Hoffman, R., Hancock, P.A., Ford, K., & Hayes, P. (2002) The triples rule. IEEE Intelligent Systems, 17 (3), 62-65

The Triples Rule

Robert R. Hoffman, Patrick Hayes, and Kenneth M. Ford, *University of West Florida*
Peter Hancock, *University of Central Florida*

A fundamental stance taken in human-centered computing is that information-processing devices must be thought of in systems terms. At first blush, this seems self-evident. However, the notion has a long history, and not just in systems engineering.

The Industrial Fatigue Research Board

During World War I, the British Health and Munitions Workers' Committee presented the results of wartime studies of fatigue and efficiency, which motivated the establishment of the Industrial Fatigue Research Board. The IFRB's mandate—marking the beginnings of ergonomics in England—was to identify the causes of worker fatigue and boredom (due to mechanization, task automation, and the need for concentrated attention) and determine ways to alleviate it.[1]

Under the IFRB's aegis and the leadership of pioneer industrial psychologists, studies were conducted of diverse jobs, covering worker selection, product design, production procedures, delivery methods, environmental effects on productivity and safety, and so on.[2] In the pages of the Board's research reports—there were 90 in all, published up through 1947—we find the first use of the phrase "the human factor."[3] We also find the first explicit notion of a man–machine system. In his 1923 "personal contribution" to the IFRB's annual report, H.G. Weston foresaw a need for a program of psychological research:

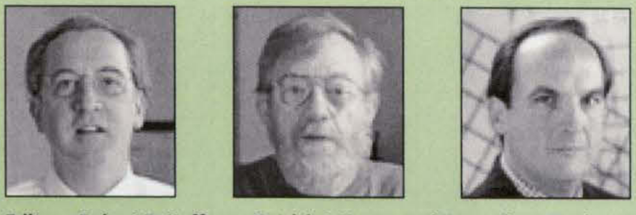

Editors: Robert R. Hoffman, Patrick J. Hayes, and Kenneth M. Ford
Institute for Human and Machine Cognition, University of West Florida
rhoffman@ai.uwf.edu

> The introduction and development of power-driven machines has effected an enormous savings of time and energy, not only by increasing the rapidity of production through substitution of mechanical power for human effort, but also by changing the character of the manipulations which remain to be performed by the operative. So great has this economy been, that is has brought with it a tendency to overlook the possibility that, while industrial machinery may be admirably adapted to the performance of its mechanical functions, it may be incompletely adapted to the needs of the human organism, upon whose efficient co-operation it depends for its productive use.[4]

Weston goes on to mention research detailing the discovery of serious design flaws in various machines that led to inefficiency and physical debilitation. For example, even commonly used lever shapes could result in serious physical debilitation if they were in an improper position. Sometimes, attempts to correct one obvious problem led to another one that the design engineers did not recognize. For example, the lightening of a roller in a laundry machine was not accompanied by a lightening of the load needed to depress the foot pedal controlling the roller, making the roller mechanism more, rather than less, difficult to use.

> It is, therefore, most important that correct design should be secured in the first place. It is difficult to see how this can be obtained to the fullest extent except as the result of definite research, undertaken with the object of determining such physiological and psychological facts as should be borne in mind when designing machines, the forms of mechanism and mechanical combinations which will conform to the needs of the operative.[4]

We could therefore say that the basic concept of the man–machine system was already entrenched in industrial psychology as early as the 1930s. To say now that we must think of humans and machines in systems terms seems perhaps more than passé. However, technology has wrought many changes since the 1930s, entailing entirely new implications for the meaning of human–machine system.

Cognitive fit: The dawn of a new age

Computers do things with symbols that reflect the meanings of those symbols, forcing us to rethink the relationship between machines and humans—the cognitive fit as well as the physical fit and how the two relate. For example, eyeglasses are an ocular prosthesis that improves vision. They must be fitted to the individual's head shape and vision; my eyeglasses won't work well for most other people. However, the use of the ocular prosthesis depends on context: If I am wearing my reading glasses and suddenly look up to examine something far away, the eyeglasses become a hindrance rather than a help. Extending this analogy to computers, we must be aware of contexts in which the human and machine collaborate and operate. This broadens the notion of context and makes it subtler; it now includes such things as goals and expectations, which are just as relevant as physical context. Also, the context can include other people in a way it hadn't before because the machine can facilitate new kinds of communication, and other peoples' concepts and ideas can be part of the context of use.

SRI International scientists in a mobile-agents research project are developing eyeglasses that do not magnify in the usual sense (see www.erg.sri.com/projects/sae/cars-poster.jpg). Rather, they can calculate where you're looking and project a virtual image onto the real scene—for example, to show you where to turn to get to the nearest convenience store. When such devices become commonplace, we will regard them as an extension of our vision and will wonder how we ever got along without them. To not have them—a failure to be able to see in new ways—will be regarded as a form of shortsightedness. However, to make such devices commonplace, we must totally reconceive what it means for a machine to have an interface.

Modern cognitive engineering has progressed through four "ages" over its 30-plus-year lifetime. During the Age of the Average Man, design was based on "one size fits all." During the Age of Adaptation, there arose a recognition of the need for flexibility. During the Age of Personalization, we realized that machine siblings could be adapted to single users. We are now entering a fourth age.

The Triples rule

In this new Age of Symbiosis, machines are made for specific humans for use in specific contexts. The unit of analysis for cognitive engineering and computer science is a triple: person, machine, and context (see Figure 1).

The Triples rule asserts that system development must take this triple as the unit of analysis, which has strong implications, including a mandate that the engineering of complex systems include detailed *cognitive work analysis*.[5–7] It also has implications for the meaning of intelligence, including artificial intelligence.

Intelligent…in a sense

It is easy, perhaps too easy, to regard humans as intelligent, but we must keep in mind that the attribution of intelligence to a human is heavily context dependent. Sitting nearly naked on the ground and poking at ants with a stick would be regarded as odd in many contexts, but not when you're in the Kalahari Desert, looking for ant trails that might lead to a source of water. The world-class expert in any particular domain will not necessarily top the charts on a standard test of general intelligence.[8] Likewise, a computational device when taken out of its intended context of application can be useless—an extreme case being when it is dropped into a swimming pool. As a physical artifact, the device will still have affordances—it could plug a leak—but it would not interact with people on the basis of its computational capabilities. Or, try entering a chess machine in a natural language processing competition. Reductio ad absurdum examples such as these perhaps make the point less well than realistic examples. A geographic information system accepts data types x, y, and z, but the user needs (context of application) to integrate data types q, r, and s, which can't be done. Or, Box A cannot communicate with Box B. These sorts of problems are daily fare in complex sociotechnical workplaces plagued by mandated and legacy systems.

To continue with the triple analysis, replace the human who knows how to work with the computational device with a human who does not, and again the device becomes essentially useless. An extreme example would be a chess machine used by someone who knows nothing about games. The machine retains some of its affordances—the human could push buttons just to see what happens or could use the machine as a doorstop. To continue the triple, change the machine (its computational and interface capabilities), and you get entirely different patterns of interactions with the human and the context.

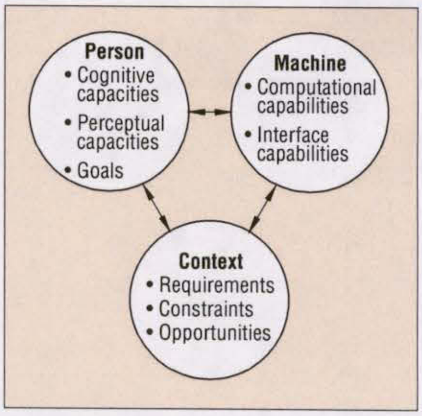

Figure 1. The triple.

Changes to the human or the context can render the machine useless or make the context inappropriate; the human is rendered less capable by changes to the context (for example, a PC user who finds himself surrounded by Macintosh users) or the machine. According to the Triples rule, intelligence emerges from the triple interaction. The rule thus runs head-on into the traditional definition of intelligence in AI: the Turing test.

Turing redux

Alan Turing's paper "Computing Machinery and Intelligence"[9] gave AI a vision and its first great challenge.[10] As Turing described it, the test is an imitation game that involves a man, a woman, and a judge, all communicating but unable to see one another. The judge's task is to decide which of the other two is the woman; the others try to persuade the judge that one of them is the woman and that the other is the man. Turing is usually understood to mean that the game should be played with the question of gender (being female) replaced by the question of species (being human). The judge is faced with the task of differentiating a human participant from a machine pretending to be human.

Borrowing from Patrick Hayes and Kenneth Ford,[11,12] we point out just three of the problems with the Turing test. First, from the standpoint of experimental design, it confounds the machine's intelligence with that of the judge. If a machine passes the test, is it demonstrated to be intelligent, or was the judge not intelligent enough to ask sufficiently telling questions? The game conditions say nothing about the judge, but the game's success depends crucially on how intelligent the judge is.

Second, the Turing test confounds intelligence with cleverness. To pass the test, a machine would have to not only give a human-like impression but also be an expert on making a good impression. It would have to avoid exhibiting any inhuman talents that it might have; it would always have to lie, cheat, and dissemble. The winner of the Loebner competition, for example, sometimes deliberately mistyped a word and then backspaced to correct it at human typing speed. This strategy is clever, but such tricks should not be central to AI.

A third difficulty with the Turing test is that the definition of intelligence keeps shifting. As AI progresses and machines increasingly perform tasks previously considered to involve human intelligence, those abilities are no longer taken to be definitive. When Eliza first appeared, some people found its conversational abilities quite human-like. No machine until then could have reacted even in a simple way to what had been said to it. But during the Loebner competition, many programs were instantly revealed as nonhuman precisely by the first hint of their behavior's resemblance to Eliza's. The ability to perform simultaneous translation could soon be reduced to the merely mechanical. Turing tests have become circular: They define the qualities for which they are claiming to be evidence.

We could argue that the Turing test should not be regarded as the defining goal for AI but as a spur to technological progress—constantly pushing us to reach for Rene Descartes' dream of the "enlightened machine."[13] But why should we take it as our only goal to build something that is just like us? A dog would never win any imitation game, but there seems to be no doubt that dogs exhibit cognition, and a machine with the cognitive and communicative abilities of a dog would be an interesting challenge for AI.[14] More importantly, our most useful computer applications (including AI programs) are often valuable by virtue of their lack of humanity. There are cameras, copiers, televisions, automobiles, battery rechargers, and laptop operating systems all incorporating algorithms that use AI ideas and techniques but are not usually advertised as "intelligent" or "expert."[12]

If we abandon the Turing test as the defining goal for AI, the goal can shift from making artificial superhumans that can replace us to making artifacts that we can use to amplify and extend our cognitive abilities. AI should play a central role in this exciting new technology of cognitive prosthetics,[15] but to do so it must turn its back on the Turing test. What, then, should be the test?

The Triples rule suggests that the problem here is with the question itself—there is no single test. In fact, there are boundless numbers of tests. In general, for any task x or y that a human or machine can do (each under the appropriate set of contextual constraints), the human–machine–context

Hoffman, R., Hancock, P.A., Ford, K., & Hayes, P. (2002) The triples rule. IEEE Intelligent Systems, 17 (3), 62-65

triple is "intelligent" if the task can be conducted better than if either the human or the machine were to conduct the task without engaging in a partnership with the other.[16] The word "better" is deliberately left open in this formulation. It might be instantiated by "more efficiently," "more rapidly," or "more economically," but it might just as well mean "more playfully." This is how it should be if the machine is human-centered. There is an appropriate fit between the human, the machine, and the context such that the human's cognitive, perceptual, and collaborative capacities are enhanced. AI thus becomes Amplified Intelligence.

This line of thinking helps put to rest the doomsaying about intelligent machines taking over the world, which is another story.[14] This is just the first of what we hope will become a series of essays that discuss candidate principles for this thing called human-centered computing. The Triples rule is one of the first such principles. We invite others to suggest additional ones. ■

References

1. S. Wyatt, J.A. Fraser, and F.G.L. Stock, *The Effects of Monotony in Work*, Report no. 56, Industrial Fatigue Research Board, His Majesty's Stationery Office, London, 1929.
2. M.S. Viteles, *Industrial Psychology*, Norton, New York, 1932.
3. E.M. Newbold, *A Contribution to the Study of the Human Factor in the Causation of Accidents*, Report no. 34, Industrial Fatigue Research Board, His Majesty's Stationery Office, London, 1926.
4. H.G. Weston, "A Note on Machine Design in Relation to the Operative," *3rd Ann. Report Industrial Fatigue Research Board, 31 December 1922*, Industrial Fatigue Research Board, His Majesty's Stationery Office, London, 1923, pp. 71–75.
5. D.D. Woods, "Coping with Complexity: The Psychology of Human Behavior in Complex Systems," *Mental Models, Tasks, and Errors*, L.P. Goodstein, H.B. Anderson, and S.E. Olsen, eds., Taylor and Francis, London, 1988, pp. 128–148.
6. D.D. Woods and D. Tinapple, "Presidential Address: Watching Human Factors People at Work," *Proc. 43rd Ann. Meeting Human Factors and Ergonomics Soc.*, Human Factors and Ergonomics Soc., 1998; http://csel.eng.ohio-state.edu/hf99.
7. K. Vicente, *Cognitive Work Analysis*, Lawrence Erlbaum, Mahwah, N.J., 2000.
8. R.R. Hoffman, "How Can Expertise Be Defined? Implications of Research from Cognitive Psychology," *Exploring Expertise*, R. Williams, W. Faulkner, and J. Fleck, eds., Macmillan, New York, 1998, pp. 81–100.
9. A.M. Turing, "Computing Machinery and Intelligence," *Mind*, vol. 59, no. 236,1950, pp. 433–460.
10. A. Hodges, *Alan Turing: The Enigma*, Simon and Schuster, New York, 1983.
11. P. Hayes and K.M. Ford, "Turing Test Considered Harmful," *Proc. 14th Int'l Joint Conf. Artificial Intelligence*, Morgan Kaufman, San Francisco, 1995, pp. 972–977.
12. K.M. Ford and P.J. Hayes, "On Computational Wings: Rethinking the Goals of Artificial Intelligence, *Scientific Am. Presents*, vol. 9, no. 4, Winter 1998, pp. 78–83.
13. R.R. Hoffman, P. Hayes, and K.M. Ford, "Human-Centered Computing: Thinking In and Out of the Box," *IEEE Intelligent Systems*, vol. 16, no. 5, Sept./Oct. 2001, pp. 76–79.
14. P.A. Hancock, *Essays on the Future of Human-Machine Systems*, Banta Information Services Group, Eden Prairie, Minn., 1997.
15. K.M. Ford, C. Glymour, and P.J. Hayes, "Cognitive Prostheses," *AI Magazine*, vol. 18, no. 3, Fall 1997, p. 104.
16. P.A. Hancock and S.F. Scallen, "Allocating Functions in Human-Machine Systems," *Viewing Psychology as a Whole*, R.R. Hoffman, M.F. Sherrick, and J.S. Warm, eds., Am. Psychological Association, Washington, D.C., 1998, pp. 509–540.

Robert R. Hoffman is a research scientist at the University of West Florida's Institute for Human and Machine Cognition and is a faculty associate at the university's Department of Psychology. He specializes in expertise studies, including knowledge elicitation and the design of human-centered systems. He received his BA, MA, and PhD in experimental psychology from the University of Cincinnati. He is a Fulbright Scholar and a fellow of the American Psychological Society. Contact him at the Inst. for Human & Machine Cognition, 40 Alcaniz St., Pensacola, FL 32501; rhoffman@ai.uwf.edu.

Patrick Hayes is a senior research scientist and the John Pace, Jr., Eminent Scholar at the University of West Florida. His research interests include knowledge representation and automatic reasoning, especially the representation of space and time and diagrammatic representations, and the philosophical foundations of AI and computer science. He received his BA in mathematics from Cambridge University and his PhD in artificial intelligence from the University of Edinburgh. Contact him at the Inst. for Human & Machine Cognition, 40 Alcaniz St., Pensacola, FL 32501; phayes@ai.uwf.edu.

Kenneth M. Ford is the founder and director of the University of West Florida's Institute for Human and Machine Cognition. He has an interdisciplinary interest in understanding cognition in both humans and other machines. He is the editor in chief of the AAAI/MIT Press, executive editor of the *International Journal of Expert Systems*, and associate editor of the *Journal of Experimental and Theoretical Artificial Intelligence*, and is a *Behavioral and Brain Sciences* associate. He received his PhD in computer science from Tulane University. Contact him at the Inst. for Human & Machine Cognition, 40 Alcaniz St., Pensacola, FL 32501; kford@ai.uwf.edu.

Peter Hancock is a professor at Central Florida University and president of the Human Factors and Ergonomics Society. His research interests include transportation safety, air traffic control, workstation design, mental workload, intelligent interfaces, simulator design, and other topics. He received a PhD in motor performance from the University of Illinois. Contact him at Dept. of Psychology, Univ. of Central Florida, Orlando, FL 32816; phancock@pegasus.cc.ucf.edu.