In Memoriam

Master of Complexity: Paul J. Feltovich

2 August 1947 – 21 January 2025

(Illustrated Version)

Robert Hoffman, Jeffrey Bradshaw, William Clancey, David Woods, Gary Klein,
Peter Hancock, Mike Prietula, Rand Spiro

Abstract
This article describes the main contributions made by the late Paul J. Feltovich to the fields of cognitive engineering and decision making.

Keywords
cognitive engineering topics, complexity topics, expert performance topics

Corresponding Author
Robert Hoffman, rhoffman@ihmc.org

Citation
Hoffman, R., Bradshaw, J., Clancey, W., Woods, D., Klein, G., Hancock, P., Prietula, M., & Spiro, R. (2025). In Memoriam Master of Complexity: Paul J. Feltovich August 2, 1947–January 21, 2025. *Journal of Cognitive Engineering and Decision Making*, *19*(3), 243–249. https://doi.org/10.1177/15553434251331788

Declaration of Conflicting Interests
The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding
The authors received no financial support for the research, authorship, and/or publication of this article.

In Memoriam

Master of Complexity: Paul J. Feltovich

2 August 1947 – 21 January 2025

Whether AI, autonomy or any other new technology "gold rush," Paul's findings—about the growth of expertise, misconceptions, and the urge to oversimplify—are real gold, a conceptual treasure chest. Once you understand that treasure, you possess a guide to navigate and outmaneuver the ever-changing complexities of this world. The phrase "master of complexity" comes to our collective minds in describing Paul Feltovich. He and his colleagues identified timeless principles that should be understood by everyone in the cognitive/learning sciences and all those who model sociotechnical and complex systems. These general findings are critical to overcome misconceptions common to technologists and managers, misconceptions that degrade the design of human-technology systems and safety management in health care, aerospace and other sectors.

Paul transcended academic boundaries. His explorations and interests evolved from a solid mathematical background to educational psychology and cognitive psychology, all influencing his seminal insights on intelligent systems. His application domains similarly evolved from psychological counseling to instructional design, to medical decision making and expertise studies. and then to collaborative cognition and cognitive systems engineering. He was a researcher and instructor at the Center for Research on Human Learning at the University of Minnesota, the Learning Research and Development Center of the University of Pittsburgh, The Southern Illinois University School of Medicine, and lastly at the Florida Institute for Human and Machine Cognition. He was named to numerous Visiting Professorships. His vita lists numerous invited presentations, technical reports, and publications in refereed professional journals. He received an extensive number of grants and awards. Paul served on multiple Editorial Boards including *The Educational Researcher,* the *Journal of the Learning Sciences,* and *Advances in Health Sciences Education.* Paul advised the NSF, the American Educational Research Association, and the Cognitive Science Society. He served as Reviewer for dozens of professional journals. He served as Editor of the journal *Teaching and Learning in Medicine*.

Amongst his contributions are some of the most widely cited papers in instructional design, cognitive psychology, expertise studies, and cognitive systems engineering. For example, his paper with Michelene Chi and Robert Glaser on problem-solving in physics (Chi, Feltovich, and Glaser, 1981) was designated a *Science Citation Classic* and marked the emergence of expertise studies. Paul was co-editor of *Expertise in Context: Human and Machine,* and co-editor of the first edition of the *Cambridge Handbook on Expertise and Expert Performance*. He contributed as a co-author to a multitude of essays on Human-Centered Computing published in *IEEE Intelligent Systems*. He was a co-author of *Accelerated Expertise: Training for High Proficiency in a Complex World*. A listing of the Paul's highly regarded and influential papers is presented at end of this essay.

Paul's text *Smart Machines in Education* (with Ken Forbus) was a landmark in intelligent tutoring systems. Here, he focused on systems for biomedicine, medical problem-solving, diagnostic reasoning, radiology, and clinical practice. In several papers on "Cognitive Flexibility Theory," Paul elucidated the causes of difficulty for learners who are trying to understand complex systems, especially those involving the trap of oversimplification. He also identified some of the key mechanisms that learners engage in their defense of misconceptions, which he called knowledge shields. That work also identified corrective strategies for mastering important conceptual complexities and assembling adaptive responses to novel real-world problems.

Based on his studies of what makes it difficult for people to understand complex dynamic systems in general (such as the circulatory system), Paul described the "dimensions of difficulty." These dimensions remain central to cognitive systems engineering since they pertain equally well to the reasoning of those who design cognitive work systems, specifically the tendency to formulate models that reduce, rather than embrace complexity in a cognitively manageable way. These dimensions are described in Table 1. These dimensions are manifest in cognitive systems engineering in the substitution myth (Christofferson and Woods, 2002), the attribution of blame, and the reductive tendency. By some accounts, scientific theories are held accountable by "Ockham's razor," which is often said to be a mandate that scientific theories be simple. Although this is not quite what William of Ockham meant, there is a tendency to fall into the trap of simplification, which is characterized by the first member in each of the contrasted dimensions.

Table 1. Application of the "dimensions of learning difficulty" to the challenges in the design of complex cognitive work systems.[1]

Are phenomena modeled as being static, or dynamic?
Are important aspects of a phenomenon captured by a fixed "snapshot," or are the critical characteristics captured only by the *changes* from frame to frame? Are phenomena modeled as being static and scalar, despite possessing dynamic characteristics?
Are factors thought of as being separable or interactive?
Separability regards processes and people as functioning in isolation/insulation, resulting in a lack of recognition of the degree of non-additive spread of effect to be expected from an intervention. Changing one component is thought to have only local effects and will not ripple throughout the operation, as in fact they will because of high interdependency. Widespread adjustments and workarounds will not take place because there will be no perceived need for them. If one substitution does not improve operation, another substitution may be tried, perhaps through many such cycles, but interactive effects will not be addressed.
Are factors or events modeled as being sequential or simultaneous?
This involves thinking of the work as a set of steps where effecting one component just has implications for the processes "down the line." This does not consider that numerous processes are going on simultaneously, for example imposing complex coordination problems, situational awareness problems, communication problems, complex and often hard to predict interactions—and so forth.
Are components modeled as being homogeneous or heterogeneous?
Homogeneity means thinking of the processes, values, ways of doing things, cultural norms, abilities, loyalties etc. as being largely of the same kind across the many diverse units of thought, collaboration or workplace.
Is a work system modeled in a mechanistic manner, or an organic manner?
Are effects traceable to simple and direct causal agents, or are they the product of more system-wide, organic functions? Can good work system design be achieved by understanding just parts of the system, or must the entire system be understood for even the parts to be understood well? Modeling cognitive work in a mechanistic set of low-level causes and effects would take no consideration of kinds of emergent, highly complex, nonlinear, interactive characteristics of the workplace as a complex system that has qualities and processes that are more than the sum of its component parts.
Are principles held to be universal or conditional?
Do principles hold without the need for substantial modification across different situations, or is there great context-sensitivity? This ignores change-over in people, individual differences in the characteristics and abilities among people, constantly changing character and demands on work, infusions of new tools, and the like. The general theme is not recognizing how much the effects of an intervention will vary (with greater or lesser desired effect) with the changing particulars of the constitution of the workplace.

Is the domain characterized in terms of regularity or irregularity?
Is a domain characterized by a high degree of regularity across cases, or do cases differ considerably from each other even when they are called by the same name? Are there strong elements of symmetry and repeatable patterns in concepts and phenomena, or is there a prevalence of asymmetry and absence of consistent pattern? The regularity stance does not take into account the highly variable nature of people, work practices, workplace norms and regulations, changing demands and pace, and so on.
Is the domain characterized in terms of linearity or nonlinearity?
Are functional relationships linear or nonlinear (i.e., are relationships between input and output variables proportional or nonproportional)? Can a single line of explanation convey a concept or account for a phenomenon, or are multiple overlapping lines of explanation required for adequate coverage? The linearity assumption is that changes, effects of interventions, perturbations of various kinds to the system will have incremental, manageable kinds of consequences. This does not account for the possibility of massive cascades and nonlinear discontinuities in a complex system.
Are domain concepts or phenomena given single or multiple interpretations?
Do elements in a situation afford single (or just a few) interpretations, functional uses, categorizations, and so on, or do they afford many? Are multiple representations required (e.g., multiple schemas, analogies, case precedents, etc.)? Are multiple perspectives required ? This reductive tendency would entail not realizing that appraisals of the success and the usability of the system can vary from lots of different perspectives, both with regard to the operation of the workplace itself (e.g., profitability, degree of stress, efficiency, quality) and from the point of view of different stakeholders (the bosses, the workers, the unions, the shareholders, etc.)—and probably many other perspectives too—e.g., the community.
Are elements or phenomena delineated as apparent on the surface, or are they described in a deeper or more abstract manner?
Superficial understanding regards the overt consequences from an intervention, the things one can easily see and measure, such as people are doing, what is getting made, how fast, and at what cost, etc. However, this approach could ignore possibly powerful forces, conditions, and consequences that are not readily apparent, such as morale, commitment, loyalty, health, satisfaction, carefulness, and workplace culture.
Are elements or phenomena delineated as discrete or continuous?
Are processes described as proceeding in discernable steps, or are they regarded as unbreakable continua? Are attributes describable by a small number of categories (e.g., dichotomous classifications like large/small), or is it necessary to recognize and utilize entire continuous dimensions (e.g., the full dimension of size) or large numbers of categorical distinctions?

The dimensions of difficulty enable people to deal with complexity, but this comes at the cost of over-simplification, misconception, and reification. Paul referred to this as the "reductive tendency."[2] The contrasts in Table 1 are fundamental to understand complex adaptive systems in the human sphere where deploying new capabilities generates complexity, but people become stuck in the surface simplifications and mis-architect systems.

Paul's most recent work was on collaborative learning. With colleagues at the Florida Institute for Human and Machine Cognition, Paul was investigating coordination, regulation, and teamwork in mixed groups of humans and intelligent software agents. The systems they developed were based on tenets for what machines can, might and must do, and were adapted from formal pertinent law, group traditions, standards of practice, and norms for acceptable everyday activity (e.g., various codes of etiquette). This laid conceptual foundations for how software agents could become acceptable to humans in complex cognitive work. It was extended to an analysis of what it would mean for a computer or robot to be "cooperative" and enable an interdependence relationship between the human and the machine. This work remains at the forefront of current trends in "Agentic AI."

To many for whom Feltovich is just a respected name in the literature, his loss would be a remote and intellectual sadness. We, his colleagues, suffer a further blow of personal diminution. Paul brought a gentle and inspiring curiosity and a unique perspective to problems. His work ethic was remarkable as evidenced by his achievements, but it was just as important to him to bring integrity and kindness to his collaborative work. Paul did not hesitate to offer criticisms but did it in a gentle way that advanced the relationship. Paul was highly respected and loved by those who knew and worked with him. He was a wonderful friend and an inspiration. It is not only an intellectual light that has been dimmed, but the presence of a gentleman and gentle man now gone. The word "wisdom" often comes to mind when people are asked about him. As so does the word, "friend."

Footnotes

1. This table is adapted from some of Paul's unpublished writings.

2. Paul initially called this the "reductive bias." The terminology was subsequently changed: "We do not see this as a bias in the sense in which cognitive science frequently uses the term. Rather, this reductive tendency is an inevitable consequence of how people learn" (Feltovich, Hoffman, Woods, & Roessler, 2004, p. 90).

References

Bradshaw, J.M., Hoffman, R.R., Johnson, M., & Woods, D.D. (2013). The seven deadly myths of "autonomous systems." *IEEE Intelligent Systems* (May/June), 54–61.

Chi, M.T.H., Feltovich, P.J., & Glaser, R. (1981). Categorization and representation of physics problems by experts and novices. *Cognitive Science 5*(2), 121–152.

Christofferson, J., & Woods, D.D. (2002). How to Make Automated Systems Team Players. In E. Salas (Ed.) *Advances in Human Performance and Cognitive Engineering Research,* Vol. 2 (pp. 1–12). Amsterdam: Elsevier Science.

Feltovich, P.J., Hoffman, R.R., Woods, D.D. & Roesler, A. (2004). Keeping it too simple: How the reductive tendency affects cognitive engineering. *IEEE Intelligent Systems* (May/June), 90–95.

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

Selected Papers by Paul Feltovich

Bradshaw, J.M., Feltovich, P.J., & Johnson, M. (2017). Human–agent interaction. In G. Boy (Ed.), *The handbook of human–machine interaction* (pp. 283–300). London: CRC Press.

Ericsson, K.A., Charness, N.B., Feltovich, P.J. & Hoffman, R.R. (Eds.) (2006). *Cambridge handbook on expertise and expert performance*. Cambridge: Cambridge University Press.

Feltovich, P.J. Bradshaw, J.M., Clancey, W.J., & Johnson, M. (2007). Toward an ontology of regulation: Socially–based support for coordination in human and machine joint activity. In G.M.P. O'Hare, A. Ricci, M.J. O'Grady, O. Dikenelli (Eds.), *Engineering Societies in the Agents World VII, 7th International Workshop* (pp. 175–192). Dublin, Ireland, September 2006. Lecture Notes in Computer Science Series, Volume 4457. Berlin: Springer.

Feltovich, P.J., Bradshaw, J.M., Clancey, W.J., Johnson, M., & Bunch, L. (2008). Progress appraisal as a challenging element in human and machine joint activity. In A. Artikis, G. O'Hare, K. Stathis, & G. Vouros (Eds.), *Engineering Societies in the Agents World VIII* (pp. 124–141). Athens, Greece, October 2007. Lecture Notes in Computer Science Series, Volume 4995. Berlin: Springer.

Feltovich, P.J., Coulson, R. L., & Spiro, R. (2001). Learners' (Mis)understandings of important and difficult concepts. In K.F. Forbus and P.J. Feltovich (Eds.), *Smart machines in education* (pp. 349–375). Menlo Park, CA: MIT Press.

Feltovich, P.J., Coulson, R. L., Spiro, R., & Adami, J.F. (1994). "Conceptual understanding and stability, and knowledge shields for fending off conceptual change." Technical Report No. 7, Southern Illinois University School of Medicine, Springfield, IL. [https://files.eric.ed.gov/fulltext/ED391807.pdf]

Feltovich, P.J., Ford, K.M. & Hoffman, R.R. (Eds.), *Expertise in context* (pp. 543–580). Cambridge, MA: MIT Press.

Feltovich P.J., Johnson P.E., Moller J.H., & Swanson D.B. (1984). LCS: The role and development of medical knowledge in diagnostic expertise. In W.J. Clancey & E.H. Shortliffe (Eds). *Readings in medical artificial intelligence: The first decade* (pp. 275–319). Reading, Mass.: Addison–Wesley.

Feltovich, P.J., Prietula, M.J., & Ericsson, K.A. (2006). Studies of expertise from psychological perspectives. In K.A. Ericsson, N. Charness, P.J. Feltovich & R.R. Hoffman (Eds.), *Cambridge handbook of expertise and expert performance* (pp. 41–61). Cambridge: Cambridge University Press.

Feltovich, P.J., Spiro, R., & Coulson, R.L. (1997). Issues of expert flexibility in contexts characterized by complexity and change. In P.J. Feltovich, K.M. Ford & R.R. Hoffman (Eds.) *Expertise in context* (pp. 125–146). Menlo Park, CA: MIT Press.

Feltovich, P.J., Spiro, R., Coulson, R.L. & Feltovich. J.(2012). Collaboration within and among minds: Mastering complexity, individually and in groups. In T. Koschmann (Ed.) *Computer Supported cooperative learning: Theory and practice of an emerging paradigm* (pp. 25 – 47). Hillsdale, NJ: Lawrence Erlbaum.

Fiore, S.M., Hancock, P.A., Janelle, C.M., Szalma, J.L., Eccles, D.W., Tenenbaum, G., Feltovich, P., & Prietula, M. (2004). From the perceptual to the organizational: How can the science of expertise better inform the practice of human performance. In: D.A. Vincenzi., M. Mouloua., and P.A. Hancock, (Eds.). *Human performance, situation awareness and automation: Current research and trends,* Vol. 2 (pp. 83–91). Mahwah, NJ.: Erlbaum.

Forbus, K.F. & Feltovich , P.J. (Eds.), *Smart machines in education*. Menlo Park, CA: MIT Press.

Hoffman, R.R., Feltovich, P.J., Fiore, S.M., Klein, G., and Ziebell, D. (2009, March/April). Accelerated learning (?). *IEEE Intelligent Systems 24*(2), 18–22.

Hoffman, R.R., Feltovich, P.J., & Ford, K.M. (1997). A general framework for conceiving of expertise and expert systems in context. In P.J. Feltovich, K.M. Ford & R.R. Hoffman (Eds.), *Expertise in context* (pp. 543–580). Cambridge, MA: MIT Press.

Hoffman, R.R., Feltovich, P.J., Ford, K. M., Woods, D.D., Klein, G. & Feltovich, A. (July/August 2002). A rose by any other name … would probably be given an acronym. *IEEE Intelligent Systems*, 72–80.

Johnson, M., Bradshaw, J.M., Feltovich, P.J., Jonker, C.M., Van Riemsdijk, B., & Sierhuis, M. (2014). Coactive design: Designing for interdependence in joint activity. *Journal of Human–Robot Interaction 3*, 43–69.

Johnson, P.E., Duran, A.S., Hassebrock, F. Moller, J. Prietula, M., Feltovich, P.J. & David B Swanson, D.B. (1981). Expertise and error in diagnostic reasoning. *Cognitive Science 5*(3), 235–283.

Klein, G. Feltovich, P.J., Bradshaw, J.M., & Woods, D.D. (2005). Common ground and coordination in joint activity. In W.B. Rouse & Kenneth K.R. Boff (Eds.) *Organizational simulation* (pp. 139–184). Hoboken, NJ: Wiley

Lesgold, A.M., Feltovich, P.J., Glaser, R., & Wang, Y. (1981). The acquisition of perceptual diagnostic skill in radiology. Technical Report No. PDS–1. Pittsburgh: University of Pittsburgh, Learning Research and Development Center.

Lesgold, A., Glaser, R., Rubinson, H., Feltovich, P.J., Klopfer, D. & Wang, Y. (1988). Expertise in a complex skill: Diagnosing X–ray pictures. In A. Lesgold, R. Glaser, H. Rubinson, P. Feltovich, D. Klopfer, & Y. Wang (Eds.) *The nature of expertis*e (pp. 311–342). Psychology Press.

Spiro, R. J, Coulson, R.L, Feltovich, P. J, & Anderson, D. K. (1988). Cognitive flexibility theory: Advanced knowledge acquisition in ill–structured domains. In *Proceedings of the 10th Annual Meeting of the Cognitive Science Society* (pp. 375–383). Mahwah, NJ: Lawrence Erlbaum.

Spiro, R.J., Feltovich, P.J., & Coulson, R.L. (1996).Two epistemic world-views: Prefigurative schemas and learning in complex domains. *Applied Cognitive Psychology 10*(7), 51–61.

Spiro, R.J., Feltovich, P.J., P. J., Gaunt, A., Hu, Y., Klautke, H., Cheng, C., Clemente, I., Leahy, S., & Ward, P. (2018). Cognitive flexibility theory and the accelerated development of adaptive readiness and adaptive response to novelty. In P. Ward, J. M., Schraagen, J. Gore, & E. Roth (Eds.), *The Oxford handbook of expertise* (pp. 951–976). Oxford: Oxford University Press.

Author Biographies

Robert R. Hoffman is Emeritus Senior Research Scientist at the Florida Institute for Human and Machine Cognition. His current interest is in the calculation of practical significance and the establishment of a general theory of macrocognitive work systems. He has authored numerous articles in the area of cognitive systems engineering, and has authored and edited numerous books including *Accelerated Expertise, Minding the Weather,* and *The Cambridge Handbook of Expertise and Expert Performance*. He holds a PhD in Experimental Psychology from the University of Cincinnati. (https://www.ihmc.us/groups/rhoffman/)

Jeffrey M. Bradshaw is an Emeritus Senior Research Scientist at the Florida Institute for Human and Machine Cognition. His current research interests include multi-agent systems, human-agent-robot teamwork, management of policy and ethical constraints in agents and cloud-based services, semantic technologies, cyber sensemaking, self-driving cars, and visualizations that exploit principles of human perception and cognition. He holds a Ph.D. in Cognitive Science from the University of Washington. (www.ihmc/groups/jbradshaw; www.jeffreymbradshaw.net).

David D. Woods is Professor Emeritus in Integrated Systems Engineering at the Ohio State University. He is Past-President of the Resilience Engineering Association and of the Human Factors and Ergonomics Society. He is a pioneer in Resilience Engineering—how to support people as they adapt to cope with complexity. His books include *Behind Human Error*, *Resilience Engineering*, and *Joint Cognitive Systems*. He started the SNAFU Catchers industry-university consortium to apply the new science to build resilience in critical digital services. He holds a Ph.D. in cognitive psychology from Purdue University. (https://www.researchgate.net/profile/David-Woods-19/research)

William J. Clancey is Emeritus Senior Research Scientist at the Florida Institute for Human and Machine Cognition. A cognitive scientist focusing on the methods and theory of AI programming, he has developed human-centered applications in medicine, education, finance, robotics, and spaceflight systems. He promotes a methodology that combines philosophical analysis, ethnography, and simulation models of work practices. He received a Ph.D. in Computer Science from Stanford University and is a Fellow of the American College of Medical Informatics, Association for Advancement of AI, Association for Psychological Science, American Association for the Advancement of Science, and the National Academy of Inventors. (https://billclancey.name/index.html)

Gary Klein is a cognitive psychologist who helped to found the Naturalistic Decision Making movement in 1989. His Recognition-Primed Decision (RPD) model has been tested and replicated several times. He also developed a Data/Frame model of sensemaking, and a Triple Path model of insight. His work relies on advanced cognitive interviewing methods that he and his team developed. He formulated the Pre-Mortem method for identifying risks and the ShadowBox method for training cognitive skills. He also led a team that developed the AIQ (Artificial Intelligence Quotient) toolkit for increasing the usability of AI/ML systems. He is the owner of MacroCognition LLC and ShadowBox LLC. (https://www.gary-klein.com/welcome)

Peter A Hancock is Provost Distinguished Research Professor at the Department of Psychology and the Institute for Simulation and Training at the University of Central Florida. His research concerns the interaction between humans and differing expressions of advanced technology. Most especially he has focused upon the dimension of trust and what factors and influences mediate a person's trust in technical systems that they use. He is the author of over eleven hundred scientific contributions including the texts; *Mind, Machine, and Morality*; *Transports of Delight*; and *Hoax Springs Eternal*. He holds a D.Sc. in human-machine systems and a Ph.D. in human motor performance with an emphasis on time perception from the University of Illinois. (http://peterhancock.ucf.edu)

Rand J. Spiro is Professor of Educational Psychology and Technology at Michigan State University. He has also been a professor at Yale, the University of Illinois at Urbana, and Harvard, in various departments at different times (psychology, computer science, cognitive science, AI, education.) He has authored numerous books and articles. His current work involves the application of Cognitive Flexibility Theory to professional learning and adaptive practice with the support of AI as a smart assistant/collaborator. His Ph.D. is in experimental cognitive psychology from Pennsylvania State University.
(https://postgutenberg.typepad.com/newgutenbergrevolution/)

Michael J. Prietula is a Professor at the Goizueta Business School at Emory University and at the Rollins School of Public Health. He holds a courtesy appointment as an External Research Scholar at the Florida Institute for Human and Machine Cognition. His research focuses on human and machine cognition, human-AI collaboration and trust, and agent-based systems. His current research involves mHealth AI design, health intervention AI models, and AI ethics. He holds a Ph.D. in Information Systems from the University of Minnesota and a Masters in Public Health from the University of Florida. (https://www.ihmc.us/groups/mprietula/)

Photographs of Paul Feltovich

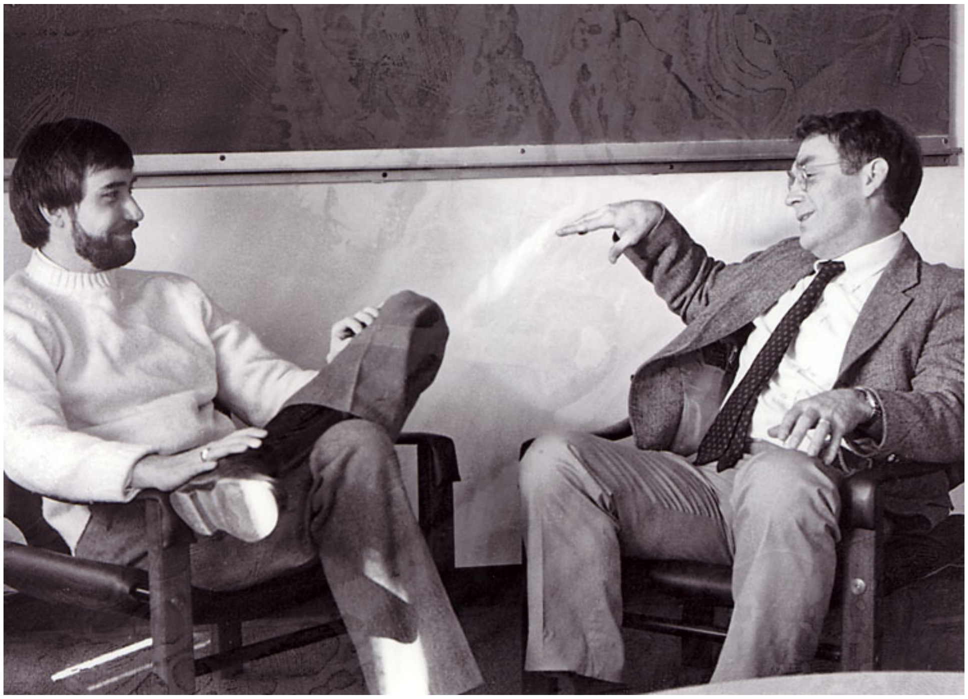

William J. Clancey and Paul J. Feltovich, October 1983, *ONR Nature of Expertise Conference*, Learning Research Development Center, Pittsburgh PA

Paul J. Feltovich, October 2007, *Engineering Societies in the Agents World Conference*, Athens, Greece

Paul J. Feltovich, February 2009, Institute for Human and Machine Cognition, Pensacola FL; office whiteboard with his notes

Paul J. Feltovich, May 2005, Seaside FL