Lessons Learned: Structuring, Abstraction, Diagnosis and Cognitive Modeling

Authors:
Mark J. Stefik, PARC, Mark.Stefik@parc.com, Phone: 650-812-4366, Fax: 650-812-4334

Hoda Eldardiry, PARC, Hoda.Eldardiry@parc.com, Phone: 650-812-4790, Fax: 650-812-4334

Peter Pirolli, PARC, Peter.Pirolli@parc.com, Phone: 650-812-4483, Fax: 650-812-4959

Christian Lebiere, Psychology Department, Carnegie Mellon University, CL@.com, Phone: 412-268-5920

Jeffrey R. Cooper, Leidos, Technical Fellow, Chief Innovation Officer, Mission Support Group, Jeffrey.Cooper@leidos.com, Phone: 571-526-7318

Danny G. Bobrow, Research Fellow, PARC, Bobrow@parc.com, Phone: 650-812-4367, Fax: 650-812-4334

Christian Fritz, PARC, Christian.Fritz@parc.com, Phone: 650-812-4876, Fax: 650-812-4334

Dave Gunning, PARC, Dave.Gunning@parc.com, Phone: 650-812-4323, Fax: 650-812-4333

Executive Summary

The goals of knowledge management are typically to improve both the quality of the work of an organization and its efficiency. For knowledge management for lessons learned, we suggest an approach where lessons learned are concise, contextualized elements abstracted from stories of an analytic process. They are rich enough to convey the relevant context but contain no more than is needed. To facilitate re-use, the lessons are generalized to convey the broader contexts in which they might apply. Social processes including curation would help in tuning lessons and provide positive reinforcement. Cognitive analysis of human analysts would provide insights into heuristics for lesson use.

The interactions and dependencies of plan steps could be inferred from AI activity representations, with causal connections identified by plan re-use technology. A key problem is the appropriate generalization of lessons. Roughly following Tenenbaum [2011], we suggest that the combination of abstract knowledge structures for representing activities as used in planning and probabilistic networks in the world models could be effective in guiding generalizations. The AI technology of plan re-use (e.g. unification) for matching and binding planning variables could enable the refinement of lessons from an abstract to an instantiated state. This approach could

both enable re-use of lessons learned and potentially provide some of bias-countering advantages of structured analytic methods.

Pulling these ideas together we imagine an analyst at work on a case using a LLKM system. We imagine relevant lessons being identified automatically, so that an analyst can tend to them as convenient. Analysts would initiate lesson creation, and discuss them with team members or experts. There would be a sandbox for testing and validating lessons and a review process for checking them over time. A lesson would have both a human-friendly presentation and an underlying machine-friendly representation. The human-friendly form would use appropriate visual representations for spatial, temporal and relational data. The machine-readable form would be used in automatic processes for generalization, testing, and so on. The lessons would be accessible to analysts, curators, and a research team designing and conducting experiments.

Overview Briefing Chart

Lessons Learned: Structuring, Abstraction, Diagnosis, and Cognitive Modeling

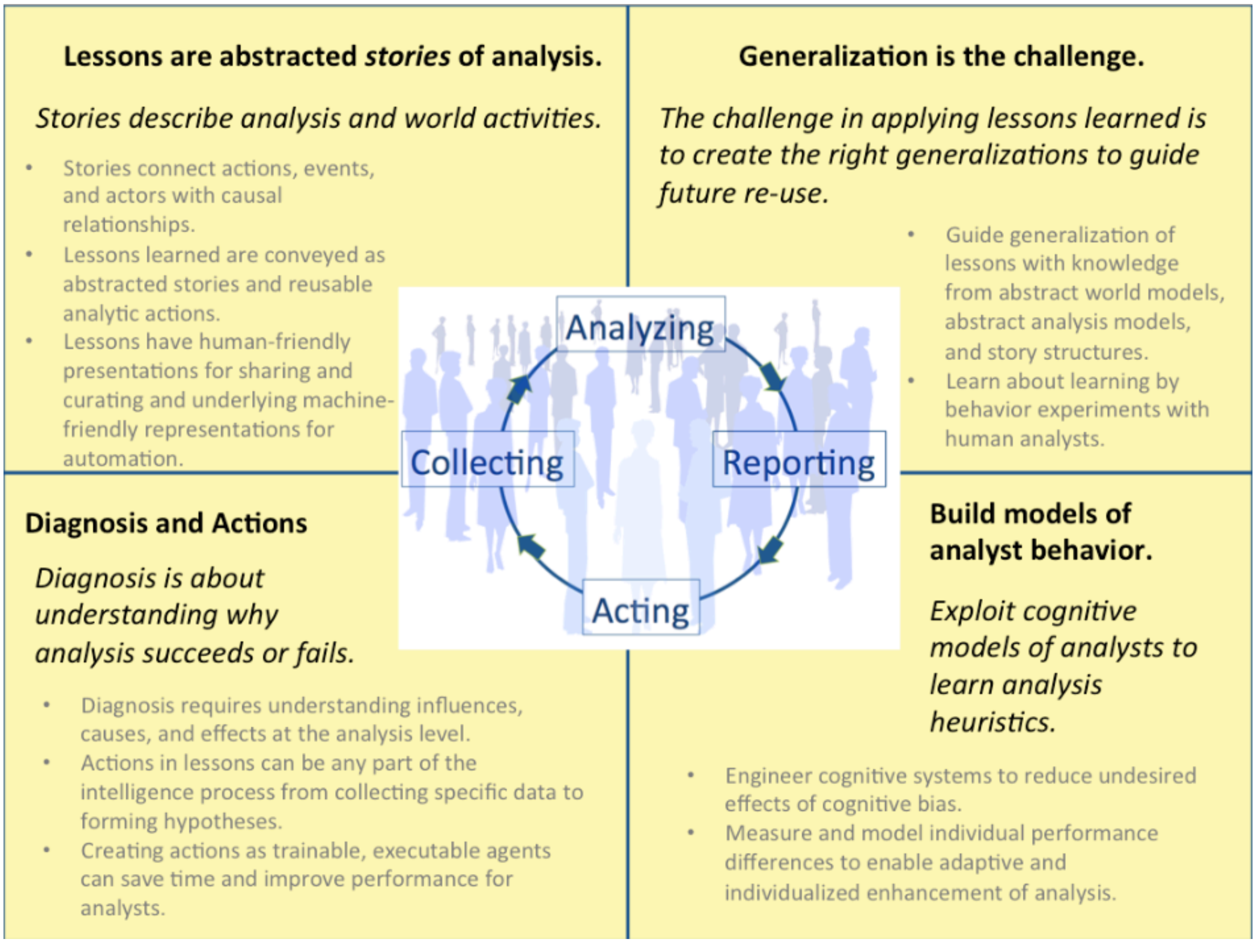

Diagnosis and Generality in Lessons Learned

The Eureka system [Bobrow & Whalen, 2002] was developed at Xerox to support the customer service engineers (CSE's) who repair the copiers and printers installed at customer sites. Eureka supported 20,000 CSE's world wide, sharing their knowledge about repair as lessons learned known as "tips". Tips were structured as stories -- organized by symptom, cause, test, and action -- and submitted by the CSE's to be curated for clarity and generality by expert field engineers. Tips covered actions that a CSE might need to do to solve a "problem" at a customer site, including tricky diagnoses, workarounds, ways of doing a job easier, and other bits of practical knowledge. A CSE would typically engage Eureka when he was stuck on a problem. He used a search engine to find relevant tips. Xerox saved money from improved efficiency in copier service because Eureka helped technicians to apply their collective experience.

Lessons as Stories

The practice of conveying *lessons learned as stories* is widespread in many settings, including intelligence analysis. The narrative structure of stories conveys information about what is important and what is unusual. A well-organized story explains how events unfold over time and space and arranges the narrative to convey an understanding of cause and effect. As a genre, intelligence stories have two levels: a world level and an analysis level. The world level describes happenings in the world. There are actors – leaders, military personnel, organizations, and so on. Actors take actions and have goals or intentions. Events happen and information about them becomes available over time. The analysis level holds a story about steps in intelligence analysis. Its actors are people in the intelligence community. Their actions include collecting data about world events, interpreting events, forming hypotheses, assessing possible actions and risks, imputing goals for actors, and other elements of the intelligence activity.

There is value in extracting lessons from an intelligence case when the analysis was difficult ("tricky diagnosis") and also when there was a notable or expensive intelligence failure. Consider the Israeli failed analysis for the 1973 surprise attack starting the Yom-Kippur War [Chorev, 1996], [Shlaim, 1976]. The figure on the right summarizes the main elements of the original analysis. The Israeli Defense Force (IDF) depended on early warning (48 hours) to mobilize its reserve forces. Israeli intelligence assumed that Egypt would not attack because it was incapable of military victory, that Egypt's military maneuvers were just practice again, and that Egypt was not ready politically. The political analysis was that Sadat continued to make threats – saying he would go to war in 1971 but did not. Israel believed that it had military superiority, that Sadat had expelled Soviet advisors, and that the Arab leadership was too divided to go to war. It was assumed that the Egyptian political goal was total defeat of Israel.

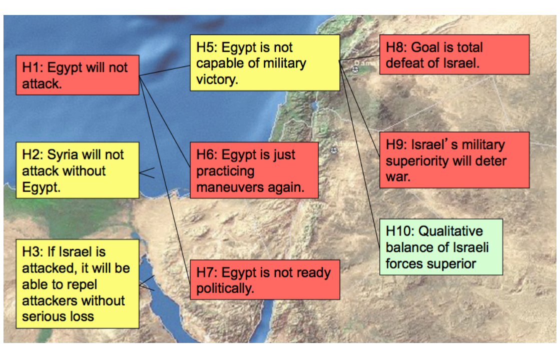

Red coloring in the figure on the left shows wrong hypotheses and yellow shows misleading ones. The challenge in learning intelligence lessons is to create analytic practices that improve performance for present and future situations. What lessons should be drawn from the case?

Diagnosis and Actions

In an intelligence failure, diagnosis is about understanding why analysis failed. This requires identifying the influences, causes, and effects at the analytic level. In the Yom Kippur case, there were multiple points of failures. The table on the right summarizes some of these points in terms of two belief sets representing (1) what the mainstream analysis concluded before the attack and (2) a post hoc analysis with revised beliefs after the attack. Judging that Israel's military superiority makes an Arab attack for defeating Israel foolish, Israeli intelligence failed to assess the possibility that Sadat would start a war of limited objectives.

	Assumption	BSet 1	BSet 2
P1	Total defeat of Israel	+	-
P2	Neutralize air force	+	-
P3	Get Sinae oil fields	0	+
P4	Limited territorial gain	-	+
P5	Sadat needs to counter internal political threats	0	+

What actions could have been taken? The tabular "surprise schema" on the left suggests actions that could be taken in the course of analysis to reduce the risk of a surprise attack. Each row of the table corresponds to a category of actions at the analysis level that would probe deeper into assumptions about enemy goals, capabilities, and readiness. The detailed actions would need to be adjusted for a situation.

Timing	Look for windows of vulnerability
Place	Look for weak places. Look for ways to soften potential targets.
Rapidity	Look at timing for response. Look at effects of slowed response.
New technologies	Check data on effectiveness of new technology.
New tactics	Look for leveraging effect of tactics. Acquire evidence of new tactics in military exercises.
New goals	Examine possible alternative goals.

Generality and Refinement

A key issue for getting value from lessons learned is to understand when to apply them. In the Eureka system, lessons are sought out and applied by CSE's when they get stuck. Intelligence analysis generally is more ongoing and less episodic. Incoming world information could signal the relevance of a lesson at any time and lessons might apply at any part of an analytic process suggested by the figure on the right. Furthermore, the description of world events is more open-ended than the world of copiers and printers. Generality is the central issue in deciding when and where lessons might apply.

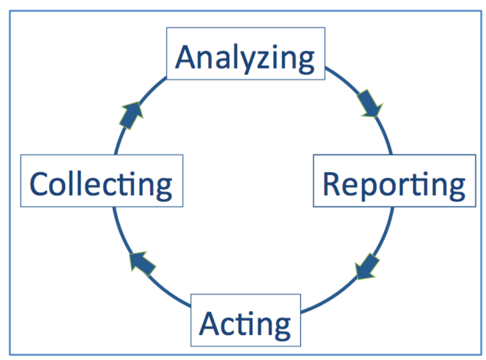

In what later situations would a lesson from the Yom Kippur case apply? In a too narrow characterization, the new situation would involve Israel, Egypt, and Syria where there is a build up of Egyptian forces to invade Israel. In a too general characterization, a lesson should be considered any time an enemy threatens an attack. The right points for lesson use fall somewhere between these extremes. The "surprise schema" table above can productively act as a set of suggestions for actions to be considered at relevant points of analysis, where the specific planning and detailed actions need to be tailored to the situation.

The representations and logic of plan monitoring and the research logic of challenges seem well-suited to adapting high-level abstract plans to appropriate actions in specific situations [Fritz, 2009], [Stefik, 1980]. These representations capture both the narrative structure and causality of stories and formal representations of abstraction. Addressing the challenges of machine learning, Tenenbaum [Tenenbaum, 2011] asks "How do human minds make inferences that go beyond the data available" and "How does the mind get so much from so little?" This is the problem of generalization. Some other source must provide the information. Tenenbaum suggests that the information comes from organized structures of abstract knowledge. At root this is the core problem of creating and retrieving lessons learned. AI representations of activities from planning research would provide a capable framework for exploring generalization and use of lessons learned in intelligence activities. Fritz's dissertation [Fritz, 2009] describes techniques for automatically identifying the rationale for a planning decision and variations on strategies that take into account uncertainties, and costs constraints.

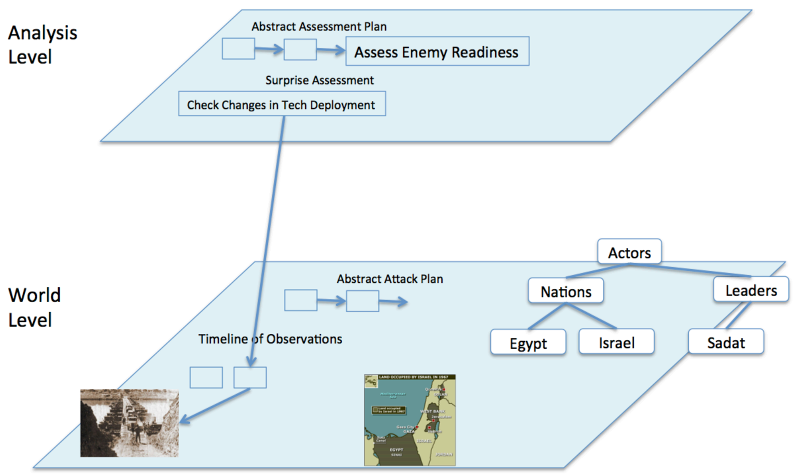

The figure on the left sketches the rough idea. At the world level, the model includes concept hierarchies, plans and abstract plans ("attack plans") with links to temporal and spatial data about unfolding events. The analysis level holds plans for analysis ("assessing enemy readiness"). The figure shows a link from an analytic step for assessing readiness to a world step to check an observation. Observations of the military "practice" exercises might have revealed technology changes that gave the Egyptians surprising advantages at the early stages of the war, specifically the use of RPG-7 rockets, RPG-43 grenades, and a novel use of high-pressure water canons to breach sand walls to undermine Israeli defenses using water from the Suez canal.

Neurocognitive Models to Improve LLKM

In the ongoing fog of uncertainty of intelligence situations, intelligence organizations rely on teams of analysts under time pressure to interpret events and to forecast likely outcomes. "They are hobbled by cognitive biases [Heuer, 1999], [Grabo, 2002], [Jervis, 2010] and exhibit analytic pathologies [Cooper, 2005]. Neurocognitive models have been developed to advance our understanding of the details of human thinking and performance involved in analytic sense making tasks [Lebiere et al., 2013]. These neurocognitive models combine symbolic processes (e.g., involving representations of the world, agents and their plans, etc.) and subsymbolic processes (e.g., associative learning; reinforcement learning; categorization; etc.). These models provide detailed explanations of the time-course of sense making processes, causal and counterfactual reasoning, individual variations in strategy and knowledge, as well as the sources of well-known heuristics and biases such as anchoring-and-adjustment heuristics, confirmation bias, probability matching, and representativeness heuristics. We propose that these recent advances in computational neurocognitive modeling can provide a foundation for novel approaches to improve the accuracy and efficacy of lesson-learned analyses, and improve the encoding and use of lessons learned knowledge. Cognitive models can be used to predict degree of surprise and unexpectedness of a case (beyond simply failure, which may have many reasons) and help guide which lessons would have the most impact on future analyses. They provide quantitative predictions as to how the structure of the analytic process and analyst knowledge impact performance in terms of speed and quality of decisions taken by both individuals and teams of analysts.

In general, these neurocognitive models simulate human performance and biases as an interaction of three interdependent factors: the structure of the analytic task+information environment, the mechanisms and limitations of cognitive processes, and the strategies on how to apply the latter to the former. The models predict how people will shape their performance and biases given their strategies+knowledge and the task+information environment. The models also predict how variations in the task+information environment (e.g., variations in knowledge management system designs) can improve performance and biases. Neurocognitive models of analytic sense making address how people induce mental models, recall and fit those mental models to data, adjust the mental models, make information foraging decisions, and update their strength of beliefs in hypotheses. The neurocognitive models explain how lack of feedback on ongoing analytic processes can lead to biases and thus provide insights and constraints on how lessons-learned knowledge might improve those biases. The models show how human working- and long-term memory, attentional focus, and pattern matching provide powerful mechanisms that perform nearly optimal in many circumstances, and the circumstances in which they fail. Neurocognitive models can be developed to predict how variations in strategy (e.g., Structured Analytic Techniques) and knowledge (expertise) can counteract these limitations and be tailored to individual differences in ability or background.

In other research, cognitive models have been used as a kind of automated agent to perform analyses. In the DARPA ACIP program [Wray et al., 2007], cognitive models were developed to automatically assemble filtered information streams into argument structures called Wigmorean trees. The cognitive models used a hybrid architecture that included aspects of massively parallel process for filtering and organizing information, recognizing patterns in data based on expertise and deliberate sequential (possibly counterfactual) reasoning processes.

Also relevant to lessons learned knowledge management system design are cognitive models of information seeking based on Information Foraging Theory [Pirolli, 2007]. These models are capable of predicting user search and navigation patterns as well as induce information needs from information seeking behavior. The models have been applied to predict and improve designs in a wide variety of domains including information retrieval, information visualization, software maintenance, and requirements analysis [e.g., Lawrance et al. 2010; Niu et al., 2013].

By analogy to previous successes in human-information interaction [Pirolli, 1999], we propose that these neurocognitive models provide a foundation for

(a) **Cognitive engineering of improved lessons-learned analytic methods and representations**, and improved KM user interaction techniques that maximize accurate and efficient knowledge gains from past cases by using a neurocognitive sense making model to predict the impact of design alternatives on user learning and performance.
(b) **User-modeling and context-modeling approaches that provide individualized adaptive support.** As for Intelligent Tutoring Systems and User Adapted Interaction systems, a model of the individual analyst can be used to select in real time the specific lessons to be learned that would have an optimal effect on the analyst's decision making.
(c) **Psychologically realistic artificial agents** capable of representing diverse strategies, knowledge, heuristics, and biases, and counterfactual reasoning. Diverse cognitive

models and counterfactual reasoning to generate a diverse set of learning material and provide opposing viewpoints in analysis.

Knowledge Management for Learning Intelligence Lessons

The goals of knowledge management are typically to improve both the quality of the work of an organization and its efficiency. We suggest an approach where lessons learned are concise, contextualized elements abstracted from stories from an analytic process. They are rich enough to convey the relevant context but contain no more than is needed. To facilitate re-use, the lessons are generalized to convey the broader contexts in which they might apply. Social processes including curation would help in tuning lessons and provide positive reinforcement. Cognitive analysis of human analysts would provide insights into heuristics for lesson use.

We also suggest various technologies and representations as candidates for the approach. The interactions and dependencies of plan steps could be inferred from the structure of plans, with causal connections identified by plan re-use technology [Fritz, 2009]. A key problem is the appropriate generalization of lessons. Roughly following Tenenbaum [2011], we suggest that the combination of abstract knowledge structures for representing activities as used in planning and probabilistic networks in the world models could be effective in guiding generalizations. Lessons as plan elements would be retrieved and refined for reuse in new situations. The AI technology of plan re-use (e.g. unification) for matching and binding planning variables could enable the refinement of lessons from an abstract to an instantiated state.

We believe that our approach could both enable re-use of lessons learned and potentially provide the bias-countering advantages of structured analytic methods. Because many of the published structured analytic techniques are manual, they are often seen by analysts under time pressure as extra work. We suggest that any approach for managing lessons learned must be experienced by its users as advantageous, preferably reducing the cost structure of sense making [Russell, 1993] while reducing biases and otherwise improving analysis. Our vision for the future of intelligence is that the future of knowledge management for lessons learned can achieve this goal when it takes on more of the concepts and approaches discussed above, enabled by both KM research and modern machine learning tools.

Pulling these ideas together we imagine an analyst at work on a case using a LLKM system. In contrast with Eureka where lessons are presented only when analysts search for them, we imagine relevant lessons being identified automatically, so that an analyst can tend to them as convenient. Analysts would initiate lesson creation, and discuss them with team members or experts. There would be a sandbox for testing and validating lessons and a review process for checking them over time. A lesson would have both human-friendly presentation and an underlying machine-friendly representation. The human-friendly form would use appropriate visual representations for spatial, temporal and relational data. The machine-readable form would be used in automatic processes for generalization, testing, and so on. The lessons would be accessible to analysts, curators, and a research team designing and conducting experiments

References[1]

Beebe, Sarah Miller, Pherson, Randolph H. *Cases in Intelligence Analysis. Structured Analytic Techniques in Action*. Los Angeles: CQ Press. 2012.

Bobrow, Daniel G. & Whalen, Jack. Community knowledge sharing in practice: The Eureka story. Reflections 4(2), p 47-59, The MIT Press.2002.

Chorev, Moni. *Surprise Attack: The Case of the Yom-Kippur War*. Washington D.C. National Defense University. 1996.

Cooper, Jeffrey R. *The CIA's Program for Improving Intelligence Analysis – Curing Analytic Pathologies*. Center for the Study of Intelligence. 2005.

Fishbein, Warren & Treverton, Gregory, *Making Sense of Transnational Threats*, Kent Center Occasional Papers, The Sherman Kent School for intelligence Analysis, CIA, October 2004.

Fishbein, Warren & Treverton, Rethinking "Alternative Analysis" to Address Transnational Threats, Kent Center Occasional Papers, The Sherman Kent School for intelligence Analysis, CIA, October 2004.

Fritz, Christian. Monitoring the Generation and Execution of Optimal Plans. Ph.D. Dissertation, University of Toronto, 2009.

Fritz, Christian & McIlwraith, Sheila A. Generating Optimal Plans in Highly Dynamic Domains. In *Proceedings of The 25th Conference on Uncertainty in Artificial Intelligence (UAI), Montreal, Canada, June 18-21*, 2009.

George, Roger Z., Bruce, James B. (Eds.) *Analyzing Intelligence: Origins, Obstacles, and Innovations.* Washington, D.C. Georgetown University Press. 2008.

Grabo, Cynthia M. *Anticipating Surprise: Analysis for Strategic Warning*. Joint Military Colleges Center for Intelligence Research. 2002

Jervis, Robert. Reports, Politics, and Intelligence Failures: The Case of Iraq. *The Journal of Strategic Studies*. Vol. 29, No. 1, 3-52, February 2006.

Jervis, Robert. *Why Intelligence Fails*. Ithaca. Cornell University Press. 2010.

Heuer Jr., Richards J. *Psychology of Intelligence Analysis*. Center for the Study of Intelligence, Central Intelligence Agency. 1999.

Heuer Jr., Richards J. & Pherson, Randolph H. *Structured Analytic Techniques for Intelligence Analysis*. Washington D.C., CQ Press. 2011.

[1] Boldface indicates when a paper author is also an author of this white paper.

Lawrance, J., Burnett, M., Bellamy, R., Bogart, C., & Swart, C. (2010). *Reactive information foraging for evolving goals*. Paper presented at the Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, Atlanta, Georgia, USA.

Lebiere, Christian, Pirolli, Peter, Thomson, Robert, Paik, Jaehyon, Rutledge-Taylor, Matthew, Staszewski, James, & Anderson, John R. *A functional model of sensemaking in a neurocognitive architecture. Computational Intelligence and Neuroscience.2013. http://dx.doi.org/10.1155/2013/921695*

Niu, N., Mahmoud, A., Chen, Z., & Bradshaw, G. (2013). *Departures from optimality: understanding human analyst's information foraging in assisted requirements tracing*. Paper presented at the Proceedings of the 2013 International Conference on Software Engineering, San Francisco, CA, USA.

Pirolli, Peter. (1999). Cognitive engineering models and cognitive architectures in human-computer interaction. In F. T. Durso, R. S. Nickerson, R. W. Schvaneveldt, S. T. Dumais, D. S. Lindsay & M. T. H. Chi (Eds.), *Handbook of applied cognition* (pp. 441-477). West Sussex, England: John Wiley & Sons.

Pirolli, Peter. (2007). *Information foraging: A theory of adaptive interaction with information*. New York: Oxford University Press.

Russell, Daniel M., **Stefik, Mark, Pirolli, Peter**, Card, Stuart K. The Cost Structure of Sensemaking. Proceedings of InterCHI, 1993.

Shlaim, Avi, Failures in National Intelligence Estimates: The Case of the Yom Kippur War, *World Politics*, Vol. 28 (April 1976).

Simon, Herbert A. "How Big Is a Chunk?" *Science*, Vol. 183 (1974), p. 487.

Stefik, Mark & Good, Lance. Design and Deployment of a Personalized News Service. *AI Magazine*. Vol. 33, No. 2. 2012.

Stefik, Mark. Planning with Constraints. (MOLGEN: Part 1). *Artificial Intelligence*. Vol 14, No 2, 111-140, 1980.

Stefik, Mark. Planning and Meta-Planning. (MOLGEN: Part 2). *Artificial Intelligence*. Vol 14, No 2, 141-169, 1980.

Tenenbaum, Joshua B., Kemp, Charles, Griffiths, Thomas L. Goodman , Noah D. How to Grow a Mind: Statistics, Structure, and Abstraction. *Science,* Vol. 331 (March 2011).

Wray, Robert, **Lebiere, Christian**, Weinstein, Peter, Jha, Krishna, Springer, Jonathan, Belding, Theodore, Best, Bradley, & Parunak, Van. Towards a complete, multi-level cognitive architecture. In *Proceedings of the 8th International Conference on Cognitive Modeling.* Ann Arbor, MI (July 2007).