Toward a Theory of Complex and Cognitive Systems

Robert R. Hoffman, *Institute for Human and Machine Cognition*
David D. Woods, *Ohio State University*

Essays in this department have presented nine propositions that we've referred to as *principles* of human-centered computing:

- *The Aretha Franklin Principle*: Do not devalue the human in order to justify the machine. Do not criticize the machine in order to rationalize the human. Advocate the human-machine system in order to amplify both.[1]
- *The Sacagawea Principle*: Human-centered computational tools need to support active organization of information, active search for information, active exploration of information, reflection on the meaning of information, and evaluation and choice among action sequence alternatives.[2]
- *The Lewis and Clark Principle*: The human user of the guidance needs to be shown the guidance in a way that is organized in terms of his or her major goals. Information needed for each particular goal should be shown in a meaningful form, and should allow the human to directly comprehend the major decisions associated with each goal.[2]
- *The Envisioned World Principle*: The introduction of new technology, including appropriately human-centered technology, will bring about changes in environmental constraints (that is, features of the sociotechnical system, or the context of practice). Even though the domain constraints might remain unchanged, and even if cognitive constraints are leveraged and amplified, changes to the environmental constraints might be negative.[3]
- *The Fort Knox Principle*: The knowledge and skills of proficient workers is gold. It must be elicited and preserved, but the gold must not simply be stored and safeguarded. It must be disseminated and utilized within the organization when needed.[4]
- *The Pleasure Principle*: Good tools provide a feeling of direct engagement. They simultaneously provide a feeling of flow and challenge.[5]
- *The Janus Principle*: Human-centered systems do not force a separation between learning and performance. They integrate them.[6]
- *The Mirror-Mirror Principle*: Every participant in a complex cognitive system will form a model of the other participant agents as well as a model of the controlled process and its environment.[7]
- *The Moving Target Principle*: The sociotechnical workplace is constantly changing, and constant change in environmental constraints might entail constant change in cognitive constraints, even if domain constraints remain constant.[3]

The term "principle" doesn't actually do much work in science. Colloquially, it's used as a tacit reference to laws, as in "This device works according to the principle of gravity." What are these so-called principles? Our answer leads to additional considerations involving the use of the principles.

Cute mnemonics?

Are the principles we've proposed simply aids to help people remember some tips from people who've grappled with issues at the intersection of humans, technology, and work? Indeed, we've deliberately given the principles names that have both mnemonic and semantic value, even though we could have given them more technical designations. And yes, they are "tips."

But they're more than that.

Cautionary tales?

Are the principles merely signposts at the fork between
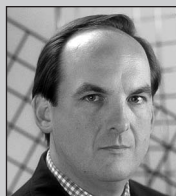
Editors: Robert R. Hoffman, Patrick J. Hayes, and Kenneth M. Ford
Institute for Human and Machine Cognition, University of West Florida
rhoffman@ihmc.us

paths to user-hostile and user-friendly systems? We and Paul Feltovich discussed the reductive tendency, which is a necessary consequence of learning: At any given time, any person's knowledge of a domain is bound to be incomplete and to some extent simplifying.[8] We pointed out that this tendency also applies to those who are creating new information technologies, especially Complex and Cognitive Systems (people working in teams, using information technology to conduct cognitive work to reach certain goals). Indeed, the people who try to create new Complex and Cognitive Systems are themselves prone to generate reductive understandings, in which complexities are simplified:

> The reductive tendency would be the assumption that a design principle has the same applicability and effects throughout the many different and changing contexts of work and practice. That is, the effects, embodied in the design principle, will hold fairly universally across differing practice environments.[8]

So, the principles are indeed important cautionary tales.

But they're more than that.

Guidelines?

Are the principles recipes that we can use to design human-centered or otherwise "good" information technologies? Take the Sacagawea Principle, for example. Can we go from that principle to a good design for a specific application? Hardly. The principles, as we've stated them, aren't entries in a cookbook that measure goodness in quarts or bytes or hours in the oven. They're not a substitute for empirical inquiry (for example, cognitive task analysis), design creativity, or proper software development. In discussing the application of human-centered computing notions to the design of intelligent systems, Axel Roesler, Brian Moon, and Robert Hoffman stated, "The principles of human-centered computing which have been discussed in essays in this Department are not entries for a cookbook; they are not axioms for design."[9]

Rather than being formulas, the principles imply design challenges. Looking back on the essays, we find one challenge expressed directly. An implication of the Fort Knox Principle is what we called the Tough Nut Problem[4]: How can we redesign jobs and processes, including workstations, computational aids, and interfaces, in such a way as to get knowledge elicitation as a "freebie" and at the same time make the usual tasks easier?

"Project managers or designers may choose to adopt [the principles] if their goal is to create good, complex cognitive systems."[9] The principles do serve as constraints on design.

But they're more than that.

Empirical generalizations?

The principles we've mentioned in this department's essays by no means exhaust the set we've generated. Consider another, for example, the Principle of Stretched Systems: Complex and Cognitive Systems are always stretched to their limits of performance and adaptability. Interventions (including innovations) will always increase the tempo and intensity of activity.

> It's certainly possible to create information technologies that don't support comprehension and navigation, don't integrate learning and performance, or fail to induce a feeling of joyful engagement.

Every system is stretched to operate at its capacity. As soon as some improvement, some new technology, exists, practitioners (and leaders, managers, and so on) will exploit it by achieving a new intensity and tempo of activity.

An example that should resonate with most readers goes as follows. "Gee, if I only had a robust voice recognition system, I could cope with all my email much better." We've heard this plea many times. But stop to consider what would really happen. Others would use the technology too, so the pace of correspondence would accelerate, and people would wind up right back where they were—in a state of mental overload. Systems always get stretched. This has happened whenever new information technologies have been introduced into, and changed, the workplace.[10]

The principles aren't just cautionary tales or design constraints; they're empirically grounded generalizations that have stood the test of time. If we tap into the literature on the philosophy of science,[11,12] we'd say that the principles are

- *Generalizations.* Referring to classes of things, not to individual things
- *Extensional generalizations.* Based on empirical or descriptive evidence

But they're more than that.

Scientific laws?

As we've stated them, the principles are what philosophers call *nomological generalizations*. That is, they're universal for the realm of discourse or for some specified boundary conditions. This criterion is important for physical law: It's literally impossible for matter to travel faster than the speed of light, for example. But it's certainly possible to create information technologies that don't support comprehension and navigation (Sacagawea, Lewis and Clark Principles), that don't integrate learning and performance (Janus Principle), or that fail to induce a feeling of joyful engagement (Pleasure Principle).

But it's impossible to create "good" human-centered systems that violate the principles. Thus, "goodness" sets a strong boundary condition and will prove, we think, to be a critical concept for cognitive engineering.

As we've stated them, the principles are what philosophers of science call *open generalizations*. That is, the evidence that's been used to induce the principles doesn't coincide with the range of application. If the evidence that is available were all the evidence there is, the science would stop. Kenneth Craik described this feature of scientific laws in the following way:

> Now all scientific prediction consists in discovering in the data of the distant past and of the immediate past (which we incorrectly call the present), laws or formulae which apply also to the future, so that if we act in accordance with those laws, our behavior will be appropriate to the future when it becomes the present.[13]

For all new applications and forms of Complex and Cognitive Systems, the principles should apply in the way Craik describes. So, what we've been calling principles are extensional, nomological generalizations whose fate is to be deter-

mined empirically. In other words, they're scientific laws.

But laws of what? This department is about making computational devices such as VCRs[14] human centered. But most of the essays have focused on technologies used in sociotechnical contexts. Complex and Cognitive Systems are systems in which multiple human and machine agents collaborate to conduct cognitive work. Cognitive work is goal-directed activity that depends on knowledge, reasoning, perceiving, and communicating. Cognitive work involves macrocognitive functions including knowledge sharing, sense making, and collaboration. Furthermore, Complex and Cognitive Systems are distributed, in that cognitive work always occurs in the context of multiple parties and interests as moments of private cognition punctuate flows of interaction and coordination. Thus, cognitive work is not private but fundamentally social and interactive.[10]

The principles—we should now say laws—are not just about HCC as a viewpoint or paradigm or community of practice; they're about Complex and Cognitive Systems in general. We do not refer to *complex cognitive systems* because that sort of adjective string would involve an ambiguity. Are they complex systems? Is it the cognition that's complex? Complex and Cognitive Systems, as we intend, uses the word "and" to express a necessary conjunction. The designation would embrace notions from "cognition in the wild"[15] and from distributed systems.[16]

So, we have a domain of discourse or subject matter. But a science needs more than that.

Steps toward a theory?

Salted throughout the essays have been statements implying that the principles hang together. An earlier essay on the Pleasure Principle stated that both the Sacagawea and the Lewis and Clark Principles

are suggestive of a state in which the practitioner is directly perceiving meanings and ongoing events, experiencing the problem they are working or the process they are controlling. The challenge is to live in and work on the problem, not to have to always fiddle with machines to achieve understanding.[5]

This suggests that the principles resonate with one another. The interplay of the principles becomes meaningful.

The Envisioned World Principle and the Moving Target Principle have a strong entailment relation. New technologies are hypotheses about how work will change, yet the context of work is itself always changing. The Envisioned World Principle involves changes to cognitive constraints (for example, task requirements) brought about by changes in environmental constraints (that is, new technologies). The Moving Target Principle asserts that cognitive constraints are also dynamic (for example, new methods of weather forecasting and new knowledge of weather dynamics involve changes in forecaster understanding). Thus, all three sources of constraint can be in constant flux whenever new technologies are introduced into the workplace.

Another criterion philosophers of sci-

> The principles—we should now say laws—are not just about HCC as a viewpoint or paradigm or community of practice; they're about Complex and Cognitive Systems in general.

ence hold out for postulates to be scientific laws is that laws must have entailment relations. They must hang together in necessary, interesting, and useful ways. Thus, what we seem to have been reaching for is a theory.

But what is it for?

Why a theory?

We see two primary motivations for a theory of Complex and Cognitive Systems. The first lurks in previous essays in this department such as the discussion of kludges and work-arounds[17] and Kim Vicente's discussion of VCRs.[14] All sorts of smart, well-intentioned people are out there building new intelligent technologies, and have been doing so for years. The notion of user-friendliness has been around for over two decades. Yet, we're all confronted daily with technologies that are

not only not user-friendly but also downright user-hostile. We're even tempted to assert this as another principle (law): *The road to user-hostile systems is paved with user-centered intentions.*[18]

Confronted with the problems that new technologies cause (apart from the problems they might solve) and new challenges entail, sponsors of systems development efforts have come to cognitive engineers crying for guidance in designing information technologies. Sponsors yearn for systems that will solve difficult problems in knowledge acquisition, collaboration, and so on, including such problems as how to enable a single person to control multiple robots or how to help weather forecasters build rich, principled visualizations of their mental models of atmospheric dynamics.

It would hardly do for cognitive engineers to reply with cute mnemonics, or cautionary tales, or cookbooks, or a disassociated collection of empirical generalizations. Cognitive engineers must present a coherent, empirically grounded, empirically testable scientific theory.

But aren't there already theories out there? Systems theory? A theory from cognitive psychology? The second motivation for a theory of Complex and Cognitive Systems is that the phenomena that occur in sociotechnical contexts are emergent and involve processes not adequately captured in either cognitive science or systems science. Explaining Complex and Cognitive Systems, and understanding their behaviors, will require more than the available perspectives and theories. Indeed, this is part of the motivation for the distinction between macrocognition and microcognition.[19]

Cognitive theory might tell us about the millisecond-to-millisecond process of attentional shift, but it doesn't say much about situational awareness. It might tell us about the processes of sentence comprehension, but it doesn't say much about sensemaking in real-world, dynamic situations.

Systems notions and notions of complexity are indeed critical for any understanding of Complex and Cognitive Systems. For instance, the Triples Rule (the unit of analysis of the human-machine-context triple)[1] and the Aretha Franklin Principle both involve systems concepts. But while systems theory can tell us about interactions and feedback, it doesn't say much about human collaboration or distributed cognition.

How do we extend the theory?

You might wonder whether the set of laws constituting a particular theory is complete. This is a high standard employed in logical or axiomatic theories, to which we might not be subject because the Theory of Complex and Cognitive Systems isn't a theory of logic or mathematics. But we might wonder, especially in light of the reductive tendency, to simply assert that the theory of Complex and Cognitive Systems is incomplete. The laws (formerly, principles) that we have mentioned in this department are certainly not all that there are—we know of some two dozen more that haven't yet been essay topics. But beyond this fuller list of laws, we assert that incompleteness is in fact a feature of the theory.

To the mathematically inclined, we might then be free to assert that the laws constituting the theory of Complex and Cognitive Systems are consistent. Rather than doing so, however, we assert that the theory's inconsistency is indeterminate. This affords one path to testability in the form of "forced inconsistency." If a Complex and Cognitive System is designed and initiated in accordance with any subset of the laws, doing so shouldn't force a violation of any other law. If that happens, the theory might need fixing.

Completeness, consistency, and testability are just three of the outstanding issues involved in creating a theory of Complex and Cognitive Systems. Obviously, more work is needed.[20] Numerous subtleties and nuances must be sorted out, involving operational definitions of key concepts and other paths to testability.

This essay aims to encourage people in the intelligent systems and cognitive engineering communities to reflect on the scientific foundations of HCC.

How do we forge a scientific foundation? Once forged, how do we use it? How do we extend, refine, and empirically test it? We invite you to correspond with this department's editors concerning more candidate laws and the challenges for a theory of Complex and Cognitive Systems. ◻

References

1. R.R. Hoffman et al., "The Triples Rule," *IEEE Intelligent Systems*, May/June 2002, pp. 62–65.

2. M. Endsley and R.R. Hoffman, "The Sacagawea Principle," *IEEE Intelligent Systems*, Nov./Dec. 2002, pp. 80–85.

3. S.W.A. Dekker, J.M. Nyce, and R.R. Hoffman, "From Contextual Inquiry to Designable Futures: What Do We Need to Get There?" *IEEE Intelligent Systems*, Mar./Apr. 2003, pp. 74–77.

4. R.R. Hoffman and L.F. Hanes, "The Boiled Frog Problem," *IEEE Intelligent Systems*, July/Aug. 2003, pp. 68–71.

5. R.R. Hoffman and P.J. Hayes, "The Pleasure Principle," *IEEE Intelligent Systems*, Jan./Feb. 2004, pp. 86–89.

6. R.R. Hoffman, G. Lintern, and S. Eitelman, "The Janus Principle," *IEEE Intelligent Systems*, Mar./Apr. 2004, pp. 78–80.

7. G. Klein et al., "Ten Challenges for Making Automation a 'Team Player' in Joint Human-Agent Activity," *IEEE Intelligent Systems*, Nov./Dec. 2004, pp. 91–95.

8. P.J. Feltovich, R.R. Hoffman, and D. Woods, "Keeping It Too Simple: How the Reductive Tendency Affects Cognitive Engineering," *IEEE Intelligent Systems*, May/June 2004, pp. 90–95.

9. R.R. Hoffman, A. Roesler, and B.M. Moon, "What Is Design in the Context of Human-Centered Computing?" *IEEE Intelligent Systems*, July/Aug. 2004, pp. 89–95.

10. D.D. Woods and R.I. Cook, "Nine Steps to Move Forward from Error," *Cognition, Technology, and Work*, vol. 4, 2002, pp. 137–144.

11. A. Kaplan, *The Conduct of Inquiry*, Chandler, 1964.

12. W. Weimer, *Notes on Methodology of Scientific Research*, Lawrence Erlbaum, 1979.

13. K.J.W. Craik, "Theory of the Operator in Control Systems: I. The Operator as an Engineering System," *British J. Psychology*, vol. 38, 1947, pp. 56–61.

14. K.J. Vicente, "Crazy Clocks: Counterintuitive Consequences of 'Intelligent' Automation," *IEEE Intelligent Systems*, Nov./Dec. 2001, pp. 74–76.

15. E. Hutchins, *Cognition in the Wild*, MIT Press, 1995.

16. G. Coulouris, J. Dollimore, and T. Kindberg, *Distributed Systems: Concepts and Design*, 3rd ed., Addison-Wesley, 2001.

17. P. Koopman and R.R. Hoffman, "Work-Arounds, Make-Work, and Kludges," *IEEE Intelligent Systems*, Nov./Dec. 2003, pp. 70–75.

18. R.R. Hoffman, G. Klein, and K.R. Laughery, "The State of Cognitive Systems Engineering," *IEEE Intelligent Systems*, Jan./Feb. 2002, pp. 73–75.

19. G. Klein et al., "Macrocognition," *IEEE Intelligent Systems*, May/June 2003, pp. 81–85.

20. D.D. Woods and R. Hoffman, *A Theory of Complex Cognitive Systems*, tech. report, Inst. Human and Machine Cognition, 2005.

Robert R. Hoffman is a senior research scientist at the Institute for Human and Machine Cognition. Contact him at IHMC, 40 So. Alcaniz St., Pensacola, FL 32502-6008; rhoffman@ihmc.us.

David D. Woods is a professor of industrial and systems engineering and the coordinator of the Cognitive Systems Engineering Laboratory at Ohio State University. Contact him at the Cognitive Systems Eng. Lab, 210 Baker Systems, Ohio State Univ., 1971 Neil Ave., Columbus, OH 43210; woods@csel.eng.ohio-state.edu.