

Cognitive Orthoses: Toward Human-Centered AI

*Kenneth M. Ford, Patrick J. Hayes,
Clark Glymour, James Allen*

■ *This introduction focuses on how human-centered computing (HCC) is changing the way that people think about information technology. The AI perspective views this HCC framework as embodying a systems view, in which human thought and action are linked and equally important in terms of analysis, design, and evaluation. This emerging technology provides a new research outlook for AI applications, with new research goals and agendas.*

As George Bernard Shaw once observed, being slandered is better than being ignored. So maybe AI researchers should be happy about what's been happening lately. After decades of pundits and philosophers arguing that AI is provably impossible, suddenly that argument has been replaced with the assertion that not only is it possible, but superhuman AI is so inevitable that it is the greatest danger ever faced by the human race. In only about a decade, the conversation has shifted from *you can't do it ...* to *you shouldn't do it!* That shift has many parallels in other domains, from vaccination, to flight, to splitting the atom, to gene manipulation. The quest for flight, as we have observed elsewhere, affords a particularly striking parallel. Even in the years the indefatigable Wright brothers were hauling their planes from Ohio to Kitty Hawk to France, ever improving them and demonstrating the improvements, scientific wags — the president of the American National Academy of Sciences no less — were arguing that their quest was impossible. By the 1920s




Figure 1. Snapshot of VirtualLabs.

possibility was not an issue, danger was. Books argued passionately that heavier-than-air flight is too dangerous to society, and should be made illegal by international agreement. In capitals around the world, including our own, lobbyists strove to pass laws forbidding attempts at flight. The parallel between AI and artificial flight (AF) is illuminating and suggests that the traditional view of the goal of AI — that is, to create a machine that can successfully imitate human behavior — is wrong.

From the very beginning, attempts at flight sought to imitate the implementation details of birds — the goal seemed to be AB (artificial birds) not AF. In particular, it was taken for granted that flying involved feathers and vigorous flapping. Even in relatively modern times, much scientific debate revolved around exactly how this flapping could be accomplished. The structural similarities and differences of birds and humans had been carefully noted and extensively studied. After all, if you take a skeletal view, people and birds seemed pretty much the same.

Although the Wright brothers were fervent bird watchers, they asked quite different questions, not about flapping, beaks, or feathers, but about lift, stability, thrust, and the physics of turning in air. AI is more abstract than AF but their histories are wonderfully analogous in that both of these strongly held human ambitions were, for a long time, focused on imitating the biological example, and this mistake, in both cases, misdirected these fields. The proper aim of AI is much larger than simply mimicking human behavior. The *scientific* goal is to provide a computational account of mental ability itself, not merely of human mentality. AI is epistemology, android epistemology. But abstract aims can be pursued apace with concrete applications, and we submit that so far, and in the foreseeable future, most of those applications have been pretty good, largely beneficial for society.

Applied AI does often give deference to the human condition, to human goals and limitations but not necessarily to human mechanisms. A principal goal of applied AI is and should be to create cognitive orthotics that can amplify and extend our cognitive abilities. That is now and near; a computational Golem is not. The articles in this special issue reflect a human-centered vision for applied AI that is less about *artificial* intelligence and more about *amplified* intelligence. From this perspective, AI systems can be usefully understood as cognitive orthoses or cognitive prostheses in some cases.¹ Eyeglasses are a simple, but compelling, example of a human-centered technology that can be regarded as a kind of ocular orthoses. Glasses leverage and extend our ability to see but would never pass a Turing test for being an eye.

Just as eyeglasses can be regarded as a kind of ocular orthoses, AI systems can be usefully construed of as a kind of *cognitive* orthoses — that is, technological systems that leverage and extend human cogni-

tion. Current work toward cognitive orthoses and prostheses reflects a fundamentally different perspective from AI's traditional Turing test ambitions. Researchers working in this framework do not set out to imitate human abilities, but to extend and amplify and provide functional substitutes for them. Turing's ghost is still with us, directing the energies of our field in certain directions and subtly discouraging others. We traditionally measure the success of AI systems by comparing them to human performance — which is rather like measuring the performance of aircraft against that of birds and complaining that aircraft do not land in trees or soil our automobiles. Pundits often talk as if our machines are engaged in a competition with the human race. One recalls the folklore story of John Henry and his race against a steam-powered hammer.

Any prosthesis or orthosis is useful only to the extent that it *fits* — in fact, the *goodness of fit* will determine system performance more than any other specific characteristic. This is true whether one considers eyeglasses, a wooden leg, or a cognitive orthosis. One can identify two broad categories of *fit* — species fit and individual fit. In some cases, a particular aspect of human function can afford a consistent fit across most of a population of interest. In many other instances, however, an *individual* fit is desirable and in these cases, relevant differences among individuals must be accommodated. In general, the design and fit of these cognitive orthoses will require a broader interdisciplinary range than has traditionally been associated with most academic units, including computer scientists, engineers, physicians, cognitive psychologists, neuroscientists, and social scientists of various other stripes.

The idea of cognitive orthotics has old and deep roots. Humanity has long recognized that the powers of mind are limited, and has always made devices to compensate for those limitations. Writing is a device for storing information outside the head so that it does not have to be remembered and the abacus was used as an arithmetical prosthesis. In more recent times, the notion of cognitive orthotics connects to Vannevar Bush and his vision of a memory orthosis. Bush also envisioned all sorts of other possible ways in which augmented cognition and perception might be possible.

A particularly compelling opportunity for cognitive orthoses is the aging population. As people age there are progressive changes in verbal skills, abstract reasoning, general intelligence, memory, and other dimensions of cognition. In addition to these changes with normal aging, microstrokes, traumatic injuries, and other physiological occurrences can also affect cognition. We envision cognitive orthoses that enhance and restore cognitive functioning.

Computers have already helped to create a number of revolutions, but as they become more capable, and more human centered, they have the potential

Take a Tutorial at AAAI-16!

The Tutorial Forum provides an opportunity for researchers and practitioners to spend two days each year freely exploring exciting advances in disciplines outside their normal focus. We believe this type of forum is essential for the cross fertilization, cohesiveness, and vitality of the AI field. We all have a lot to learn from each other; the Tutorial Forum promotes the continuing education of each member of AAAI.

Tutorial Schedule for Friday, February 12

9:00 AM - 1:00 PM

FA1: CP-Nets

FA2: Organ Exchanges:

A Success Story of AI in Healthcare

FA3: Recent Directions in Heuristic-Search

FA4: Symbolic Methods for Hybrid Inference,
Optimization, and Decision-making

2:00 PM - 6:00 PM

FP1: AI for Disasters (1 hour, 45 minutes)

FP2: Answer Set Programming

Modulo Theories (1 hour, 45 minutes)

FP3: CogSketch

FP4: Deep Learning: from Foundations
to Implementation

FP5: Type-Based Methods for Interaction in
Multiagent Systems

Tutorial Schedule for Saturday, February 13

9:00 AM - 1:00 PM

SA1: AI Planning and Scheduling for
Real-World Applications

SA2: Constraint (Logic) Programming

SA3: Diffusion in Social Networks

SA4: How to Automatically Machine Read the Web

2:00 - 6:00 PM

SP1: Algorithm Configuration: A Hands on Tutorial

SP2: Algorithms for Maximum Satisfiability
with Applications to AI

SP3: Computational Epidemiology and
Public Health Policy Planning

SP4: Learning and Inference in
Structured Prediction Models

(All tutorials are 4 hours, including breaks, unless otherwise noted.)

to allow us to continue to revolutionize ourselves. Some futurist thinkers take this competition idea very seriously and worry that these mechanical rivals for intellectual dominance will soon take over our planet and treat us like domestic pets or worse. Even in fictional accounts of superhuman AI run amok, the source of the hazard was often not that the machine was “too intelligent” but that it was “too human.” For example, HAL’s design reflects AI’s old ambition to create an artificial human. However, simpler, more reliable and cost-effective methods exist for creating humans and they are not in short supply and arguably in excess. Rather than intelligent computers becoming our rivals or doing our thinking for us . . . they will (and have already) become our amplifiers and teammates.

Note

1. An orthosis is a device that helps correct or amplify an

existing capability, whereas a prosthesis is a device that replaces a missing capability.

Kenneth M. Ford is founder and chief executive officer of the Florida Institute for Human and Machine Cognition (IHMC) — a not-for-profit research institute with locations in Pensacola, and Ocala, Florida.

Patrick J. Hayes is a senior research scientist at the Florida Institute for Human and Machine Cognition (IHMC).

Clark Glymour is a senior research scientist at the Florida Institute for Human and Machine Cognition (IHMC). He also is the Alumni University Professor of Philosophy at Carnegie Mellon University.

James Allen is the associate director and a senior researcher at the Institute for Human and Machine Cognition (IHMC) in Pensacola, Florida. He is also the John H. Dessauer Distinguished Professor of Computer Science at the University of Rochester.