Artificial Intelligence Community and Self

Peter Tu

■ *The means by which artificial intelligence systems will both interact and influence individuals and community are considered, with a description of current capabilities such as video analytics, agency, natural language processing, and online data analysis. Open research questions such as the roles of analogies, associative memory, and the grounding problem are discussed, with speculation presented regarding the possible ramification of this research agenda. Such topics will include how artificial intelligence and humanity might co-evolve; artificial creativity; and life without work.*

Increasingly, artificial intelligence (AI)-enabled agents are interacting with individuals and communities via a myriad of platforms and devices. Physical instantiations include various forms of robots as well as avatars presented via flat screen displays. Commercial products such Amazon's Alexa and Apple Corporation's Siri provide for spoken interactions using mobile and in-home devices, while artificial agents and other online entities are deployed routinely across social media ecosystems. Such mechanisms may require the ability to gather social awareness based on nonverbal cues, interact via physical gestures, understand and produce spoken as well as written speech, manipulate context relevant knowledge, acquire personal information, and develop strategies that can be used to influence both

individuals and communities. These agents can be used to supply information, act as a user interface, and provide companionship; and they can also be used as a means for persuasion, and possibly, manipulation. As an introduction to this transformative topic, we start by describing how AI technology is currently used to enable social interaction. We start with visual awareness, agency, natural language, and individual profiling. The next section will focus on topics that must be addressed to construct the next generation of AI — also known as the Third Wave. This section will include a review of how human language may have formed; the role of analogies; development of AI–community interactions; the roles of associative memory; and the grounding problem. The discussion concludes with speculation regarding the possible ramifications of this technological path. Topics will include co-evolution, coming to grips with our evolutionary past, AI to assist creativity, a human life without work, and the temptation to improve the species itself.

Visual Awareness

In addition to the detection of objects and scene elements, advances in computer vision have allowed for the capture of nonverbal social cues. Capabilities include tracking individuals through crowds, collection of high-resolution imagery via pan-tilt-zoom devices, construction of articulated skeleton models from visible imagery, interpretation of body language (affective pose), estimation of gaze direction, measurement of pupil motion and dilation, and classification of facial expression. Systems like those described by Tu et al. (2015) have been used to capture such information to infer group-level rapport and hostility, among other things.

Physical Agency

Robotic systems such as the one described by McGinn et al. (2019) have demonstrated their utility with respect to social engagement with elderly populations. Interactions include playing trivia games and bingo. An interactive advertising system, described by Tu et al (2019), is comprised of a flat screen display, a collocated sensing camera, and an outcome feedback camera. The display presents an Avatar that has the goal of convincing people to enter a store. The Avatar has a repertoire of actions such as offering a coupon or presenting a piece of merchandise. A sensing camera is used to represent the state of a possible customer. State variables include affective pose, facial expressions, gaze direction, and eyeball movement. An outcome feedback camera is used to detect if and when a customer elects to enter the store. Given access to such feedback, a reinforcement learning algorithm is then used to construct a policy that indicates which actions should be taken as a function of observed customer state.

Natural Language Processing

Several technologies have allowed for new forms of interaction via spoken and written speech. Deep learning methods enable the transformation of spoken language into text. Deep learning methods can also be used to model language in terms of statistical distributions. Given a sequence of words, a prediction of the next word in the sequence can now be produced. Such capabilities have led to the construction of embedding methods where documents are transformed into numerical representations such that Euclidian distance is equivalent to semantic distance. This allows for reasoning over the semantic content of such documents. These methods, coupled with representations such as knowledge graphs, allow agents to participate in various forms of human–machine interactions such as those that are supported by mobile and in-home devices.

Online Agents

Increasingly, artificial agents are attempting to interact with individuals via social media. These agents may focus on exposing individuals to various forms of persuasive information or to make a social connection between individuals so that one can be influenced by another. To increase the efficacy of such methods, various forms of automatic profiling have been investigated. An accepted tool for producing a psychologic profile is to administer an extensive questionnaire. Questions such as *do you like new experiences?* or *do you find such-and-such irritating?* are answered by the subject. It turns out that in general people can be clustered based on their responses to such questionnaires. This results in a set of psychologic dimensions known as OCEAN (openness to experience, conscientiousness, extroversion, agreeableness, neuroticism). For example, *O* stands for openness to new experiences. On one extreme of the O-scale, people are attracted to new forms of art and different types of food. Taking the questionnaire provides for accurate estimates of one's OCEAN coordinates and this can be viewed as a psychograph. Cambridge Analytica proposed the construction of psychographs using AI methods as a proxy for administering a questionnaire. They started by building an online application that, after receiving consent, administers the questionnaire to a small number of Facebook users. They then made use of the subject's Facebook user profile (what they liked to eat, where they hung out, and so on) to determine if an AI-based regression focused on predicting the answers of the questionnaire could be produced. They were successful. They then applied the AI regression to the friends of everyone that took the questionnaire, generating estimated psychographs for all such individuals. This resulted on the order of 50 million estimated psychographs, but also raised questions regarding privacy concerns. See Kroll (2018) for additional background.

Going Forward

Based on technologies such as computer vision, natural language processing, knowledge representation, and data mining, we have already seen the rise of social robots, avatars that learn over time as well as chatbots that inhabit our homes, phones, and social media networks. However, such interactions remain limited. Open research questions remain regarding how to construct more meaningful relationships as would be needed to instantiate a familiar, a guardian, or even a friend. To begin this discussion, consider Michael Tomasello's (2008) work on the origins of human communication. The great apes are able to use intentional gestures to request or demand actions from others. In contrast, early humans were able to use gestures (pointing and pantomime) to accomplish three basic things: make requests for actions from others; give information that might be of assistance; and express one's feelings regarding the environment. The ability to both request and assist is a form of cooperative communication that involves mechanisms such as common conceptual ground, joint intentionality, and various forms of recursive mind reading. The drive toward expressing one's opinions can be viewed as a mechanism for building common conceptual ground with the goal of minimizing within-tribe differences and maximizing between-tribe differences. You might be willing to cooperate with a non-tribe member, but you generally only share your admiration for a pleasant sunset with a member of the tribe. While *pointing* allows for identifying objects within the current field of view, *pantomime* allows for identification of nonvisual referents. As these pantomimes become conventionalized, they can become increasingly arbitrary in nature. This allows for development of spoken language that is based on completely arbitrary but socially accepted symbols. Arbitrary languages (both spoken and sign) have the complexity required to construct narratives that can describe events and actions of multiple agents over both space and time. Various frameworks can then emerge so that these sequences of events can appear to be causal and even make sense. Through this analysis we see that cognitive skills evolve phylogenetically (over the lifetime of the species), enabling the creation of cultural products (history), which then provide developing children with the biologic and cultural tools they need to develop ontogenetically (over the lifetime of the organism).

Based on the previous analysis of human cooperation, the construction of more meaningful relationships between AI agents and their human counterparts may require several technological leaps. Such topics may include the use of analogies and metaphors, artificial community, associative memory, and the grounding problem. These advancements will be challenging from both computational and philosophical perspectives. If progress can be made along these lines of investigation, such efforts may lead to the construction of what may be known as *grounded AI agents*. The following is a brief description of these AI questions.

Analogies

In the book by Hofstadter and Sander (2013) we see the argument that analogies should be viewed as the *currency of thought*. They are ubiquitous and are more than simple graph isomorphisms. As we develop analogy landscapes, it can be argued that the meaning of symbolic concepts can be established via their analogies to other symbolic concepts — this may be one approach to the grounding problem (Deacon 1998). Jaynes (1990) argues that understanding a thing is to arrive at a metaphor for that thing by substituting something more familiar to us. The feeling of familiarity is the feeling of understanding. A theory is a metaphor between a model and data. Understanding in science is the feeling of similarity between complicated data and a familiar model. The instantiation of a wide variety of fluid analogy engines will be a cornerstone for third-wave AI research. Tasks such as recognition, memory, and the ability to operate seamlessly across multiple domains will require a deeper understanding of how analogies are formed and manipulated.

Artificial Community

The *Flynn effect* refers to the observation that standardized intelligence quotient results are continuously increasing. Such rapid developments are difficult to attribute to slow-moving evolutionary processes. However, it can be argued that cultural artifacts such as language, which encapsulates increasingly sophisticated conceptual spaces, has the capacity to support such remarkable progress. Thus, new forms of intelligence may emerge from the concept of community.

Associative Memory

There seems to be an extraordinary number of cognitive mechanisms that do not rely on consciousness. The conscious mind might only oversee the *identification* of problems, and not be responsible for solving them. The illusion of consciousness may be a result of integrating disparate inputs of the subconscious into a consistent narrative Eagleman (2011). In Kahneman (2011), we see the argument that what we think of as our conscious selves should be thought of as the *B* actor that thinks it is the star. The real star of the show is associative memory. Given unforeseen circumstances, relevant memories seem to simply come to mind. How are memories encoded? Should we view analogies as exotic memory-indexing functions? How are disparate memories combined, resulting in an interpretation of novel circumstances? Significant progress in AI research hinges on the ability to address these intriguing questions. The mechanisms behind associative memory will dictate how data are ultimately transformed into experience.

Grounding

John Searle's *Chinese Room problem* was used as an argument against the Turing Test. In this scenario, a person is situated in a box and is handed questions written in Chinese characters. He or she has a set of rules written in various books that allows him or her to produce a set of Chinese characters that provide a suitable answer to such queries. The person does not speak or read Chinese and at no point are the Chinese characters translated into their native language. Searle's argument is that even though such a system seems to satisfy the Turing test, we would be hard-pressed to call such a system *intelligent* because the *grounding problem* (the ability to establish symbolic meaning) is not addressed. It can be argued that systems must start to learn using mathematics that are both continuous and discrete or symbolic in nature (Steels and Hild 2012). If successful, AI will transition from problems of state estimation to the interpretation of meaning.

Ramifications

The potential long-term effects of advanced AI on society are now considered, to include the possible co-evolution of humanity and AI; the unshackling of our perspectives from our evolutionary past; new forms of creativity; the ability to build meaningful lives in the absence of work; and whether we choose to identify and neutralize the predators that have always been part of the human experience. Brief speculation regarding these intriguing topics is then given.

Co-evolution

It has been hypothesized that the co-habitation of humans and canines may have resulted in significant cognitive changes for both species. During this period, the prefrontal cortex of the human brain increased significantly, at the possible expense of our senses of hearing and smell. Dogs, in turn, gave up their wolf-like cunning in trade for more stable living conditions. Could exposure to advanced social AI allow for another jump in human cognitive evolution? Like the reliance on the global position system, which has diminished our ability to navigate, online forms of social interaction seem to reduce our capacity to recognize social cues such as facial expression and eye contact. On the positive side, it may be the case that co-evolution with sophisticated AI Agents may provide us with new coping mechanisms to help alleviate the effects of loneliness, depression, and even, perhaps, autism.

Escaping Our Past

Robert Wright's (2017) analysis of Buddhism makes the argument that we are still firmly in the grip of our evolutionary past. For example, for the last 10,000 generations most humans would have gone through life encountering at most 50 people. The opinion that those 50 people have of you might very well determine whether your genes make it into the next generation. So, we are all the product of 10,000 fathers and 10,000 mothers who managed to win their popularity contests. Flash-forward to today. For many people, the prospect of attending a cocktail party, where they don't know anyone and hence there is a reasonable likelihood of being ignored or even rejected, is a great source of anxiety. This is odd when one considers the fact that the opinions of the other partygoers will have almost no bearing on one's survival. Could an AI agent in the form of a familiar presence help us to see the world as it is, as opposed to through the lens of our evolutionary past?

Creativity

Charles Fernyhough (2016) makes the point that the back-and-forth nature of dyadic conversations between diverse individuals often leads to creative solutions. He goes further to assert that via theory-of-mind mechanisms, we are often able to hold such conversations within our own minds. He makes the argument that John Lennon and Paul McCartney may have spent relatively little time directly collaborating on pop songs, but because both had strong models of each other, they each might have been able to hold an internal form of co-evolving musical creation. It may be the case that interactions between humans and social AI agents may result in new forms of creative conversation.

Life without Work

Rifkin (2014) heralds a future where marginal costs (the cost of producing additional units of a product or service, once initial fixed costs have been accounted for) shrink asymptotically to zero. The cost of living for workers will be significantly reduced if the ability to produce anything, anywhere, at almost no cost, can be established on a locality-by-locality basis. Whole communities will be able to decouple themselves from reliance on 21st century global manufacturing. Key zero-marginal cost technologies will include the cost of raw materials, such as the recycling of waste produced by zero-marginal-cost communities; the cost of physical labor, such as robots gifted with the ability to learn and perform any physical task; the cost of specialized manufacturing, such as 3D printing capable of producing any type of object; the cost of energy, which would include renewables such as solar and wind; and the cost of research and management, performed by AI. A possible consequence of such developments is that work plays an increasingly smaller role in people's lives. It may be the case that social AI agents will be needed to help humanity achieve meaningful existences in the absence of work.

Snakes in Suits

In Ronson's (2012) book, we find that persons with psychopathic tendencies occur in the general population with a frequency of approximately one in 100.

In corporate America, the statistics are closer to three in 100. Bob Hare has established the following checklist for the purpose of identifying such individuals: glib and superficial charm; grandiose self-worth; proneness to boredom; pathological lying; display of cunning and manipulation; lack of remorse or guilt; shallow affect; lack of empathy; parasitic lifestyle; poor behavioral controls; promiscuity; early behavioral problems; lack of realistic long-term goals; impulsivity; irresponsible behaviors; inability to accept responsibility for own actions; numerous short-term relationships; juvenile delinquency; and criminal versatility. Throughout history, these villains have been the cause of great sorrow and pain. Modern AI could be used to actively identify and neutralize such predators. However, the ethical question that AI protectors of society must address is: *should they?* The human condition is a spectrum, and as the poet René Char put it, we must each "develop [our] legitimate strangeness."

References

Deacon, T. W. 1998. *The Symbolic Species: The Coevolution of Language and the Brain*. New York: W. W. Norton and Company.

Eagleman, D. 2011. *Incognito: The Secret Lives of the Brain*. New York: Pantheon Books.

Fernyhough, C. 2016. *The Voices Within: The History and Science of How We Talk to Ourselves*. New York: Basic Books.

Hofstadter, D., and Sander, E. 2013. *Surfaces and Essences: Analogy as the Fuel and Fire of Thinking*. New York: Basic Books.

Jaynes, J. 1990. *The Origin of Consciousness in the Break Down of the Bicameral Mind*. Boston, MA: Houghton Mifflin Company.

Kahneman, D. 2011. *Thinking Fast and Slow*. New York: Farrar, Straus, and Giroux.

Kroll, A. 2018. Cloak and Data: The Real Story Behind Cambridge Analytica's Rise and Fall. *Mother Jones* May/June. www.motherjones.com/politics/2018/03/cloak-and-data-cambridge-analytica-robert-mercer.

McGinn, C.; Bourke, E.; Murtagh, A.; Donovan, C.; and Cullinan, M. F. 2019. Meeting Stevie: Perceptions of a Socially Assistive Robot by Residents and Staff in a Long-Term Care Facility. In *Proceedings of the 14th ACM/IEEE International Conference on Human-Robot Interaction*, 602–3. New York: Association for Computing Machinery. https://doi.org/10.1109/HRI.2019.8673161

Rifkin, J. 2014. *The Zero Marginal Cost Society: The Internet of Things, the Collaborative Commons, and the Eclipse of Capitalism*. New York: Palgrave Macmillan.

Ronson, J. 2012. *The Psychopath Test: A Journey Through the Madness Industry*. New York: Riverhead Books.

Steels, L., and Hild, M., editors. 2012. *Language Grounding in Robots*. Berlin, Germany: Springer. https://doi.org/10.1007/978-1-4614-3064-3

Tomasello, M. 2008. *Origins of Human Communication*. Cambridge, MA: The MIT Press. https://doi.org/10.7551/mitpress/7551.001.0001

Tu, P.; Chen, J.; Chang, M. C.; Yu, T.; Tian, T. P.; Rubin, G.; and Logan-Terry, A. 2015. Cross-Cultural Training Analysis via Social Science and Computer Vision Methods. *Procedia Manufacturing* 3: 4068–75. https://doi.org/10.1016/j.promfg.2015.07.978

Tu, P.; Kubricht, J.; Pang, A. S.; Chen, A.; Gao, T.; and Foehner, J. 2019. A Learning Approach to Interactive Advertising. In *16th Institute of Electrical and Electronics Engineers (IEEE) International Conference on Advanced Video and Signal Based Surveillance (AVSS)*. Piscataway, N.J.: IEEE.

Wright, R. 2017. *Why Buddhism Is True: The Science and Philosophy of Meditation and Enlightenment*. New York: Simon & Schuster.


Courtesy, iStock

Join Us for AAAI-22 in Vancouver, British Columbia, Canada !

As many know, AAAI-21 was intended to be in Vancouver, Canada, but was held virtually due to the current pandemic. With hope for a brighter future, we invite you to join us in Vancouver, British Columbia, Canada, February 22 through March 1, 2022 for the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI-22), the Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence (IAAI-22), and the Twelfth Symposium on Educational Advances in Artificial Intelligence. We have learned a lot this past year, and hope to incorporate some of the virtual features of 2021 into 2022.

AAAI-22 will be held at the Vancouver Convention Centre. Located on Vancouver's downtown waterfront with a dramatic mountain backdrop, the award-winning Vancouver Convention Centre offers one of the most beautiful settings in the world. The center is steps away from hotels, restaurants, and all the entertainment, retail, and nightlife that the city offers. Voted as one the '"Best City in the Americas" year after year, Vancouver is a vibrant, cosmopolitan city nestled between the Pacific Ocean and the soaring peaks of the North Shore Mountains. With the convention centre and Vancouver as home base, you will be well positioned to explore this unique region. Seaplanes, helicopters, ferries and cruise ships are just minutes away. Discover attractions like Victoria – British Columbia's capital – and the scenic Gulf Islands. Take a leisurely drive up the Sea-to-Sky Highway to Whistler, one of North America's top ski resorts, sample the award-winning vintages of the Okanagan wine country, or travel the Alaskan coastline.

For complete information about options in Vancouver, please visit www.tourismvancouver.com. The 2022 Call for Papers will be available in early summer at the following url:

www.aaai.org/aaai22

Peter Tu is GE's chief scientist for AI. His research revolves around video analytics in the form of multi-camera situational awareness and behavior recognition. His current interests also include advanced AI topics such as the grounding problem and multiagent cooperation.