Thanassis Rikakis, Virginia Tech
Aisling Kelliher, Virginia Tech
Jia-Bin Huang, Virginia Tech
Hari Sundaram, University of Illinois at Urbana-Champaign

Progressive Cyber-Human Intelligence for Social Good

Insights

- The use of AI to advance human intelligence and the overall human condition is not yet widespread.
- Systems where humans and AI agents can enhance their abilities together are important for a healthy society.
- New forms of cyber-human intelligence could require a radical rethinking of the social contract.

Imagine a movie-trailer voice intoning, "In a world where AI has learned to partner with humans to peacefully advance society…" And then forget it! That movie, you see, is never getting made—no explosions, no antagonism, no killer robots. However, outside Hollywood, the co-adaptive development of humans and artificial intelligence (AI) may be worth a bit more consideration. There is little doubt at this point that the growth and maturation of AI will be a major influence on our economy and society overall. Significant work is underway both on advancing AI and on combining human and artificial intelligence to improve the functionality and user experience of AI-based methods, tools, and services. Advanced AI is successfully reshaping many transactional contexts such as image search and purchase recommendations, as well as contexts that involve repetitive activity, such as manufacturing. However, AI is progressing much more slowly in contexts that involve rich experiences aimed at advancing human intelligence and the overall human condition—for example, in education. A potentially unintended consequence of this is increased emphasis on the lower-hanging fruit of

IMAGE BY METAMORWORKS / SHUTTERSTOCK

transactional and repetitive contexts, and less emphasis on the more complex human-development contexts that are critical for a healthy society. This article proposes a design approach for tackling the integration of AI into human-development contexts while promoting the development of new forms of cyber-human intelligence.

Building on some of the growing literature exploring design and AI [1,2], our team proposes the concept of *progressive cyber-human intelligence* (PCHI). We define PCHI as a design approach for interactive cyber-human systems. The approach aims to continuously evolve the intelligence, abilities, and experience of the key classes of human actors (e.g., system designer, content expert, end user) in addition to enhancing the abilities of the AI components. PCHI design aims to build on related theories including phenomenology, constructivism, and distributed cognition, while also being differentiated from the types of human-in-the loop AI approaches that focus primarily on the use of human intelligence for the improvement of machine-learning (ML) algorithms. By continuously expanding the ability of humans, together with the abilities of the cyber components, PCHI design focuses on the evolution of the human experience, avoids zero-sum situations such as AI replacing people, and promotes continuously advancing horizons for the development of technology.

PCHI design is aimed at contexts where the ultimate goal is human development through training and/or learning, and where both the human and the machine are learning together. It is specifically aimed at complex learning and training contexts in which the AI needs continuous assistance from human expertise due to a number of key conditions. These include the ongoing evolution of the environment in which the person is learning and evolving, and the high dimensionality of the decision space, as multitudes of cognitive, sensorimotor, and affect parameters need to be addressed. Additional conditions include latent variables, meaning that the parameters affecting the learning are not all observable, and sparse and costly data, for example capturing and tagging data relating to people learning complex tasks in the wild.

PCHI design accepts that humans and machines may have different approaches to decision making. As Herbert Simon suggests in *Models of Man*, humans may focus on decision strategies that satisfy bounded human aspirations rather than rational algorithms that search for optimal decisions, which is the case with computational agents. It is a well-established philosophical principle that humans are not just decision-making machines. Rather, humans are also actors that can seek deep experiences for the sake of the embodied experience itself, without a decision-type outcome. A recent survey of designers who use AI techniques also highlighted the need to recognize the different structures and goals of human intelligence and the statistical intelligence of computing [1]. Ultimately, PCHI design aims to expose the learning process of the human and cyber agents to each other, as well as to evolve the agents toward improving the unique expertise of both humans and cyber agents in synergy. This in turn supports the emergence of powerful forms of cyber-human intelligence in contexts such as evidence-based healthcare, adaptive education, and cultural creativity, which are uniquely positioned to contribute to the major societal challenges of the 21st century.

We believe there are four key principles for practicing PCHI design:

• *The context must call for meaningful human-machine co-development.* PCHI design works best in contexts that focus on complex and aspirational human development. These must be developmental contexts that cannot be handled adequately by improved human ability alone, or through the development of standalone technologies.

• *The basic materials of the learning space are low-level features.* Many current human learning and training structures exist within a context of human sensorimotor and/or intelligence limitations, which often include undiscovered human bias or blind spots [3]. We believe that the basic material for PCHI design should be low-level features of the learning space that can be combined in different ways to form novel cyber-human learning paths that are not limited by human ability or bias. We define low-level features as fundamental building blocks that can be combined in modular ways to reproduce all functions of the system.

• *Evidence-based modeling of the learning space.* Complex interactive learning is still at an early stage of development and therefore not fully discovered or optimized. The design of a PCHI system starts with in-situ interactive learning experiments that produce robust data [4]. This data can then be mined by human experts and/or algorithms to model the interactive learning and support evidence-based development and improvement of the system. An interactive context where the users are performing well and continuously improving will incentivize participation and the contribution of good data. Understanding the learning process also facilitates the appropriate choice of machine intelligence for particular learning contexts.

• *Participatory, iterative design drives the emergence of cyber-human intelligence.* As PCHI systems are developed and refined, later iterations progressively elevate the role of the human agents in the learning, in tandem with enhancing the machine intelligence and improving the system design. This is achieved by leveraging both the data and improved learning models that emerge from earlier cycles, and the progressively enhanced abilities of the human actors from exposure to the system. In addition, outcomes from the customization of computational intelligence to the learning model and the informed offloading of subtasks to the appropriate human or AI agents assist in the integrated learning cycle.

PCHI design requires long-term, genuine commitment by many human actors. Such commitment can be motivated by working in contexts of societal significance. Healthcare, education, and the arts provide three primary contexts for exploring the potential of our approach. In the following section, we give an example of how our team has applied PCHI design to develop systems for interactive neurorehabilitation. We then discuss adaptive curricula and innovation in creative practice as additional areas where PCHI can have significant impact.

The basic material for PCHI design should be low-level features of the learning space that can be combined in different ways to form novel cyber-human learning paths not limited by human ability or bias.

APPLYING PCHI PRINCIPLES TO HOME-BASED INTERACTIVE NEUROREHABILITATION

Neurorehabilitation is a complex and multifaceted therapeutic approach aimed at helping patients recover from nervous system injuries such as brain trauma, Parkinson's disease, or stroke. Here are some of the key elements of the PCHI system we developed to support neurorehabilitation in stroke patients.

Appropriate context. Stroke survivors can benefit significantly from long-term rehabilitation, but the costs and logistics in providing this service at scale are prohibitive. Home-based therapy offers a promising but challenging alternative [5]. There are expansive efforts to develop low-cost systems that are easily embedded in the home and allow patients to practice different tasks by themselves that generalize to activities of daily living. The key outstanding issue to solve is reproducing a complex adaptive rehabilitation experience that advances patient self-efficacy and quality of life, without the continuous presence of the therapist. Much of the expert input of the therapist cannot be fully replicated by computational agents due to sparse and inconsistent data, high dimensionality of the decision space, latent variables, and continuous adaptation. Our own work aims for

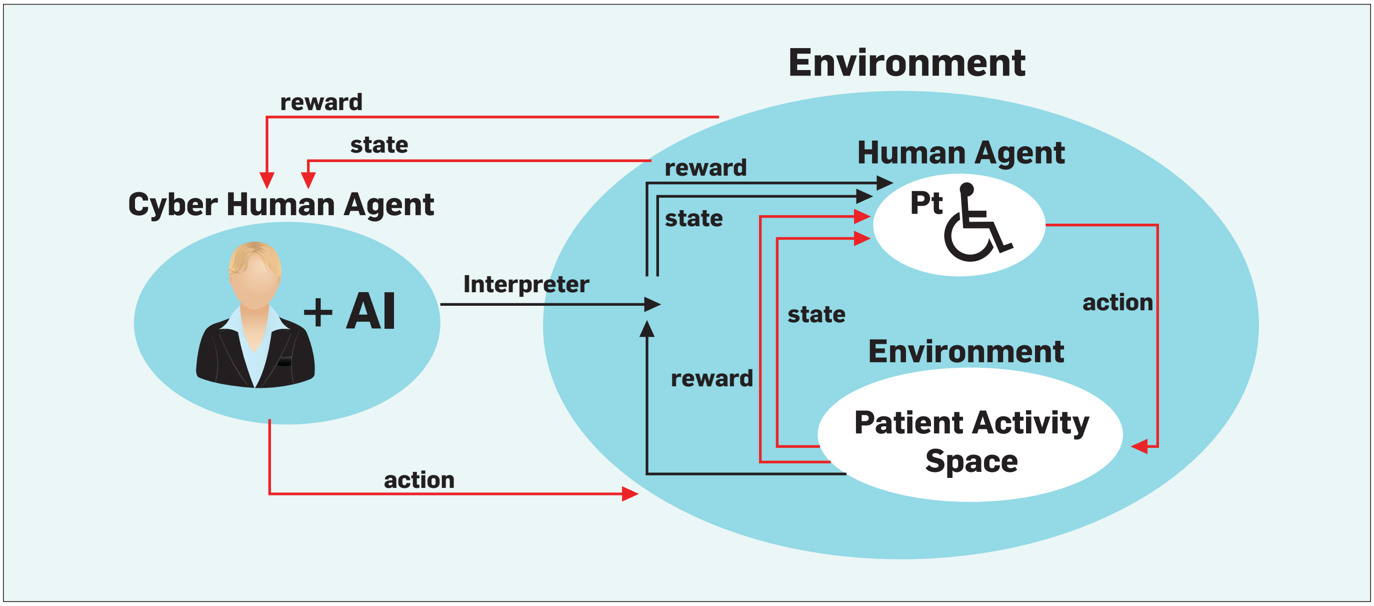

Figure 1. Modeling the cyber-human reinforcement learning space in interactive neurorehabilitation.

lightly supervised therapy for the upper extremity of stroke survivors through the development of progressive cyber-human intelligence between the patient, the therapist, appropriate AI agents, and the design team [6].

Basic low-level features. Working with rehabilitation experts, we are developing a library of low-level therapy features for our system that can be combined in many different configurations and sequences to create adaptive exercise protocols. They include modular objects and tasks for training, interrelated kinematic features and daily-activity parameters for monitoring improvement, and modular multimodal feedback components for the patient.

Modeling the learning space. Extensive related work [3] and an intensive sequence of in-situ experiments by our team facilitate the data modeling of the interactive learning, which is continuously improving. At a high level, it can be presented as a nested reinforcement learning model. The patient's reinforcement learning loop also serves as the environment component of a larger PCHI-reinforcement learning loop. In that larger loop, the therapist and the computational agent act together as a cyber-physical agent that aims to optimize the adaptive training strategy. The agent also assists the patient in interpreting action and state feedback, as those abilities may have been compromised by the patient's stroke (Figure 1).

Iterative design for PCHI. Our system supports information flows that promote continuous learning by the human actors in tandem with learning by the AI agent. The therapist needs to learn to make decisions using the information captured by the system, both raw data presented through appropriate visualizations and patterns extracted by the system. The therapist also needs to help the system learn by rating the performance of the tasks and annotating the key events they use to inform the adaptation decisions. As the patient enhances the quantity and quality of their participation, they learn more and produce better data, which then informs the system and the therapist. The computational agent needs to first learn to imitate the therapist-rating and adaptation decisions, and then make suggestions to the therapist before it takes over some of the rating and decision making. By interacting with the system observations, the therapist also improves their understanding, especially in areas where human ability is limited (e.g., detailed observation of kinematics, extraction of statistical properties over multiple sessions). Iterative participatory design leverages the emerging intelligence of the cyber and human actors to improve the system.

In this project, we use AI to enhance the ability of the therapist to ultimately deliver effective therapy to more patients and at lower costs. We believe the reinforcement learning components of our PCHI approach can be useful in developing systems for other neuro and musculoskeletal contexts, such as spinal cord injury and Parkinson's disease.

PCHI APPLICATIONS IN EDUCATION—TOWARD ADAPTIVE CURRICULA

In the past few decades, the field of higher education and the employment landscape have undergone considerable change. Many of the jobs that students will do once they graduate are currently unknown, as are the skills that will be needed for those jobs. A great deal of the breakthroughs in this century will be at the intersection of areas of knowledge, thus making knowledge segregation harder and even less productive. We now understand that people learn differently and have different types of intelligences, and that all of these intelligences are needed if we are to solve truly complex problems. In addition, the means of delivery have expanded rapidly, offering students even more potential combinations of learning paths. To facilitate adaptive learning, education experts are beginning to break down the knowledge space into low-level features or modules.

The curricula of the future could take the form of as-needed, customizable combinations of modules, of different lengths and credit hours, with each module delivered in a number of different ways (from synchronous person-to-person, to asynchronous interactive) and through different approaches (from theoretical to applied and everything in between). The goal state of *What does the student want to do?* will emerge as the student matures and the faculty and advisors learn more about each student. This progressive learning helps inform the adaptive learning path.

Customizing the curriculum path of each student by hand in such a large possibility space is an overwhelming job for any involved human actor. We thus propose that this problem and opportunity space is ideally fitted to a PCHI design approach. There are increasing amounts of quantitative data on the higher-education learning process, including data from massive online learning courses, from interactive learning (e.g., CMU's Simon Initiative), and from national consortia on paths to student success (EAB, Gallup-Purdue index). This data can drive models, mined by experts working with computational agents that can evolve the expertise of both the human (e.g., student, faculty, advisor) and cyber agents, so they can together support successful adaptive education at scale. As in our earlier example, here PCHI can be presented as four interrelated but different intelligent agents that are co-mastering relevant versions of self-paced curriculum-learning models.

PCHI APPLICATIONS IN THE CREATIVE ARTS

We look to some of the greatest artists of the past century for insights into how the creative arts can maintain their role as an innovation force by accelerating their ability for originality. Composers Claude Debussy and Iannis Xenakis suggested that inspiration for continuous innovation in music can come from creative explorations outside of music, such as from observing the movement of hay in the fields or using stochastic processes to create open forms. Of course, progress has been and continues to be made in this direction, as various facets of the arts form partnerships with diverse fields of expertise and practice to explore novel ideas (e.g., Nelly Ben Hayoun Studios, Atonation).

Searching for interaction possibilities with other areas of emerging knowledge, or across different modes of experiencing, and then translating the findings into meaningful creative innovations is a huge endeavor—one that may not be best served by human intelligence alone. This is another area where PCHI design can be applied to evolve combinations of human-embodied intelligence and computational statistical intelligence in a large and fast-moving search space. PCHI can be used to search for emerging patterns of sense-making in the modern world, and to identify paths for embedding these patterns into creative outcomes that promote reflection on the evolving human condition. The work being presented at SIGGRAPH, Ars Electronica, and ISEA highlights how the computational arts have matured. In addition, industry partners are developing exciting tools, such as Autodesk's Dreamcatcher project, that integrate AI into the creative process, allowing for open-ended searches within and across different forms of expression.

CONCLUSION

As a concluding example, at a recent Autodesk Ideas Meeting Change summit on "Co-creating with Machine Intelligence," the industry and academic discussants alike wondered if current technological changes were disrupting the economy and society to the point where a new societal contract was needed. Such concerns are to some extent legitimized by current and impending job displacement due to automation and the emergence of new skills and training required to function in an economy partly driven by major technological advances like big data analytics and AI. Perhaps, the participants suggested, a new social contract could emerge and evolve through integrative, loosely connected actions that explore the evolution of AI in partnership with the development of a rich and productive society. We are excited to join with the summit participants in working on this vision, and hope others will also join in helping make unlikely Hollywood fantasy a potential reality.

ENDNOTES

1. Dove, G., Halskov, K., Forlizzi, J., and Zimmerman, J. UX design innovation: Challenges for working with machine learning as a design material. *Proc. of the 2017 CHI Conference on Human Factors in Computing Systems.* ACM, New York, 2017, 278–288.
2. Holmquist, L.E. Intelligence on tap: Artificial intelligence as a new design material. *Interactions 24*, 4 (Jul.–Aug. 2017), 28–33.
3. Reinkensmeyer, D.J., Burdet, E., Casadio, M., Krakauer, J.W., Kwakkel, G., Lang, C.E., Swinnen, S.P., Ward, N.S., and Schweighofer, N. Computational neurorehabilitation: Modeling plasticity and learning to predict recovery. *Journal of Neuroengineering and Rehabilitation 13*, 1 (2016), 42.
4. Kartoun, U. Text nailing: An efficient human-in-the-loop text-processing method. *Interactions 24*, 6 (Nov.–Dec. 2017), 44–49.
5. Baran, M., Lehrer, N., Duff, M., Wolf, S., Rymer, Z., and Rikakis, T. Interdisciplinary concepts for design and implementation of mixed reality interactive neurorehabilitation systems for stroke. *Physical Therapy 95* (2015), 449–460.
6. Kelliher, A., Choi, J., Huang, J-B., Rikakis, T., and Kitani, K. HOMER: An interactive system for home-based stroke rehabilitation. *Proc. of the 19th International ACM SIGACCESS.* ACM, New York, 2017, 379–380.

Thanassis Rikakis is professor of bioengineering and of performing arts at Virginia Tech. His research spans systems design, arts, and technology, focusing on interactive neurorehabilitation, adaptive learning, and cyber-human intelligence. He is the founding chair of the Calhoun Honors Discovery Program and director of the Calhoun Center for Higher Education Innovation.
→ rikakis@vt.edu

Aisling Kelliher is an associate professor of computer science with a joint appointment in the Institute for Creativity, Arts, and Technology at Virginia Tech. She co-leads the Interactive Neurorehabilitation Lab (http://vtinr.com/), where her team designs and develops interactive systems for stroke rehabilitation.
→ aislingk@vt.edu

Jia-Bin Huang is an assistant professor in the Bradley Electrical and Computer Engineering Department at Virginia Tech. His main research interests are in the areas of computer vision, computer graphics, and machine learning. He is particularly interested in exploiting physically grounded constraints to enable efficient visual learning.
→ jbhuang@vt.edu

Hari Sundaram is an associate professor in the computer science and advertising departments at the University of Illinois at Urbana-Champaign. His research lies at the intersection of social network analysis and computational advertising and aims to persuade large groups of people to adopt healthy behaviors.
→ hs1@illinois.edu

DOI: 10.1145/3231559