The Macro Architecture Hypothesis: Applications to Modeling Teamwork, Conflict Resolution, and Literary Analysis

Robert L. West (robert_west@carleton.ca),
Emmanuelle Hancock (emme.hancock@gmail.com),
Sterling Somers (sterlingsomers@gmail.com),
Korey MacDougall (koreymacdougall@gmail.com),
Francis Jeanson (fjeanson@connect.carleton.ca)
Institute of Cognitive Science, Carleton University, Ottawa, Ontario, Canada

Abstract

ICCM was founded to create a scientific forum for cognitive modeling results. The gold standard for cognitive models has long been the good fit of a model to experimental data. In this paper we argue that although this type of work is essential, we also need *explicit* theory and methods for creating and evaluating cognitive models of real world tasks. This is discussed in terms of the relationship between macro and micro cognition and our own theory and methods for bridging the two. These are illustrated with three related research projects on teamwork, conflict resolution, and literary analysis.

Keywords: Macro Cognition; Cognitive Modeling; Cognitive Architectures; Teamwork; Mediation; Conflict Resolution; Literary Analysis

Introduction

In theory, cognitive architectures can provide a scientific basis for modeling complex, real world behaviors and for creating more efficient, safer systems for humans to work in. However, to a large extent, this laudable goal remains unfulfilled. Cognitive Modeling in systems design (if it is used at all) is mostly limited to the use of human factors models of specific, isolated elements of the task (e.g., work load, attention, etc.). There are a number of possible reasons for this; cognitive architectures are still limited and unproven in some ways, and they are difficult to connect to real world simulation engines. However, another important reason is that many are not convinced that cognitive architectures are useful or valid for modeling complex, real world behaviors (e.g., Klein et al., 2002).

Related to this, the study of cognition has been divided into micro and macro cognition (Cacciabue & Hollnagel, 1995; Klein et al., 2003), where micro cognition refers to cognition as it is studied in cognitive psychology experiments and macro refers to cognition in real world tasks. The question is then whether or not architectures built to model micro cognitive tasks can scale up to usefully model macro cognitive behavior. Some have argued that they cannot (Klein et al., 2002), whereas we have argued that they can (West & Nagy, 2007). More specifically, we have argued that the difference between micro and macro cognition is analogous to the difference between neural models and (micro) cognitive models (West & Pronovost, 2009). Therefore what is needed is a theoretical account of how micro cognition produces macro cognition. However, in addition to theory, novel approaches to methodology and testing are required as this endeavor involves moving beyond the experimental paradigm of micro cognition into the complex real world tasks of macro cognition. In this paper we put outline a theoretical framework and discuss methodologies for evaluation with examples.

The Macro Architecture Hypothesis

The macro architecture hypothesis is one way of framing the relationship between micro and macro cognition. The hypothesis proposes there is a is a macro level architecture that is built on a micro level architecture (which is built on a neural architecture). The macro cognitive architecture is hypothesized to exist in the brains of individuals and to enable us to apply our information processing abilities (micro cognition) to complex, dynamic, multi agent, real world tasks (macro cognition).

The idea of a cognitive architecture was proposed by Newell (Newell, 1990). It is based on a systems level view of intelligence. The assertion that the human brain is designed in a systems level way is, itself, a hypothesis. The test of this hypothesis is whether or not a unified cognitive architecture can be created. Difficult philosophical questions can be asked about the reality of systems levels and architectures but we will ignore them in this paper (although see West & Leibovitz, 2012). The point of this paper is to outline a practical framework for linking macro and micro cognition and for evaluating such models on macro cognitive tasks..

Micro and Macro architectures

When building a model of a task it is common practice to first construct a unit task model (i.e., a high level model of the components of the task) and then to work out how the micro cognitive architecture accomplishes these. In contrast, the macro cognitive architecture is concerned with how the unit tasks are organized, interrupted, re-organized, and re-interpreted in response to external events. So the first step in building a model in a macro cognitive architecture is the same as building any other cognitive model, except that the unit task structure is constrained by the macro architecture.

Evaluating macro models is both similar and different from evaluating micro models. The big difference is in terms of the types of measurements and analyses that can be used. Macro cognition is generally at a time scale that is too long and too noisy to use reaction time measures. Also, macro tasks are not randomized so it is often the case that the order of events will not be the same across trials, and this means that averaging is of limited use. Finally, macro cognition is often concerned with higher level constructs that can only be evaluated using qualitative methods.

Here, it is important to note the distinction between system levels and levels of analysis. Level of analysis refers to analyzing something at a particular level of measurement (e.g., fMRI, RT, questionnaire, discourse analysis, etc.). Systems level means that all behavior at a particular level can be explained by one unified architecture. Level of analysis and systems level tend to be correlated in that higher systems levels are usually studied at a higher levels of analysis, but this is not a rule.

In terms of similarities, macro models and micro models are both models and certain truths about modeling apply equally to both. In particular, we argue that Newell's (1973) critique of modeling in cognitive psychology applies equally to the study of macro cognition. In his well known paper, You can't play 20 questions with nature and win, Newell pointed out that creating different models for each cognitive phenomena is of limited use. The models provide insight into individual phenomena but the practice, unchecked, results in a bewildering plethora of unrelated models. This criticism can be applied to macro cognition as well. Newell's solution for cognitive psychology was to work towards creating a unified cognitive architecture. Potentially, this approach can also be usefully applied to macro cognition.

Scientific Evaluation

In terms of evaluation, like micro cognitive architectures, macro cognitive architectures cannot be evaluated within a strict Popperian scientific framework. This is because, although a specific model built within an architecture can be disproven, this does not mean that the architecture is false. It is also possible that the model was built in the wrong way or that the architecture needs an additional component or a minor adjustment, or that the task was misunderstood. Because of this, Newell (1990) noted that cognitive architectures should be evaluated within a Lakatosian scientific framework (Lakatos, 1970; also see Cooper, 2007 for discussion of this framework applied to cognitive architectures). In the Lakatosian framework theories are evaluated across time in terms of progress and usefulness. Therefore, in this framework, an architecture is considered scientific as long as it continues to further unify different phenomena and produce parsimonious explanations. For example, under this framework the theory that planets orbit in circles was initially a valid scientific theory as it produced significant progress in understanding our solar system, but it became less valid as more and more epicycles were required to describe the orbits, which were actually elliptical and only approximated by circles.

An important concept for applying the Lakatosian approach to cognitive architectures is Lakatos's idea that a theory (or architecture) can be understood in terms of core and peripheral commitments. When a model fails, the first line of defense are the peripheral commitments whereas the core concepts are only challenged if it is not possible to make progress by altering the peripheral concepts (ACT-R is a good example of an architecture that is developed in this way, see Cooper, 2007, for discussion).

Cooper (2007) also notes that, in addition to core and peripheral commitments models built in (micro) cognitive architectures have to be evaluated in terms of the accuracy of the task model (i.e., the knowledge added into the architecture to allow it to do the task). He also points out the scientific desirability of validating the task model separately so it does not become an added source of variance for evaluating the architecture. Tasks used in cognitive psychology experiments are kept very simple so the task model is reasonably obvious. However, this is not the case for macro cognitive tasks. A lot of assumptions about the task knowledge are needed to get a micro cognitive architecture to model macro cognitive tasks.

A macro cognitive architecture would ameliorate this problem. A macro cognitive architecture would have core and peripheral components that constrain how task knowledge is organized. Figure 1 illustrates how a macro cognitive architecture could be combined with micro cognitive and neural architectures within a Lakatosian framework adapted to accommodate systems levels. Note that in this scheme the core mechanisms of the architecture above are used to challenge the peripheral mechanisms of the architecture below. If the core mechanisms of different levels are at odds then the system would be considered incommensurate. In terms of upward constraints, it must be possible to build the architecture above on the architecture below (possibly with some modifications to peripheral components). Essentially, the core components define the core functionality of the each systems level.

Using this framework, a macro cognitive architecture can be evaluated in two different ways. The first is whether it can reasonably and efficiently model human behavior across a diverse set of macro cognitive tasks. The way to evaluate this is to build the architecture and test it across a diverse set of macro cognitive tasks. The second way is whether the macro cognitive architecture can be reasonably and efficiently produced by a micro cognitive architecture. The way to evaluate this is to build the macro cognitive architecture on top of a micro cognitive architecture.

Note that this also changes how the micro cognitive architecture is evaluated. Instead of being evaluated directly on its ability to perform the task it is evaluated based on its ability to provide the core and peripheral functionality of the macro cognitive architecture. If it can, and the macro architecture can model the task, then it should all work (although the whole thing should be run to check it).

Neuroscience

Some people are uncomfortable with a systems level approach and find it hard to think about functions that are divorced from neurons. Currently the dominant way of thinking about (micro) cognitive functions is that that they are produced by dedicated groups of neurons, or neural modules. However, this is not the only possibility. The alternative is neural reuse (Anderson, 2010) whereby cognitive functions are created through the interaction of different neural groups. We do not take a stand on this but neural reuse is probably how the macro cognitive architecture is produced. That is, it is produced through the interaction of basic cognitive functions, which are produced through neural activity. Therefore, neural evidence for the macro cognitive architecture would come from the pattern of interactions across the brain and not from localization. Neural imaging techniques are now able to identify the network of brain areas involved across tasks with reasonable accuracy (e.g., Varela et al, 2001) so it is theoretically possible to link the macro architecture directly to neural activity.

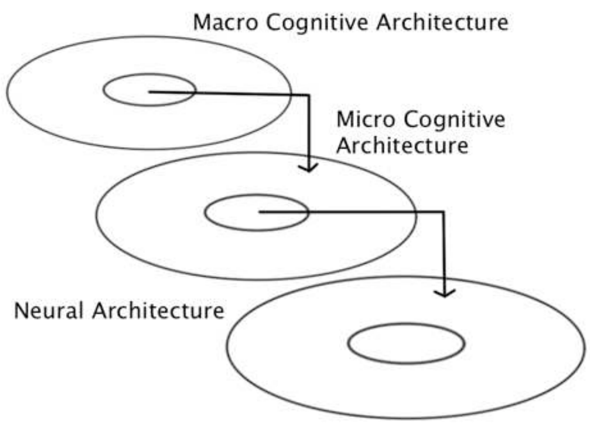

Figure 1. Lakatosian framework adapted for use with systems levels

In terms of the neural origins of the macro cognitive architecture, there are two possibilities. One is that the macro cognitive architecture is genetically hardwired. That is, brains come into the world designed to get tasks done in the world. The other possibility is that it is learned. According to the Rationality Principle (Card, Moran, & Newell, 1983), if people are exposed to similar problems, over time they will converge on learning similar ways to deal with the problem. Therefore, the macro cognitive architecture could be created through the developmental process of learning how to do macro tasks in the real world. However, in either case the result is the same - a system for processing macro cognitive tasks that is similar across individuals. It is also possible that the two processes work together as in language development.

SGOMS

Most studies of macro cognition can be considered studies of experts operating in the real world. This is probably due to the practical need to produce applied research for different areas, but whatever reason, the dominant approach in the study of expertise (regardless if the focus is micro or macro) is to treat each domain of expertise separately (see Ericsson et al, 2006). This is also the dominant approach for Cognitive Engineering within Systems Engineering (Kirlik, 2012). Most work on expert learning does not consider the possibility of a task independent way of organizing expert knowledge. In contrast, we used the SGOMS architecture which is based on the hypothesis that all expert behavior is mediated by a fixed set of interacting cognitive structures.

SGOMS (West & Nagy, 2007) is essentially a control system that allows expertise to occur in chaotic, multi agent environments with interruptions and re-planning. SGOMS is an extension of the GOMS architecture (Card, Moran, & Newell, 1983), which is accurate for modeling isolated experts in controlled conditions. In the GOMS architecture, the highest construct is the *unit task*, which is a control structure that mediates between the micro cognitive demands of the task and the micro cognitive limitations of the individual. SGOMS adds an additional control structure, called *planning units*, above unit tasks. Planning units mediate between the need of the expert to execute their unit tasks in an effective way given a constantly shifting context, interruptions, and unexpected events.

SGOMS:ACT-R

SGOMS:ACT-R is a version of SGOMS implemented in ACT-R (Anderson & Lebiere, 1998) which is a well studied and well tested micro cognitive architecture. From an ACT-R perspective, SGOMS:ACT-R represents the hypothesis that the right way is to build a model of an expert in ACT-R is by using the SGOMS macro cognitive architecture (see West, & Pronovost, 2009, West & Somner, 2011, for a discussion). Doing this requires every ACT-R model to have a fixed set of dedicated productions that mediate the productions and declarative memory content related to specific tasks. As outlined above, these productions could be hard wired or they could be arrived at through developmental learning and the Rationality Principle, or both. It also requires minor alterations to some of the peripheral components of ACT-R.

Note, though, that this is not the first attempt to create systems for generating models of macro level tasks using micro level architectures (see Ritter et al 2006 for a review). In our opinion, this type of endeavor implicitly presupposes some sort of macro architecture. At a minimum, the concept of a macro architecture is a candidate for framing this type of research and making the goals and commitments clearer.

However, in this paper we will focus on SGOMS macro cognitive models. That is, models constructed from the unit task level up. These models rely only on the SGOMS architecture and can be evaluated independently, without reference to the ACT-R implementation. This is because one of the theoretical commitments of SGOMS is that the macro architecture functions to protect micro cognition within unit tasks from interruptions and unexpected events. However, the second phase of evaluating an SGOMS model would be to implement the micro cognitive details and evaluate it on that basis. Below we describe three projects that illustrate macro cognitive approaches to creating and evaluating our models. The important point is that the same architecture must work for all tasks.

Mediation

Mediation is an important macro cognitive activity. Whenever multiple agents must agree on a decision or a plan, differences can arise and must be resolved. To some degree all people have strategies for mediation, although some may be more effective than others. To study this we constructed a model of professional insight mediation. An insight mediator focuses on helping the parties involved in the mediation gain deeper insight into the problem and hopefully resolve it (Melchin & Picard, 2008). Other schools of mediation also exist and one of the goals of this project is to use the model to understand the ways in which different approaches to professional mediation are the same and different. The other thing that is good about studying professional, trained mediators is that they are experts in dealing with problematic mediation styles that clients bring to the table. Therefore, by modeling the mediation process we also gain insight into common problems that prevent people from resolving conflict on their own.

The model building process involved using textbooks, followed by interviews with experts, more adjustments to the model, more interviews, more adjustments, and finally a model tracing session using a video of a simulated mediation (the simulation was not scripted, the clients were played by trained mediators who re-created common communication problems for the mediator to work with). Model tracking produced remarkably good results. All actions (at the unit task level and above) were consistent with the model. In two cases there was some ambiguity but these were satisfactorily resolved by interviewing the mediator from the video. This result is important because it shows that SGOMS is able to model mediation as it occurs in real world tasks

In terms of communication we noticed that the mediator would pause noticeably before the beginning a new planning unit and wait to see if the parties were ready to move on. This is consistent with the theory that planning units represent conceptually coherent information (West & Nagy, 2007). According to the SGOMS architecture only one planning unit can be active at a time and there is a cost (in terms of time and effort) for changing from one to another. Signaling the change to other agents and waiting to make sure the the previous planning unit is actually over therefore makes sense. Based on this observation and reasoning we will add the pause as a signal for planning unit change to the theory and evaluate it in future studies.

Another finding came out of discussions with the experts who noticed that important goals and values were not represented in the model. This highlighted an important issue - the difference between static knowledge represented in the model and dynamic effects that occur as a result of running the model, which we will refer to as emergent effects. Emergent effects are often the goal of the process being studied so it is important to represent them as part of the model in order to create a complete picture. For example, one intended effect of the process of insight mediation is to create insight, but there is no planning unit or unit task for creating insight. To represent goals and values that arise out of the process of running an SGOMS model we added emergent effects to the elements of SGOMS. Although emergent effects are not involved in running the model they are important for evaluating the behavior of the model and for understanding the relationship between the process and the goals and values driving the task.

Currently we are working on translating the mediation model into a board game, which we hope will facilitate training by taking students through the process in a realistic way (if the model is right). Feedback on the usefulness of the game will provide further opportunities to evaluate both the model, the SGOMS architecture, and the implementation of the SGOMS architecture in ACT-R. We have also investigated the relationship between SGOMS and board games in a pilot study run by one of the authors (E.H.) who hosts a board game night. After analyzing the different games in terms of planning units, unit tasks, and constraints, E.H. reported a significant decrease in the time to teach a new game using the SGOMS descriptions. We plan to follow this up with an experimental study to confirm it.

Team Play

Team work is an important area of macro cognition that has received a lot of attention. SGOMS has strong implications for how team work occurs. West & Nagy (2007) characterized the macro level in terms of the tension between reacting to interruptions and unexpected events (i.e., real time, real world chaos, uncooperative agents) and maintaining coordination with other people (i.e., team work, planning, coordination). The SGOMS architecture was created as a control system that mediates between these two factors. SGOMS extends the GOMS system for modeling expert behavior to allow the same structures to be used to model expertise in complex, messy environments, where standard GOMS models break down (Kieras & Santoro, 2004; West & Nagy, 2007).

In terms of communication between team members, SGOMS predicts that the nature of discussion between agents should be different depending on whether the information communicated is meant to inform planning unit choice or unit task choice. In SGOMS, planning units represent acting according to a plan so discussions about selecting, modifying, or creating a new planning unit to deal with a novel situation can only occur when switching between planning units. Within a planning unit communication should be limited to factors relevant to choosing the next unit task or interrupting and exiting the planning unit. Moreover, because unit tasks are theorized to be units that are not expected to be interrupted or to produce downtime an agent cannot stop in the middle of a unit task to communicate. Communication, if it occurs, would be interleaved with performing the unit task and could be a source of distraction, stress and error (as in errors caused by driving while talking on a cell phone, which can be modeled accurately in ACT-R, see Salvucci, 2006).

The modeling process for this project was somewhat different since it is a game and not real life. To investigate this, two of the authors (S.S. and F.J.) spent countless hours developing expertise on team play in the X-Box video game, Gears of War (Psycho Level). We started by defining the operators as all of the actions that can be achieved through the controller. We then recorded protocols from the players during game play. We identified an initial set of planning units and unit tasks and wrote each on a sticky note and put the sticky notes hierarchically (planning units on top, unit tasks on bottom) on the wall next to the game. We then iterated between playing the game and rearranging the sticky notes as well as adding and deleting them. We also used a method we call *expert substitution*, in which a less experienced player (R.W.) took the place of one of the experts (S.S.) who then acted as a coach. This method was very helpful for uncovering aspects of the task that were implicit for the experts. Eventually it was realized that most of the game was covered by three planning units: find cover, defend cover, and take ground. Others included: get ammo, revive comrade, select new weapon, and survive (for when things have gone badly wrong).

In terms of communication we found a very clear pattern. Discussions about what to do next occurred only when the players were between planning units and both in a secure, safe area. Therefore, discussing what to do next was added as a planning unit that could fire only under the constraints mentioned. An important prediction of the SGOMS architecture is that there is no representation of task knowledge higher than planning units. Therefore, discussions about which planning unit to do next have to be conducted through the use of planning units. In theory, if there was a disagreement, planning units similar to those in the mediation model would come into play (note- if there is no discussion or contemplation needed the next planning can be chosen through a production rule that triggers the planning unit, therefore it is also possible to switch very quickly to an appropriate planning unit if required, so if a hidden alien were to attack during a discussion it would interrupt the discussion planning unit and immediately trigger a switch to the survive planning unit).

When the players were engaged in a planning unit involving fighting the communication was very different. Instead of discussion they would call out relevant information as it was observed. Mostly this involved the movement of enemy combatants. Through this system team members gathered location information through their auditory system (i.e., information from their teammate) as well as from their own visual system. By sharing visual information they also created a common ground (Klein et al, 2004) for decision making allowing them to act in coordination without actually planning the coordination. The coordination arose through both players applying the same expert knowledge to the same, common ground data set. During less time pressured planning units there was casual conversation, sometimes related to the task. This was modeled in SGOMS as interleaving two planning units, the current one and the discussion planning unit. This can happen in the model if the current planning unit has downtime between unit tasks.

Taking the game theme one step further, based on this study, the pilot study by E, and other work showing that SGOMS is particularly good at describing games (West & Pronovost, 2009), we have hypothesized that the appeal of a good game is that it allows us to to exercise a natural urge to use our macro cognitive architecture to display expertise. This hypothesis is testable as it predicts that games that are consistent with the SGOMS architecture will be more enjoyable and easier to learn than games that are not. Further research is planned to evaluate this hypothesis

Literary Analysis

If we all share a common architecture for expressing expertise then the form of this architecture should be apparent in descriptions of expertise that people find compelling. Writers and story-tellers often depict experts who share their (the author's) skill sets, or do extensive research to make their characters psychologically believable. Portrayals of experts in literature and other story-telling media are often highly accurate, and tend to be holistic and situated. This is a source of more complete and complex pictures of what constitutes expertise in various domains (particularly those that involve competition, social interaction, language, and complex or multi-part tasks). Many forms of expertise involve fluidly (and often idiosyncratically) adapting to people and situations, and these facets of expertise are more difficult to address in standard experimental protocols.

If the macro architecture hypothesis is true then we should also expect to see regularities across literary descriptions of expertise. Furthermore, if the SGOMS architecture is valid then we should find regularities consistent with the architecture. In particular, we hypothesize that compelling literary portrayals of expertise are created (at least in part) by leveraging aspects of the narrative that reveal the underlying architecture. In this way, a reader with no knowledge of an expert domain can still judge and relate to a character portrayed as an expert in that domain.

Expert-oriented and expert-produced literature also provides a valuable window into the importance and function of specialized vocabularies. Many, if not most, expert domains, require the development and use of highly specific vocabularies and dialects that allow the rapid and efficient compression and transmission of information between experts. As noted in the Mediation and Team Play sections, SGOMS can be used to analyze communication patterns. Therefore, we hypothesize that more compelling portrayals of expertise will use specialized vocabulary to communicate unit tasks and planning units in a manner consistent with SGOMS (currently we are developing a literary analysis to systematically investigate these hypotheses).

Conclusion

We have presented the idea of a macro architecture and described SGOMS as an example of a macro architecture. We have also described three very different research projects and how they are related to evaluating the SGOMS architecture. This approach to research has implications for both macro and micro cognitive modeling.

For Macro Cognition, the introduction of an architecture deals with the problem of proliferating an endless stream of unrelated (or loosely related) one-off models. Furthermore, the requirement that the architecture must work across knowledge domains creates a much stronger scientific framework for evaluation.

For Micro Cognition the introduction of a macro architecture creates the opportunity for a divide and conquer strategy. Instead of scaling up micro architectures in different ways to model different macro level tasks, micro architectures are scaled up to have the functionality needed to model all macro level tasks across all knowledge domains. One way of thinking about this is as a strategy to more efficiently achieve Newell's (1990) goal of a complete, fully functional cognitive architecture.

References

Anderson, J. R. & Lebiere, C. (1998). The atomic components of thought. Mahwah, NJ:Erlbaum.

Anderson, M. L. (2010). Neural reuse: A fundamental organizational principle of the brain. Behavioral and brain sciences , 33 (4), 245-66.

Card, S., Moran, T., and Newell, A. (1983) The Psychology of Human-Computer Interaction, Lawrence Erlbaum Associates, Hillsdale, NJ.

Cooper, R. P. (2007). The Role of Falsification in the Development of Cognitive Architectures: Insights from a Lakatosian Analysis. Cognitive Science 31 (2007) 509–533

Ericsson, K. A., Charness, N., Hoffman, R. R., & Feltovich, P. J. (Eds.). (2006). The Cambridge handbook of expertise and expert performance. New York: Cambridge University Press.

Kieras, D., and Santoro P. (2004) Computational GOMS Modeling of a Complex Team Task: Lessons Learned. In Proceedings of CHI 2004 (pp. 97-104), April 24–29, 2004, Vienna, Austria.

Klein (2002). Year One Final Report: High-Level Cognitive Processes In Field Research. Prepared through participation in the Advanced Decision Architectures Collaborative Technology Alliance sponsored by the U.S. Army Research Laboratory under Cooperative Agreement DAAD19-01-2-0009

Klein, G., Woods, D. D., Bradshaw, J. D., Hoffman, R. R., and Feltovich, P. J. (November/December 2004). Ten challenges for making automation a "team player" in joint human-agent activity. IEEE: Intelligent Systems, pp. 91-95.

Kirlik, A. (2012), The Emerging Toolbox of Cognitive Engineering Models. Proceedings of the International Conference on Cognitive Modeling.

Lakatos, I. (1970). Falsification and the methodology of scientific research programs. In I. Lakatos & A. Musgrave (eds.) Criticism and the Growth of Knowledge (pp. 91–196). Cambridge, UK: Cambridge University Press.

Melchin, K., & Picard, C. (2008). Transforming Conflict through Insight, University of Toronto Press, Scholarly Publishing Division, Toronto, ON.

Newell, A. (1973). You Can't Play 20 Questions with Nature and Win: Projective Comments on the Papers of this Symposium. In W. G. Chase (Ed.), Visual Information Processing: Proceedings of the 8th Symposium on Cognition (pp. 283-308). New York:Academic Press.

Newell, A. (1990). Unified theories of cognition. Cambridge: Harvard University Press.

Ritter, F. E., Haynes, S. R., Cohen, M. A., Howes, A., John, B. E., Best, B., Lebiere, C., Jones, R. M., Lewis, R. L., St. Amant, R., McBride, S. P., Urbas, L., Leuchter, S., & Vera, A. (2006). High-level behavior representation languages revisited. In Proceedings of the Seventh International Conference on Cognitive Modeling, 404-407. Trieste, Italy: Edizioni Goliardiche.

Salvucci, D. (2006). Modeling Driver Behavior in a Cognitive Architecture. Human Factors, 48 (2), 362-380.

West, R. L. & Leibovitz, D. P. (2012). Understanding each other: Defining a conceptual space for cognitive modeling. The proceedings of Cognitive Science.

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

West, R. L., & Pronovost, S. (2009). Modeling SGOMS in ACT-R: Linking Macro- and Microcognition. Journal of Cognitive Engineering and Decision Making, 3 (2), 194-207.

Varela F, Lachaux JP, Rodriguez E, Martinerie J. (2001). The brainweb: phase synchronization and large-scale integration. Nat Rev Neurosci. 2001 Apr;2(4):229-39.

West, R. L., & Somner, S. (2011) Scaling up from Micro Cognition to Macro Cognition: Using SGOMS to build Macro Cognitive Models of Sociotechnical Work in ACT-R. The proceedings of Cognitive Science