Barriers, Bridges, and Progress in Cognitive Modeling for Military Applications

Kevin A. Gluck, PhD
Air Force Research Laboratory

The role of the Air Force Research Laboratory (AFRL), like the other service laboratories, is to conduct the basic and applied research and advanced technology development necessary to create new future technology options for the Department of Defense. At the Warfighter Readiness Research Division of AFRL's Human Effectiveness Directorate we have initiated a research program focused on mathematical and computational cognitive process modeling for replicating, understanding, and predicting human performance and learning. This research will lead to new technology options in the form of human-level synthetic teammates, cognitive readiness analysis tools, and predictive and prescriptive knowledge tracing algorithms. Creating a future in which these objectives become realities requires tightly-coupled, multi-disciplinary, collaborative interaction among scientists and engineers dedicated to overcoming the myriad challenges standing between current reality and our future vision.

Barriers and Bridges

There are many barriers to progress in cognitive science in general and to computational cognitive process modeling in particular. I will emphasize just two of them here. The first is a *domain barrier*. There exists an infinite variety of domains in which humans learn and perform and in order to simulate human performance and learning in a particular domain we must provide relevant domain knowledge to the simulated human. Transfer from one domain to the next is largely a function of the degree to which the knowledge in the two domains overlaps. The reason this is problematic for scientific progress is that the domains typically used to study human cognitive functioning in the laboratory are very different than the domains of application in the real world. Laboratory domains are mostly simple, abstract, and of short duration, whereas real world application domains are complex, situated, and of long duration. Thus, in the field of cognitive science we must look for ways to build bridges between laboratory and applied contexts.

The other barrier I will emphasize here is a ***disciplinary barrier***. Cognitive science is a field of study comprised of seven sub-disciplines: anthropology, artificial intelligence, education, linguistics, neuroscience, philosophy, and psychology. These sub-disciplines involve very different methods, frameworks, and theories, and it is challenging to make progress at disciplinary intersections. For instance, there is a powerful zeitgeist currently associated with neuroscience-based explanations of phenomena in various fields ranging from the more obvious, such as psychology (leading to the creation of a field known as neuropsychology) to the less obvious, such as economics (leading to the creation of a field known as neuroeconomics). This has led some to begin speculating that there ought to be ways to improve the readiness of our military personnel by capitalizing on the tools, methods, empirical results and theories of neuroscience. Simultaneously, there is interest in bringing together the sub-disciplines of anthropology, artificial intelligence, and psychology in order to better understand and prepare for multi-cultural interaction. Making scientific progress across these disciplinary boundaries requires that we build bridges among the neural, cognitive, and social bands of human experience (Newell, 1990). Anderson and Gluck (2000) noted that the same challenge exists in connecting neuroscience and educational practice and proposed that cognitive architectures are an appropriate formalism for building such bridges. I propose that cognitive architectures also are an appropriate formalism for building bridges from neuroscience to the military's cognitive readiness applications, via cognitive phenomena and models.

The Solution: Cognitive Architectures

The purpose of all scientific disciplines is to identify invariant features and explanatory mechanisms for the purpose of understanding the phenomena of interest in the respective disciplines. Within the cognitive science community there is an approximately 50-year history of empirical research that involves using carefully constructed (usually simple and abstract) laboratory tests to isolate components of the human cognitive system in order to model and understand them. Sometimes optimistically referred to as "divide and conquer," this approach has led to comprehensive empirical documentation and sophisticated theories of hundreds of phenomena (e.g., fan effect, framing effect, Stroop effect) and functional components (e.g., attention, perception, memory, cognition, motor movement). A subset of the cognitive science community have become concerned that this divide and conquer approach is not leading to a unified understanding of human cognitive functioning, and have proposed cognitive architectures as the solution to that problem. Thus, cognitive architectures are intended to serve an integrative, cumulative role within the cognitive science community. They are where the fractionated theories come together in a unifying account not only of the computational functionality of the component processes, but also of the architectural control structures that define the relationships among those components, and of the representation of knowledge content that is used by cognition. Gray (2007) explains how these three theoretical spaces (components, control structures, and knowledge) interact and provides numerous case studies of each. Ultimately, it is at the intersection of these theories that cognitive architectures exist.

Ongoing Cognitive Modeling Research

Our cognitive modeling research program at the Air Force Research Laboratory's Mesa Research Site is organized around a set of methodological strategies with associated benefits. First, we are using and improving on the ACT-R (Anderson et al., 2004) cognitive architecture because it provides a priori theoretical constraints on the models we develop, facilitates model re-use among members of the ACT-R research community, and serves the integrating, unifying role described earlier. Second, we use the architecture, or equations and algorithms inspired by it, to make quantitative predictions in order to facilitate eventual transition to applications that make accurate, precise predictions about human performance and learning. Third, we develop models in both abstract, simplified laboratory tasks and in more realistic, complex synthetic task environments in order to begin constructing those bridges between the laboratory and the real world. Lastly, we compare the predictions of our models to human subjects data, in order to evaluate the necessity and sufficiency of the computational mechanisms and parameters that are driving those predictions and in order to evaluate the validity of the models. We are pursuing this research strategy in several lines of research, which I briefly describe next.

Knowledge Tracing. This is our only research line that is entirely mathematical modeling and does not involve a computational modeling component. The current approach is an extension and (we think) improvement to the General Performance Equation proposed by Anderson & Schunn (2000); thus, it derives from the computational implementation of learning and forgetting processes in ACT-R. The new equation allows us to make performance predictions or prescribe the timing and frequency of training, both of which will enable tailored training experiences at individual and team levels of analysis (Jastrzembski, Gluck, & Gunzelmann, 2006).

Communication. One of the barriers standing between us and human-level synthetic teammates is that we don't have a valid computational implementation of natural language – verbal or otherwise. This is critical because good teammates adapt their communications in order to facilitate accomplishing the shared mission. Our research in natural language modeling involves extending the Double R computational cognitive linguistic theory to knowledge-rich, time-pressured, team performance environments similar to those encountered in real-world situations, such as unmanned air vehicle reconnaissance missions (Ball, Heiberg, & Silber, 2007).

Spatial Competence. Spatial cognition has long been a sub-specialization within the cognitive science community, but typically individual scientists or research groups adopt particular phenomena to study without worrying about how the pieces of the spatial cognitive system come back together to create a more general competence. Reflecting this state of affairs, it turns out there is no comprehensive theory of the mechanisms and processes that allow for spatial competence. Our research in this area is pushing the field and the ACT-R architecture in the direction of a neurofunctional and architectural view of how spatial competence is realized in the brain and the mind (Gunzelmann & Lyon, 2006).

Fatigue. There is a rich history of sleep-related fatigue research conducted in and sponsored by the military laboratories. We are adding a new twist to that tradition by implementing new architectural mechanisms and processes that allow us to replicate the effects of sleepiness on the cognitive system. The process models are then combined with biomathematical models of the circadian and sleep homeostat systems to create the capacity to predict what the precise effects of sleep deprivation or long-term sleep restriction will be in a given performance context (Gunzelmann, Gluck, Kershner, Van Dongen, & Dinges, 2007).

High Performance and Volunteer Computing. As our cognitive modeling research expanded in breadth and depth and our scientific and technical objectives grew more ambitious we began to exceed the capacity of our local computing resources. In the search first for more resources and subsequently for more intelligent and efficient use of available resources, we have begun to use both high performance computing and volunteer computing as platforms for processor horsepower. We have demonstrated that such platforms can indeed be used productively for faster progress in cognitive modeling (Gluck, Scheutz, Gunzelmann, Harris, & Kershner, 2007) and are investing in additional software improvements for facilitating the use of these resources.

An Important Direction for the Research Community

Finally, I close by mentioning an important research direction for the cognitive modeling community: overcoming the knowledge engineering bottleneck. The key here is not the development of tools for doing manual knowledge engineering more efficiently, although that is a perfectly fine idea in the interim. Instead, I believe it is critical that we develop the ability for our modeling architectures to acquire their own knowledge without direct human assistance. This will require a variety of learning mechanisms, based on a combination of cognitive psychology, machine learning, and internet search algorithms.

References

Anderson, J. R., Bothell, D., Byrne, M. D., Douglass, S., Lebiere, C., & Qin, Y. (2004). An integrated theory of the mind. *Psychological Review 111*, 1036-1060.

Anderson, J. R., & Gluck, K. A. (2001). What role do cognitive architectures play in intelligent tutoring systems? In D. Klahr & S. M. Carver (Eds.), *Cognition and instruction: 25 years of progress* (pp. 227-261). Mahwah, NJ: Lawrence Erlbaum Associates.

Anderson, J. R., & Schunn, C. D. (2000). Implications of the ACT-R learning theory: No magic bullets. In R. Glaser (Ed.), *Advances in instructional psychology: Educational design and cognitive science*, *Vol. 5*. Mahwah, NJ: Erlbaum.

Ball, J., Heiberg, A. & Silber, R. (in press). Toward a large-scale model of language comprehension in ACT-R 6. Proceedings of the 8th International Conference on Cognitive Modeling.

Gray, W. D.(Ed.) (2007). *Integrated models of cognitive systems***.** New York: Oxford University Press.

Gluck, K. A., Scheutz, M., Gunzelmann, G., Harris, J., Kershner, J. (2007). Combinatorics meets processing power: Large-scale computational resources for BRIMS. In *Proceedings of the Sixteenth Conference on Behavior Representation in Modeling and Simulation* (pp. 73-83). Orlando, FL: Simulation Interoperability Standards Organization. (Selected for Recommended Reading List)

Gunzelmann, G., Gluck, K. A., Kershner, J., Van Dongen, H. P. A., & Dinges, D. F. (2007). Understanding decrements in knowledge access resulting from increased fatigue. In D. S. McNamara & J. G. Trafton (Eds.), *Proceedings of the 29th Annual Meeting of the Cognitive Science Society*. Austin, TX: Cognitive Science Society

Gunzelmann, G., & Lyon, D. R. (in press). Mechanisms of human spatial competence. In T. Barkowsky, C. Freksa, M. Knauff, B. Krieg-Bruckner, and B. Nebel (Eds.) Lecture Notes in Artificial Intelligence: Proceedings of Spatial Cognition (V) 2006. Berlin, Germany: Springer-Verlag.

Jastrzembski, T. S., Gluck, K. A., & Gunzelmann, G. (2006). Knowledge tracing and prediction of future trainee performance. In the *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference* (pp. 1498-1508). Orlando, FL: National Training Systems Association. (*Winner of the Best Paper Award*)

Newell, A. (1990). *Unified Theories of Cognition*. Harvard University Press, Cambridge, MA.