive models by which to test and extend the ACT theory. The trailing “R” in ACT-R stands for Rational. ACT-R 6.0 is the current manifestation of this theory represented in architecture (Anderson, 2007).

This section will briefly examine the two main points of interest that are addressed through the ACT-R model that is presented later in the chapter: interruptions and communication in expert cognition.

Interruptions

Dealing with interruptions effectively is crucially important in many forms of expertise. In sports, the role of defense is precisely to interrupt the offensive strategies of the opponent. In rhetorically heavy forms of expertise, such as law, politics, or journalism, gracefully creating or recovering from interruptions may be decisive in emerging the victor in an engagement. In scientific research, if work is released that invalidates (or perhaps more damningly, confirms) the particular question an investigator was working on, the study or programme underway may have to be interrupted. In aviation, inclement weather may necessitate a modification to the flight plan. And in complex, chaotic environments, like those characteristic of the video games

analyzed in the present work, interruptions are often built-in: the situation changes so rapidly that plans must be frequently readjusted.

These different forms of interruptions place a variety of cognitive demands upon the experts, at various time-scales and incurring variable degrees of impairment. Two cognitive impacts are particularly important in the context of this chapter: task-switching without system failure, and recovery without lost progress. Concerning the first, human beings are quite robust in the face of unexpected events. We do not typically freeze like deer in headlights when presented with something unexpected (though of course this does happen at times). Even in cases where we *do* freeze up, it tends not to last forever; we are generally able to recover eventually. In modern expert systems and cognitive models, however, the situation is different: unexpected events tend to crash the system (Holland, 1986). This is not realistic, and resuming after an interruption ought to be possible, particularly if we are attempting to describe and model naturalistic cognition. The first variant of the model below addresses this issue in the context of Python ACT-R models of expert cognition.

Concerning the second impact, human beings are able to manage a number of activities simultaneously. If we are writing an e-mail and receive a phone call mid-way through, we can return to writing the e-mail after completing the call. We can pick up where we left off, rather than being forced to start from the beginning. As another example: we can stop stirring a pot of soup to put the chicken in the oven, and then return to the soup; we do not have to cook each dish until it is finished. While novices often must fully complete a single task before turning attention to another, experts tend

to be more skillful in managing parallel, domain-specific activities (Bainbridge, 1997; Zylberberg et al., 2011). A professor may be supervising 10 graduate and undergraduate students at a given time, and a high-ranking military officer may have dozens of units under their command; these experts must be capable of juggling the demands of the various projects underway.

When switching attention from one project to another, we rarely have to start from the beginning. If that were required, complex projects such as writing a book or building a house would be next to impossible. Humans have developed a wide range of tools to make such task switching and resumption less costly, such as to-do lists, progress reports, and automated reminders on smart phones. Our memory and perceptual systems also provide an enormously powerful platform to do this naturally. We can simply remember where we left off, or we can survey the scene and use cues to get us back on track. As an example of the latter, bartenders often set out glasses along the bar in order to more easily recall which drinks to make for a complex order (this is often used as an illustration of the extended mind hypothesis, (Clark, 2009)). This ability to get back on track is an important part of naturalistic cognition, and ought to be accounted for in cognitive models of such. The first model variant is used to discuss how this capacity could be implemented in computational models.

Communication

As discussed at length in the previous chapter, communication is centrally important in many forms of expertise. The ability to pass information back and forth can make a team much greater than the sum of its parts. While a complete analysis of this capacity is currently impossible (it would likely require significant advances in linguistics,

psychology, and media studies), a simplified model of the ability may be quite useful in macro-cognitive modeling. The model presented in this chapter includes a rudimentary communication system, which functions almost like telepathy: once an agent has found the goal location, it is able to broadcast this location to all other agents. The duration of this communication is one production cycle (50 ms of simulation time), though this cost could be varied in alternative scenarios.

The 50 ms timing is not meant to reflect real communication time costs; obviously it takes longer than 50 ms for people to communicate with one another, especially if the communication involves plans of any degree of complexity. The single production cycle time cost currently serves as a placeholder. More realistic timings may be explored in future work, particularly if there is an attempt to match simulation data to human performance timing data. Importantly, however, modifying the duration of the communication process has no effect on the performance of the model, other than prolonging the average time to completion in the model by the duration of the communication process. In other words, if the communication requires 3 seconds, the average time to completion will increase by 3 seconds, but in all other respects, the pattern of behavior in the model remains unchanged.

The model serves as a demonstration of efficiency gains from communication as a function of group size, and of a useful bottom-up mechanism for modeling it in Python ACT-R (Stewart & West, 2007). As will be shown below, efficiency gains increased in proportion to the number of cooperative agents in the simulation. In the real world, the situation is of course more complicated. As groups grow in size, direct and immediate

communication becomes more costly, and is eventually supplemented by more complex mechanisms, such as newspaper distribution and encyclopedias. Future work could use this initial implementation to more carefully explore the relationship between the cost and benefit of communication, and potential cognitive mechanisms underlying communication behavior. One possible route of investigation would be using the production learning mechanism of ACT-R to simulate learning utility values of the communication productions. Another might involve modeling cognitive load in connection with the decision to communicate or not. A third line of inquiry could explore communication with a limited range, to examine spatial network effects of communication.

Practical difficulties of modeling interruptions

In the current implementation of Python ACT-R, unexpected interruptions cause the system to crash. More precisely, changing the contents of a central buffer before a production fires will cause the previously readied production to be abandoned. If a set of modeled actions is dependent on being executed in the proper sequence, changing these buffer contents will cause the expert system to fail.

One possible mechanism by which to create such interruptions is to include higher-than-average-utility productions that will “hi-jack” the production system. This does not, however, allow for a graceful response; it will typically cause the model to terminate or hang. This is because there is no way to “buffer”, or hold in reserve, a suspended action, if the same buffer is being used to handle top-down, deliberate actions and simulated bottom up information.

The lack of such a graceful interrupt mechanism has 3 principal weaknesses in the context of the current work. First, it means that the model is not capable of context-sensitive responses to information from the environment, particularly concerning threat or interruption. A great deal of macro-cognition research is concerned with understanding the relationship of agent to environment, and how this coupling permits or prohibits certain actions, facilitates cognitive processing, and so on (Klein et al., 2003; Schraagen et al., 2008). Because so much of expertise is context sensitive, a usable mechanism for modeling intelligent use of bottom up information from the environment is seen here as desirable. This principle also applies to evaluating one's own actions in the context of a current plan. In the example discussed below, the significance of a failed or interrupted action is different according to whether that action was executed as part of a complex plan, or was performed as a stand-alone action.

Second, the risk of interruptions is continuous. Any bit of information, entering the cognitive system at any point or through any channel, could in principle be an interruption. This is true regardless of whether the agent is familiar with the incoming information. To illustrate, consider a boxer engaged in a match: if he intends to throw a combination, and is "stuffed" after the first punch (i.e., his opponent gets closer than striking distance, making it impossible to continue punching) he will quickly modify his plan of attack. This is the sort of interruption for which the athlete is well trained. On the other hand, if the sprinkler system goes off during the match, this is the sort of interruption for which it is highly unlikely that he has specifically prepared. He will still be able to react, however, due to his general knowledge about how the world works.

There are two points to be made from this illustration. First, we can differentiate between expected interruptions and unexpected interruptions, with the former being those that are domain specific and familiar to the expert, perhaps with an associated, well-practiced response. The latter could be anything for which a response does not fall properly within the realm of the type of expertise in question. The second point is that because *anything* could serve as an interruption, it must be possible to tell a generalized story about bottom up monitoring of the environment. It seems implausible that after every single micro-action (which could be of arbitrarily fine grain size), an individual could actively check whether an interruption has occurred. The present chapter proceeds from the view that interruptions are frequently handled by a bottom-up system.

Finally, there is evidence from neuroscience that the brain may possess a “neural alarm system”, which is able to evaluate potential sources of threat without “higher” or conscious processing (Liddell et al., 2005). This is hypothesized to be a brainstem-amygdala-cortical loop. This research suggests that real brains may use fast, bottom-up evaluation of stimuli in determining whether to orient or react, bypassing the need for conscious appraisal. If this is an accurate assessment, it suggests a candidate mechanism by which the problem of threat/interruption assessment is handled without overloading the top-down functionality of the mind (i.e., top-down checking the environment after every micro-action).

The next section addresses the proposed solution to this bottom-up interrupt problem in the context of Python ACT-R and SGOMS.

Methodology

Efficiency of the group was quantified by the duration of simulated time required to reach the goal location. Communication is quick and global, requiring only one production cycle (50 ms). In such a design, efficiency gains as a result of the ability to communicate are inversely proportional to the size of the group, as will be shown in the analysis in the second half of this chapter.

The method used to address the above issues is mainly computational. Before attempting to build and test detailed models of the cognitive processes underlying interruptions, the ACT-R architecture must be less limited in terms of the types of interruptions it can tolerate. In the remainder of this chapter, two potential modifications to the ACT-R architecture are presented and tested. The goal is to demonstrate that the modifications meet the goal of giving the architecture a reasonably stable, versatile, and intelligent interrupt mechanism that can serve as a platform for future work, refining and evaluating models of interruptions.

Architecture Changes

ACT is often referred to as a production system (Anderson, 1996), but it is actually a hybrid system that uses a production system as a central organizer (note that this is different than the construct of a central executive as generally understood in psychology; see Miyake & Shah, 1999). It is a “hybrid” because it also incorporates modules in the periphery that can perform various operations that are not necessarily represented as productions. These modules are typically, but not necessarily, under the control of the central production systems.

The ACT production system represents the actions of procedural memory. It is composed of production rules (if/then rules) that fire based on the contents of the buffers. Buffers in ACT represent neural connections between modules. In some cases there is neural evidence for these buffers and their particular connections (Lenk, Möbus, Özyurt, Thiel, & Claassen, 2011); in other cases, the buffers and modules represent abstract functionality. The buffers hold chunks that represent the content of the communication between modules. Chunks contain predicate information. As suggested by their name, the buffers can hold a chunk until the receiving module is ready to take it. Modules operate in parallel and represent key cognitive abilities, such as declarative memory, vision, motor, goal tracking, etc.

As noted in Cooper (2007) ACT-R follows a Lakatosian framework for science. In particular, there is a core ACT-R architecture that represents the official version of ACT, and there is a periphery which contains various modules and modifications that represent proposed additions or changes that still need broader evaluation by the ACT-R community. What is being proposed in this chapter would fall in the periphery.

Figure 8 is a schematic of a standard ACTR model with 2 arbitrary modules. It can be seen from this image that the architectural structure places the procedural memory production system at the center of communication. When a production fires it can deliver requests to the modules for information or direct the module to execute an action. When modules with buffers have completed a request, the information is placed in their buffer, where it can then be accessed by the central production system. Information placed in the buffers by modules is used for two things: deciding which

production to fire next by the CPS, and providing information to be used by the production that fires (an in the context or unit task buffer in the model presented here; see below). Modules cannot communicate between each other, except by going through the procedural memory production system. Therefore the production system plays a central coordinating role and also forms a bottleneck as only one production can fire at a time. In this sense, the production system can be seen as a serial control system for a parallel architecture. Note that the CPS has a single associated buffer, typically referred to as the goal buffer or focus buffer. In brief, a production cycle consists of the CPS executing some action internally (e.g., retrieving a chunk from declarative memory) or making a call to a module.

Figure 9 sketches the model presented in this chapter, with the changes made to the standard ACT-R structure. V1 and V2 represent the two vision modules (top-down/search and bottom-up/entity recognition); Mo represents the motor module which moves the agent through the environment; and B's represent buffers. It can be seen from the diagram that the two vision modules share a buffer. Note also that V2, the bottom-up visual module described above (entity recognition) can directly affect the CPS, but is never directly called by the CPS. Finally, it can be seen that the CPS has multiple buffers, unlike in Figure 8. This represents the SGOMS task structure model, which involves a separate buffer for planning units, unit tasks, and context. These three buffers, along with the standard focus buffer, control an agent's execution of sequential actions.

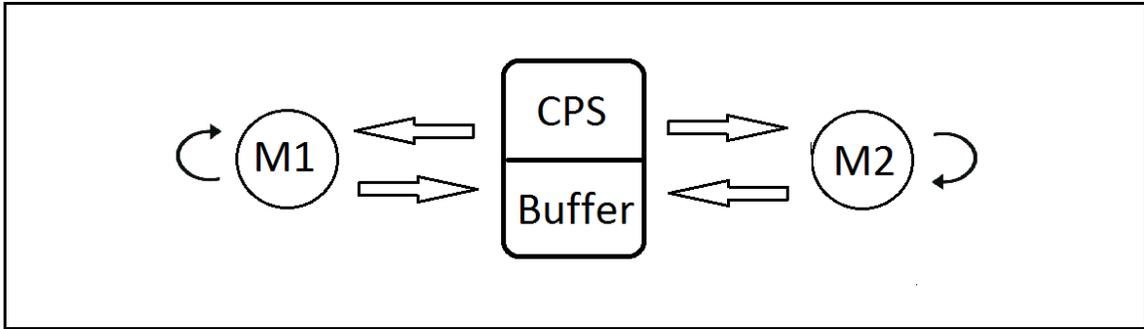


Figure 8. An ACT-R model with two arbitrary modules (M1, M2). The central production system makes calls to modules which can return information to the focus/goal buffer of the CPS. Modules can interact directly with the environment, as indicated by the looping arrows on the periphery, but not with each other. The CPS cannot interact directly with the environment.

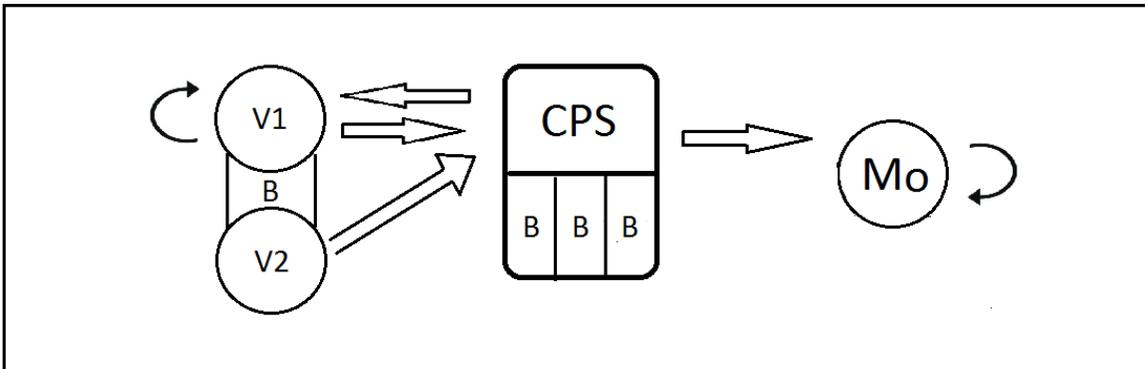


Figure 9. Schematic of the ACT-R expert cooperation model presented in this chapter. V1 and V2 represent the top-down and bottom-up vision modules (search and entity recognition, respectively), Mo represents the motor module, which. The multiple buffers associated with the CPS represent the SGOMS task cycling mechanism, relying on a separate buffer for each unit task, planning units, and context tracking. Note that the bottom-up vision module can directly drive the CPS.

The term “production system” can be confusing. The computer code underlying the ACT-R modules uses production rules, and Python ACT-R explicitly uses production systems in the construction of modules. What defines the ACT-R *procedural memory production system* is its positioning in the overall architecture: it functions as a top-down control system, directing all other subsystems. The problems discussed above are inherent to this architectural organization but can be solved through some minor but functionally significant modifications. These are described below.

Bottom-up Buffering - As argued above, modeling bottom-up information processing is critical for creating models of interruptions, especially unexpected interruptions. However, the ACT-R architecture is limited in its ability to deal with bottom-up information. The current standard solution is to have a “vision module” which can “notice” objects or events in the environment as long as it is not preoccupied with top-down driven search. This mechanism causes everything that is noticed bottom-up to be placed in the visual buffer, a process known as “buffer stuffing” (Hussain & Wood, 2009). The model presented here uses a version of buffer stuffing for vision, and a more direct, non-buffered mechanism for audition.

Buffer stuffing is usually employed only for vision, and relies on an (at least implicit) notion of low fidelity environments, as discussed above. Using buffer stuffing to model the bottom-up flow of visual information as experienced by a human would be virtually impossible without some simplification or intermediary; pure, raw, or unmediated bottom-up processing would be continuous. Under such conditions, the

visual buffer would be continuously stuffed, and the vision system would therefore prevented from executing any top-down driven actions, such as visual search.

Furthermore, any such model of this process must account for the transformation of raw sensory input to object or pattern recognition. While a human participant may report simply “seeing an enemy” in the game, the process of constructing that visual gestalt is known to involve enormously complicated neural/cognitive machinery behind the scenes (DiCarlo & Cox, 2007; Isabel Gauthier & Tarr, 2002; Lowe, 1999). Arguably, ACT has sidestepped this issue by implicitly assuming that the brain has filtered out irrelevant bottom up chunks before buffer stuffing occurs, relying on the technique of low-fidelity environments to model this. This is one way of justifying the common practice of only putting visual information that is relevant to the task in the environment.

The buffer stuffing approach falls short, however, if we wish to model the processing of unexpected information or interruptions. This is particularly problematic in cases where the agent is in the process of serially executing a series of steps, as such as a unit task. In these cases, if a piece of bottom up piece of information were to hijack the top-down production system, it could either crash the system or cause it to hang if unfinished productions were interrupted. One possible solution to deal with this problem is to use a secondary buffer for the vision module that responds only to bottom up information. An alternative, and the one implemented here, is to have two separate modules which both use the same buffer (i.e., the visual buffer, in this model), but to grant a degree of intelligence to the modules such that they can direct agent behavior

directly, without being forced to interrupt the central production system in the middle of executing a production.

Figure 9 schematically illustrates how this mechanism is implemented. The *visual search* module is implemented as a top-down visual module. It is driven directly by the central production system. It has access to the visual buffer, which will buffer the “line of sight” that is constructed by the module’s productions. Operating in parallel is the *entity recognition*, which is implemented as a bottom-up visual module. This cycles in parallel to the central production system, and executes buffer stuffing when it notices either an enemy or a bunny in the line of sight. These visual recognitions cause the agent, upon completion of the unit task they are currently engaged in, to either engage the enemy or ignore the bunny. The key here is that the bottom up productions cause the agent to respond appropriately by triggering the relevant planning unit. This idea was drawn from observation of the professional Counter Strike players, discussed in Chapter 4, who were able to make incredibly quick motor responses to visual information. The intent was to implement a system whereby the bottom up entity recognition system could direct behavior intelligently, rather than cueing a decision-making process in the central production system to decide what to do next based on context.

Failing to integrate these modules and buffers correctly creates a serious problem for the procedural production system when interruptions occur, as the buffer contents may change before the correct procedural production has a chance to fire (or to retrieve the chunk requested by a previous production). This is a computational issue.

Production systems are often fragile, especially when multiple sequential productions must fire in a specific pattern. Changing the information used for matching (i.e., the information to drive the firing of a production) will disrupt the sequence and cause a crash. This can be managed for small amounts of bottom-up information, but becomes problematic as modeled environments scale in complexity. Using a low-fidelity environment allowed this issue to be managed in the present work, but implementing a more flexible and robust mechanism for this ability, in order to handle more complex information flows, is an area for future development.

Note, finally, that current thinking in neuroscience is that there are multiple upward and downward flows of information in the brain (Delorme, Rousselet, Macé, & Fabre-Thorpe, 2004; Karakaş, Başar-Eroğlu, Özesmi, Kafadar, & Erzençin, 2000; Latinus, VanRullen, & Taylor, 2010; Theeuwes, 2010), possibly with substantial recursion (Pollack, 1990). Accounting for this duality is thus desirable if one wishes to construct neutrally plausible models of cognition; the dual flow of information in the present model is motivated partially by this issue.

Bottom-up Production Systems and Interruption - In ACT there is only one route to interrupt a top-down production-driven sequence via bottom-up information: have a unique production for each bottom-up chunk that indicates the need for an interrupt. To do this, these productions need to have a higher utility value than normal productions in the procedural production system. Such productions will be triggered immediately when the bottom up chunk appears and will interrupt any ongoing sequence of productions. However, this is problematic because one cannot predict

when these productions will fire, and it is difficult to write complex code that can be immediately interrupted at any point and still be capable of recovering. From the other side of the issue, experts do not react immediately to every interruption. Instead, as noted above, they react intelligently based on the nature of the interruption and the current context. These issues have been addressed with the use of bottom-up processing modules in conjunction with the SGOMS task structure representation. Together, this permits an agent to be sensitive to where it is in the execution of a task, and to only respond to an interruption when it can do so gracefully.

Several custom modules have been used in this model. A motor module controls movement of the agent, and thus serves as the lowest level interface between the environment and the agent. This motor module is atheoretical. The vision system is modeled as two linked modules, the top down *search* module, and the bottom up *entity recognition* module. Finally, there is a bottom up communication module, described next.

The communication module functions as a mechanism by which one agent can directly transmit information to another, and the recipient agent can store that information in memory. This allows an agent who has spotted the goal location to tell others how to reach it. The module runs in parallel to the central production system, and is triggered when the visual system detects the goal location. Once the goal is spotted, the agent immediately communicates to all other agents. The information that is transmitted is *not the location of the goal*; it is instead the *planning unit* that will allow the agent to find the goal. This implementation is based on the SGOMS hypothesis that

communication between experts often occurs at the level of shared planning units (see Chapter 3), and that by representing planning units in the same way, experts enhance cooperative capacity. Implementing the generalized SGOMS task execution mechanism (presented below) was thus a key part of modeling this communication process.

Integration with SGOMS

SGOMS implemented in ACT-R can be understood as a set of generic productions used to manage task specific productions. Although there are different types of interruptions and different responses to them, according to SGOMS theory the most common way that experts experience interruptions is the following: an interruption occurs, the expert finishes the unit task that they are engaged in, creates a chunk representing where they were interrupted, exits the current planning unit, and evaluates the interruption (West et al., 2013; West & Nagy, 2007). Again according to SGOMS, unit tasks should be short enough that it is unlikely they will be interrupted. In other words, unit tasks are meant to be finished. This implies that there is no automatic way to gracefully recover from an unexpected interruption that prevents the completion of a unit task. This type of interruption is where serious errors are predicted to occur; experts would be expected to avoid these situations by quickly finishing the current unit task. This particular type of interruption is not addressed in the model presented here, but could be investigated further using the same system.

An important caveat is that the above applies only to unexpected interruptions. By definition, when someone is an expert, they have effective ways of dealing with expected interruptions. For minor, expected interruptions, the unit task would contain appropriate quick fixes. To return to the boxer example used above: if the fighter is

“stuffed” while executing a combination, he might immediately switch to a clinch position. On the other hand, if there is a potentially serious interruption that requires the unit task to be abandoned, a path to quickly and effectively abandon the unit task might exist if the expert has trained for this (presumably training would be needed, as such events are hypothesized to be rare in practice). However, in SGOMS theory, unit tasks represent islands of activity carefully crafted in terms of length so as to avoid disruption and allow for a graceful exit if an interruption occurs. Theoretically then, in most cases experts should deal with interruptions by first completing the current unit task, and only then responding to the interruption. This is how the interrupt mechanism functions in the present model: the current unit task is completed before responding to the interruption. In contrast, a planning unit may be interrupted part way through completion without issue. This is demonstrated in the behavior of the model presented here.

In the model, the SGOMS task structure is reflected by the planning unit and unit-task performance system. This is a content-agnostic production cycling mechanism that causes the agent to:

- 1) Execute a unit task to completion once it has been begun,
- 2) Execute the elements of a planning unit in the proper order, and
- 3) *Avoid crashing* if the agent is interrupted.

To accomplish the second function, agents rely on a context buffer, which operates as a simplified situational awareness model similar to that presented in Chapter 3.

The code for this model is presented in appendix D and is hosted at <https://github.com/KoreyMac54/SGOMS> .

Description of Model Functioning

The following section presents a schematic overview of the model and details of how the interruption and communication mechanisms are implemented.

At the beginning of each run of the simulation, the environment is populated with agents, monsters, bunnies, and the goal location, all randomly assigned to unique locations on the grid world. The world is a square grid, and data was generated using world sizes of 5, 7, 9, 11, and 13 units squared (i.e., 5x5, 7x7, etc.). The number of agents was also varied, for the range 1-9 agents.

Data was collected for two variations of the model. In the first, the agents all attempt to locate the goal square independently. Each agents moves randomly until it reaches the goal square, and then waits for the remaining agents to locate the goal. This variant will hereafter be referred to as the “wander model”. In the second variant, agents are able to communicate. Once the first agent spots the goal square, it informs the other agents about how to reach the goal. This variant will be referred to as the “communication variant”. In both variants, once all agents have reached the goal, the simulation ends.

This set-up yielded a 5 x 9 x 2 (world size by agents by model variant) design with 90 conditions. One thousand (1000) trials were collected for each condition, yielding 90,000 runs of the model. A pilot model, involving far fewer runs, is also discussed briefly.

Functioning of the Model – Interruptions: Conceptually, the functioning of an interruption in the model is handled in the following way: while the agent is attempting to navigate to the goal location (the primary goal of the agents), if the bottom-up visual system (the *entity detection* module) detects an enemy in a location ahead of the agent, the navigation planning unit is abandoned upon completion of the current unit task, and the “engage enemy” planning unit is initiated. This is made possible by a production in the bottom-up visual module that instructs the agent to abandon the current planning unit and respond to the noticed entity. Once the “engage enemy” planning unit is completed, the agent resumes the navigation planning unit. There are also “bunny” characters in the environment, which are noticed by the bottom-up visual system but ignored (the agent carries on navigating). These are meant to illustrate cases when the bottom-up visual system notices something in the environment, temporarily interrupting the agent, but not requiring further action by the agent.

Functioning of the Model – Communication: Agents begin the simulation without knowledge of how to navigate to the goal location. Once one of the agents spots the goal, this agent communicates to the others. The broadcasting agent uses a production to transmit to the other agents a planning unit, which is a chunk that the receiving agents store in declarative memory. Once this information is in the declarative memory of the remaining agents, they stop moving randomly, and instead use the new navigation planning unit to move toward the goal.

Evaluation of the Model

As noted above, the present chapter is concerned with computational issues in the Python ACT-R architecture, namely how to effectively implement interruption and communication behaviors. The intent was not to compare model behavior to human data and determine goodness of fit of the model. While this is valuable, such an effort would be premature in the present case. Instead, the goal was to demonstrate through implementation the viability of a bottom-up interruption and communication system in an SGOMS ACT-R framework. The evaluation of the two components is presented below.

Evaluation of Interruption Mechanism

The challenge of implementing an interrupt mechanism in Python ACT-R is a feature of the ACT architecture: it is meant to model a top-down driven system that is tailored to routine serialized behavior. Inherent in the design of the architecture are limitations on how sequences of action can be executed and interrupted. There is no default mechanism in ACT by which information from the environment can gracefully interrupt an agent mid-task, and attempting to do will normally terminate the model prematurely.

To model the capacity to handle interruptions, bottom-up driven modules were used, in conjunction with the SGOMS task framework based on unit tasks and planning units. By representing task information in the agent according to SGOMS theory, and providing the agent with the bottom-up visual buffering capacity, it was possible to demonstrate that interruptions could be handled without crashing the model.

Evaluating the interruption mechanism thus amounted to building a model that could execute the desired behavior without crashing. The model is, in effect, an existence proof. Successfully building the model suggested that the proposed mechanism is computationally viable.

An example run from the model is presented in Appendix E, where it can be seen by reviewing the output that the agent is able to interrupt their current planning unit when necessary without the model crashing (see, for example). Once the interruption has been dealt with, the agent can resume carrying out their previous planning unit, and continues this cycle until completion of the main task.

Evaluation of the Communication Mechanism

The communication behavior of the model served two roles: to test the implementation of the bottom-up buffering system in a second context, and to explore the notion of group efficiency gains as a function of communication (see Chapter 4). As with the interrupt mechanism, the successful functioning of the model supports the viability of the implemented buffer mechanism. As for the second question, concerning gains in group efficiency, the simulation output data suggests that the SGOMS ontology concerning expert communication, derived in the previous chapter, is a promising approach to modeling this capacity in experts. The following section will examine the timing data from the model, which suggests that using the SGOMS ontology in conjunction with both computational and anthropologically-inspired methods may be a fruitful framework for building cognitive models of expert teams.

Quantitative Analysis of Single Agent Pilot

The first analysis investigated the average time taken for a single agent to reach the goal square in the Wander Model, as a function of world size and delta. Initial delta is defined as the number of squares between the agent's starting position and the position of the goal. These two values, world size and delta, are related but not identical. The maximum delta of any world size is given by $(2W_s - 3)$, where W_s = size of the world on one side (i.e., length or width of world). For example, a world size of 3 x 3 gives $2(3) - 3 = 3$, and a world size of 7 x 7 gives $2(7) - 3 = 11$. This is illustrated below in Figure 10.

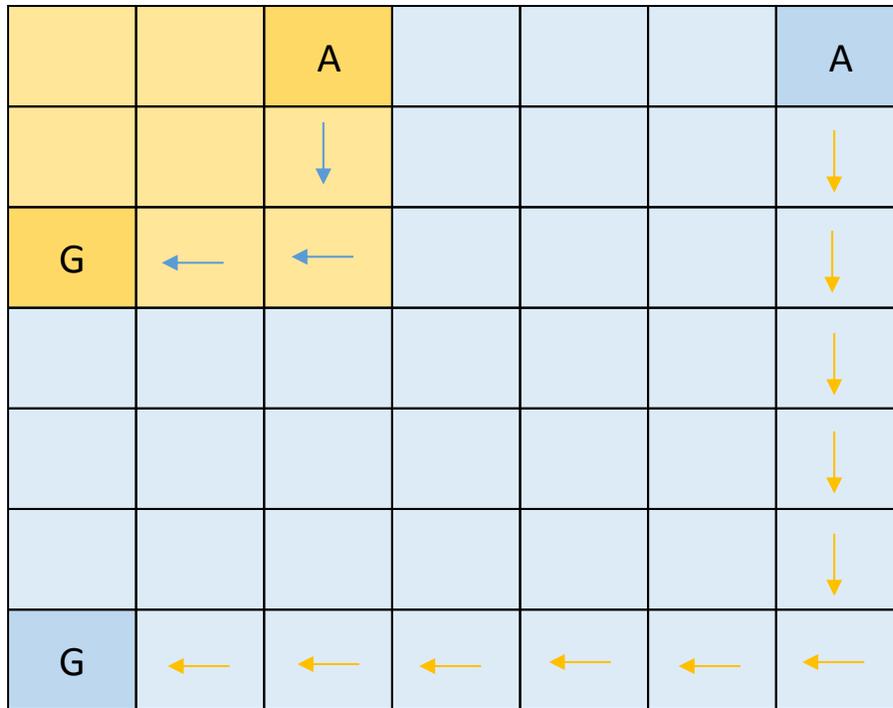


Figure 10. Grid world with maximum initial delta. World size determines maximum distance an agent can spawn from a goal. Shown are size 3x3 (yellow) and 7x7 (blue).

The data show that both increasing world size and increasing delta lead to an increase in average time for the agent to reach the goal. This was expected, of course, and is reported simply to highlight that larger worlds require, on average, more time to navigate. The results of 1000 trials with a single agent in worlds of size 5, 7, 9, and 11, and 13 are shown below (total of 5000 runs). Figure 11 shows average time by world size, and Figure 12 shows average time by initial delta.

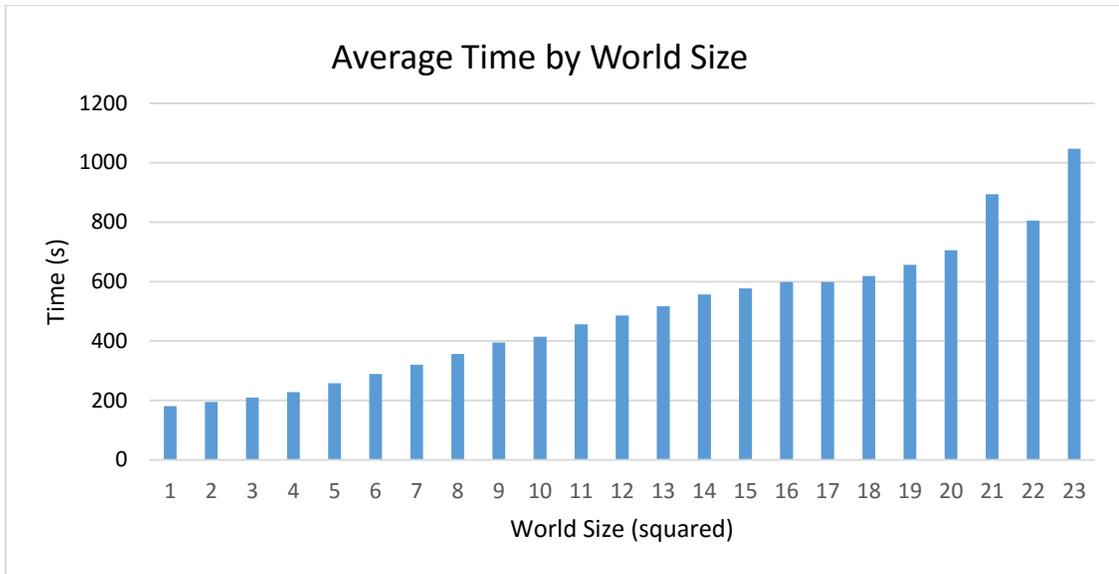


Figure 11. Average time to completion for a single agent in Wander Model, based on size of world. Number on x-axis represents the root of total map-squares (i.e., world size 5 = 5², or 25 squares).

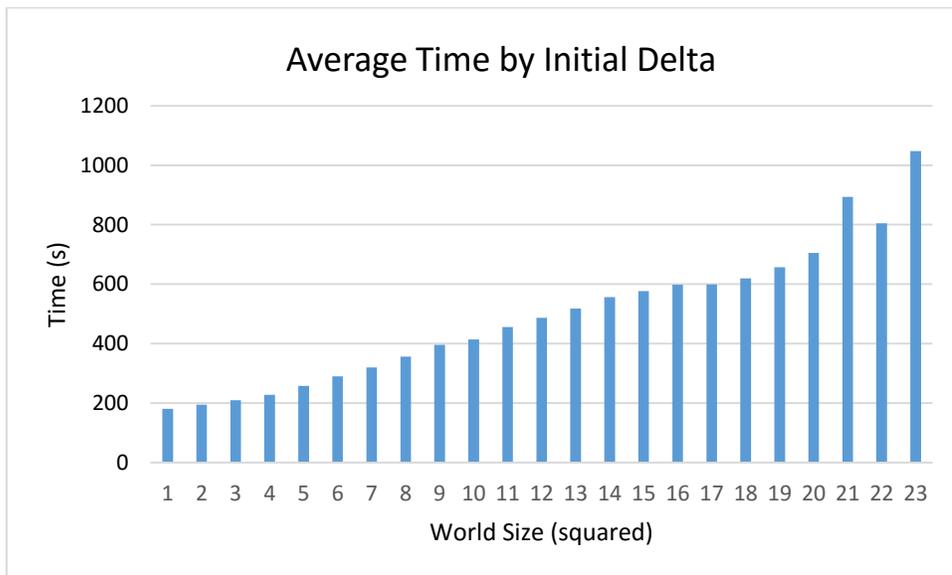


Figure 12. Average time to completion for a single agent in Wander Model, based on initial delta.

Quantitative Analysis of Wander Model

As mentioned above, each variant of the model was run under 45 conditions (5 world sizes by 9 agent counts), for 1000 runs per condition, yielding 45,000 total runs. The wander model, in which agents each had to find their own way to the goal by moving randomly, is presented here.

Data showed that as world size increased, so too did average time for completion, reproducing the result from the single-agent pilot presented above. This is presented in Figure 13. All times below are presented in seconds of simulation time.

As the number of agents in the simulation increased, the average time to completion for the group also increased. Agents could not occupy the same square on the map. The additional time required to complete the simulation was thus the result of agents being forced to navigate around more “obstacles”, i.e., other agents. This increase in time to completion is shown in Figure 14.

Figure 15 is an alternative representation of the simulation data presented in Figures 13 and 14. Figure 15 breaks down average timing according to both world size and agent count.

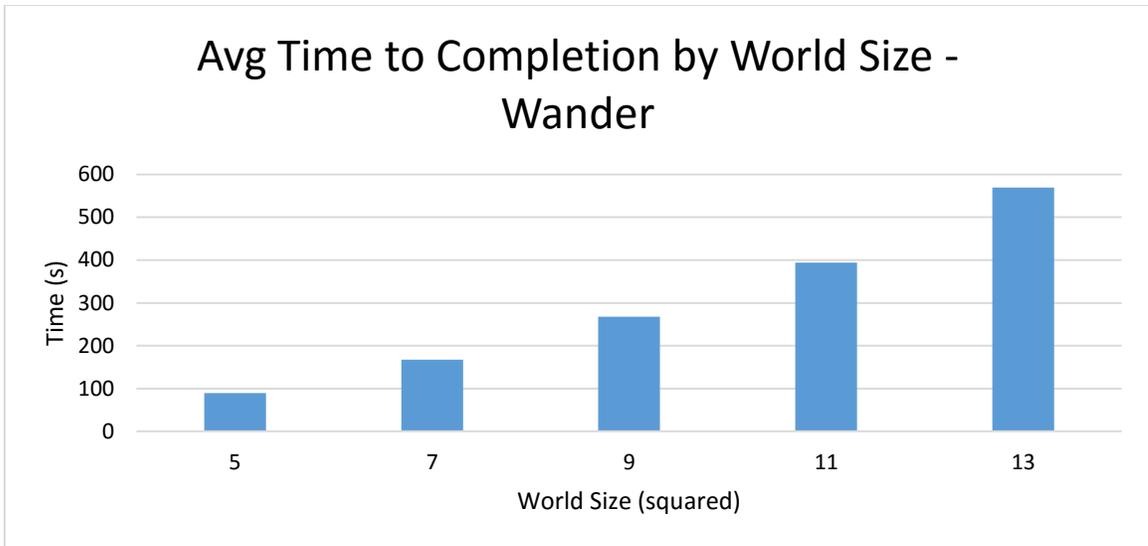


Figure 13. Average time to completion in multi-agent wander model, based on world size.

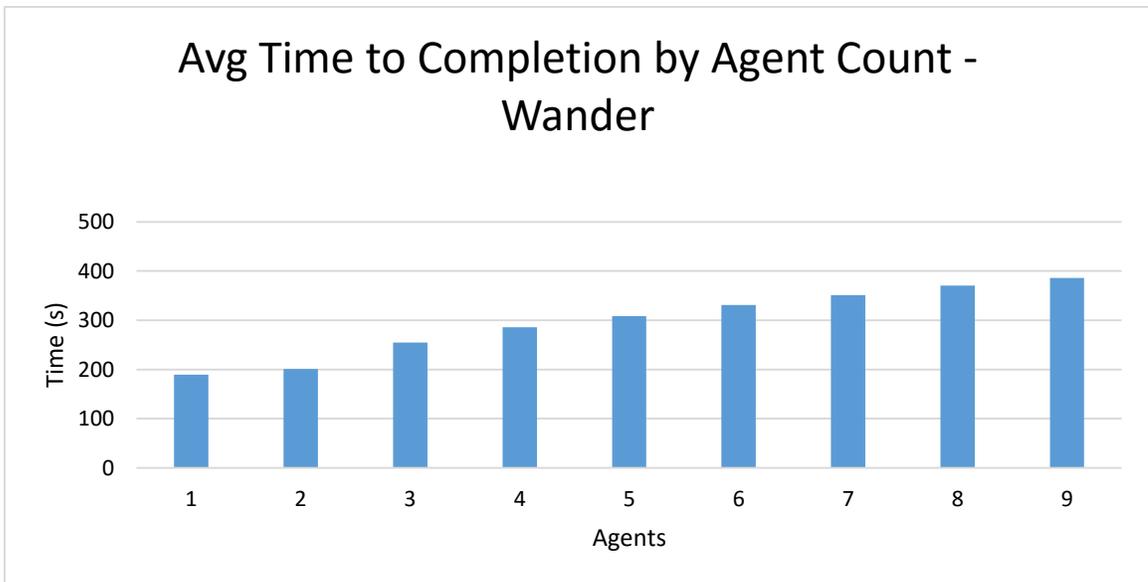


Figure 14. Average time to completion in multi-agent wander model, based on agent count.

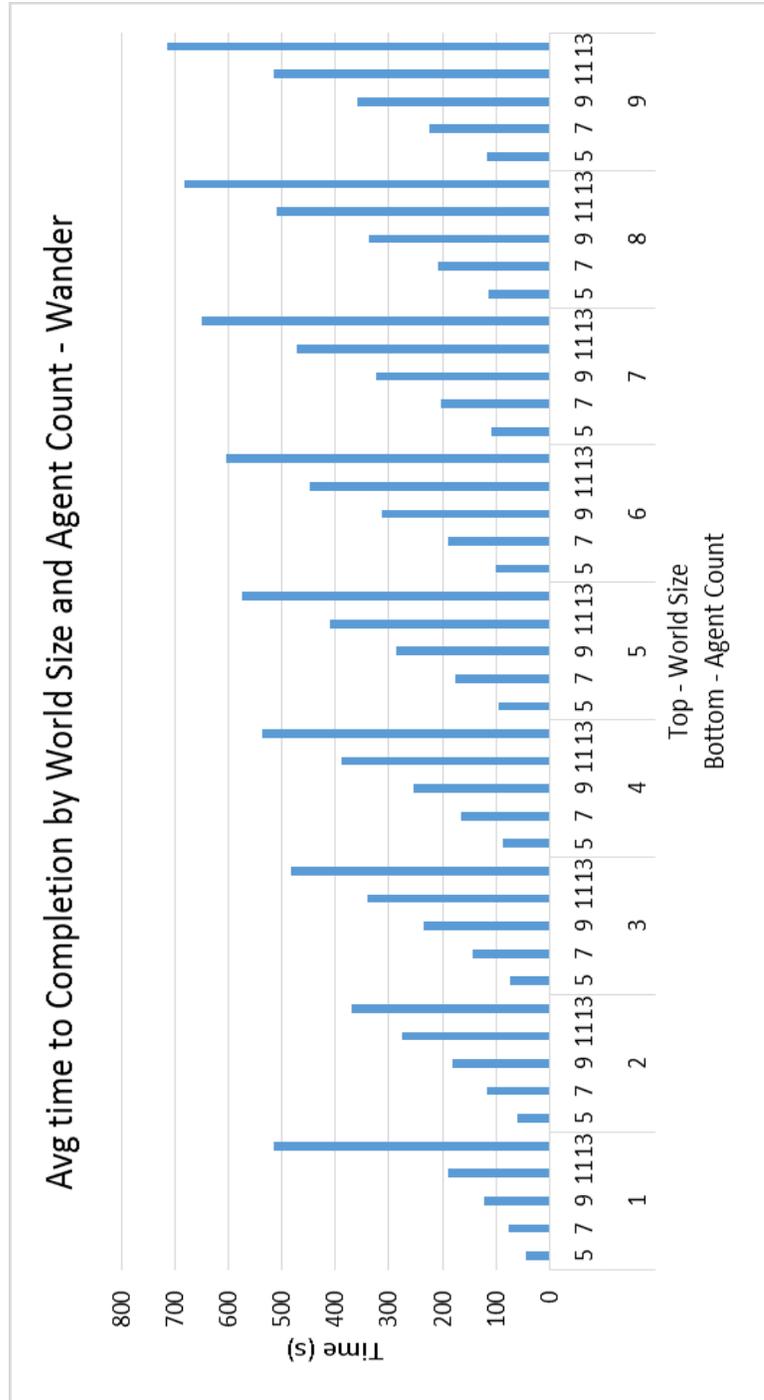


Figure 15. Average time to completion in multi-agent wander model, based on world size and agent count.

Quantitative Analysis of Communication Model

The communication model was also run in 45 conditions (5 world sizes by 9 agent counts) for 1000 trials each, totaling 45,000 runs. As with the wander model, average time to completion increased as world size increased, averaged over all agent counts. This is presented below in Figure 16.

Figure 17 presents average time to completion according to agent count in the communication model. Note that the pattern is reversed in the communication model as compared to the wander model: instead of average time to completion increasing as the agent count increases, time to completion decreases as more agents are added to the simulation. The reason is simple: as soon as the first agent locates the goal, the rest of the agents are informed about how to reach it, and converge quickly upon it. The more agents there are in the simulation, the sooner one of them will find the goal. This was the key hypothesis to be tested in implementing the communication module in the model.

Figure 18 presents the combined data from Figures 16 and 17. It is clear from the image that adding more agents to the simulation more than offsets the increase in time to completion seen in the wander model. This is taken as support for a successful implementation of an SGOMS ACT-R mechanism that can produce “greater than additive gains” through group cooperation.

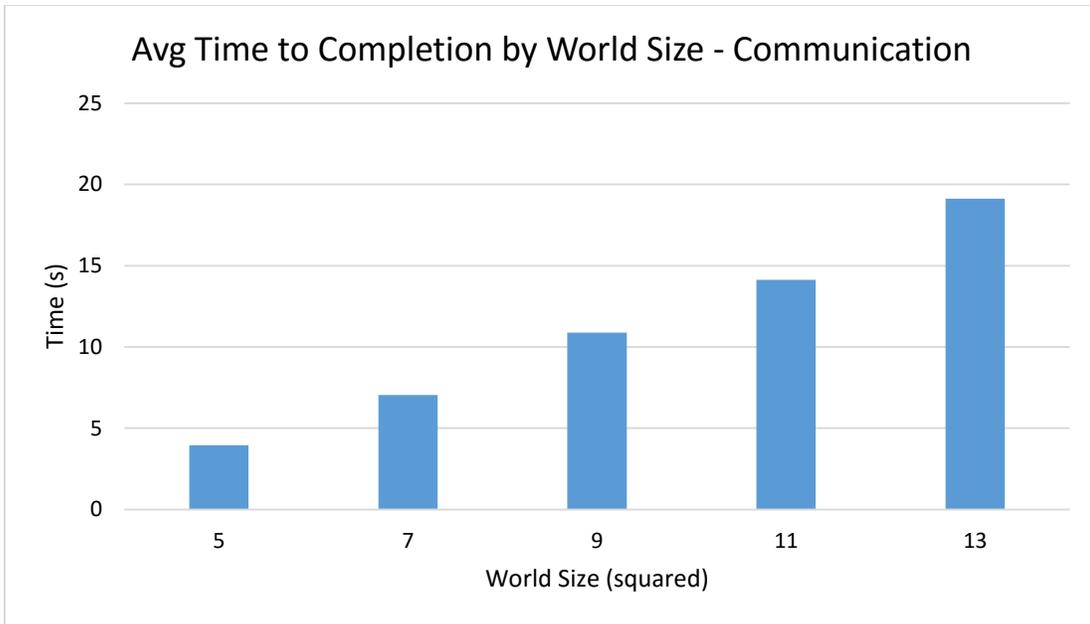


Figure 16. Average time to completion in multi-agent communication model, based on world size

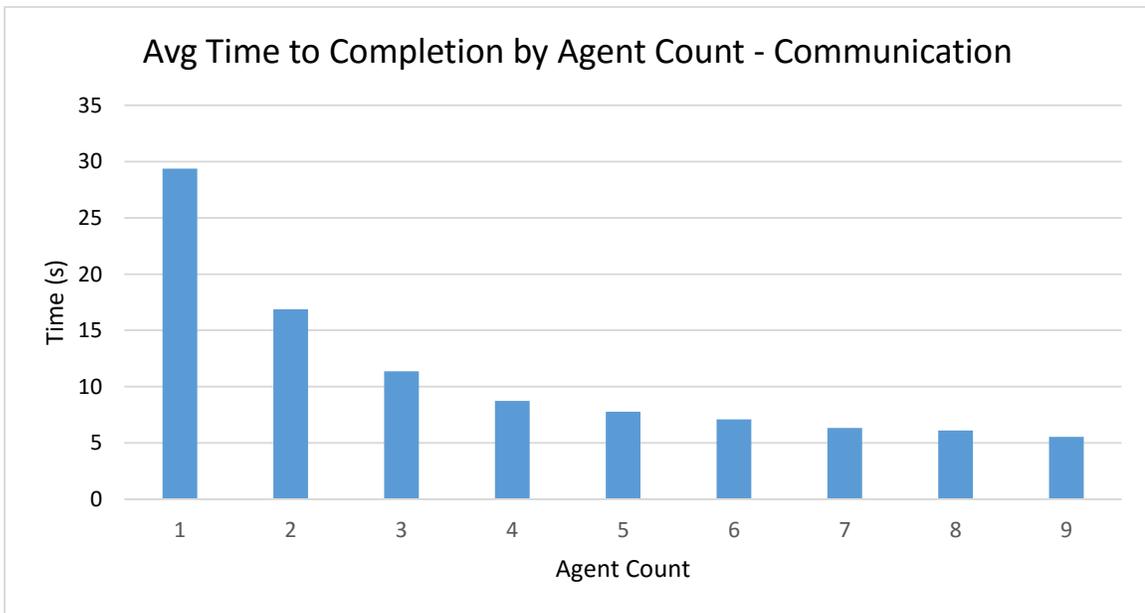


Figure 17. Average time to completion in multi-agent communication model, based on number of agents, averaged across world size.

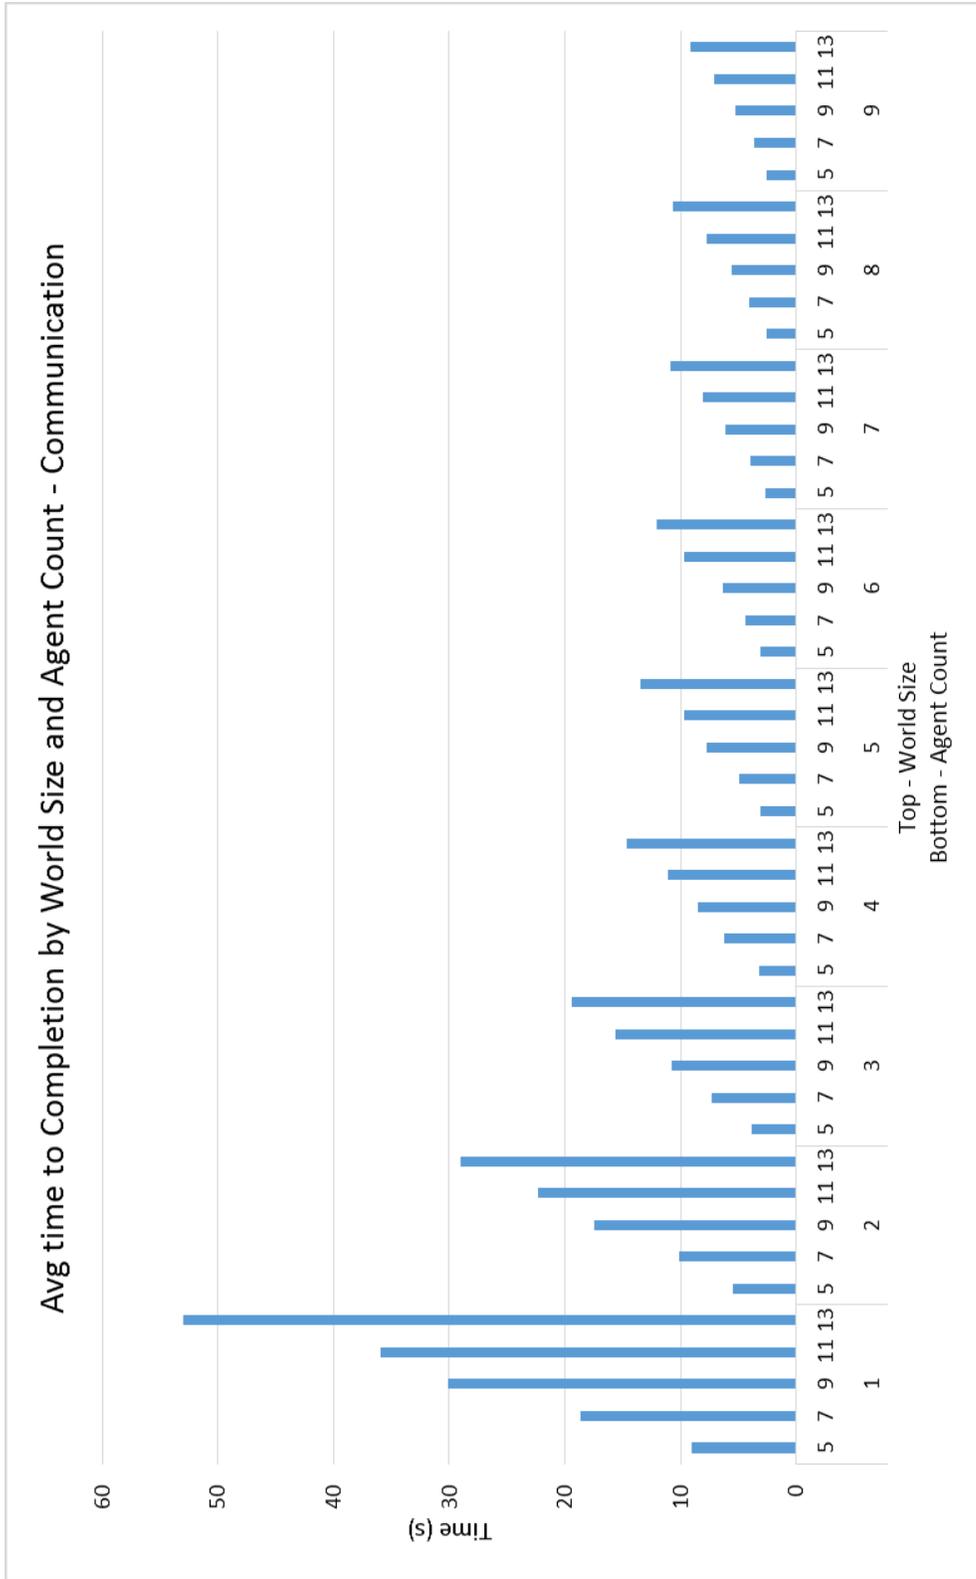


Figure 18. Average time to completion in multi-agent communication model, based on world size and number of agents.

Conclusions

Dealing effectively with interruptions and communication are two key components of expertise in the real world. The models presented in this chapter have demonstrated an implementation of these two elements in an SGOMS cognitive model built in Python ACT-R.

Modeling interruptions in Python ACT-R is difficult due to the strongly top-down nature of the architecture. By using the bottom-up buffering system presented above, and the SGOMS ontology of task structure, it was possible to model cognitive agents that are capable of dealing gracefully with interruptions while completing complex tasks.

Communication is a complex component of expertise and human behavior more generally. The communication model discussed above is intended to address a tightly circumscribed aspect of expert communication, informed by the SGOMS modeling framework: the communication of planning units by cooperative agents. Data collected from the presented implementation shows a clear trend of improved performance by the group when communication of planning units is enabled.

Chapter 6 – Summary of Contributions, Conclusions, and Future Directions

This work has attempted to demonstrate the viability of a particular multi-method approach to the problem of constructing a unified theory of expertise. I have tried to argue that the notion can be effectively tackled at the level of macro cognition, and that a cognitive architectural perspective (the Universal Architecture of Expertise hypothesis), an approach of Methodological Pipelining, and a contextualizing and evaluating of the work in a broader scientific community according to principles of Lakatosian Analysis can provide early purchase on the construction of such a theory. The work thus far suggests that such a research programme has legs. This chapter will summarize the contributions presented in this thesis, and relate them back to the larger theoretical and methodological goals outlined in the beginning of the document.

Summary of Contributions

The first study presented a novel methodology for creating macro cognitive models of performance in naturalistic contexts. Using a combination of SGOMS macro-architecture ideas, video annotation, and iterative model tracing, we demonstrated a viable method for capturing and predicting behavior in complex scenarios. This approach could usefully be applied to modelling a range of other expert tasks, and offers a principled way of creating evaluating such models.

The second study produced an early-stage typology of expert communication behaviors, inspired by the Universal Architecture of Expertise hypothesis and the SGOMS macro-cognitive modeling framework. Using concepts from these two positions, I was

able to extract usable insights into expert coordination that are valuable for the effort of scaling up single agent macro cognitive models to multi agent, coordinated task performance. Specifically, I identified two patterns in communication, related to planning units and interruptions, which were later incorporated into the computational implementation of Study 3. Given the real-world importance of coordination and communication in expert activities, such analyses are a valuable component of a modeling paradigm directed at multi agent expert modeling.

The third study demonstrated a macro-architecture inspired mechanism for graceful interruptions in Python ACT-R. This allowed agents to quickly change tasks in light of unexpected information without crashing the simulation. The ability to effectively handle interruptions is an important component of expertise in many domains, and therefore an approach aimed at modeling naturalistic expert cognition ought to be capable of handling such interruptions. Traditional expert systems and simulation architectures have difficulty coping with such conditions. The SGOMS task structure of unit tasks and planning units provided a theoretical scaffolding around which to test ideas about graceful interruptions, and the mechanism was successfully implemented.

The third study also demonstrated a communication system implemented in Python ACT-R that provided a cognitively plausible account of the basis for greater-than-additive gains in cooperative, multi-agent scenarios. Using the concept of shared planning units, I was able to analyze and computationally implement ideas generated through the communication analyses from the second study. Cooperation and the

division of labour are features of real-world expertise than have been discussed at length in many disciplines, but we do not yet have a fully satisfactory theory of the cognitive capacities that underlie this ability. This implementation is a step in that direction.

Finally, the series of 3 studies taken together demonstrated the viability of the notion of Methodological Pipelining. By chaining together methods, I was able to generate novel research questions, and effectively combine ideas and techniques from disparate fields. I have argued through this thesis that such an approach is necessary for achieving the goal of synthesizing the vast literature on expertise, and for constructing more comprehensive theories of expert cognition and behavior. I hope that the conceptual unity that runs through the research programme presented here is taken as evidence for the value of such a multi-modal approach.

Conclusions and Future Directions

Expertise is an enormously important dimension of human life. It impacts our ability to survive, our social hierarchies, our economies, the careers and leisure activities we pursue, and the links between generations. The suite of cognitive and cultural factors that are involved in the development and deployment of expertise make the topic an ideal candidate for interdisciplinary research. Many fields have turned their attention to the problem, and a wealth of information has been accumulated. There remains a great deal of targeted work to be done, by memory and personality psychologists, educational policy makers, cultural anthropologists, economists, narrative theorists, artificial intelligence researchers, and computational cognitive modelers. But this targeted work

must be complemented by efforts at synthesis, at creating unified models of expertise that incorporate information from all of these disciplines. This is a tall order, as it requires a cooperated community of analysts that is fluent in many different languages. While computers and the internet have made it easier than ever before to collaborate, it has also reduced the cost of publishing and distributing to next to nothing, driving down the signal-to-noise ratio in scientific writing, flooding the marketplace of ideas, and forcing even narrower specialization if the goal is mastery.

I view the work presented here as a prototype for what an interdisciplinary research program into expertise might look like. It does not offer a final answer concerning what expertise is, how it can be augmented, or how it ought to be studied. It has aimed at the more modest goal of testing a combination of methodologies to see whether such chaining can be fruitful. There are dozens of research methods that could have been chosen for such a pipeline; hundreds if the grain size is made smaller. I make no claim as to the superiority of the methodological mix presented here.

While the work presented in a dissertation is in some important sense meant to be that of an individual, I don't see unification as the purview of an individual. I don't believe, for example, that in an increasingly interconnected world, theoretical synthesis is the last bastion of individualized genius; that a writer alone in a room can pull it all together. I suspect that the recent paper that emerged from Large Hadron Collider in Geneva (ATLAS Collaboration & CMS Collaboration, 2015), announcing a refined measurement of the mass of the Higgs-Boson and listing 5,124 authors, is indicative of where science is heading. In other words, the "methodological pipeline" that I have

presented here is not offered as *my* solution to unification. I think it is more interesting to view it from a distance, as a potential model for the activity of communities of researchers that I have attempted to channel. In this sense it is not a paradigm or set of steps that I would follow in my own work, but a way of framing a relationship between different research communities and their output. The problem is not “how can I make sense of all this?”, but “how ought the work be structured so that we can make sense of it together?” Perhaps research and writing that draws on multiple sources in its construction is also more inviting to readers from diverse backgrounds, and serves therefore to link communities. This is a relatively simple point, but I think an important one. I suspect there is significant value in conducting further work toward normalizing and systematizing the generation and communication of scientific output such that it can be more easily incorporated into inter-/multi-disciplinary work.

Moving forward, the topic of expertise seems an inexhaustible source of interesting questions, several of which I hope to tackle in the not-too-distant future. Four of the most intriguing involve fictional experts, economic narratives, grey market expertise, and emerging technologies scaffolding expertise. Concerning the first, I am intrigued by the possibility that experts portrayed in fiction, particularly novels which provide simulated access to the mind of another individual, may support the social distribution of labour as discussed in Chapter 4. I think that the principle advantage of having well-defined expert groups, such as doctor and mechanic, is that it reduces the cognitive load on individuals living in complex societies. Perhaps fiction aids in propagating these relationships-to-experts through society, instructing readers about

where to turn for various needs. I believe computational linguistic analysis of corpora could offer insight into this question.

By economic narratives I refer to the stories that frame the way that members of a society feel about their job or career. I suspect evidence can be found that changes in national economic policy are likely reflected in the forms of narratives that are presented to the public. As an example, in America there is currently an upswing in money being invested in technology startups, mimicking the dot com boom of the 90s. At the same time, there have been a number of television programs, both fiction and non, that lionize entrepreneurs and technologists (e.g., Shark Tank, Silicon Valley). There are many ways to mobilize a work force and develop the expertise of a people; stories are likely one of the most potent.

Grey market expertise draws on the notion of gray markets from economics. These are constituted by economic activities which pass through unofficial channels. Though not exactly illegal, such activities are also not quite legitimate; hence, gray. An example is the import and unofficial resale of electronics. In the case of expertise, the Internet has enabled individuals to access information through many new channels, and these channels are often unregulated. A simple example is medical information and the so-called “Doctor Google” effect (Bouwman, Teunissen, Wijburg, & Linthorst, 2010; Leune & Nizard, 2012). People with access to the Internet can find hundreds or thousands of pages describing their symptoms, and may use this to make inferences about what may be ailing them. The problem is that most people are not doctors, and are not trained to differentiate sound medical advice from quackery. Doctors and

hospitals exist precisely to take that responsibility out of our hands, and a key function of medical licensing is to reassure the patient that the doctor possesses the expertise she claims to. The same sort of uncertainty is now present in essentially all domains, because anyone who wants to pay for webhosting can pass themselves off as an expert online. What I would like to investigate is the variety of mechanisms by which expertise can be accurately signaled in these new channels of information distribution. As mentioned in Chapter 4, many websites implement thumbs-up/down voting systems to allow the best answer to a question to rise to the top, and various other mechanisms are also in use. I believe people will become increasingly savvy concerning information consumption, in the same way that most internet and email users can now immediately recognize spam for what it is. This makes for a second, related topic of interest: the expertise required to effectively extract information from a web filled with noise and false signals.

Finally, I would like to explore the role that technology plays in enhancing and transmitting expertise. The last 3 topics point toward this, referring to novels, television, and the internet. Each new medium has great potential for training and collaboration, thereby impacting the shape of expertise. I believe that virtual reality and augmented reality, two forms of immersive storytelling and data visualization (among many other uses) that appear to be on the verge of mass public adoption, are going to have an enormous impact on human abilities. These technologies will find extensive use in training people on a host of different skills, possibly also giving rise to novel forms of

expertise that have yet to be imagined. This particular intersection between technology and expertise is one that I will be tracking very closely.

References

- Abbott, A. (1988). *The system of professions: An essay on the division of expert labor* (1st ed.). Chicago: University of Chicago Press.
- Adams, D. (2002). *The salmon of doubt: Hitchhiking the universe one last time*. New York: Harmony.
- Aikins, J. S. (1983). Prototypical knowledge for expert systems. *Artificial Intelligence*, 20(2), 163–210. [http://doi.org/10.1016/0004-3702\(83\)90017-6](http://doi.org/10.1016/0004-3702(83)90017-6)
- Alexander, P. (2003). The development of expertise: The journey from acclimation to proficiency. *Educational Researcher*, 32(8), 10–14. <http://doi.org/10.3102/0013189X032008010>
- Amant, R. S., Freed, A. R., & Ritter, F. E. (2005). Specifying ACT-R models of user interaction with a GOMS language. *Cognitive Systems Research*, 6(1 SPEC.ISS.), 71–88. <http://doi.org/10.1016/j.cogsys.2004.09.008>
- Anderson, C. A., & Bushman, B. J. (1997). External validity of “trivial” experiments: The case of laboratory aggression. *Review of General Psychology*, 1(1), 19–41. <http://doi.org/10.1037/1089-2680.1.1.19>
- Anderson, J. R. (1996). ACT: A simple theory of complex cognition. *American Psychologist*, 51(4), 355–365. <http://doi.org/10.1037/0003-066X.51.4.355>
- Anderson, J. R. (1998). Social stimuli and social rewards in primate learning and cognition. *Behavioural Processes*, 42(2-3), 159–175. [http://doi.org/10.1016/S0376-6357\(97\)00074-0](http://doi.org/10.1016/S0376-6357(97)00074-0)
- Anderson, J. R. (2005). Human symbol manipulation within an integrated cognitive architecture. *Cognitive Science*, 29(3), 313–341.

- Anderson, J. R. (2007). *How can the human mind occur in the physical universe?*
<http://doi.org/10.1093/acprof:oso/9780195324259.001.0001>
- Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review*, *111*(4), 1036–60. <http://doi.org/10.1037/0033-295X.111.4.1036>
- Anderson, M. L. (2010). Neural reuse: A fundamental organizational principle of the brain. *The Behavioral and Brain Sciences*, *33*(4), 245–266; discussion 266–313.
<http://doi.org/10.1017/S0140525X10000853>
- Anjaneyulu, K. S. R. (1998). Expert systems: An introduction. *Resonance*, *3*(3), 46–58.
<http://doi.org/10.1007/BF02837614>
- Argyris, C. (1973). Some limits of rational man organizational theory. *Public Administration Review*, *33*(3), 253–267. <http://doi.org/10.2307/974803>
- Artman, H., & Garbis, C. (1998). Team communication and coordination as distributed cognition. *Proceedings of 9th Conference of Cognitive Ergonomics: Cognition and Cooperation*, 151–156. Retrieved from <http://kth.diva-portal.org/smash/record.jsf?pid=diva2:479674>
- Atkinson, P. (1988). Ethnomethodology: A critical review. *Annual Review of Sociology*, *14*(1), 441–465. <http://doi.org/doi:10.1146/annurev.so.14.080188.002301>
- ATLAS Collaboration, & CMS Collaboration. (2015). Combined measurement of the Higgs Boson mass in pp collisions at $\sqrt{s}=7$ and 8 TeV with the ATLAS and CMS experiments. *Phys. Rev. Lett.*, *114*(19), 191803. <http://doi.org/10.1103/PhysRevLett.114.191803>
- Baas, N., & Emmenche, C. (1997). On emergence and explanation. *Intellectica*, *2*, 67–83.

- Baddeley, A. (2003). Working memory and language: An overview. In *Journal of Communication Disorders* (Vol. 36, pp. 189–208). [http://doi.org/10.1016/S0021-9924\(03\)00019-4](http://doi.org/10.1016/S0021-9924(03)00019-4)
- Baer, J. (2012). Domain specificity and the limits of creativity theory. *Journal of Creative Behavior*. <http://doi.org/10.1002/jocb.002>
- Bainbridge, L. (1997). The change in concepts needed to account for human behavior in complex dynamic tasks. *IEEE Transactions on Systems, Man, and Cybernetics Part A: Systems and Humans*. <http://doi.org/10.1109/3468.568743>
- Balthazard, P. A., Potter, R. E., & Warren, J. (2002). The effects of extraversion and expertise on virtual team interaction and performance. In *Proceedings of the 35th Annual Hawaii International Conference on System Sciences* (Vol. January, p. 10). <http://doi.org/10.1109/HICSS.2002.994426>
- Bandura, A. (1971). *Social learning theory. Social Learning Theory*. New York: General Learning Press. <http://doi.org/10.1111/j.1460-2466.1978.tb01621.x>
- Barrett, H. C., & Kurzban, R. (2006). Modularity in cognition: Framing the debate. *Psychological Review*, 113(3), 628–647. <http://doi.org/10.1037/0033-295X.113.3.628>
- Bednarz, J. C. (1988). Cooperative hunting Harris' hawks (*Parabuteo unicinctus*). *Science*, 239(4847), 1525–7. <http://doi.org/10.1126/science.239.4847.1525>
- Beilock, S. L., Wierenga, S. A., & Carr, T. H. (2003). Memory and expertise: What do experienced athletes remember? In *Expert performance in sports: Advances in research on sport expertise* (pp. 295–320). Champaign, IL: Human Kinetics.
- Bender, A., Hutchins, E., & Medin, D. (2010). Anthropology in cognitive science. *Topics in Cognitive Science*, 2, 374–385. <http://doi.org/10.1111/j.1756-8765.2010.01082.x>

- Bennett, L. J. (1990). Modularity of mind revisited. *British Journal for the Philosophy of Science*.
<http://doi.org/10.1093/bjps/41.3.429>
- Berkowitz, A. L., & Ansari, D. (2010). Expertise-related deactivation of the right temporoparietal junction during musical improvisation. *NeuroImage*, *49*(1), 712–719.
<http://doi.org/10.1016/j.neuroimage.2009.08.042>
- Berkowitz, L., & Donnerstein, E. (1982). External validity is more than skin deep: Some answers to criticisms of laboratory experiments. *American Psychologist*, *37*(3), 245–257.
<http://doi.org/10.1037/0003-066X.37.3.245>
- Bernard, A. B., Redding, S. J., & Schott, P. K. (2007). Comparative advantage and heterogeneous firms. *Review of Economic Studies*, *74*(1), 31–66. <http://doi.org/10.1111/j.1467-937X.2007.00413.x>
- Berube, M. (2011). The science wars redux. *Democracy*, *19*(Winter), 64–74.
<http://doi.org/10.1119/1.19040>
- Beshers, S. N., & Fewell, J. H. (2001). Models of division of labor in social insects. *Annual Review of Entomology*, *46*, 413–440. <http://doi.org/10.1146/annurev.ento.46.1.413>
- Bizzell, P. (1992). *Academic discourse and critical consciousness. Pittsburgh series in composition, literacy, and culture*. Pittsburgh: University of Pittsburgh Press.
<http://doi.org/10.2307/358592>
- Bialic, M., Turella, L., Campitelli, G., Erb, M., & Grodd, W. (2012). Expertise modulates the neural basis of context dependent recognition of objects and their relations. *Human Brain Mapping*, *33*(11), 2728–2740. <http://doi.org/10.1002/hbm.21396>
- Bloom, B. S. (1985). Generalizations about talent development. In B. S. Bloom (Ed.), *Developing*

- talent in young people* (pp. 507–549). New York: Ballantine Books.
- Bloor, D. (1991). The strong programme in the sociology of knowledge. *Knowledge and Social Imagery*, 2, 3–23. <http://doi.org/10.1016/j.shpsa.2003.12.011>
- Bloor, D. (2004). Sociology of scientific knowledge. In *Handbook of epistemology* (pp. 919–962). Boston: Kluwer Academic Publishers. <http://doi.org/10.1007/BF00993464>
- Blum, B. I. (1996). *Beyond programming: To a new era of design*. New York: Oxford University Press.
- Bogner, A., Littig, B., & Menz, W. (Eds.). (2009). *Interviewing Experts*. Basingstoke: Palgrave Macmillan. <http://doi.org/10.1057/9780230244276>
- Böhme, G. (1975). The social function of cognitive structures: A concept of the scientific community within a theory of action. In K. D. Knorr, H. Straasser, & H. G. Zilian (Eds.), *Determinants and controls of scientific development* (pp. 205–225). Dordrecht: Reidel Publishing Company.
- Boldrin, M., & Levine, D. K. (2009). Does intellectual monopoly help innovation? *Review of Law Economics*, 5(3), 991–1024. <http://doi.org/10.2202/1555-5879.1438>
- Bolger, F., & Wright, G. (1992). Reliability and validity in expert judgment. In *Expertise and decision support* (pp. 47–76). New York: Springer US.
- Bonabeau, E., Sobkowski, A., Theraulaz, G., & Deneubourg, J. (1997). Adaptive task allocation inspired by a model of division of labor in social insects. In *Biocomputing and Emergent Computation* (pp. 36–45). <http://doi.org/10.1.1.25.8016>
- Bouilloud, J.-P. (2003). Review essay: The reception of the Sokal affair in France--“Pomo”

hunting or intellectual McCarthyism? A propos of impostures intellectuelles by A. Sokal and J. Bricmont. *Philosophy of the Social Sciences*, 33(1), 122–137.

<http://doi.org/10.1177/0048393102250309>

Bouwman, M., Teunissen, Q., Wijburg, F., & Linthorst, G. (2010). “Doctor Google” ending the diagnostic odyssey in lysosomal storage disorders: Parents using internet search engines as an efficient diagnostic strategy in rare diseases. *Archives of Disease in Childhood*, 95(8), 642–644.

Boyke, J., Driemeyer, J., Gaser, C., Büchel, C., & May, A. (2008). Training-induced brain structure changes in the elderly. *J Neurosci*, 28(28), 7031–7035.

<http://doi.org/10.1523/JNEUROSCI.0742-08.2008>

Bozeman, B. (1993). A Theory Of Government “Red Tape.” *Journal of Public Administration Research and Theory*, 3(3), 273–303. <http://doi.org/10.2307/1181785>

Brajnik, G., Guida, G., & Tasso, C. (1990). User modeling in expert man-machine interfaces: a case study in intelligent information retrieval. *IEEE Transactions on Systems, Man, and Cybernetics*, 20(1), 166–185. <http://doi.org/10.1109/21.47819>

Brandenburger, A. M., & Nalebuff, B. J. (1995). The right game: Use game theory to shape strategy. *Long Range Planning*, 28(5), 128. [http://doi.org/10.1016/0024-6301\(95\)90326-7](http://doi.org/10.1016/0024-6301(95)90326-7)

Brand-Gruwe, S., Wopereis, I. G. J. H., & Vermetten, Y. (2005). Information problem solving by experts and novices: Analysis of a complex cognitive skill. *Computers in Human Behavior*, 21, 487–508. <http://doi.org/10.1016/j.chb.2004.10.005>

Brante, T. (1988). Sociological Approaches to the Professions. *Acta Sociologica*, 31(2), 119–142.

<http://doi.org/10.1177/000169938803100202>

- Bresnahan, T. F., & Reiss, P. C. (1990). Entry in monopoly markets. *The Review of Economic Studies*, 57(4), 531–553. <http://doi.org/10.2307/2298085>
- Bressler, S. L., & Menon, V. (2010). Large-scale brain networks in cognition: Emerging methods and principles. *Trends in Cognitive Sciences*. <http://doi.org/10.1016/j.tics.2010.04.004>
- Brewer, G. A., & Walker, R. M. (2010). The impact of red tape on governmental performance: An empirical analysis. *Journal of Public Administration Research and Theory*, 20(1), 233–257. <http://doi.org/10.1093/jopart/mun040>
- Bricker, L. A., & Bell, P. (2008). Conceptualizations of argumentation from science studies and the learning sciences and their implications for the practices of science education. *Science Education*, 92(3), 473–498. <http://doi.org/10.1002/sce.20278>
- Brint, S. (1994). *In an age of experts: The changing role of professionals in politics and public life*. Princeton, NJ: Princeton University Press.
- Brown, A. D., Coman, A., & Hirst, W. (2009). The role of narratorship and expertise in social remembering. *Social Psychology*, 40(3), 119–129. <http://doi.org/10.1027/1864-9335.40.3.119>
- Brown, J. S., Collins, A., & Duguid, P. (1989). Situated cognition and the culture of learning. *Educational Researcher*, 18(1), 32–42. <http://doi.org/10.3102/0013189X018001032>
- Bruyer, R., & Scailquin, J. C. (1998). The visuospatial sketchpad for mental images: testing the multicomponent model of working memory. *Acta Psychologica*, 98(1), 17–36. [http://doi.org/10.1016/S0001-6918\(97\)00053-X](http://doi.org/10.1016/S0001-6918(97)00053-X)
- Buchanan, B. G., & Shortliffe, E. H. (1984). *Rule-based expert systems: The MYCIN experiments of the Stanford Heuristic Programming Project*. Reading, MA: Addison Wesley.

- Bullinaria, J. A. (2007). Understanding the emergence of modularity in neural systems. *Cognitive Science*, 31(787188060), 673–695. <http://doi.org/10.1080/15326900701399939>
- Button, G., & Sharrock, W. (2009). Studies of work and the workplace in HCI: Concepts and techniques. *Synthesis Lectures on Human-Centered Informatics*, 2(1), 1–96. <http://doi.org/10.2200/S00177ED1V01Y200903HCI003>
- Byrne, R. W., & Whiten, A. (1989). Machiavellian intelligence: Social expertise and the evolution of intellect in monkeys, apes, and humans. *American Scientist*. [http://doi.org/10.1016/0003-3472\(89\)90054-7](http://doi.org/10.1016/0003-3472(89)90054-7)
- Calder, B. J., Phillips, L. W., & Tybout, A. M. (1982). The concept of external validity. *Journal of Consumer Research*, 9(3), 240. <http://doi.org/10.1086/208920>
- Camerer, C. (1999). Behavioral economics: Reunifying psychology and economics. *Proceedings of the National Academy of Sciences*, 96(19), 10575–10577. <http://doi.org/10.1073/pnas.96.19.10575>
- Cammelli, A., & Socci, F. (1992). Informatics for the law: LEXIS -- A legal expert system on Italian family law. *Expert Systems With Applications*, 4(4), 409–414. [http://doi.org/10.1016/0957-4174\(92\)90133-D](http://doi.org/10.1016/0957-4174(92)90133-D)
- Campbell, J. (1949). *The hero with a thousand faces*. <http://doi.org/10.2307/537371>
- Carbonell, J. G. (1986). Derivational analogy: A theory of reconstructive problem solving and expertise acquisition. In R. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach* (Vol. 2, pp. 371–392).
- Card, S. K., Moran, T. P., & Newell, A. (1986). The model human processor: An engineering model for human performance. *Handbook of Perception and Human Performance*.

<http://doi.org/10.1177/107118138102500180>

Chalmers, D. J. (2011). Strong and weak emergence. In *The re-emergence of emergence: The emergentist hypothesis from science to religion*.

<http://doi.org/10.1093/acprof:oso/9780199544318.003.0011>

Chandrashekara, K., & Gadagkar, R. (1991). Behavioural castes, dominance and division of labour in a primitively eusocial wasp. *Ethology*, *87*, 269–283.

<http://doi.org/10.1111/j.1439-0310.1991.tb00252.x>

Chant, J. F. (1963). Irrational behavior and economic theory: A comment. *Journal of Political Economy*, *71*(5), 505. <http://doi.org/10.1086/258800>

Chase, W. G., & Simon, H. (1973a). Perception in chess. *Cognitive Psychology*, *4*, 55–81.

Chase, W. G., & Simon, H. (1973b). The mind's eye in chess. In W. G. Chase (Ed.), *Visual Information Processing* (pp. 215–281). New York: Academic Press.

Cheatham, D., & Lane, S. (2002). Differential access hypothesis: The effects of task and information type on the validity of knowledge acquisition methods. In *Proceedings of the Forty-Sixth Meeting of the Human Factors & Ergonomic Society* (pp. 487–491). Baltimore, MD: Human Factors and Ergonomics Society.

Chen, Y., Hsu, C. Y., Liu, L., & Yang, S. (2012). Constructing a nutrition diagnosis expert system. In *Expert Systems with Applications* (Vol. 39, pp. 2132–2156).

<http://doi.org/10.1016/j.eswa.2011.07.069>

Cheung, O. S., & Bar, M. (2012). Visual prediction and perceptual expertise. *International Journal of Psychophysiology*. <http://doi.org/10.1016/j.ijpsycho.2011.11.002>

- Chi, M. T. (2006). Laboratory methods for assessing experts' and novices' knowledge. In *The Cambridge Handbook of Expertise and Expert Performance* (pp. 167–184).
- Chi, M. T., Feltovich, P. J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science*, *5*, 121–152.
http://doi.org/10.1207/s15516709cog0502_2
- Chi, M. T., Glaser, R., & Rees, E. (1982). Expertise in problem solving. *Advances in the Psychology of Human Intelligence*, *1*, 7–75. <http://doi.org/10.1002/wcs.47>
- Chipman, S. F., Schraagen, J. M., & Shalin, V. L. (2000). Introduction to cognitive task analysis. In J. M. Schraagen, S. F. Chipman, & V. J. Shute (Eds.), *Cognitive task analysis*. Mahwah, NJ: Lawrence Erlbaum Associates. Retrieved from
<http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Introduction+to+Cognitive+Task+Analysis+?#0>
- Clancey, W. J. (2006). Observation of work practices in natural settings. In *The Cambridge handbook of expertise and expert ...* (pp. 127–145). <http://doi.org/10.2277/052184097X>
- Clark, A. (2009). *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*. *Supersizing the Mind: Embodiment, Action, and Cognitive Extension*.
<http://doi.org/10.1093/acprof:oso/9780195333213.001.0001>
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, *58*(January), 7–19.
<http://doi.org/10.1111/j.1467-9744.2009.01021.x>
- Clark, R. E., Feldon, D. F., Van Merriënboer, J. J. G., Yates, K., & Early, S. (2007). Cognitive task analysis. *Handbook of Research on Educational Communications and Technology, 2006*, 577–593. Retrieved from

<http://books.google.com/books?hl=en&lr=&id=gyt5AgAAQBAJ&oi=fnd&pg=PP1&dq=Cognitive+Task+Analysis&ots=cfB0yBQby8&sig=FYrCXZS7Ufem9INWIPkKg9n-wrc>

Cleary, T. J., & Zimmerman, B. J. (2005). Self-regulation differences during athletic practice by experts, non-experts, and novices. *Journal of Applied Sport Psychology, 13*(2), 185–206.
<http://doi.org/10.1080/104132001753149883>

Collins, H. M. (1983). The sociology of scientific knowledge: Studies of contemporary science. *Annual Review of Sociology*. <http://doi.org/10.1146/annurev.so.09.080183.001405>

Collins, H. M. (2001). Scientific knowledge, sociology of. In *International Encyclopedia of the Social & Behavioral Sciences* (pp. 13741–13746). <http://doi.org/10.1016/B978-0-08-097086-8.85029-1>

Collins, H. M., & Evans, R. (2002). The third wave of science studies: Studies of expertise and experience. *Social Studies of Science, 32*(2), 235–296.
<http://doi.org/10.1177/0306312702032002003>

Coltheart, M. (1999). Modularity and cognition. *Trends in Cognitive Sciences, 3*(3), 115–120.
[http://doi.org/10.1016/S1364-6613\(99\)01289-9](http://doi.org/10.1016/S1364-6613(99)01289-9)

Cook, V. J., & Newson, M. (1996). *Chomsky's universal grammar: an introduction*. *Chomsky's universal grammar: An introduction* (2nd ed.). Cambridge, MA: Blackwell Publishers.

Cooke, N. J., Gorman, J. C., & Rowe, L. J. (2009). An ecological perspective on team cognition. In *Team effectiveness in complex organizations: Cross-disciplinary perspectives and approaches*. (pp. 157–182). <http://doi.org/10.1037/001790>

Cooper, R. P. (2007). The role of falsification in the development of cognitive architectures: Insights from a lakatosian analysis. *Cognitive Science, 31*, 509–533.

<http://doi.org/10.1080/15326900701326592>

Corbet, P. S. (1980). Biology of Odonata. *Annual Review of Entomology*, 25, 189–217.

<http://doi.org/10.1146/annurev.en.25.010180.001201>

Costinot, A. (2009). On the origins of comparative advantage. *Journal of International Economics*, 77(2), 255–264. <http://doi.org/10.1016/j.jinteco.2009.01.007>

Curby, K. M., & Gauthier, I. (2010). To the trained eye: Perceptual expertise alters visual processing. *Topics in Cognitive Science*, 2(2), 189–201. <http://doi.org/10.1111/j.1756-8765.2009.01058.x>

Dawson, M. (2013). *Mind, body, world: Foundations of cognitive science*. Edmonton, AB: AU Press.

de Groot, A. (1965). *Thought and choice in chess*. *The American Journal of Psychology* (Vol. 79). The Hague: Mouton Publishers. <http://doi.org/10.5117/9789053569986>

De Langhe, R. (2010). The division of labour in science: The tradeoff between specialisation and diversity. *Journal of Economic Methodology*, 17(1), 37–51. <http://doi.org/10.1080/13501780903528960>

Delorme, A., Rousselet, G. A., Macé, M. J. M., & Fabre-Thorpe, M. (2004). Interaction of top-down and bottom-up processing in the fast visual analysis of natural scenes. *Cognitive Brain Research*, 19(2), 103–113. <http://doi.org/10.1016/j.cogbrainres.2003.11.010>

Derry, S., Pea, R., Barron, B., Engle, R., Erickson, F., Goldman, R., ... Sherin, B. (2010). Conducting video research in the learning sciences: Guidance on selection, analysis, technology, and ethics. *Journal of the Learning Sciences*, 19(1), 3–53. <http://doi.org/10.1080/10508400903452884>

- DeWitt, I., & Rauschecker, J. P. (2013). Wernicke's area revisited: Parallel streams and word processing. *Brain and Language*, *127*(2), 181–191.
<http://doi.org/10.1016/j.bandl.2013.09.014>
- Diaper, D. (2002). Scenarios and task analysis. In *Interacting with Computers* (Vol. 14, pp. 379–395). [http://doi.org/10.1016/S0953-5438\(02\)00005-X](http://doi.org/10.1016/S0953-5438(02)00005-X)
- DiCarlo, J. J., & Cox, D. D. (2007). Untangling invariant object recognition. *Trends in Cognitive Sciences*, *11*(8), 333–341. <http://doi.org/10.1016/j.tics.2007.06.010>
- Díez, F. J., Mira, J., Iturralde, E., & Zubillaga, S. (1997). DIAVAL, a Bayesian expert system for echocardiography. *Artificial Intelligence in Medicine*. [http://doi.org/10.1016/S0933-3657\(97\)00384-9](http://doi.org/10.1016/S0933-3657(97)00384-9)
- Donald, D. L. (1998). A tutorial on ergonomic and process modeling using QUEST and IGRIP. In *1998 Winter Simulation Conference. Proceedings (Cat. No.98CH36274)* (Vol. 1, pp. 297–302). <http://doi.org/10.1109/WSC.1998.744982>
- Doppelmayr, M., Finkenzeller, T., & Sauseng, P. (2008). Frontal midline theta in the pre-shot phase of rifle shooting: Differences between experts and novices. *Neuropsychologia*, *46*(5), 1463–1467. <http://doi.org/10.1016/j.neuropsychologia.2007.12.026>
- Dorussen, H., Lenz, H., & Blavoukos, S. (2005). Assessing the reliability and validity of expert interviews. *European Union Politics*, *6*(3), 315–337.
<http://doi.org/10.1177/1465116505054835>
- Dreyfus, H. L., & Dreyfus, S. E. (1986). *Mind over machine: The power of human intuition and expertise in the era of the computer*. New York: Free Press.
<http://doi.org/10.1109/MEX.1987.4307079>

- Dreyfus, H. L., & Dreyfus, S. E. (1991). Towards a phenomenology of ethical expertise. *Human Studies*, 14(4), 229–250. <http://doi.org/10.1007/BF02205607>
- Dreyfus, S. E., & Dreyfus, H. L. (1980). A five-stage model of the mental activities involved in directed skill acquisition. *Operations Research Center*, (February), 1–18.
<http://doi.org/ADA084551>
- Dror, I. E., & Harnad, S. (2008). Offloading cognition onto cognitive technology. *Cognition Distributed: How Cognitive Technology Extends Our Minds.*, 1–23.
- Drori, G. S. (2003). *Science in the modern world polity: Institutionalization and globalization*. Stanford, CA: Stanford University Press.
- Duarte, A., Pen, I., Keller, L., & Weissing, F. J. (2012). Evolution of self-organized division of labor in a response threshold model. *Behavioral Ecology and Sociobiology*, 66(6), 947–957.
<http://doi.org/10.1007/s00265-012-1343-2>
- Duff, P. A. (2010). Language socialization into academic discourse communities. *Annual Review of Applied Linguistics*, 30, 169–192. <http://doi.org/10.1017/S0267190510000048>
- Duncker, K. (1945). On problem solving. *Psychological Monographs*, 58, 270.
<http://doi.org/10.1037/h0093599>
- Durkheim, E. (1893). *The division of labor in society* (2014th ed.). New York: Free Press.
<http://doi.org/10.2307/1116093>
- Ebersbach, S., & Peng, Z. (2008). Expert system development for vibration analysis in machine condition monitoring. *Expert Systems with Applications*, 34(1), 291–299.
<http://doi.org/10.1016/j.eswa.2006.09.029>

- Eccles, D., & Feltovitch, P. J. (2014). Implications of domain-general “psychological support skills” for transfer of skill and acquisition of expertise. *The Internet and Higher Education*, 42(2), 43–60. <http://doi.org/10.1002/piq>
- Echterhoff, G., & Hirst, W. (2009). Social influence on memory. *Social Psychology*. <http://doi.org/10.1027/1864-9335.40.3.106>
- Elgin, M. (2008). Theory-laden observation and incommensurability. *Organon F*, 15(1), 3–19.
- Ellingsen, T. (1992). Strategic buyers and the social cost of monopoly. *American Economic Review*, 81(3), 648–657.
- Emery, F. E. (1959). *The emergence of a new paradigm of work*. Canberra, Australia: Australian National University, Centre for Continuing Education.
- Epic Games. (2011). Gears of War 3.
- Erickson, K. I., Colcombe, S. J., Wadhwa, R., Bherer, L., Peterson, M. S., Scalf, P. E., ... Kramer, A. F. (2007). Training-induced functional activation changes in dual-task processing: An fMRI study. *Cerebral Cortex*, 17(1), 192–204. <http://doi.org/10.1093/cercor/bhj137>
- Ericsson, K. A. (2009). Development of professional expertise: Toward measurement of expert performance and design of optimal learning environments. In *Development of professional expertise: Toward measurement of expert performance and design of optimal learning environments* (pp. xv, 552).
- Ericsson, K. A., Charness, N., Feltovitch, P. J., & Hoffman, R. R. (2006). *The Cambridge handbook of expertise and expert performance*. (K. A. Ericson, N. Charness, P. J. Feltovich, & R. R. Hoffman, Eds.) (1st ed.). New York: Cambridge University Press. <http://doi.org/10.2277/0521600812>

- Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, *102*(2), 211–245. <http://doi.org/10.1037/0033-295X.102.2.211>
- Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, *100*(3), 363–406. <http://doi.org/10.1037/0033-295X.100.3.363>
- Ericsson, K. A., & Smith, J. (1991). *Toward a general theory of expertise: Prospects and limits*. Cambridge, MA: Cambridge University Press.
- Ernst, G. W., & Newell, A. (1969). *GPS: A case study in generality and problem solving*. New York: Academic Press.
- Evans, J. S. B. T., & Over, D. E. (2008). Whole mind theory: Massive modularity meets dual processes. *Thinking & Reasoning*, *14*(2), 200–208. <http://doi.org/10.1080/13546780701652563>
- Farrington-Darby, T., & Wilson, J. R. (2006). The nature of expertise: A review. *Applied Ergonomics*. <http://doi.org/10.1016/j.apergo.2005.09.001>
- Feldman, J. (1999). Is there a post-modern science? *Philosophy and Social Action*, *25*(3), 19–30.
- Ferguson-Gow, H., Sumner, S., Bourke, A. F. G., & Jones, K. E. (2014). Colony size predicts division of labour in attine ants. *Proceedings of the Royal Society B: Biological Sciences*, *281*(1793), 20141411–20141411. <http://doi.org/10.1098/rspb.2014.1411>
- Ferster, D., & Miller, K. D. (2000). Neural mechanisms of orientation selectivity in the visual cortex. *Annual Review of Neuroscience*, *23*, 441–471. <http://doi.org/10.1146/annurev.neuro.23.1.441>

- Fiore, S. M., Rosen, M. A., Smith-Jentsch, K. A., Salas, E., Letsky, M., & Warner, N. (2010). Toward an understanding of macrocognition in teams: predicting processes in complex collaborative contexts. *Human Factors*, 52(2), 203–224.
<http://doi.org/10.1177/0018720810369807>
- Fodor, J. A. (1989). *The modularity of mind : an essay on faculty psychology*. *Reasoning: Studies of human inference and its foundations*. <http://doi.org/10.2307/413815>
- Fontana, A., & Wolcott, H. F. (1999). The art of fieldwork. *Contemporary Sociology*.
<http://doi.org/10.2307/2655068>
- Franzmann, A., Jansen, A., & Münte, P. (Eds.). (2015). *Legitimizing Science: National and global public (1800-2010)*. Chicago, IL: The University of Chicago Press.
- Freedman, D. (2010). *Wrong: Why experts keep failing us--and how to know when not to trust them*. New York: Little, Brown and Company.
- Frith, C. D. (2008). Social cognition. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 363(1499), 2033–9. <http://doi.org/10.1098/rstb.2008.0005>
- Frith, C. D., & Frith, U. (2012). Mechanisms of social cognition. *Annual Review of Psychology*, 63, 287–313. <http://doi.org/10.1146/annurev-psych-120710-100449>
- Fuller, S. (2001). Thomas Kuhn and the science wars. *British Journal for the History of Science*, 34(120), 99.
- Gallant, S. I. (1988). Connectionist expert systems. *Communications of the ACM*.
<http://doi.org/10.1145/42372.42377>
- Galton, F. (1869). *Hereditary genius: An inquiry into its laws and consequences*. London: Julian

Friedman Publishers.

Gardner, D. (2011). *Future babble: Why expert predictions fail - and why we believe them anyway*. Toronto: McLelland & Stewart.

Garfinkel, H. (1986). *Ethnomethodological studies of work*. (H. Garfinkel, Ed.) *Studies in ethnomethodology* (1st ed.). London: Routledge.

Gauthier, I., & Tarr, M. J. (2002). Unraveling mechanisms for expert object recognition: bridging brain activity and behavior. *Journal of Experimental Psychology. Human Perception and Performance*, 28(2), 431–446. <http://doi.org/10.1037/0096-1523.28.2.431>

Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform “face area” increases with expertise in recognizing novel objects. *Nature Neuroscience*, 2(6), 568–73. <http://doi.org/10.1038/9224>

Gazda, S. K., Connor, R. C., Edgar, R. K., & Cox, F. (2005). A division of labour with role specialization in group-hunting bottlenose dolphins (*Tursiops truncatus*) off Cedar Key, Florida. *Proceedings. Biological Sciences / The Royal Society*, 272(1559), 135–40. <http://doi.org/10.1098/rspb.2004.2937>

Geary, D. C., & Huffman, K. J. (2002). Brain and cognitive evolution: forms of modularity and functions of mind. *Psychological Bulletin*, 128(5), 667–698. <http://doi.org/10.1037/0033-2909.128.5.667>

Geertz, C. (1973). *The interpretation of cultures: Selected essays*. New York: Basic Books.

Gelder, T. Van, & Port, R. F. (1995). It's about time: An overview of the dynamical approach to cognition. In T. Van Gelder & R. F. Port (Eds.), *Mind as motion: Explorations in the dynamics of cognition* (pp. 1–43). Cambridge, MA: MIT Press.

- Germain, M. I. (2006). Stages of psychometric measure development: the example of the generalized expertise measure (GEM). *Academy of Human Resource Development*, (1993), 893–898.
- Gick, M. L., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology*.
- Giere, R. (2002). Scientific cognition as distributed cognition. *The Cognitive Basis of Science*, 285–299. <http://doi.org/10.1017/cbo9780511613517.016>
- Gilmore, D. J., & Green, T. R. G. (1988). Programming plans and programming expertise. *The Quarterly Journal of Experimental Psychology Section A*, 40(3), 423–442. <http://doi.org/10.1080/02724988843000005>
- Gintis, H. (2000). *Game theory evolving: A problem-centered introduction to modeling strategic interaction*. Princeton, NJ: Princeton University Press. [http://doi.org/10.1016/S1090-5138\(00\)00072-6](http://doi.org/10.1016/S1090-5138(00)00072-6)
- Glaser, R. (1996). Changing the agency for learning: Acquiring expert performance. In *The road to excellence: The acquisition of expert performance in the arts and sciences, sports and games* (pp. 303–312). New York: Psychology Press.
- Glimcher, P. W., Dorris, M. C., & Bayer, H. M. (2005). Physiological utility theory and the neuroeconomics of choice. *Games and Economic Behavior*. <http://doi.org/10.1016/j.geb.2004.06.011>
- Gobet, F., Lane, P. C. R., Croker, S., Cheng, P. C.-H., Jones, G., Oliver, I., & Pine, J. M. (2001). Chunking mechanisms in human learning. *Trends in Cognitive Sciences*, 5(6), 236–243. [http://doi.org/10.1016/S1364-6613\(00\)01662-4](http://doi.org/10.1016/S1364-6613(00)01662-4)
- Gordon, D. M. (1996). The organization of work in social insect colonies. *Nature*, 380(6570),

121–124. <http://doi.org/10.1038/380121a0>

Government of Canada - Population of the federal public service. (2016). Retrieved April 10, 2016, from <http://www.tbs-sct.gc.ca/psm-fpfm/modernizing-modernisation/stats/ssa-pop-eng.asp>

Grabner, R. H. (2014). The role of intelligence for performance in the prototypical expertise domain of chess. *Intelligence*, *45*(1), 26–33. <http://doi.org/10.1016/j.intell.2013.07.023>

Gray, W., John, B., & Atwood, M. (1993). Project Ernestine: Validating a GOMS analysis for predicting and explaining real-world task performance. *Human-Computer Interaction*, *8*(3), 237–309. http://doi.org/10.1207/s15327051hci0803_3

Greenbaum, J., & Kyng, M. (1991). *Design at work: Cooperative design of computer systems*. Mahwah, NJ: Lawrence Erlbaum Associates. <http://doi.org/10.1007/BF02298175>

Gregory, J., & Nussbaum, K. (1982). Race against time: Automation of the office. *Office Technology and People*, *1*(2/3), 197–236. <http://doi.org/10.1108/eb022612>

Groenewegen, P. P. (2006). Trust and the sociology of the professions. *European Journal of Public Health*. <http://doi.org/10.1093/eurpub/ckl003>

Guida, A., Gobet, F., Tardieu, H., & Nicolas, S. (2012). How chunks, long-term working memory and templates offer a cognitive explanation for neuroimaging data on expertise acquisition: A two-stage framework. *Brain and Cognition*, *79*(3), 221–244. <http://doi.org/10.1016/j.bandc.2012.01.010>

Haddock, A. (2004). Rethinking the “strong programme” in the sociology of knowledge. *Studies in History and Philosophy of Science Part A*, *35*(1), 19–40. <http://doi.org/10.1016/j.shpsa.2003.12.011>

- Harari, Y. N. (2014). *Sapiens: A brief history of humankind* (1st ed.). Oxford: Signal.
- Harsanyi, J. C. (1966). A general theory of rational behavior in game situations. *Econometrica*, 34(3), 613–634. <http://doi.org/10.2307/1909772>
- Hartman, M., Bolton, E., & Fehnel, S. E. (2001). Accounting for age differences on the Wisconsin Card Sorting Test: decreased working memory, not inflexibility. *Psychology and Aging*, 16(3), 385–399. <http://doi.org/10.1037/0882-7974.16.3.385>
- Hauw, D., & Durand, M. (2007). Situated analysis of elite trampolinists' problems in competition using retrospective interviews. *Journal of Sports Sciences*, 25(2), 173–83. <http://doi.org/10.1080/02640410600624269>
- Henderson, P. (1997). The rise and fall of anarcho-syndicalism in South America, 1880-1930. In *Ideologues & Ideologies in Latin America* (pp. 20–26).
- Henrich, J. (2015). *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. Princeton, NJ: Princeton University Press.
- Heritage, J. (1984). *Garfinkel and Ethnomethodology*. Cambridge, UK: Polity Press. <http://doi.org/10.2307/2069993>
- Herman, D. (2000). Narratology as a cognitive science. *Image & Narrative*, 1(1), 1–26.
- Herman, D. (2013). *Storytelling and the sciences of mind*. Cambridge: MIT Press.
- Herzmann, G., & Curran, T. (2011). Experts' memory: an ERP study of perceptual expertise effects on encoding and recognition. *Memory & Cognition*, 39(3), 412–432. <http://doi.org/10.3758/s13421-010-0036-1>

- Heyck, H. (2008). Defining the computer: Herbert Simon and the bureaucratic mind--Part 2. *Annals of the History of Computing, IEEE, 30*(2), 52–63.
<http://doi.org/10.1109/MAHC.2008.19>
- Hidden Path Entertainment, & Valve Corporation. (2012). Counter Strike: Global Offensive.
- Higgins, E. T. (2000). Social cognition: Learning about what matters in the social world. *European Journal of Social Psychology, 30*(1), 3–39. [http://doi.org/10.1002/\(SICI\)1099-0992\(200001/02\)30:1<3::AID-EJSP987>3.0.CO;2-I](http://doi.org/10.1002/(SICI)1099-0992(200001/02)30:1<3::AID-EJSP987>3.0.CO;2-I)
- Hilgartner, S. (1997). The Sokal affair in context. *Science, Technology & Human Values, 22*(4), 506 – 522. <http://doi.org/10.1177/016224399702200404>
- Hirst, R. (2003). Scientific jargon, good and bad. *Journal of Technical Writing and Communication, 33*(3), 201–229. <http://doi.org/10.2190/J8JJ-4YD0-4R00-G5N0>
- Hoffman, R. R. (1987). The problem of extracting the knowledge of experts from the perspective of experimental psychology. *AI Magazine, 8*(2), 53. <http://doi.org/10.1609/aimag.v8i2.583>
- Hoffman, R. R., Shadbolt, N. R., Burton, A. M., & Klein, G. (1995). Eliciting knowledge from experts: A methodological analysis. *Organizational Behavior and Human Decision Processes, 62*(2), 129–158. <http://doi.org/10.1006/obhd.1995.1039>
- Hoffman, R. R., Trafton, G., & Roebber, P. (2006). *Minding the weather: How expert forecasters think*. Cambridge, MA: MIT Press.
- Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Trans. Comput.-Hum. Interact., 7*, 174–196.
<http://doi.org/http://doi.acm.org/10.1145/353485.353487>

- Holland, J. H. (1986). Escaping brittleness: The possibilities of general-purpose learning algorithms applied to parallel rule-based systems. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning, an artificial intelligence approach. Volume II* (pp. 593–623). San Francisco, CA: Morgan Kaufmann Publishers Inc.
- Holland, J. H. (2012). *Signals and boundaries: Building blocks for complex adaptive systems*. Cambridge, MA: MIT Press. <http://doi.org/10.1007/s10710-013-9180-2>
- Horgan, J. (2012). What Thomas Kuhn really thought about scientific “truth.” *Scientific American Blog*, 1–4.
- Horn, J., & Masunaga, H. (2006). A merging theory of expertise and intelligence. In *The Cambridge handbook of expertise and expert performance*. (pp. 587–611).
- Howe, M. J., Davidson, J. W., & Sloboda, J. A. (1998). Innate talents: Reality or myth? *The Behavioral and Brain Sciences*, 21, 399–407. <http://doi.org/10.1017/S0140525X9800123X>
- Howorth, J. (2004). Discourse, ideas, and epistemic communities in european security and defence policy. *West European Politics*, 27(2), 211–234.
<http://doi.org/10.1080/0140238042000214883>
- Hoy, W. K. (1990). Organizational climate and culture: A conceptual analysis of the school workplace. *Journal of Educational and Psychological Consultation*, 1(2), 149–168.
http://doi.org/10.1207/s1532768xjepc0102_4
- Hunt, S. D., & Morgan, R. M. (1995). The comparative advantage theory of competition. *Journal of Marketing*, 59(2), 1–15. <http://doi.org/10.2307/1252069>
- Hussain, F., & Wood, S. (2009). Modelling the efficiencies and interactions of attentional networks. In *Lecture Notes in Computer Science* (Vol. 5395 LNAI, pp. 139–152).

http://doi.org/10.1007/978-3-642-00582-4_11

Hutchins, E. (1995). *Cognition in the wild*. MIT Press. <http://doi.org/10.1023/A:1008642111457>

Hutchins, E. (2000). *Distributed cognition*. IESBS *Distributed Cognition*. Retrieved from <http://files.meetup.com/410989/DistributedCognition.pdf>

Hutchins, E. (2010). Cognitive ecology. *Topics in Cognitive Science*, 2(4), 705–715. <http://doi.org/10.1111/j.1756-8765.2010.01089.x>

Imm, K., & Stylianou, D. A. (2012). Talking mathematically: An analysis of discourse communities. *Journal of Mathematical Behavior*, 31(1), 130–148. <http://doi.org/10.1016/j.jmathb.2011.10.001>

Jackson, R. W., & Parasuraman, A. (1986). The Yellow Pages as an advertising tool for small businesses. *American Journal of Small Business*, 10(4), 29–35. Retrieved from <http://search.ebscohost.com/login.aspx?direct=true&db=buh&AN=5749932&site=ehost-live>

Jacobson, M. J. (2001). Problem solving, cognition, and complex systems: Differences between experts and novices. *Complexity*, 6(3), 41–49. <http://doi.org/10.1002/cplx.1027>

Jeanson, R., Kukuk, P. F., & Fewell, J. H. (2005). Emergence of division of labour in halictine bees: Contributions of social interactions and behavioural variance. *Animal Behaviour*, 70(5), 1183–1193. <http://doi.org/10.1016/j.anbehav.2005.03.004>

John, B. E., & Kieras, D. E. (1996). The GOMS family of user interface analysis techniques: comparison and contrast. *ACM Transactions on Computer-Human Interaction*. <http://doi.org/10.1145/235833.236054>

- Johnson, K. E., & Mervis, C. B. (1997). Effects of varying levels of expertise on the basic level of categorization. *Journal of Experimental Psychology: General*, *126*(3), 248–277.
<http://doi.org/10.1037/0096-3445.126.3.248>
- Jones, C. (2009). Interdisciplinary approach - Advantages, disadvantages, and the future benefits of interdisciplinary studies. *Essai*, *7*(2009), 75–81.
- Jonna, R. J. (2015). Monopoly capital and labor. *Labor Studies Journal*, *40*(3), 262–274.
http://doi.org/10.14452/MR-026-03-1974-07_1
- Judd, T. (2000). Division of labour in colony defence against vertebrate predators by the social wasp *Polistes fuscatus*. *Animal Behaviour*, *60*(1), 55–61.
<http://doi.org/10.1006/anbe.2000.1449>
- Julian, G. E., & Fewell, J. H. (2004). Genetic variation and task specialization in the desert leaf-cutter ant, *Acromyrmex versicolor*. *Animal Behaviour*, *68*(1), 1–8.
<http://doi.org/10.1016/j.anbehav.2003.06.023>
- Jung, C. G. (1980). The archetypes and the collective unconscious. *Chemistry, Volume 9*, 550.
Retrieved from http://books.google.co.uk/books?id=_gDvgAEACAAJ
- K. Ruhleder, B. J. (1997). Capturing complex, distributed activities: Video-based interaction analysis as a component of workplace ethnography. In *Information Systems and Qualitative Research* (pp. 246–275). <http://doi.org/10.1007/978-0-387-35309-8>
- Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *American Economic Review*. <http://doi.org/10.1257/000282803322655392>
- Kalyuga, S. (2011). Cognitive load theory: How many types of load does it really need? *Educational Psychology Review*, *23*(1), 1–19. <http://doi.org/10.1007/s10648-010-9150-7>

- Karakaş, S., Başar-Eroğlu, C., Özesmi, Ç., Kafadar, H., & Erzenin, Ö. Ü. (2000). Gamma response of the brain: A multifunctional oscillation that represents bottom-up with top-down processing. *International Journal of Psychophysiology*. [http://doi.org/10.1016/S0167-8760\(00\)00137-9](http://doi.org/10.1016/S0167-8760(00)00137-9)
- Kassubek, J., Schmidtke, K., Kimmig, H., Lücking, C. H., & Greenlee, M. W. (2001). Changes in cortical activation during mirror reading before and after training: An fMRI study of procedural learning. *Cognitive Brain Research*, *10*(3), 207–217. [http://doi.org/10.1016/S0926-6410\(00\)00037-9](http://doi.org/10.1016/S0926-6410(00)00037-9)
- Kayaoglu, M. N. (2013). Impact of extroversion and introversion on language-learning behaviors. *Social Behavior and Personality*, *41*(5), 819–826. <http://doi.org/10.2224/sbp.2013.41.5.819>
- Keil, F. C., Stein, C., Webb, L., Billings, V. D., & Rozenblit, L. (2008). Discerning the division of cognitive labor: An emerging understanding of how knowledge is clustered in other minds. *Cognitive Science*, *32*(2), 259–300. <http://doi.org/10.1080/03640210701863339>
- Keller, S. S., Crow, T., Foundas, A., Amunts, K., & Roberts, N. (2009). Broca's area: Nomenclature, anatomy, typology and asymmetry. *Brain and Language*. <http://doi.org/10.1016/j.bandl.2008.11.005>
- Kellman, P. J., & Garrigan, P. (2009). Perceptual learning and human expertise. *Physics of Life Reviews*. <http://doi.org/10.1016/j.plrev.2008.12.001>
- Kheirbek, M. A., & Hen, R. (2015). (Radio)active neurogenesis in the human hippocampus. *Cell*, *153*(6), 1183–1184. <http://doi.org/10.1016/j.cell.2013.05.033>
- Kieras, D. E., & Meyer, D. E. (1994). The EPIC architecture for modeling human information-processing and performance: A brief introduction. *Security*, *1*(1).

- Kieras, D. E., & Meyer, D. E. (2000). The role of cognitive task analysis in the application of predictive models of human performance. In *Cognitive task analysis* (pp. 237–260).
- Kimball, D. R., & Holyoak, K. J. (2000). Transfer and expertise. *The Oxford Handbook of Memory*, 109–122. Retrieved from http://users.phhp.ufl.edu/rbauer/cognitive/Articles/kimball_holyoak_2000.pdf
- Kipp, M. (2014). ANVIL: The video annotation research tool. In J. Durand, U. Gut, & G. Kristoffersen (Eds.), *Handbook of Corpus Phonology* (pp. 420–436). Oxford: Oxford University Press.
- Kirlik, A., Plamondon, B. D., & Jagacinski, R. J. (1993). Supervisory control in a dynamic and uncertain environment: Laboratory task and crew performance. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(4), 1130–1138.
- Kirwan, B., & Ainsworth, L. K. (1992). *A guide to task analysis*. Washington, DC: Taylor & Francis. <http://doi.org/doi:10.4324/9780203221457>
- Klegon, D. (1978). The sociology of professions: An emerging perspective. *Work and Occupations*, 5(3), 259–283. <http://doi.org/10.1177/073088847800500301>
- Klein, G., Feltovich, P. J., Bradshaw, J. M., & Woods, D. D. (2005). Common ground and coordination in joint activity. *Organizational Simulation*, (53), 139–184. <http://doi.org/10.1002/0471739448.ch6>
- Klein, G., & Hoffman, R. R. (2008). Macrocognition, mental models, and cognitive task analysis methodology. In L. G. Militello, T. Ormerod, J. M. Schraagen, & R. Lipshitz (Eds.), *Naturalistic Decision Making and Macrocognition* (pp. 57–80). Aldershot: Ashgate Publishing Ltd.

- Klein, G., Ross, K. G., Moon, B. M., Klein, D. E., Hoffman, R. R., & Hollnagel, E. (2003).
Macrocognition. *IEEE Intelligent Systems*, *18*(3), 81–85.
<http://doi.org/10.1109/MIS.2003.1200735>
- Kleinmintz, O. M., Goldstein, P., Mayselless, N., Abecasis, D., & Shamay-Tsoory, S. G. (2014).
Expertise in musical improvisation and creativity: The mediation of idea evaluation. *PLoS ONE*, *9*(7). <http://doi.org/10.1371/journal.pone.0101568>
- Kling, R., & Star, S. L. (1998). Human centered systems in the perspective of organizational and social informatics. *Computers and Society*, *28*(1), 22–29.
<http://doi.org/10.1145/277351.277356>
- Knapp, S. J. (1999). Analyzing narratives of expertise: Toward the development of a burkeian pentadic scheme. *The Sociological Quarterly*, *40*(4), 587–612.
<http://doi.org/10.1111/j.1533-8525.1999.tb00569.x>
- Kochan, J. (2010). Contrastive explanation and the “strong programme” in the sociology of scientific knowledge. *Social Studies of Science*, *40*(1), 127–144.
<http://doi.org/10.1177/0306312709104780>
- Koedinger, K. R., & Anderson, J. R. (1990). Abstract planning and perceptual chunks: Elements of expertise in geometry. *Cognitive Science*, *14*(4), 511–550. [http://doi.org/10.1016/0364-0213\(90\)90008-K](http://doi.org/10.1016/0364-0213(90)90008-K)
- Koester, A. (2006). Investigating workplace discourse. *Investigating Workplace Discourse*, 1–183.
<http://doi.org/10.4324/9780203015742>
- Kolodner, J. L. (1983). Maintaining organization in a dynamic long-term memory. *Cognitive Science*, *7*(4), 243–280. [http://doi.org/10.1016/S0364-0213\(83\)80001-9](http://doi.org/10.1016/S0364-0213(83)80001-9)

- Kuhn, T. S. (1962). *The structure of scientific revolutions*. Chicago: University of Chicago Press.
- Lahti, E. (2016). Valve puts in \$1 million for all future major CS:GO tournaments. Retrieved April 10, 2016, from <http://www.pcgamer.com/valve-all-major-csgo-tournaments-will-be-1-million/>
- Laird, J. E. (2012). The Soar cognitive architecture. *AISB Quarterly*, 171(134), 224–235.
<http://doi.org/http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=12784>
- Lakatos, I. (1970). Falsification and the methodology of scientific research programmes. In *Criticism and the Growth of Knowledge* (Vol. 4, pp. 91–196). <http://doi.org/10.2277/>
- Lapadat, J. C., & Lindsay, A. C. (1998). Examining transcription : A theory-laden methodology. *Annual Meeting of the American Educational Research Association*, 1–20.
- Latinus, M., VanRullen, R., & Taylor, M. J. (2010). Top-down and bottom-up modulation in processing bimodal face/voice stimuli. *BMC Neuroscience*, 11, 36.
<http://doi.org/10.1186/1471-2202-11-36>
- Latour, B. (1999). *Pandora's hope: Essays on the reality of science studies. Essays on the Reality of Science Studies*. <http://doi.org/10.2307/2655272>
- Latour, B., & Noor, A. (2002). The science wars: A dialogue. *Common Knowledge*, 8(1), 71–79.
<http://doi.org/10.1215/0961754X-8-1-71>
- Laverick, D. M. (2007). Motivation, metacognition, mentors, and money: Ingredients that support teaching expertise. *Early Childhood Education Journal*, 34(4), 247–249.
<http://doi.org/10.1007/s10643-006-0144-5>
- Laws, K. R. (1999). A meta-analytic review of Wisconsin Card Sort studies in schizophrenia:

- general intellectual deficit in disguise? *Cognitive Neuropsychiatry*, 4(1), 1–30; discussion 31–35. <http://doi.org/10.1080/135468099396025>
- Lazaric, N. (2008). Routines and routinization: An exploration of some micro-cognitive foundations. In M. C. Becker (Ed.), *Handbook of Organizational Routines* (pp. 205–227). Northampton, MA: Edward Elgar.
- Lee, H., Voss, M. W., Prakash, R. S., Boot, W. R., Vo, L. T. K., Basak, C., ... Kramer, A. F. (2012). Videogame training strategy-induced change in brain function during a complex visuomotor task. *Behavioural Brain Research*, 232(2), 348–357. <http://doi.org/10.1016/j.bbr.2012.03.043>
- Leismuller, G., & Schimpf, E. J. (2012). *Economics of competition*. New York, NY: Nova Science Publishers. Retrieved from <http://www.scopus.com/inward/record.url?eid=2-s2.0-84896570502&partnerID=40&md5=ce2da708487b1eeae8fede4aff8ea728>
- Leith, P. (2010). The rise and fall of the legal expert system. *European Journal of Law and Technology*, 1, 179.
- Lenk, J. C., Möbus, C., Özyurt, J., Thiel, C. M., & Claassen, A. (2011). From fMRI data to cognitive models: Testing the ACT-R brain mapping hypothesis with an ex-post model. *Cognitive Systems Research*, 12(3-4), 321–335. Retrieved from http://www.thinkmind.org/index.php?view=article&articleid=cognitive_2011_1_30_40098
- Lesgold, A., Rubinson, H., Feltovich, P., Glaser, R., Klopfer, D., & Wang, Y. (1988). Expertise in a complex skill: Diagnosing x-ray pictures. In *The Nature of Expertise* (pp. 311–342).
- Leune, A., & Nizard, J. (2012). Doctor Google: Use of Internet during pregnancy in France in 2009. *Journal de Gynecologie, Obstetrique et Biologie de La Reproduction*, 41(3), 243–254.

- Levy, K. N., & Scala, J. W. (2012). Transference, transference interpretations, and transference-focused psychotherapies. *Psychotherapy, 49*(3), 391–403.
<http://doi.org/10.1037/a0029371>
- Lewis, M. (2010). *The big short: Inside the doomsday machine*. London: Penguin Books Ltd.
- Ley, T., Ulbrich, A., Scheir, P., Lindstaedt, S. N., Kump, B., & Albert, D. (2008). Modeling competencies for supporting work-integrated learning in knowledge work. *Journal of Knowledge Management, 12*(6), 31–47. <http://doi.org/10.1108/13673270810913603>
- Li, D., Fu, Z., & Duan, Y. (2002). Fish-Expert: A web-based expert system for fish disease diagnosis. *Expert Systems with Applications, 23*(3), 311–320.
[http://doi.org/10.1016/S0957-4174\(02\)00050-7](http://doi.org/10.1016/S0957-4174(02)00050-7)
- Liddell, B. J., Brown, K. J., Kemp, A. H., Barton, M. J., Das, P., Peduto, A., ... Williams, L. M. (2005). A direct brainstem-amygdala-cortical “alarm” system for subliminal signals of fear. *NeuroImage, 24*(1), 235–243. <http://doi.org/10.1016/j.neuroimage.2004.08.016>
- Little, M., Jordens, C. F., & Sayers, E.-J. (2003). Discourse communities and the discourse of experience. *Health, 7*(1), 73–86. <http://doi.org/10.1177/1363459303007001619>
- Litzinger, T. A., Lattuca, L. R., Hadgraft, R. G., & Newstetter, W. C. (2011). Engineering education and the development of expertise. *Journal of Engineering Education, 100*(1), 123–150.
<http://doi.org/10.1002/j.2168-9830.2011.tb00006.x>
- Liu, M., Schallert, D. L., & Carroll, P. J. (2004). Working memory and expertise. *Interpreting, 6*(1), 19–42. <http://doi.org/10.1075/intp.6.1.04liu>
- Loveday, T., Wiggins, M. W., Festa, M., Schell, D., & Twigg, D. (2013). Pattern recognition as an indicator of diagnostic expertise. *Pattern Recognition - Applications and Methods, 204*, 1–

11. <http://doi.org/10.1007/978-3-642-36530-0>

Lowe, D. G. (1999). Object recognition from local scale-invariant features. In *Proceedings of the Seventh IEEE International Conference on Computer Vision* (Vol. 2, pp. 1150–1157).

<http://doi.org/10.1109/ICCV.1999.790410>

Luff, P., & Heath, C. (2012). Some “technical challenges” of video analysis: Social actions, objects, material realities and the problems of perspective. *Qualitative Research*, 12(3), 255–279. <http://doi.org/10.1177/1468794112436655>

Luff, P., Hindmarsh, J., & Heath, C. (2000). *Workplace studies: Recovering work practice and informing system Design. Management Science*.

<http://doi.org/10.1017/CBO9780511628122>

Lutz, A., Greischar, L. L., Perlman, D. M., & Davidson, R. J. (2009). BOLD signal in insula is differentially related to cardiac function during compassion meditation in experts vs. novices. *NeuroImage*, 47(3), 1038–1046.

<http://doi.org/10.1016/j.neuroimage.2009.04.081>

MacDougall, W. K., Martin, M., Nagy, N., & West, R. L. (2015). A method for building models of expert cognition in naturalistic environments. In *Proceedings of the International Conference on Cognitive Modeling 2015* (pp. 25–30).

MacDougall, W. K., West, R. L., & Hancock, E. (2014). Modeling multi-agent chaos: Killing aliens and managing difficult people. In *36th annual meeting of the Cognitive Science Society* (pp. 2603–2608). Quebec, QC. Retrieved from

<https://mindmodeling.org/cogsci2014/papers/451/paper451.pdf>

Machery, E. (2007). Massive modularity and brain evolution. *Philosophy of Science*, 74(5), 825–

838. <http://doi.org/10.1086/525624>

Magnus, P. D. (2007). Distributed cognition and the task of science. *Social Studies of Science*, 37(2), 297–310. <http://doi.org/10.1177/0306312706072177>

Manfredini, F., Lucas, C., Nicolas, M., Keller, L., Shoemaker, D., & Grozinger, C. M. (2014). Molecular and social regulation of worker division of labour in fire ants. *Molecular Ecology*, 23(3), 660–672. <http://doi.org/10.1111/mec.12626>

Marchant, G., Robinson, J., Anderson, U., & Schadewald, M. (1991). Analogical transfer and expertise in legal reasoning. *Organizational Behavior and Human Decision Processes*, 48(2), 272–290. [http://doi.org/10.1016/0749-5978\(91\)90015-L](http://doi.org/10.1016/0749-5978(91)90015-L)

Martin, B. R., Nightingale, P., & Yegros-Yegros, A. (2012). Science and technology studies: Exploring the knowledge base. *Research Policy*. <http://doi.org/10.1016/j.respol.2012.03.010>

Martin, M. (2014). A set of code for SGOMS modelling, including a GUI interface. Retrieved from <https://github.com/mattmartin256/SGOMS>

Marx, K. (1978). Communist manifesto. *The Marx-Engels Reader*, 1(November), 473–500.

Mason, A., & Singh, C. (2011). Assessing expertise in introductory physics using categorization task. *Physical Review Special Topics - Physics Education Research*, 7(2). <http://doi.org/10.1103/PhysRevSTPER.7.020110>

Mason, J. R., & Rosner, J. (2007). Where did the risk go? How misapplied bond ratings cause mortgage backed securities and collateralized debt obligation market disruptions. *SSRN Electronic Journal*, (May), 1–87. <http://doi.org/10.2139/ssrn.1027475>

- Maurer, U., Zevin, J. D., & McCandliss, B. D. (2008). Left-lateralized N170 effects of visual expertise in reading: evidence from Japanese syllabic and logographic scripts. *Journal of Cognitive Neuroscience, 20*(10), 1878–1891. <http://doi.org/10.1162/jocn.2008.20125>
- Mayfield, W. a., Kardash, C. M., & Kivlighan, D. M. . J. (1999). Differences in experienced and novice counselors' knowledge structures about clients: Implications for case conceptualization. *Journal of Counseling Psychology, 46*(4), 504–514. <http://doi.org/10.1037/0022-0167.46.4.504>
- Maynard, D. (1991). The diversity of ethnomethodology. *Annual Review of Sociology, 17*(1), 385–418. <http://doi.org/10.1146/annurev.soc.17.1.385>
- McCoy, L. P. (2005). Effect of demographic and personal variables on achievement in eighth-grade algebra. *The Journal of Educational Research, 98*(3), 131–135. <http://doi.org/10.3200/JOER.98.3.131-135>
- McGraw, K. M., & Pinney, N. (1990). The effects of general and domain-specific expertise on political memory and judgment. *Social Cognition, 8*(1), 9–30. <http://doi.org/10.1521/soco.1990.8.1.9>
- McLuhan, M. (1964). *Understanding media: The extensions of man*. New York: McGraw-Hill. <http://doi.org/10.2307/2711172>
- McMahan, R. P., Bowman, D. A., Zielinski, D. J., & Brady, R. B. (2012). Evaluating display fidelity and interaction fidelity in a virtual reality game. *IEEE Transactions on Visualization and Computer Graphics, 18*(4), 626–633. <http://doi.org/10.1109/TVCG.2012.43>
- Memmert, D. (2009). Pay attention! A review of visual attentional expertise in sport. *International Review of Sport and Exercise Psychology, 2*(2), 119–138.

<http://doi.org/10.1080/17509840802641372>

Memmert, D. (2011). Creativity, expertise, and attention: exploring their development and their relationships. *Journal of Sports Sciences*, 29(1), 93–102.

<http://doi.org/10.1080/02640414.2010.528014>

Memmert, D., Simons, D. J., & Grimme, T. (2009). The relationship between visual attention and expertise in sports. *Psychology of Sport and Exercise*, 10(1), 146–151.

<http://doi.org/10.1016/j.psychsport.2008.06.002>

Meyer, M., Elmer, S., & Jäncke, L. (2012). Musical expertise induces neuroplasticity of the planum temporale. *Annals of the New York Academy of Sciences*, 1252(1), 116–123.

<http://doi.org/10.1111/j.1749-6632.2012.06450.x>

Michaelian, K., & Sutton, J. (2013). Distributed cognition and memory research: History and current directions. *Review of Philosophy and Psychology*, 4(1), 1–24.

<http://doi.org/10.1007/s13164-013-0131-x>

Mieg, H. A. (2001). *The social psychology of expertise: Case studies in research, professional domains, and expert roles*. Mahwah, NJ: Lawrence Erlbaum Associates.

Miljkovic, D. (2005). Rational choice and irrational individuals or simply an irrational theory: A critical review of the hypothesis of perfect rationality. *Journal of Socio-Economics*, 34(5),

621–634. <http://doi.org/10.1016/j.socec.2003.12.031>

Milton, J., Solodkin, A., Hlustik, P., & Small, S. L. (2007). The mind of expert motor performance is cool and focused. *NeuroImage*, 35(2), 804–813.

<http://doi.org/10.1016/j.neuroimage.2007.01.003>

Minsky, M. (1986). *Society of mind*. New York: Simon & Schuster.

- Mishler, E. G. (1995). Models of narrative analysis: A typology. *Journal of Narrative & Life History*, 5(2), 87–123. <http://doi.org/10.1075/jnlh.5.2.01mod>
- Miyake, A., & Shah, P. (Eds.). (1999). *Models of working memory: Mechanisms of active maintenance and executive control*. Cambridge, MA: Cambridge University Press.
- Moore, C. D., Cohen, M. X., & Ranganath, C. (2006). Neural mechanisms of expert skills in visual working memory. *The Journal of Neuroscience*, 26(43), 11187–11196. <http://doi.org/10.1523/JNEUROSCI.1873-06.2006>
- Morrow, D. G., Menard, W. E., Stine-Morrow, E. a, Teller, T., & Bryant, D. (2001). The influence of expertise and task factors on age differences in pilot communication. *Psychology and Aging*, 16(1), 31–46. <http://doi.org/10.1037/0882-7974.16.1.31>
- Münste, T. F., Altenmüller, E., & Jäncke, L. (2002). The musician's brain as a model of neuroplasticity. *Nature Reviews. Neuroscience*, 3(6), 473–478. <http://doi.org/10.1038/nrn843>
- Mussa, M., & Rosen, S. (1978). Monopoly and product quality. *Journal of Economic Theory*, 18(2), 301–317. [http://doi.org/10.1016/0022-0531\(78\)90085-6](http://doi.org/10.1016/0022-0531(78)90085-6)
- Mutton, T., Hagger, H., & Burn, K. (2011). Learning to plan, planning to learn: the developing expertise of beginning teachers. *Teachers and Teaching*, 17(4), 399–416. <http://doi.org/10.1080/13540602.2011.580516>
- Naiman, R. J. (1999). A perspective on interdisciplinary science. *Ecosystems*, 2(4), 292–295. <http://doi.org/10.1007/s100219900078>
- Neishabouri, A., & Faisal, a A. (2014). Saltatory conduction in unmyelinated axons: clustering of Na(+) channels on lipid rafts enables micro-saltatory conduction in C-fibers. *Frontiers in*

- Neuroanatomy*, 8(October), 109. <http://doi.org/10.3389/fnana.2014.00109>
- Neumann, C., Duboscq, J., Dubuc, C., Ginting, A., Irwan, A. M., Agil, M., ... Engelhardt, A. (2011). Assessing dominance hierarchies: Validation and advantages of progressive evaluation with Elo-rating. *Animal Behaviour*. <http://doi.org/10.1016/j.anbehav.2011.07.016>
- Newell, A. (1973). You can't play 20 questions with nature and win: Projective comments on the papers of this symposium. In *Visual Information Processing* (pp. 283–308). <http://doi.org/10.1021/bi047419c>
- Newth, D. (2009). Asynchronous Iterated Prisoner's Dilemma. *Adaptive Behavior*, 17(2), 175–183. <http://doi.org/10.1177/1059712309104313>
- Newzoo.com. (2015). The eSports economy will generate at least \$465 million in 2017. Retrieved from <https://newzoo.com/insights/articles/esports-economy-will-generate-least-465-million-2017/>
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, 84(3), 231–259. <http://doi.org/10.1037/0033-295X.84.3.231>
- Noble, K. G., Wolmetz, M. E., Ochs, L. G., Farah, M. J., & McCandliss, B. D. (2006). Brain-behavior relationships in reading acquisition are modulated by socioeconomic factors. *Developmental Science*, 9(6), 642–654. <http://doi.org/10.1111/j.1467-7687.2006.00542.x>
- Nokes, T. J., Schunn, C. D., & Chi, M. T. (2010). Problem solving and human expertise. In *International Encyclopedia of Education* (pp. 265–272). <http://doi.org/10.1016/B978-0-08-044894-7.00486-3>
- Noss, R., Hoyles, C., & Pozzi, S. (2002). Abstraction in expertise: A study of nurses' conceptions

- of concentration. *Journal for Research in Mathematics Education*, 33(3), 204.
<http://doi.org/10.2307/749725>
- Novick, L. R. (1988). Analogical transfer, problem similarity, and expertise. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 14(3), 510–520.
<http://doi.org/10.1037/0278-7393.14.3.510>
- Nowotny, H. (2003). Democratising expertise and socially robust knowledge. *Science and Public Policy*, 30(3), 151–156. <http://doi.org/10.3152/147154303781780461>
- Nunn, R. (2008). A network model of expertise. *Bulletin of Science, Technology & Society*, 28, 414–427. <http://doi.org/10.1177/0270467608322587>
- Nyhus, E., & Barceló, F. (2009). The Wisconsin Card Sorting Test and the cognitive assessment of prefrontal executive functions: A critical update. *Brain and Cognition*, 71(3), 437–451.
<http://doi.org/10.1016/j.bandc.2009.03.005>
- Olick, J. K., & Robbins, J. (1998). Social memory studies: From “collective memory” to the historical sociology of mnemonic practices. *Annual Review of Sociology*, 24(1), 105–140.
<http://doi.org/10.1146/annurev.soc.24.1.105>
- Olson, M. H., Lucas Jr., H. C., & Kling, R. (1982). The impact of office automation on the organization: Some implications for research and practice. *Communications of the ACM*, 25(11), 838–847. <http://doi.org/10.1145/358690.358720>
- Orr, M. R. (1992). Parasitic flies (Diptera: Phoridae) influence foraging rhythms and caste division of labor in the leaf-cutter ant, *Atta cephalotes* (Hymenoptera: Formicidae). *Behavioral Ecology and Sociobiology*, 30(6), 395–402. <http://doi.org/10.1007/BF00176174>
- Osborne, M. J. (2000). *An introduction to game theory*. Book.

<http://doi.org/10.1192/pb.26.12.476-a>

Osuagwu, C. C., & Okafor, E. C. (2010). Framework for eliciting knowledge for a medical laboratory diagnostic expert system. *Expert Systems with Applications*, 37(7), 5009–5016.

<http://doi.org/10.1016/j.eswa.2009.12.012>

Owen, A. M., Hampshire, A., Grahn, J. A., Stenton, R., Dajani, S., Burns, A. S., ... Ballard, C. G. (2010). Putting brain training to the test. *Nature*, 465(7299), 775–8.

<http://doi.org/10.1038/nature09042>

Paletz, S. B. F., Kim, K. H., Schunn, C. D., Tollinger, I., & Vera, A. (2013). Reuse and recycle: The development of adaptive expertise, routine expertise, and novelty in a large research team. *Applied Cognitive Psychology*, 27(4), 415–428. <http://doi.org/10.1002/acp.2928>

Pandey, S. K., & Scott, P. G. (2002). Red rape: A review and assessment of concepts and measures. *Journal of Public Administration Research & Theory*, 12(4), 553.

<http://doi.org/Article>

Parry, A. (1996). Scientific practice and ordinary action: ethnomethodology and social studies of science. *Physiotherapy*. [http://doi.org/10.1016/S0031-9406\(05\)66899-X](http://doi.org/10.1016/S0031-9406(05)66899-X)

Pastells, A. A. I., & Roca, D. S. (2003). A comparative analysis of the phonological loop versus the visuo-spatial sketchpad in mental arithmetic tasks in 7-8 y.o. children. *Psicothema*, 15(2), 241–246.

PC Games News. (2015). How Counter-Strike: Global Offensive is turning into the world's most exciting eSport. Retrieved April 10, 2016, from <http://www.pcgamesn.com/counter-strike-global-offensive/how-counter-strike-global-offensive-is-turning-into-the-worlds-most-exciting-esport>

- Pinho, A. L., de Manzano, Ö., Fransson, P., Eriksson, H., & Ullén, F. (2014). Connecting to create: expertise in musical improvisation is associated with increased functional connectivity between premotor and prefrontal areas. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, *34*(18), 6156–63. <http://doi.org/10.1523/JNEUROSCI.4769-13.2014>
- Pitt, M. a, Myung, I. J., & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review*, *109*(3), 472–491. <http://doi.org/10.1037/0033-295X.109.3.472>
- Plato. (1966). *Plato's Republic*. (I. A. Richards, Ed.). Cambridge, MA: Cambridge University Press. <http://doi.org/10.2307/623766>
- Plato. (2003). *Plato's Phaedrus*. (S. Scully, Ed.). Newburyport, MA: Focus Publishing/R. Pullins Co.
- Pollack, J. B. (1990). Recursive distributed representations. *Artificial Intelligence*, *46*(1-2), 77–105. [http://doi.org/10.1016/0004-3702\(90\)90005-K](http://doi.org/10.1016/0004-3702(90)90005-K)
- Popper, K. (1963). Science as Falsification. *Conjectures and Refutations, Readings in the Philosophy of Science*, 33–39.
- Posner, R. (2011). *A failure of capitalism: The crisis of '08 and the descent into depression*. Cambridge, MA: Harvard University Press.
- Press, W. H., & Dyson, F. J. (2012). Iterated Prisoner's Dilemma contains strategies that dominate any evolutionary opponent. *Proceedings of the National Academy of Sciences of the United States of America*, *109*(26), 10409–13. <http://doi.org/10.1073/pnas.1206569109>
- Prinz, J. (2006). Is the mind really modular? In R. J. Stainton (Ed.), *Contemporary debates in*

- cognitive science* (pp. 22–36). Malden, MA: Blackwell Publishers.
- Propp, V. Y. (2012). *The Russian folktale*. (S. Forrester, Ed.). Detroit: Wayne State University Press. <http://doi.org/10.1353/mat.2014.0023>
- Rampton, S., & Stauber, J. (2001). *Trust us, we're experts! How industry manipulates science and gambles with your future*. New York, NY: Putnam.
- Raskin, E. (1935). Comparison of scientific and literary ability: a biographical study of eminent scientists and men of letters of the nineteenth century. *The Journal of Abnormal and Social Psychology*, 31, 20–35. <http://doi.org/10.1037/h0060483>
- Redies, C., & Puelles, L. (2001). Modularity in vertebrate brain development and evolution. *BioEssays*. <http://doi.org/10.1002/bies.10014>
- Reeves, S., Brown, B., & Laurier, E. (2009). Experts at play: Understanding skilled expertise. *Games and Culture*, 4, 205–227. <http://doi.org/10.1177/1555412009339730>
- Repko, A. F. (2008). *Interdisciplinary research: Process and theory*. Los Angeles: SAGE.
- Restif, O., Hayman, D. T. S., Pulliam, J. R. C., Plowright, R. K., George, D. B., Luis, A. D., ... Webb, C. T. (2012). Model-guided fieldwork: Practical guidelines for multidisciplinary research on wildlife ecological and epidemiological dynamics. *Ecology Letters*. <http://doi.org/10.1111/j.1461-0248.2012.01836.x>
- Rezlescu, C., Barton, J. J. S., Pitcher, D., & Duchaine, B. (2014). Normal acquisition of expertise with greebles in two cases of acquired prosopagnosia. *Proceedings of the National Academy of Sciences of the United States of America*, 111(14), 5123–8. <http://doi.org/10.1073/pnas.1317125111>

- Ries, E. (2011). *The lean startup: How today's entrepreneurs use continuous innovation to create radically successful businesses*. New York, NY: Crown Business. <http://doi.org/23>
- Ringach, D. L., Bredfeldt, C. E., Shapley, R. M., & Hawken, M. J. (2002). Suppression of neural responses to nonoptimal stimuli correlates with tuning selectivity in macaque V1. *Journal of Neurophysiology*, *87*(2), 1018–27. <http://doi.org/10.1152/jn.00614.2001>
- Risko, E. F., Laidlaw, K., Freeth, M., Foulsham, T., & Kingstone, A. (2012). Social attention with real versus reel stimuli: toward an empirical approach to concerns about ecological validity. *Frontiers in Human Neuroscience*, *6*(May), 1–11. <http://doi.org/10.3389/fnhum.2012.00143>
- Ritter, F., & Larkin, J. H. (1994). Developing process models as summaries of HCI action sequences. *Human Computer Interaction*, *9*(4), 345–383.
- Rivas, T., Paz, M., Martinn, J. E., Matias, J. M., Garcia, J. F., & Taboada, J. (2011). Explaining and predicting workplace accidents using data-mining techniques. *Reliability Engineering and System Safety*, *96*(7), 739–747. <http://doi.org/10.1016/j.res.2011.03.006>
- Robben, A. C. G. M., & Sluka, J. A. (Eds.). (2012). *Ethnographic fieldwork: An anthropological reader* (2nd ed.). Malden, MA: Wiley-Blackwell.
- Robbins, P. (2013). Modularity and mental architecture. *Wiley Interdisciplinary Reviews: Cognitive Science*. <http://doi.org/10.1002/wcs.1255>
- Roberts, R. E., Bain, P. G., Day, B. L., & Husain, M. (2013). Individual differences in expert motor coordination associated with white matter microstructure in the cerebellum. *Cerebral Cortex*, *23*(10), 2282–2292. <http://doi.org/10.1093/cercor/bhs219>
- Rogers, Y. (2006). Distributed cognition and communication. *Library and Information Science*,

- (1990), 51–66. <http://doi.org/10.1016/B0-08-044854-2/00862-2>
- Rogers, Y., & Ellis, J. (1994). Distributed cognition: An alternative framework for analysing and explaining collaborative working. *Journal of Information Technology*, 9(2), 119–128.
<http://doi.org/10.1057/jit.1994.12>
- Rosano, M. J. (2000). Citizenship and Socrates in Plato's Crito. *The Review of Politics*, 62(03), 451.
<http://doi.org/10.1017/S0034670500041656>
- Rossion, B., Kung, C.-C., & Tarr, M. J. (2004). Visual expertise with nonface objects leads to competition with the early perceptual processing of faces in the human occipitotemporal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 101(40), 14521–14526. <http://doi.org/10.1073/pnas.0405613101>
- Rouncefield, M., & Tolmie, P. (Eds.). (2011). *Ethnomethodology at work: Directions in ethnomethodology and conversation Analysis*. Farnham, UK: Ashgate Publishing Ltd.
- Rowe, A. J., & Watkins, P. R. (1992). Beyond expert systems: Reasoning, judgment, and wisdom. *Expert Systems With Applications*, 4(1), 1–10. [http://doi.org/10.1016/0957-4174\(92\)90035-Q](http://doi.org/10.1016/0957-4174(92)90035-Q)
- Rubin, C. T. (1998). Science Wars. *Academic Questions*, 11(3), 85–89.
<http://doi.org/10.1007/s12129-998-1033-8>
- Rumelhart, D. E., McClelland, J. L., & Williams, R. J. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. Cambridge, MA: MIT Press.
<http://doi.org/10.1017/CBO9781107415324.004>
- Sabers, D. S., Cushing, K. S., & Berliner, D. C. (1991). Differences among teachers in a task characterized by simultaneity, multidimensional, and immediacy. *American Educational*

Research Journal, 28(1), 63–88. <http://doi.org/10.3102/00028312028001063>

Sahin, S., Tolun, M. R., & Hassanpour, R. (2012). Hybrid expert systems: A survey of current approaches and applications. *Expert Systems with Applications*, 39(4), 4609–4617. <http://doi.org/10.1016/j.eswa.2011.08.130>

Samuels, R. (1998). Evolutionary psychology and the massive modularity hypothesis. *The British Journal for the Philosophy of Science*, 49(4), 575–602. <http://doi.org/10.1093/bjps/49.4.575>

Sanbonmatsu, D. M., Strayer, D. L., Medeiros-Ward, N., & Watson, J. M. (2013). Who multi-tasks and why? Multi-tasking ability, perceived multi-tasking ability, impulsivity, and sensation seeking. *PLoS ONE*, 8(1). <http://doi.org/10.1371/journal.pone.0054402>

Sanfey, A. G., Loewenstein, G., McClure, S. M., & Cohen, J. D. (2006). Neuroeconomics: Cross-currents in research on decision-making. *Trends in Cognitive Sciences*, 10(3), 108–116. <http://doi.org/10.1016/j.tics.2006.01.009>

Schack, T., & Mechsner, F. (2006). Representation of motor skills in human long-term memory. *Neuroscience Letters*, 391(3), 77–81. <http://doi.org/10.1016/j.neulet.2005.10.009>

Schectman, J. M., Schorling, J. B., Nadkarni, M. M., & Voss, J. D. (2005). Determinants of physician use of an ambulatory prescription expert system. *International Journal of Medical Informatics*, 74(9), 711–717. <http://doi.org/10.1016/j.ijmedinf.2005.05.011>

Schindler, S. (2013). Theory-laden experimentation. *Studies in History and Philosophy of Science Part A*, 44(1), 89. <http://doi.org/10.1016/j.shpsa.2012.07.010>

Schnotz, W., & Kürschner, C. (2007). A reconsideration of cognitive load theory. *Educational Psychology Review*, 19(4), 469–508. <http://doi.org/10.1007/s10648-007-9053-4>

- Schraagen, J. M., Chipman, S. F., & Shalin, V. L. (2000). *Cognitive task analysis*. Retrieved from <http://books.google.com/books?hl=en&lr=&id=gyt5AgAAQBAJ&oi=fnd&pg=PP1&dq=Cognitive+Task+Analysis&ots=cfC2rtL7xa&sig=km1krdtRdaFLfuRalbc1Mflqtkl>
- Schraagen, J. M., Militello, L. G., Ormerod, T., & Lipshitz, R. (2008). *Naturalistic decision making and macrocognition*. Burlington, VT: Ashgate. <http://doi.org/2007039710>
- Seamster, T. L., Redding, R. E., & Kaempff, G. L. (1997). *Applied cognitive task analysis in aviation*. London: Ashgate Publishing Ltd.
- Selinger, E., & Crease, R. (2007). The philosophy of expertise. *Ethics: An International Journal of Social, Political, and Legal Philosophy*, 117(2), 377–381.
- Shafto, P., & Coley, J. D. (2003). Development of categorization and reasoning in the natural world: Novices to experts, naive similarity to ecological knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(4), 641–649. <http://doi.org/10.1037/0278-7393.29.4.641>
- Shapin, S. (1995). Here and everywhere: Sociology of scientific knowledge. *Annual Review of Sociology*, 21(1), 289–321. <http://doi.org/10.1146/annurev.soc.21.1.289>
- Shinn, T., Sörlin, S., & Crawford, E. T. (1993). Denationalizing science: The contexts of international scientific practice. In *Sociology of the sciences* (Vol. 16). Dordrecht, Boston: Kluwer Academic Publishers.
- Shipp, S. (2004). The brain circuitry of attention. *Trends in Cognitive Sciences*. <http://doi.org/10.1016/j.tics.2004.03.004>
- Shoham, Y. (2008). Computer science and game theory. *Communications of the ACM*, 51, 74. <http://doi.org/10.1145/1378704.1378721>

- Siegel, A. J., Kaftanoglu, O., Fondrk, M. K., Smith, N. R., & Page, R. E. (2012). Ovarian regulation of foraging division of labour in Africanized backcross and pollen-hoarding honeybees. *Animal Behaviour*, *83*(3), 653–658. <http://doi.org/10.1016/j.anbehav.2011.12.009>
- Sieloff, C. G. (1999). “If only HP knew what HP knows”: The roots of knowledge management at Hewlett-Packard. *Journal of Knowledge Management*, *3*(1), 47–53. <http://doi.org/10.1108/13673279910259385>
- Silva, P., Gago, P., Ribeiro, J. C. B., Santos, M. F., Portela, F., Abelha, A., ... Pinto, F. (2014). An expert system for supporting traditional Chinese medicine diagnosis and treatment. *Procedia Technology*, *16*, 1487–1492. <http://doi.org/10.1016/j.protcy.2014.10.169>
- Silvia, P. J., Kaufman, J. C., & Pretz, J. E. (2009). Is creativity domain-specific? Latent class models of creative accomplishments and creative self-descriptions. *Psychology of Aesthetics, Creativity, and the Arts*, *3*(3), 139–148. <http://doi.org/10.1037/a0014940>
- Simon, H. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, *69*(1), 99–118. <http://doi.org/10.2307/1884852>
- Simon, H. (1996a). The architecture of complexity: Hierarchic systems. *The Sciences of the Artificial*, *106*(6), 467–482. <http://doi.org/10.2307/985254>
- Simon, H. (1996b). *The sciences of the artificial* (3rd ed.). Cambridge, MA: MIT Press. [http://doi.org/10.1016/S0898-1221\(97\)82941-0](http://doi.org/10.1016/S0898-1221(97)82941-0)
- Simonton, D. K. (1990). *Psychology, science, and history: An introduction to historiometry*. New Haven, CT: Yale University Press.
- Simonton, D. K. (2006). Historiometric methods. In K. A. Ericsson, N. Charness, P. J. Feltovitch, & R. R. Hoffman (Eds.), *The Cambridge Handbook of Expertise and Expert Performance* (pp.

- 319–335). New York, NY: Cambridge University Press.
- Skogh, G. (1976). The social costs of monopoly and regulation: Some comments. *Journal of Political Economy*, 84(6), 1319. <http://doi.org/10.1086/260516>
- Slack, J. D., & Semati, M. M. (1997). Intellectual and political hygiene: The “Sokal affair.” *Critical Studies in Media Communication*, 14(3), 201–227.
- Smagorinsky, P. (1998). Thinking and speech and protocol analysis. *Mind, Culture, and Activity*, 5(3), 157–177. http://doi.org/10.1207/s15327884mca0503_2
- Smith, A. (1976). *An inquiry into the nature and causes of the wealth of nations*. (A. S. Skinner & R. H. Cambell, Eds.). New York, NY: Clarendon Press.
<http://doi.org/10.7208/chicago/9780226763750.001.0001>
- Smith, V. L. (2005). Behavioral economics research and the foundations of economics. *Journal of Socio-Economics*, 34(2), 135–150. <http://doi.org/10.1016/j.socec.2004.09.003>
- Smith, V. L., & Ellsworth, P. C. (1987). The social psychology of eyewitness accuracy: Misleading questions and communicator expertise. *Journal of Applied Psychology*, 72(2), 294–300.
<http://doi.org/10.1037/0021-9010.72.2.294>
- Smith-Jentsch, K. A., Johnston, J. H., & Payne, S. C. (1998). Measuring team-related expertise in complex environments. In *Making decisions under stress: Implications for individual and team training*. (pp. 61–87).
- Sohn, Y. W., & Doane, S. M. (2004). Memory processes of flight situation awareness: Interactive roles of working memory capacity, long-term working memory, and expertise. *Human Factors*, 46(3), 461–475. <http://doi.org/10.1518/hfes.46.3.461.50392>

- Sokal, A. (1996a). A physicist experiments with cultural studies. *Lingua Franca*, May/June, 62–64. Retrieved from <http://linguafranca.mirror.theinfo.org/9605/sokal.html>
- Sokal, A. (1996b). Transgressing the boundaries: An afterword. *Philosophy and Literature*, 20(2), 338–346. <http://doi.org/10.1353/phl.1996.0078>
- Sokal, A. (1998). What the Social Text affair does and does not prove. *Critical Quarterly*, 40(2), 3–18. <http://doi.org/10.1111/1467-8705.00151>
- Sosniak, L. A. (2006). Retrospective interviews in the study of expertise and expert performance. In *The Cambridge handbook of expertise and expert performance* (pp. 287–301). Cambridge, MA: Cambridge University Press.
- Spalding, K. L., Bergmann, O., Alkass, K., Bernard, S., Salehpour, M., Huttner, H. B., ... Frisé, J. (2013). Dynamics of hippocampal neurogenesis in adult humans. *Cell*, 153(6), 1219–1227. <http://doi.org/10.1016/j.cell.2013.05.002>
- Spence, A. M. (1975). Monopoly, quality, and regulation. *The Bell Journal of Economics*, 6, 417–429. <http://doi.org/10.2307/3003237>
- Spence, T. M., & Brucks, M. (1997). The moderating effects of problem characteristics on experts' and novices' judgments. *Journal of Marketing Research*, 34(May), 233. <http://doi.org/10.2307/3151861>
- Sperber, D. (2001). In defense of massive modularity. *Language, Brain and Cognitive Development: Essays in Honor of Jacques Mehler*, 7(October), 47–57.
- Stander, P. E. (1992). Cooperative hunting in lions: The role of the individual. *Behavioral Ecology and Sociobiology*, 29(6), 445–454. <http://doi.org/10.1007/BF00170175>

Starkes, J. L., & Allard, F. (Eds.). (1993). *Cognitive issues in motor expertise*. Amsterdam: Elsevier.

[http://doi.org/10.1016/S0166-4115\(08\)61463-6](http://doi.org/10.1016/S0166-4115(08)61463-6)

Sternberg, R. J., & Grigorenko, E. L. (Ed.). (2003). *Perspectives on the psychology of abilities, competencies, and expertise*. Cambridge, MA: Cambridge University Press.

Stewart, T. C., & West, R. L. (2007). Deconstructing and reconstructing ACT-R: Exploring the architectural space. *Cognitive Systems Research*, 8(3), 227–236.

<http://doi.org/10.1016/j.cogsys.2007.06.006>

Stich, S. (1985). Could man be an irrational animal?: Some notes on the epistemology of rationality. *Synthese*, 64(1), 115–135.

<http://doi.org/10.1093/acprof:oso/9780199733477.003.0003>

Stoet, G., O'Connor, D. B., Conner, M., & Laws, K. R. (2013). Are women better than men at multi-tasking? *BMC Psychology*, 1(1), 1–18. <http://doi.org/10.1186/2050-7283-1-18>

Stubbe, M., Lane, C., Hilder, J., Vine, E., Vine, B., Marra, M., ... Weatherall, A. (2003). Multiple discourse analyses of a workplace interaction. *Discourse Studies*, 5(3), 351–388.

<http://doi.org/10.1177/14614456030053004>

Sun, R. (2008). Introduction to computational cognitive modeling. In *Cambridge handbook of computational psychology* (pp. 1–36). Cambridge, MA: Cambridge University Press.

<http://doi.org/http://dx.doi.org/10.1017/CBO9780511816772.003>

Sun, R. (2009). Theoretical status of computational cognitive modeling. *Cognitive Systems Research*, 10(2), 124–140. <http://doi.org/10.1016/j.cogsys.2008.07.002>

Sutton, J., Harris, C. B., Keil, P. G., & Barnier, A. J. (2010). The psychology of memory, extended cognition, and socially distributed remembering. *Phenomenology and the Cognitive*

- Sciences*, 9(4), 521–560. <http://doi.org/10.1007/s11097-010-9182-y>
- Sweller, J., Ayres, P., & Kalyuga, S. (2011). *Cognitive Load Theory*. New York, NY: Springer.
http://doi.org/10.1007/978-1-4419-8126-4_6
- Tan, Y. H., & Thoen, W. (2000). INCAS: A legal expert system for contract terms in electronic commerce. *Decision Support Systems*, 29(4), 389–411. [http://doi.org/10.1016/S0167-9236\(00\)00085-3](http://doi.org/10.1016/S0167-9236(00)00085-3)
- Tanaka, J. W., & Taylor, M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology*, 23(3), 457–482. [http://doi.org/10.1016/0010-0285\(91\)90016-H](http://doi.org/10.1016/0010-0285(91)90016-H)
- Theeuwes, J. (2010). Top-down and bottom-up control of visual selection. *Acta Psychologica*, 135(2), 77–99. <http://doi.org/10.1016/j.actpsy.2010.02.006>
- Thiel, P., & Masters, B. (2014). *Zero to one: Notes on startups, or how to build the future*. New York, NY: Crown Publishing Group.
- Todorov, T. (1971). The 2 principles of narrative. *Diacritics*, 1(1), 37–44.
- Toth, A. L., & Robinson, G. E. (2005). Worker nutrition and division of labour in honeybees. *Animal Behaviour*, 69(2), 427–435. <http://doi.org/10.1016/j.anbehav.2004.03.017>
- Trickett, S. B., & Trafton, J. G. (2009). A primer on verbal protocol analysis. In D. Schmorow, J. Cohn, & D. Nicolson (Eds.), *The PSI Handbook of Virtual Environments for Training and Education* (pp. 332–346).
- Tsang, P. S., & Shaner, T. L. (1998). Age, attention, expertise, and time-sharing performance. *Psychology and Aging*, 13(2), 323–47. <http://doi.org/10.1037/0882-7974.13.2.323>

- Turk, D. J., Handy, T. C., & Gazzaniga, M. S. (2005). Can perceptual expertise account for the own-race bias in face recognition? A split-brain study. *Cognitive Neuropsychology*, *22*(7), 877–883. <http://doi.org/10.1080/02643290442000383>
- Turvey, M. T., & Carello, C. (2012). On intelligence from first principles: Guidelines for inquiry into the hypothesis of physical intelligence (PI). *Ecological Psychology*, *24*(1), 3–32. <http://doi.org/10.1080/10407413.2012.645757>
- Valkanova, V., Eguia Rodriguez, R., & Ebmeier, K. P. (2014). Mind over matter - what do we know about neuroplasticity in adults? *International Psychogeriatrics*, *26*(6), 891–909. <http://doi.org/10.1017/S1041610213002482>
- Van de Wiel, M. W. J., & Van den Bossche, P. (2013). Deliberate practice in medicine: The motivation to engage in work-related learning and its contribution to expertise. *Vocations and Learning*, *6*(1), 135–158. <http://doi.org/10.1007/s12186-012-9085-x>
- Van Der Maas, H. L. J., & Wagenmakers, E. J. (2005). A psychometric analysis of chess expertise. *American Journal of Psychology*, *118*(1), 29–60. <http://doi.org/10.2307/30039042>
- Van der Velde, F. (1993). Is the brain an effective Turing machine or a finite-state machine? *Psychological Research*, *55*(1), 71–79. <http://doi.org/10.1007/BF00419895>
- Van Dijk, T. A. (2006). Discourse, context and cognition. *Discourse Studies*, *8*(1), 159–177. <http://doi.org/10.1177/1461445606059565>
- Van Stijn, E., & Wensley, A. (2001). Organizational memory and the completeness of process modeling in ERP systems. *Business Process Management Journal*, *7*(3), 181–194. <http://doi.org/10.1108/14637150110392647>
- Vayalakkara, J., Backhaus, S. D., Bradley, J. D., Simco, E. R., & Golden, C. J. (2000). Abbreviated

- form of the Wisconsin Card Sort Test. *International Journal of Neuroscience*, *103*(1-4), 131–137. <http://doi.org/1148T991056> [pii]
- Velicer, W. F., & Prochaska, J. O. (1999). An expert system intervention for smoking cessation. *Patient Education and Counseling*, *36*(2), 119–129. [http://doi.org/10.1016/S0738-3991\(98\)00129-3](http://doi.org/10.1016/S0738-3991(98)00129-3)
- Vera, D., & Crossan, M. (2005). Improvisation and innovative performance in teams. *Organization Science*, *16*(3), 203–224. <http://doi.org/10.1287/orsc.1050.0126>
- Vicente, K. J. (1988). Adapting the memory recall paradigm to evaluate interfaces. *Acta Psychologica*, *69*(3), 249–278. [http://doi.org/10.1016/0001-6918\(88\)90035-2](http://doi.org/10.1016/0001-6918(88)90035-2)
- Vicente, K. J. (1999). *Cognitive work analysis: Toward safe, productive, and healthy computer-based work*. Mahwah, NJ: Erlbaum.
- Vicente, K. J., & Wang, J. H. (1998). An ecological theory of expertise effects in memory recall. *Psychological Review*, *105*, 33–57. <http://doi.org/10.1037//0033-295X.105.1.33>
- Vogel, G. (2009). Do chimps have culture? *Science*. Retrieved from <Go to ISI>://000248288700019
- Von Stengel, B. (1996). Efficient computation of behavior strategies. *Games and Economic Behavior*, *14*(2), 220–246. <http://doi.org/10.3390/s111009327>
- Vuust, P., Pallesen, K. J., Bailey, C., Van Zuijen, T. L., Gjedde, A., Roepstorff, A., & ??stergaard, L. (2005). To musicians, the message is in the meter: Pre-attentive neuronal responses to incongruent rhythm are left-lateralized in musicians. *NeuroImage*, *24*(2), 560–564. <http://doi.org/10.1016/j.neuroimage.2004.08.039>

- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Cambridge, MA: Harvard University Press. <http://doi.org/10.2307/3726822>
- Wagner, W. P., Otto, J., & Chung, Q. B. (2002). Knowledge acquisition for expert systems in accounting and financial problem domains. *Knowledge-Based Systems*, 15(8), 439–447. [http://doi.org/10.1016/S0950-7051\(02\)00026-6](http://doi.org/10.1016/S0950-7051(02)00026-6)
- Waibel, M., Floreano, D., Magnenat, S., & Keller, L. (2006). Division of labour and colony efficiency in social insects: effects of interactions between genetic architecture, colony kin structure and rate of perturbations. *Proceedings of the Royal Society of London B: Biological Sciences*, 273(1595), 1815–1823. <http://doi.org/10.1098/rspb.2006.3513>
- Walker, G. (2005). Verbal protocol analysis. In S. Neville, A. Hedge, K. Brookhuis, E. Salas, & H. Hendrick (Eds.), *Handbook of Human Factors and Ergonomics Methods* (pp. 301–311). Boca Raton: CRC Press. <http://doi.org/doi:10.1201/9780203489925.ch30>
- Walker, M., Takayama, L., & Landay, J. A. (2002). High-fidelity or low-fidelity, paper or computer? Choosing attributes when testing web prototypes. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 46(5), 661–665. <http://doi.org/10.1177/154193120204600513>
- Weber, M. (1978). *Economy and society: An outline of interpretative sociology* (1978th ed.). Oakland, CA: University of California Press.
- Weber, N., & Brewer, N. (2003). Expert memory: The interaction of stimulus structure, attention, and expertise. *Applied Cognitive Psychology*, 17(3), 295–308. <http://doi.org/10.1002/acp.865>
- Webster, C. A., & Schempp, P. (2007). Self-monitoring: Demystifying the wonder of expert

- teaching. *Journal of Physical Education, Recreation & Dance*, 1(79), 1–58.
- Weisberg, M., & Muldoon, R. (2009). Epistemic landscapes and the division of cognitive labor. *Philosophy of Science*, 76, 225–252. <http://doi.org/10.1086/644786>
- Weiser, M., & Shertz, J. (1983). Programming problem representation in novice and expert programmers. *International Journal of Man-Machine Studies*, 19, 391–398.
[http://doi.org/10.1016/S0020-7373\(83\)80061-3](http://doi.org/10.1016/S0020-7373(83)80061-3)
- Weiss, E. M. (1999). Perceived workplace conditions and first-year teachers' morale, career choice commitment, and planned retention: a secondary analysis. *Teaching and Teacher Education*, 15(8), 861–879. [http://doi.org/10.1016/S0742-051X\(99\)00040-2](http://doi.org/10.1016/S0742-051X(99)00040-2)
- West, R. L., Hancock, E., Somers, S., MacDougall, W. K., & Jeanson, F. (2013). The macro architecture hypothesis: Applications to modeling teamwork, conflict resolution, and literary Analysis. In *Proceedings of the 12th International Conference on Cognitive Modeling* (pp. 427–432).
- West, R. L., & MacDougall, W. K. (2014). The macro-architecture hypothesis: Modifying Newell's system levels to include macro-cognition. *Biologically Inspired Cognitive Architectures*, 8, 140–149. <http://doi.org/10.1016/j.bica.2014.03.009>
- West, R. L., & Nagy, G. (2007). Using GOMS for modeling routine tasks within complex sociotechnical systems: Connecting macrocognitive models to microcognition. *Journal of Cognitive Engineering and Decision Making*, 1(2), 186–211.
<http://doi.org/10.1518/155534307X232848>
- White, R. K. (1931). The versatility of genius. *Journal of Social Psychology*, 2(4), 460–489.
<http://doi.org/10.1080/00224545.1931.9918987>

- Williams, R. (2013). Cognitive anthropology: Distributed cognition and gesture. In C. Müller, A. Cienki, E. Fricke, S. H. Ladewig, D. McNeill, & S. Tebendorf (Eds.), *Body - Language - Communication: An international handbook on multimodality in human interaction* (pp. 240–258). Berlin: Walter de Gruyter GmbH.
- Williams, W. M., & Sternberg, R. J. (1988). Group intelligence: Why some groups are better than others. *Intelligence*, 12(4), 351–377. [http://doi.org/10.1016/0160-2896\(88\)90002-5](http://doi.org/10.1016/0160-2896(88)90002-5)
- Wilson, E. O. (1980). Caste and division of labor in leaf-cutter ants (Hymenoptera: Formicidae: Atta) - I. The overall pattern in *A. sexdens*. *Behavioral Ecology and Sociobiology*, 7(2), 143–156. <http://doi.org/10.1007/BF00299520>
- Wood, E., Zivcakova, L., Gentile, P., Archer, K., De Pasquale, D., & Nosko, A. (2012). Examining the impact of off-task multi-tasking with technology on real-time classroom learning. *Computers and Education*, 58(1), 365–374. <http://doi.org/10.1016/j.compedu.2011.08.029>
- Woolley, A. W., Chabris, C. F., Pentland, A., Hashmi, N., & Malone, T. W. (2010). Evidence for a collective intelligence factor in the performance of human groups. *Science*, 330(6004), 686–8. <http://doi.org/10.1126/science.1193147>
- Wynne, B. (2015). Seasick on the third wave? Subverting the hegemony of propositionalism : Response to Collins & Evans (2002). *Social Studies of Science*, 33(3), 401–417.
- Xiong, J., Ma, L., Wang, B., Narayana, S., Duff, E. P., Egan, G. F., & Fox, P. T. (2009). Long-term motor training induced changes in regional cerebral blood flow in both task and resting states. *NeuroImage*, 45(1), 75–82. <http://doi.org/10.1016/j.neuroimage.2008.11.016>
- Yao, Y., Zeng, Y., Zhong, N., & Huang, X. (2007). Knowledge retrieval (KR). In *Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence, WI 2007* (pp. 729–735).

<http://doi.org/10.1109/WI.2007.4427181>

Yoerg, S. I. (1991). Ecological frames of mind: the role of cognition in behavioral ecology. *The Quarterly Review of Biology*, 66(3), 287–301. <http://doi.org/10.1086/417243>

Yudes, C., Macizo, P., & Bajo, T. (2011). The influence of expertise in simultaneous interpreting on non-verbal executive processes. *Frontiers in Psychology*, 2(OCT).

<http://doi.org/10.3389/fpsyg.2011.00309>

Zetzel, E. R. (1956). Current concepts of transference. *The International Journal of Psychoanalysis*, 37, 369–375.

Zhang, J., Ackerman, M. S., & Adamic, L. (2007). Expertise networks in online communities: structure and algorithms. *Forum American Bar Association*, 2005, 221–230.

<http://doi.org/10.1145/1242572.1242603>

Zhang, J., & Patel, V. L. (2006). Distributed cognition, representation, and affordance. *Pragmatics & Cognition*, 14(2), 333–341. <http://doi.org/10.1075/pc.14.2.12zha>

Zimmerman, B. J. (2006). Development and adaptation of expertise: The role of self-regulatory process and beliefs. In *The Cambridge handbook of expertise and expert performance* (Vol. 186, pp. 705–722). Cambridge, MA: Cambridge University Press.

<http://doi.org/10.1097/ACM.0b013e31828b0007>

Zizzo, D. J. (2005). Economic man: Self-interest and rational choice. *Behavioral and Brain Sciences*, 28(6), 837–838. <http://doi.org/10.1017/S0140525X05450141>

Zylberberg, A., Dehaene, S., Roelfsema, P. R., & Sigman, M. (2011). The human Turing machine: A neural framework for mental programs. *Trends in Cognitive Sciences*.

<http://doi.org/10.1016/j.tics.2011.05.007>

Appendix A – Example ANVIL annotation XML

The following is the raw output of an ANVIL annotation file in XML (extended markup language). This file is read into anvil to display the annotation tracks in sync with the video file. This block of text is an annotation of the first block of time on a video (i.e., first planning unit annotated), lasting 5.52 seconds. This can be seen from the line - `<el index="0" start="1.08" end="6.6">`.

All of these variables are encoded for each frame of video data, spanning over demarcated blocks of time. The combination of all these variables represents the player's situation awareness model at any given time

```
<?xml version="1.0" encoding="UTF-16"?>
<annotation>
  <head>
    <specification src="../../../../GitHub/SGOMS_Annotations/SGOMS SPEC model v7i
Failure5.xml" />
    <video src="../../Converted Videos/Korey,Checkout,Wave5 - 2015.01.07 - MVI_0383(H263) -
Attempt2.mov" master="true" />
    <info key="coder" type="String">Matthew</info>
    <info key="encoding" type="String">UTF-16</info>
  </head>
  <body>
    <track name="Context" type="primary">
      <el index="0" start="1.08" end="6.6">
        <attribute name="token">grab ammo</attribute>
        <comment>==== Environment Variables ====

buffer_context = environment:game_started
buffer_context = environment:between_rounds
buffer_context = environment:calm
buffer_context = environment:enemies_not_coming #NIM
buffer_context = environment:explosive_not_dangerously_close #NIM
buffer_context = environment:many_far-mid_enemies_not_present #NIM
```

```
buffer_context = environment:many_close_quarter_enemies_not_present
```

```
===== Environment Object Variables =====
```

```
buffer_context = environment_object:command_post_bought
```

```
buffer_context = environment_object:free_not_ammo_accessible # Not sure  
# Agent likely to die getting ammo
```

```
buffer_context = environment_object:free_ammo_not_available # not sure  
# Simply not present in environment
```

```
buffer_context = environment_object:expensive_ammo_accessible  
# Expensive ammo is always available
```

```
buffer_context = environment_object:ammo_accessible # Either free or expensive ammo
```

```
buffer_context = environment_object:engagement_position_not_located #NIM
```

```
buffer_context = environment_object:engagement_position_not_taken #NIM
```

```
===== Agent Variables =====
```

```
*** Location ***
```

```
buffer_context = agent:not_at_suitable_place_to_engage
```

```
buffer_context = agent:defensible_position_not_located #NIM
```

```
buffer_context = agent:defensible_base_position_located
```

```
buffer_context = agent:not_at_base
```

```
buffer_context = agent:not_at_defensible_position
```

```
buffer_context = agent:not_at_ideal_defensible_position
```

```
*** Movement ***
```

buffer_context = agent:moving #Just prior to this action, he was moving
Moving includes running, rolling, shifting and walking
buffer_context = agent:walk

*** Actions ***

buffer_context = agent:not_aiming #NIM
buffer_context = agent:not_firing
buffer_context = agent:not_cq_engage_enemy_target #NIM
buffer_context = agent:scanning
buffer_context = agent:not_grab_ammo #NIM
buffer_context = agent:not_reloading #NIM

*** Status ***

buffer_context = agent:not_behind_cover
buffer_context = agent:not_safe_behind_cover #NIM
buffer_context = agent:feels_defensive #unknown
buffer_context = agent:not_being_mobbed #NIM
buffer_context = agent:not_caught_in_crossfire #NIM
buffer_context = agent:health_high

*** Ammo ***

buffer_context = agent:stored_ammo_low
need ammo / # could get ammo, but no need / # don't need ammo (basically full)
buffer_context = agent:gun_ammo_high # should reload / don't need to reload
buffer_context = agent:ammo_located

*** Gun ***

buffer_context = agent:current_gun_shotgun

buffer_context = agent:current_gun_appropriate # arbitrary choice

===== Enemy variables =====

No enemies present, so none of these apply

buffer_context = enemy:alive / enemy:injured / enemy:dead

The base state / # If enemy crawling around, i.e. critically injured / # Completely dead

buffer_context = enemy:located / location_not_known

buffer_context = enemy:target / enemy:not_target #NIM

buffer_context = enemy:presence_known / enemy:presence_unknown #NIM

buffer_context = enemy:vulnerable / enemy:not_vulnerable

buffer_context = enemy:visible / enemy:not_visible

buffer_context = enemy:near_agent / not_near_agent #NIM

buffer_context = enemy:distance_dangerously_close / distance_not_dangerously_close #NIM

buffer_context = enemy:within_close_quarters_range / not_within_close_quarters_range

buffer_context = enemy:suitable_to_kill_close_quarters / not_suitable_to_kill_close_quarters

Usually big enemies are not suitable to kill cq (for example)</comment>

</el>

<el index="1" start="6.6" end="8.28">

<attribute name="token">move and scan</attribute>

</el>

<el index="2" start="8.32" end="11.6">

<attribute name="token">switch gun</attribute>

<comment>===== Environment Variables =====

buffer_context = environment:game_started
buffer_context = environment:between_rounds
buffer_context = environment:calm
buffer_context = environment:enemies_not_coming #NIM
buffer_context = environment:explosive_not_dangerously_close #NIM
buffer_context = environment:many_far-mid_enemies_not_present #NIM
buffer_context = environment:many_close_quarter_enemies_not_present

Appendix B – Sample Gears of War Team Communication Transcriptions

The following are 4 examples of transcribed audio from cooperative Gears of War 3 gameplay.

Game 1

1 - Force them to come to us.

2 - That is true

2 - I'll take the right side then?

1 - Yeah

1 - So they drop ... some of them drop the hammerburst, right?

2 - Yeah

1 - I'm going to switch to the hammerburst when I see one, not as a strategic thing

2 - There's one on your side

1 - Can we destroy that middle segment ...

1 - There's a skitter coming there. Those little skitter bugs

2 - There's a boomer on your side

2 - I'll try to get a better position

1 - Yeah this isn't a great position either. I'm going to move.

1 - Are you still in the back there? No, you're in the front.

2 - Yeah.

1 - (inaudible) bugs

1 - bug

1 - Is that a boomer back there? No, guy with a sword.

1 - Guy behind you.

1 - I need a new weapon, I'm almost out of ammo. (inaudible)

2 - You can probably take the next

1 - Oh I thought we were (inaudible)

2 - Four left

1 - Hunt them down

Game 2

1 - So where are you?

2 - I'm...

1 - Okay.

2 - Dark ring.

1 - I say let's start back here. Um...

1 - Maybe those back corners aren't the best place but like...here

2 - But the thing is if someone throws a grenade.

1 - Well you can get away though, if you roll out of the way.

1 - There's one on the left.

1 - I'll cover the left you cover the right.

2 - Got it.

2 - Moving up.

2 - Flame time. Right side.

1 - I'll let you get it. There's a flamethrower on the left here.

2 - I'll flank him.

2 - Got him.

1 - Is the boomer dead?

2 - I can't see him.

1 - Oh, he's alive somewhere.

1 - Soldier on the left. Got him.

1 - Boomer on the left. Got him.

1 - These big things... they're called Kantis?

2 - Kanthis.

1 - Kanthis, yeah. There's one in the center.

2 - They can revive other smaller guys, so that would be my first pick.

1 - There's uh...

2 - There's three on my side.

1 - They can revive... even if they're dead?

1 - If they're down or if they're dead?

2 - Down.

2 - Kanthis in back.

1 - Somebody on the left.

2 - Moving up

1 - theres one running across the left here.

1 - in back

1 - one in the center

2 - there's one one on the rigt

1 - Right or left? Well ...

1 - Your back is in the way.

2 - (Inaudible)

1 - We got a...

1 - I think he's a grenadier.

1 - Whoa... where did that come from?

2 - Yo, my side.

1 - What side?

1 - He's in the center. A soldier.

Game 3

1 - One on the left here.

2 - One on the right.

1 - Two coming down the center.

2 - Kanthis on the right.

1 - Okay.

2 - I'm going to stand right here, on the (inaudible).

1 - There's a boomer on the left.

2 - Got it.

2 - Boomer on the right.

2 - They're on my side.

1 - Okay.

1 - Yeah. Boomer on the right.

2 - I had a grinder on my side, but ...

1 - There's another ... mulcher I think. Coming out of the bathroom.

1 - Boomer on the left.

2 - Flamethrower on the right.

1 - Yeah, boomer's in the back.

1 - This... I'm down. Coming towards you.

1 - Good.

2 - I'm low on ammo.

1 - One on the right. Okay.

2 - I'm pretty much out of ammo, actually.

1 - Okay, well I could cover you while you go replenish.

1 - Boomer coming in.

2 - (Inaudible)

2 - Probably want to go..

2 - Oh crap, another boomer...somewhere.

1 - Okay.

2 - And a mulcher.

1 - Okay I'll try to get the mulcher while you go get the ammo.

1 - They're down.

2 - There's a boomer on the left. On your side. Same side.

2 - Got the ammo.

1 - Only three left.

2 - Boomer on the left.

2 - (inaudible)

2 - One left.

1 - He's over here.

2 - Oh, it's that guy.

Game 4

1 - Set this?

2 - Yeah, let's start in the same place.

1 - All right.

2 - One in the center.

1 - One on the right.

2 - There's a skitter coming in the middle.

1 - Are you taking him?

2 - I can't see him.

1 - Flamethrower on my side.

1 - I'm down.

2 - Shoot. Oh you're down?

1 - Yeah.

2 - Okay.

1 - Crawling towards you.

Appendix C – Sample Counter Strike: Global Offensive Transcript

31

00:02:56,400 --> 00:02:58,400

Guys don't forget what i just told you in the debriefing

32

00:02:58,400 --> 00:03:03,400

You guys listen me and just don't panic when we are in the lead

33

00:03:12,400 --> 00:03:18,400

u remember PTR shox? he was always hitting me through mid doors

34

00:03:22,400 --> 00:03:24,400

If I remember they love splitting B

35

00:03:34,400 --> 00:03:36,400

I saw nothing mid I go B

36

00:03:43,400 --> 00:03:45,400

They got the A short, Can you please grenade long NBK?

37

00:03:48,400 --> 00:03:49,400

1 long, dead

38

00:03:50,400 --> 00:03:52,400

They all running short very fast

39

00:03:53,400 --> 00:03:55,400

All 4 short, im very low hp

40

00:03:57,400 --> 00:03:58,400

last short

41

00:04:02,400 --> 00:04:04,400

They didn't even kill someone, that's ridiculous

42

00:04:08,400 --> 00:04:10,400

I think we got the worst of the scene

43

00:04:12,400 --> 00:04:14,400

Nade long please with me smithz

44

00:04:18,399 --> 00:04:19,399

They bought a scout

45

00:04:28,400 --> 00:04:29,400

I smoke A doors

46

00:04:40,400 --> 00:04:43,400

No 1 is holding mid? No informations for A short?

47

00:04:44,400 --> 00:04:45,400

No I don't have any info im not watching short from mid

48

00:04:51,400 --> 00:04:52,400

Ready to flash if I ask you?

49

00:04:53,400 --> 00:04:54,400

yeah yeah no worries

50

00:05:04,400 --> 00:05:05,400

1 is B with a Deagle

51

00:05:12,400 --> 00:05:14,400

They running long, many people long

52

00:05:16,400 --> 00:05:18,400

At least 2 more long with 1 scout

53

00:05:22,400 --> 00:05:25,400

Short 1, that was the guy who was lurking tunnels

54

00:05:26,400 --> 00:05:28,400

im in the corner, they probably running back to you guys

55

00:05:30,400 --> 00:05:32,400

They can't rotate I got this, 2 upper mid

56

00:05:33,400 --> 00:05:35,400

1 more, ok i got him bomb is mid

57

00:05:36,400 --> 00:05:38,400

3v1 don't show yourself smithz play calm

58

00:05:39,400 --> 00:05:40,400

Good job

59

00:05:44,400 --> 00:05:52,400

buy grenades, stay focused, nade long with us smithz

60

00:05:58,400 --> 00:06:00,400

Play safe for the team, don't aggro

61

00:06:13,400 --> 00:06:14,400

I think I heard steps short

62

00:06:21,400 --> 00:06:22,400

They all are on mid

63

00:06:25,400 --> 00:06:26,400

Last B

64

00:06:34,400 --> 00:06:38,20

Let's do the lemondogs strat

65

00:06:41,400 --> 00:06:44,400

If they rush really fast that's gonna be hard

66

00:06:46,400 --> 00:06:50,340

I have 2 flashes i'll stay on site

67

00:06:51,400 --> 00:06:54,380

I flash

68

00:06:54,440 --> 00:06:56,380

Im blind

69

00:07:08,60 --> 00:07:08,820

Im blind mid

70

00:07:21,140 --> 00:07:23,159

1 is tunnels 100%

71

00:07:24,320 --> 00:07:25,219

Tell me when I resmoke

72

00:07:25,220 --> 00:07:26,320

Yeah go right now

73

00:07:47,500 --> 00:07:48,180

Yeah

74

00:07:51,620 --> 00:07:52,320

I flash

75

00:07:58,900 --> 00:08:00,599

1 hit leg

76

00:08:01,920 --> 00:08:02,980

I need backup

77

00:08:06,380 --> 00:08:07,520

I'll stay mid guys

78

00:08:09,620 --> 00:08:10,800

I'll flash twice on the smoke

79

00:08:12,700 --> 00:08:13,200

2nd time

80

00:08:14,100 --> 00:08:15,320

1 more A site with a sniper

81

00:08:15,320 --> 00:08:17,320

Just right behind

82

00:08:17,320 --> 00:08:19,320

Low hp, 1 hp

83

00:08:23,420 --> 00:08:24,360

Dead

84

00:08:25,260 --> 00:08:26,360

1 short maybe

85

00:08:26,360 --> 00:08:28,360

What is he doing?

86

00:08:28,360 --> 00:08:30,360

1 was close short aswell

87

00:08:35,799 --> 00:08:37,380

Nice try

88

00:08:38,380 --> 00:08:39,400

There was 2 left

Appendix D – ACT-R Model Code

```
world_size=0
world_x_range=world_size
world_y_range=world_size
number_of_agents=0

import sys
import os

sys.path.append(os.getcwd() + '\\CCMSuite')

import ccm

log=ccm.log(html=False)

from ccm.lib.actr import *

import random

goal_square = str(random.choice(range(world_x_range))),
str(random.choice(range(world_y_range)))

pass#print goal_square

agent_list = []

class Environment(ccm.Model):    # items in the environment look and act like chunks - but
note the syntactic differences

    prepping_world = True

    while prepping_world:

        occupied_tally = 0

        squares = [ccm.Model(isa='square', x=x, y=y, occupied=0, occupant='nil') \
            for x in range(world_x_range) for y in range(world_y_range)]

        for square in squares:

            if (str(square.x), str(square.y)) == goal_square:

                square.occupied = 1

                square.occupant = 'goal'
```

```

for square in squares:
    x = random.choice(range(10))
    if x == 3 and \
        (str(square.x), str(square.y)) != goal_square:
        square.occupant = 'monster'
        square.occupied = 1

    elif x == 2 and \
        (str(square.x), str(square.y)) != goal_square:
        square.occupant = 'bunny'
        square.occupied = 1

for s in squares:
    if s.occupied == 1:
        occupied_tally += 1
if occupied_tally + number_of_agents >= world_x_range * world_y_range:
    continue
else:
    prepping_world = False

agent_list = []

```

```

class MotorModule(ccm.ProductionSystem): # create a motor module do the actions
    #directions: 0 = up, 1 = right, 2 = down, 3 = left
    production_time = 0.05

def moveForward(self):
    pass#print self.parent.instance_name, 'is trying to move forward'
    if self.parent.facing == 0:
        target_value = self.parent.y_coordinate - 1

```

```

target_x_y = self.parent.x_coordinate, target_value
if self.parent.y_coordinate > 0:
    for s in self.parent.parent.squares:
        if (s.x, s.y) == target_x_y:
            target_square = s
            if s.occupied == '1':
                pass#print 'cannot pass: square occupied'
                pass
            else:
                pass#print 'target square is unoccupied'
                pass#print self.parent.instance_name, ' moved forward to target'
                pass#print '-1y'
                self.parent.y_coordinate -= 1
        else:
            pass#print self.parent.instance_name, ' at edge of map; cannot move forward'
            pass

elif self.parent.facing == 1:
    target_value = self.parent.x_coordinate + 1
    target_x_y = target_value, self.parent.y_coordinate
    if self.parent.x_coordinate < world_x_range:
        for s in self.parent.parent.squares:
            if (s.x, s.y) == target_x_y:
                target_square = s
                if s.occupied == '1':
                    pass#print 'cannot pass: must engage monster first'
                    pass
                else:
                    pass#print 'target square is unoccupied'
                    pass#print self.parent.instance_name, ' moved forward to target'

```

```

        pass#print '+1x'
        self.parent.x_coordinate += 1
else:
    pass#print self.parent.instance_name, ' at edge of map; cannot move forward'
    pass

elif self.parent.facing == 2:
    target_value = self.parent.y_coordinate + 1
    target_x_y = self.parent.x_coordinate, target_value
    if self.parent.y_coordinate < world_y_range:
        for s in self.parent.parent.squares:
            if (s.x, s.y) == target_x_y:
                target_square = s
                if s.occupied == '1':
                    pass#print 'cannot pass: must engage monster first'
                    pass
                else:
                    pass#print 'target square is unoccupied'
                    pass#print self.parent.instance_name, ' moved forward to target'
                    pass#print '-1y'
                    self.parent.y_coordinate += 1
            else:
                pass#print self.parent.instance_name, ' at edge of map; cannot move forward'
                pass

elif self.parent.facing == 3:
    target_value = self.parent.x_coordinate - 1
    target_x_y = target_value, self.parent.y_coordinate
    if self.parent.x_coordinate > 0:
        for s in self.parent.parent.squares:
            if (s.x, s.y) == target_x_y:

```

```

        target_square = s
        if s.occupied == '1':
            pass#print 'cannot pass: must engage monster first'
            pass
        else:
            pass#print 'target square is unoccupied'
            pass#print self.parent.instance_name, ' moved forward to target'
            pass#print '-1x'
            self.parent.x_coordinate -= 1
    else:
        pass#print self.parent.instance_name, ' at edge of map; cannot move forward'
        pass

def turnLeft(self):
    pass#print self.parent.instance_name, ' is turning left'
    if self.parent.facing == 0:
        self.parent.facing = 3
    else:
        self.parent.facing -= 1
    pass#print 'turned left'
def turnRight(self):
    pass#print self.parent.instance_name, ' is turning right'
    if self.parent.facing == 3:
        self.parent.facing = 0
    else:
        self.parent.facing += 1
    pass#print 'turned right'

##navigation methods to approach known target
def faceGoalY(self):

```

```

pass#print 'running faceGoalY'
if self.parent.y_coordinate == int(goal_square[1]):
    pass
elif self.parent.y_coordinate < int(goal_square[1]):
    if self.parent.facing == 2:
        pass
    elif self.parent.facing == 1:
        self.parent.motor.turnRight()
    elif self.parent.facing == 3:
        self.parent.motor.turnLeft()
    elif self.parent.facing == 0:
        self.parent.motor.turnRight()
        self.parent.motor.turnRight()
elif self.parent.y_coordinate > int(goal_square[1]):
    if self.parent.facing == 2:
        self.parent.motor.turnRight()
        self.parent.motor.turnRight()
    elif self.parent.facing == 1:
        self.parent.motor.turnLeft()
    elif self.parent.facing == 3:
        self.parent.motor.turnRight()
    elif self.parent.facing == 0:
        pass

```

```

def approachGoal_Y(self):
    pass#print 'running approachGoal Y'

    if self.parent.y_coordinate == int(goal_square[1]):
        pass
    else:

```

```

    delta_y = self.parent.y_coordinate - int(goal_square[1])
    for i in range(abs(delta_y)):
        self.parent.motor.moveForward()
        yield 0.05

def faceGoalX(self):
    pass#print 'running faceGoalX'
    if self.parent.x_coordinate == int(goal_square[0]):
        pass
    elif self.parent.x_coordinate < int(goal_square[0]):
        if self.parent.facing == 0:
            self.parent.motor.turnRight()
        elif self.parent.facing == 1:
            pass
        elif self.parent.facing == 2:
            self.parent.motor.turnLeft()
        elif self.parent.facing == 3:
            self.parent.motor.turnRight()
            self.parent.motor.turnRight()
    elif self.parent.x_coordinate > int(goal_square[0]):
        if self.parent.facing == 0:
            self.parent.motor.turnLeft()
        elif self.parent.facing == 1:
            self.parent.motor.turnRight()
            self.parent.motor.turnRight()
        elif self.parent.facing == 2:
            self.parent.motor.turnRight()
        elif self.parent.facing == 3:
            pass

```

```

def approachGoal_X(self):
    pass#print 'running approachGoal X'
    if self.parent.x_coordinate == int(goal_square[0]):
        pass
    else:
        delta_x = self.parent.x_coordinate - int(goal_square[0])
        for i in range(abs(delta_x)):
            self.parent.motor.moveForward()
        yield 0.05

class Visual_Search_Module(ccm.ProductionSystem):
    production_time = 0.05

    def check_self_location(self):
        visual_string_temp = str('x_loc:' + str(self.parent.x_coordinate) \
            + ' y_loc:' + str(self.parent.y_coordinate) + ' facing:' + str(self.parent.facing))
        self.parent.visual_buffer.set(visual_string_temp)
        pass#print 'in ', self.parent.instance_name, '\s visual buffer is',
        self.parent.visual_buffer.chunk

    def line_of_sight_constructor(self):
        #line_of_sight_vector = squares in front of the agent, to edge of map
        pass#print self.parent.instance_name, ' is at: ', self.parent.x_coordinate,
        self.parent.y_coordinate

        if self.parent.facing == 0:
            line_of_sight_vector = [s for s in env.squares if \
                s.x == self.parent.x_coordinate and s.y < self.parent.y_coordinate]
        elif self.parent.facing == 1:
            line_of_sight_vector = [s for s in env.squares if \
                s.y == self.parent.y_coordinate and s.x > self.parent.x_coordinate]

```

```

elif self.parent.facing == 2:
    line_of_sight_vector = [s for s in env.squares if \
        s.x == self.parent.x_coordinate and s.y > self.parent.y_coordinate]
elif self.parent.facing == 3:
    line_of_sight_vector = [s for s in env.squares if \
        s.y == self.parent.y_coordinate and s.x < self.parent.x_coordinate]

for i in range(len(line_of_sight_vector)):
    if line_of_sight_vector[i].occupant == 'monster':
        line_of_sight_vector = line_of_sight_vector[0:i+1]
        ##^^ this causes the agent to see only up to the first obstruction (ie, monster)
        break
pass#print self.parent.instance_name, ' LoS vector is ...', line_of_sight_vector
for s in line_of_sight_vector:
    pass#print s.x, s.y
    pass

self.parent.line_of_sight = line_of_sight_vector

def update_top_down_vision(self):
    for square in self.parent.line_of_sight:
        if square.occupant == 'goal':
            self.parent.visual_buffer.set('goal:visible')
        elif square.occupant == 'monster' or square.occupant == 'bunny':
            self.parent.visual_buffer.set(square)

def check_goal(self):
    pass#print self.parent.instance_name, ' checking goal'
    if str(self.parent.x_coordinate) == goal_square[0] and \

```

```

    str(self.parent.y_coordinate) == goal_square[1]:
        pass#print self.parent.instance_name, ' goal conditions match'
        pass#print self.parent.instance_name, ' reached the goal!!!!'
        self.parent.focus_buffer.set('at_goal')
else:
    pass#print 'not yet at goal'

def vision_update(self): #wraps all vision methods, calls them all together

    self.parent.top_down_vision.check_self_location()
    self.parent.top_down_vision.line_of_sight_constructor()
    self.parent.top_down_vision.update_top_down_vision()
    self.parent.top_down_vision.check_goal()

class Entity_Recognition_Module(ccm.ProductionSystem):
    #production_time = 0.01

    def monsterSpotting(DMbuffer='planning_unit:!kill_monster',
visual_buffer='occupant:monster', DM='busy:False', unit_task_buffer=""):
        pass#print self.parent.instance_name, 'initiating monster slaying protocol'
        DM.request('planning_unit:kill_monster')
        context_buffer.set('last_action:none')
        unit_task_buffer.clear()

    def bunnySpotting(DMbuffer='planning_unit:!kill_monster', visual_buffer='occupant:bunny',
DM='busy:False', unit_task_buffer=""):
        pass#print 'bunny spotted; ignoring it'
        pass#print self.parent.instance_name, 'ignoring bunny'
        DM.request('planning_unit:wander')
        visual_buffer.clear()
        context_buffer.set('last_action:none')
        unit_task_buffer.clear()

```

```

#this does not exist in the Wander Model
class Bottom_Up_Communication_Module(ccm.ProductionSystem):
    callOutMade = False
    def call_out_goal(visual_buffer='goal:visible'):
        if callOutMade == True:
            pass#print 'call out already made, not making it again'
            pass
        elif callOutMade == False:
            pass#print 'goal has been spotted: communicating planning unit to teammates'
            for agent in self.parent.parent.agent_list:
                agent.DM.add('planning_unit:navigation_sequential UnitTask1:approach_Y
UnitTask2:approach_X UnitTask3:conclude')
                agent.DM.add('planning_unit:navigation_sequential2 UnitTask1:approach_X
UnitTask2:approach_Y UnitTask3:conclude')
                callOutMade = True
                visual_buffer.clear()

#define agent properties/methods
class MyAgent(ACTR):

##class attributes: buffers, modules, variables
    focus_buffer=Buffer()
    visual_buffer=Buffer()
    context_buffer=Buffer()
    unit_task_buffer=Buffer()

    DMbuffer=Buffer() # create a buffer for the declarative memory (henceforth DM)
    DM=Memory(DMbuffer)#,latency=0.3, threshold=0) # create DM and connect it to its buffer
    #dm_bl=DMBaseLevel(DM,decay=0.5,limit=None)

```

```

motor=MotorModule()
communicationModule=Bottom_Up_Communication_Module()
top_down_vision=Visual_Search_Module()
bottom_up_vision = Entity_Recognition_Module()
moveList = [motor.moveForward, motor.turnLeft, motor.turnRight]
line_of_sight = []

#these set the class variable to the global variable; allows for agent positioning to be tied to
world size

world_x_range = world_x_range
world_y_range = world_y_range
#need a global reference to goal square for delta calc
goal_coord = goal_square
finished = False
placed = False

def init():
    DM.add('planning_unit:wander UnitTask1:random_movement UnitTask2:conclude')
    DM.add('planning_unit:kill_monster UnitTask1:ready UnitTask2:aim UnitTask3:fire
UnitTask4:conclude')
    #find an open square to place agent on
    while not self.placed:
        x = random.choice(range(world_x_range))
        y = random.choice(range(world_y_range))
        init_x_y = (x,y)
        for s in self.parent.squares:
            if (s.x, s.y) == (x,y):
                init_square = s
                if init_square.occupied == 1:
                    continue
                elif init_square.occupied == 0:

```

```

        self.x_coordinate = x
        self.y_coordinate = y
        self.placed = True

    self.facing = random.randint(0,3)

    initial_delta=(abs(self.x_coordinate - int(goal_coord[0])) + abs(self.y_coordinate -
int(goal_coord[1])))

    self.log.delta=initial_delta

    top_down_vision.check_self_location()

    DM.request('planning_unit:!kill_monster')

    context_buffer.set('last_action:none')

#planning unit selection productions

    def start_planning_unit(DMbuffer='planning_unit:?planning_unit',
context_buffer='last_action:none'):

        x = sorted(DMbuffer.chunk.keys())

        y = DMbuffer.chunk[x[0]]

        context_buffer.clear()

        unit_task_buffer.set(y)

    def continue_planning_unit(DMbuffer='planning_unit:?planning_unit',
context_buffer='last_action:?last_action!none!conclude', DM='busy:False'):

        x = sorted(DMbuffer.chunk.keys())

        a = [key for key,value in DMbuffer.chunk.items() if value==last_action]

        nextUnitTask = x.index(a[0]) + 1

        y = DMbuffer.chunk[x[nextUnitTask]]

        context_buffer.clear()

        unit_task_buffer.set(y)

    def finish_planning_unit(DMbuffer='planning_unit:?planning_unit',
context_buffer='last_action:conclude'):

        pass#print 'planning unit', planning_unit, 'finished'

        context_buffer.set('last_action:none')

```

```

    unit_task_buffer.clear()
    DM.request('planning_unit:?.')
    ##fix

#sequential navigation unit tasks
def approach_Y(unit_task_buffer='approach_Y'):
    motor.faceGoalY()
    motor.approachGoal_Y()
    top_down_vision.vision_update()
    unit_task_buffer.clear()
    context_buffer.set('last_action:approach_Y')

def approach_X(unit_task_buffer='approach_X'):
    motor.faceGoalX()
    motor.approachGoal_X()
    top_down_vision.vision_update()
    unit_task_buffer.clear()
    context_buffer.set('last_action:approach_X')

#random movement navigation unit task
def wander(unit_task_buffer='random_movement'):
    random.choice(moveList)()
    top_down_vision.vision_update()
    unit_task_buffer.clear()
    context_buffer.set('last_action:random_movement')

#monster killing unit tasks
def ready(unit_task_buffer='ready'):
    if self.line_of_sight == []:
        pass#print self.instance_name, 'nothing in sight, passing on ready'

```

```

DM.request('planning_unit:!kill_monster')
context_buffer.set('last_action:none')
unit_task_buffer.clear()
visual_buffer.clear()
else:
    for s in self.line_of_sight:
        if s.occupant == 'monster':
            pass#print self.instance_name, 'preparing to engage enemey'
            unit_task_buffer.clear()
            context_buffer.set('last_action:ready')
            continue
    if all(s.occupant != 'monster' for s in self.line_of_sight):
        pass#print 'no monsters here; passing on ready'
        DM.request('planning_unit:!kill_monster')
        context_buffer.set('last_action:none')
        unit_task_buffer.clear()
        visual_buffer.clear()

def aim(unit_task_buffer='aim'):
    pass#print 'aiming weapon'
    unit_task_buffer.clear()
    context_buffer.set('last_action:aim')

def fire(unit_task_buffer='fire'):
    pass#print 'fire!!!!!!'
    for i in self.parent.squares:
        if (i.x, i.y) == (visual_buffer.chunk['x'], visual_buffer.chunk['y']):
            i.occupant='nil'
            i.occupied=0
            pass#print 'monster slain x.x'

```

```

    unit_task_buffer.clear()

    context_buffer.set('last_action:fire')

#generic planning unit conclude unit task
def conclude(unit_task_buffer='conclude', DMbuffer='planning_unit:?planning_unit'):
    pass#print 'concluding the planning unit', planning_unit
    top_down_vision.vision_update()
    unit_task_buffer.clear()
    context_buffer.set('last_action:conclude')

#goal confirmation productions
def reached_goal(focus_buffer='at_goal'):
    pass#print self.instance_name, 'ending from focus buffer'
    unit_task_buffer.clear()
    DMbuffer.clear()
    context_buffer.clear()
    visual_buffer.clear()
    self.finished = True
    focus_buffer.set('waiting_for_teammate')

def wait_for_teammate_to_finish(focus_buffer='waiting_for_teammate'):
    production_time = 0.1
    if all(a.finished == True for a in self.parent.agent_list):
        pass#print 'all agents have finished'
        self.stop()
    else:
        pass#print self.instance_name, 'is waiting'
        pass#print self.parent.agent_list
        for a in self.parent.agent_list: pass#print a.finished

```

```
        focus_buffer.set('waiting_for_teammate')
        pass#print 'xxxxxxxxxxxxxxxxxxxxxxxxxxx'

env=Environment()

for a,i in enumerate(range(number_of_agents)):
    a = MyAgent()
    a.instance_name = str('agent' + str(i))
    env.agent = a
    env.agent.log = log
    env.agent_list.append(a)

#ccm.log_everything(env)
env.run()
# for s in env.squares:
#     pass#print s.occupant
# exit()
ccm.finished()
```

Appendix E – Sample Output of the Model with Debugging Print Statements

Parameters of run:

5x5 world; 5 agents

```
0.000 production_threshold None
0.000 production_time_sd None
0.000 finished False
0.000 production_time 0.05
0.000 production_match_delay 0
0.000 world_y_range 5
0.000 instance_name agent0
0.000 placed False
0.000 world_x_range 5
0.000 production_threshold None
0.000 production_time_sd None
0.000 finished False
0.000 production_time 0.05
0.000 production_match_delay 0
0.000 world_y_range 5
0.000 instance_name agent1
0.000 placed False
0.000 world_x_range 5
0.000 production_threshold None
0.000 production_time_sd None
0.000 finished False
0.000 production_time 0.05
0.000 production_match_delay 0
0.000 world_y_range 5
0.000 instance_name agent2
0.000 placed False
```

0.000 world_x_range 5
0.000 x_coordinate 4
0.000 y_coordinate 4
0.000 placed True
0.000 facing 0
0.000 delta 4
in agent0 's visual buffer is facing0 x_loc4 y_loc4
0.000 x_coordinate 0
0.000 y_coordinate 1
0.000 placed True
0.000 facing 3
0.000 delta 5
in agent2 's visual buffer is facing3 x_loc0 y_loc1
0.000 x_coordinate 3
0.000 y_coordinate 3
0.000 placed True
0.000 facing 3
0.000 delta 4
in agent1 's visual buffer is facing3 x_loc3 y_loc3
0.050 production start_planning_unit
0.050 production start_planning_unit
0.050 production start_planning_unit
0.100 production None
0.100 production None
0.100 production None
0.100 production wander
0.100 production wander
0.100 production wander
0.150 production None
agent0 is trying to move forward

target square is unoccupied

agent0 moved forward to target

-1y

0.150 y_coordinate 3

in agent0 's visual buffer is facing0 x_loc4 y_loc3

agent0 is at 4 3

agent0 LoS vector is ... [ccm.model.Model instance at 0x02832DC8, ccm.model.Model instance at 0x02832DF0, ccm.model.Model instance at 0x02832E18]

0.150 line_of_sight [ccm.model.Model instance at 0x02832DC8, ccm.model.Model instance at 0x02832DF0, ccm.model.Model instance at 0x02832E18]

agent0 checking goal

not yet at goal

0.150 production None

agent2 is turning left

0.150 facing 2

turned left

in agent2 's visual buffer is facing2 x_loc0 y_loc1

agent2 is at 0 1

agent2 LoS vector is ... [ccm.model.Model instance at 0x02832AA8, ccm.model.Model instance at 0x02832B20, ccm.model.Model instance at 0x02832B48]

0.150 line_of_sight [ccm.model.Model instance at 0x02832AA8, ccm.model.Model instance at 0x02832B20, ccm.model.Model instance at 0x02832B48]

agent2 checking goal

not yet at goal

0.150 production continue_planning_unit

0.150 production None

agent1 is trying to move forward

target square is unoccupied

agent1 moved forward to target

-1x

0.150 x_coordinate 2

in agent1 's visual buffer is facing3 x_loc2 y_loc3

agent1 is at 2 3

agent1 LoS vector is ... [ccm.model.Model instance at 0x02832B20, ccm.model.Model instance at 0x02832BE8]

0.150 line_of_sight [ccm.model.Model instance at 0x02832B20, ccm.model.Model instance at 0x02832BE8]

agent1 checking goal

not yet at goal

0.150 production continue_planning_unit

0.150 production continue_planning_unit

agent0 initiating monster slaying protocol

0.200 production (changed before firing)

0.200 production start_planning_unit

0.200 production None

0.200 production None

0.200 production conclude

0.200 production conclude

0.250 production None

concluding the planning unit wander

in agent2 's visual buffer is facing2 x_loc0 y_loc1

agent2 is at 0 1

agent2 LoS vector is ... [ccm.model.Model instance at 0x02832AA8, ccm.model.Model instance at 0x02832B20, ccm.model.Model instance at 0x02832B48]

0.250 line_of_sight [ccm.model.Model instance at 0x02832AA8, ccm.model.Model instance at 0x02832B20, ccm.model.Model instance at 0x02832B48]

agent2 checking goal

not yet at goal

0.250 production finish_planning_unit

0.250 production None

concluding the planning unit wander

in agent1 's visual buffer is facing3 x_loc2 y_loc3

agent1 is at 2 3

agent1 LoS vector is ... [ccm.model.Model instance at 0x02832B20, ccm.model.Model instance at 0x02832BE8]

0.250 line_of_sight [ccm.model.Model instance at 0x02832B20, ccm.model.Model instance at 0x02832BE8]

agent1 checking goal

not yet at goal

0.250 production finish_planning_unit

0.250 production None

0.250 production ready

0.300 production None

planning unit wander finished

0.300 production None

agent0 preparing to engage enemey

0.300 production continue_planning_unit

0.300 production start_planning_unit

0.300 production None

planning unit wander finished

0.300 production start_planning_unit

0.350 production None

0.350 production ready

0.350 production None

0.350 production ready

0.350 production None

0.350 production aim

0.400 production None

no monsters here; passing on ready

0.400 production start_planning_unit

0.400 production None

no monsters here; passing on ready

0.400 production start_planning_unit

0.400 production None

aiming weapon

0.400 production continue_planning_unit

0.450 production None

0.450 production fire

0.450 production None

0.450 production wander

0.450 production None

0.450 production wander

0.500 production None

agent2 is turning left

0.500 facing 1

turned left

in agent2 's visual buffer is facing1 x_loc0 y_loc1

agent2 is at 0 1

agent2 LoS vector is ... [ccm.model.Model instance at 0x02832B98, ccm.model.Model instance at 0x02832C60, ccm.model.Model instance at 0x02832D28, ccm.model.Model instance at 0x02832DF0]

0.500 line_of_sight [ccm.model.Model instance at 0x02832B98, ccm.model.Model instance at 0x02832C60, ccm.model.Model instance at 0x02832D28, ccm.model.Model instance at 0x02832DF0]

agent2 checking goal

not yet at goal

0.500 production continue_planning_unit

0.500 production None

agent1 is turning right

0.500 facing 0

turned right

in agent1 's visual buffer is facing0 x_loc2 y_loc3

agent1 is at 2 3

agent1 LoS vector is ... [ccm.model.Model instance at 0x02832C38, ccm.model.Model instance at 0x02832C60, ccm.model.Model instance at 0x02832C88]

0.500 line_of_sight [ccm.model.Model instance at 0x02832C38, ccm.model.Model instance at 0x02832C60, ccm.model.Model instance at 0x02832C88]

agent1 checking goal

not yet at goal

0.500 production continue_planning_unit

0.500 production None

fire!!!!!!

monster slain x.x

0.500 production continue_planning_unit

0.550 production None

0.550 production conclude

0.550 production None

0.550 production conclude

0.550 production None

0.550 production conclude

0.600 production None

concluding the planning unit kill_monster

in agent0 's visual buffer is facing0 x_loc4 y_loc3

agent0 is at 4 3

agent0 LoS vector is ... [ccm.model.Model instance at 0x02832DC8, ccm.model.Model instance at 0x02832DF0, ccm.model.Model instance at 0x02832E18]

0.600 line_of_sight [ccm.model.Model instance at 0x02832DC8, ccm.model.Model instance at 0x02832DF0, ccm.model.Model instance at 0x02832E18]

agent0 checking goal

not yet at goal

0.600 production finish_planning_unit

0.600 production None

concluding the planning unit wander

in agent2 's visual buffer is facing1 x_loc0 y_loc1

agent2 is at 0 1

agent2 LoS vector is ... [ccm.model.Model instance at 0x02832B98, ccm.model.Model instance at 0x02832C60, ccm.model.Model instance at 0x02832D28, ccm.model.Model instance at 0x02832DF0]

0.600 line_of_sight [ccm.model.Model instance at 0x02832B98, ccm.model.Model instance at 0x02832C60, ccm.model.Model instance at 0x02832D28, ccm.model.Model instance at 0x02832DF0]

agent2 checking goal

not yet at goal

0.600 production finish_planning_unit

0.600 production None

concluding the planning unit wander

in agent1 's visual buffer is facing0 x_loc2 y_loc3

agent1 is at 2 3

agent1 LoS vector is ... [ccm.model.Model instance at 0x02832C38, ccm.model.Model instance at 0x02832C60, ccm.model.Model instance at 0x02832C88]

0.600 line_of_sight [ccm.model.Model instance at 0x02832C38, ccm.model.Model instance at 0x02832C60, ccm.model.Model instance at 0x02832C88]

agent1 checking goal

not yet at goal

0.600 production finish_planning_unit

goal has been spotted communicating planning unit to teammates

0.650 production None

planning unit wander finished

0.650 production None

planning unit wander finished

0.650 production start_planning_unit

0.650 production start_planning_unit

0.650 production None

planning unit kill_monster finished

0.650 production start_planning_unit

0.700 production None

0.700 production None

0.700 production wander
0.700 production approach_X
0.700 production None
0.700 production approach_X
0.750 production None
agent2 is turning left
0.750 facing 0
turned left
in agent2 's visual buffer is facing0 x_loc0 y_loc1
agent2 is at 0 1
agent2 LoS vector is ... [ccm.model.Model instance at 0x02832AD0]
0 0
0.750 line_of_sight [ccm.model.Model instance at 0x02832AD0]
agent2 checking goal
not yet at goal
0.750 production None
running faceGoalX
agent1 is turning right
0.750 facing 1
turned right
in agent1 's visual buffer is facing1 x_loc2 y_loc3
agent1 is at 2 3
agent1 LoS vector is ... [ccm.model.Model instance at 0x02832D78, ccm.model.Model instance at 0x02832E40]
0.750 line_of_sight [ccm.model.Model instance at 0x02832D78, ccm.model.Model instance at 0x02832E40]
agent1 checking goal
not yet at goal
running approachGoal X
agent1 is trying to move forward
target square is unoccupied

agent1 moved forward to target

+1x

0.750 x_coordinate 3

0.750 production None

running faceGoalX

in agent0 's visual buffer is facing0 x_loc4 y_loc3

agent0 is at 4 3

agent0 LoS vector is ... [ccm.model.Model instance at 0x02832DC8, ccm.model.Model instance at 0x02832DF0, ccm.model.Model instance at 0x02832E18]

0.750 line_of_sight [ccm.model.Model instance at 0x02832DC8, ccm.model.Model instance at 0x02832DF0, ccm.model.Model instance at 0x02832E18]

agent0 checking goal

not yet at goal

running approachGoal X

0.750 production continue_planning_unit

0.750 production continue_planning_unit

0.750 production continue_planning_unit

agent1 is trying to move forward

target square is unoccupied

agent1 moved forward to target

+1x

0.800 x_coordinate 4

goal has been spotted communicating planning unit to teammates

0.800 production None

0.800 production approach_Y

0.800 production None

0.800 production approach_Y

0.800 production None

0.800 production conclude

0.850 production None

running faceGoalY

agent1 is turning left

0.850 facing 0

turned left

in agent1 's visual buffer is facing0 x_loc4 y_loc3

agent1 is at 4 3

agent1 LoS vector is ... [ccm.model.Model instance at 0x02832DC8, ccm.model.Model instance at 0x02832DF0, ccm.model.Model instance at 0x02832E18]

0.850 line_of_sight [ccm.model.Model instance at 0x02832DC8, ccm.model.Model instance at 0x02832DF0, ccm.model.Model instance at 0x02832E18]

agent1 checking goal

not yet at goal

running approachGoal Y

agent1 is trying to move forward

target square is unoccupied

agent1 moved forward to target

-1y

0.850 y_coordinate 2

0.850 production None

running faceGoalY

in agent0 's visual buffer is facing0 x_loc4 y_loc3

agent0 is at 4 3

agent0 LoS vector is ... [ccm.model.Model instance at 0x02832DC8, ccm.model.Model instance at 0x02832DF0, ccm.model.Model instance at 0x02832E18]

0.850 line_of_sight [ccm.model.Model instance at 0x02832DC8, ccm.model.Model instance at 0x02832DF0, ccm.model.Model instance at 0x02832E18]

agent0 checking goal

not yet at goal

running approachGoal Y

agent0 is trying to move forward

target square is unoccupied

agent0 moved forward to target

-1y

0.850 y_coordinate 2
0.850 production continue_planning_unit
0.850 production continue_planning_unit
0.850 production None
concluding the planning unit wander
in agent2 's visual buffer is facing0 x_loc0 y_loc1
agent2 is at 0 1
agent2 LoS vector is ... [ccm.model.Model instance at 0x02832AD0]
0.850 line_of_sight [ccm.model.Model instance at 0x02832AD0]
agent2 checking goal
not yet at goal
0.850 production finish_planning_unit
agent1 is trying to move forward
target square is unoccupied
agent1 moved forward to target
-1y
0.900 y_coordinate 1
agent0 is trying to move forward
target square is unoccupied
agent0 moved forward to target
-1y
0.900 y_coordinate 1
goal has been spotted communicating planning unit to teammates
goal has been spotted communicating planning unit to teammates
0.900 production None
0.900 production conclude
0.900 production None
0.900 production conclude
0.900 production None
planning unit wander finished

0.900 production start_planning_unit
agent1 is trying to move forward
target square is unoccupied
agent1 moved forward to target
-1y
0.950 y_coordinate 0
.....
1.400 production wait_for_teammate_to_finish
1.450 production None
aiming weapon
1.450 production continue_planning_unit
1.450 production None
agent1 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.450 production wait_for_teammate_to_finish
1.450 production None
agent0 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.450 production wait_for_teammate_to_finish
1.500 production None
1.500 production fire
1.500 production None
agent1 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.500 production wait_for_teammate_to_finish
1.500 production None
agent0 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]

1.500 production wait_for_teammate_to_finish
1.550 production None
fire!!!!!!
monster slain x.x
1.550 production continue_planning_unit
1.550 production None
agent1 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.550 production wait_for_teammate_to_finish
1.550 production None
agent0 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.550 production wait_for_teammate_to_finish
1.600 production None
1.600 production conclude
1.600 production None
agent1 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.600 production wait_for_teammate_to_finish
1.600 production None
agent0 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.600 production wait_for_teammate_to_finish
1.650 production None
concluding the planning unit kill_monster
in agent2 's visual buffer is facing0 x_loc3 y_loc0
agent2 is at 3 0
agent2 LoS vector is ... []

1.650 line_of_sight []
agent2 checking goal
not yet at goal
1.650 production finish_planning_unit
1.650 production None
agent1 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.650 production wait_for_teammate_to_finish
1.650 production None
agent0 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.650 production wait_for_teammate_to_finish
1.700 production None
planning unit kill_monster finished
1.700 production start_planning_unit
1.700 production None
agent1 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.700 production wait_for_teammate_to_finish
1.700 production None
agent0 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.700 production wait_for_teammate_to_finish
1.750 production None
1.750 production ready
1.750 production None
agent1 is waiting

[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]

1.750 production wait_for_teammate_to_finish

1.750 production None

agent0 is waiting

[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]

1.750 production wait_for_teammate_to_finish

1.800 production None

agent2 nothing in sight, passing on ready

no monsters here; passing on ready

1.800 production None

agent1 is waiting

[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]

1.800 production None

agent0 is waiting

[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]

1.800 production start_planning_unit

1.800 production wait_for_teammate_to_finish

1.800 production wait_for_teammate_to_finish

1.850 production None

1.850 production None

agent1 is waiting

[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]

1.850 production None

agent0 is waiting

[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]

1.850 production approach_X

1.850 production wait_for_teammate_to_finish
1.850 production wait_for_teammate_to_finish
1.900 production None
running faceGoalX
agent2 is turning right
1.900 facing 1
turned right
in agent2 's visual buffer is facing1 x_loc3 y_loc0
agent2 is at 3 0
agent2 LoS vector is ... [ccm.model.Model instance at 0x02832DC8]
4 0
1.900 line_of_sight [ccm.model.Model instance at 0x02832DC8]
agent2 checking goal
not yet at goal
running approachGoal X
agent2 is trying to move forward
target square is unoccupied
agent2 moved forward to target
+1x
1.900 x_coordinate 4
1.900 production None
agent1 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.900 production None
agent0 is waiting
[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80,
__main__.MyAgent instance at 0x0289EE40]
1.900 production continue_planning_unit
1.900 production wait_for_teammate_to_finish
1.900 production wait_for_teammate_to_finish

goal has been spotted communicating planning unit to teammates

1.950 production None

1.950 production None

agent1 is waiting

[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80, __main__.MyAgent instance at 0x0289EE40]

1.950 production None

agent0 is waiting

[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80, __main__.MyAgent instance at 0x0289EE40]

1.950 production approach_Y

1.950 production wait_for_teammate_to_finish

1.950 production wait_for_teammate_to_finish

2.000 production None

running faceGoalY

in agent2 's visual buffer is facing1 x_loc4 y_loc0

agent2 is at 4 0

agent2 LoS vector is ... []

2.000 line_of_sight []

agent2 checking goal

agent2 goal conditions match

agent2 reached the goal!!!!

running approachGoal Y

2.000 production None

agent1 is waiting

[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80, __main__.MyAgent instance at 0x0289EE40]

2.000 production None

agent0 is waiting

[__main__.MyAgent instance at 0x02832A08, __main__.MyAgent instance at 0x02893F80, __main__.MyAgent instance at 0x0289EE40]

2.000 production reached_goal

2.000 production wait_for_teammate_to_finish
2.000 production wait_for_teammate_to_finish
2.050 production None
agent2 ending from focus buffer
2.050 finished True
2.050 production None
all agents have finished
[Finished in 0.3s]