

# Artificial Intelligence and the “Barrier” of Meaning

## Workshop Participants

**Alan Mackworth** is a Professor of Computer Science at the University of British Columbia. He was educated at Toronto (B.A.Sc.), Harvard (A.M.) and Sussex (D.Phil.). He works on constraint-based artificial intelligence with applications in vision, robotics, situated agents, assistive technology and sustainability. He is known as a pioneer in the areas of constraint satisfaction, robot soccer, hybrid systems and constraint-based agents. He has authored over 130 papers and co-authored two books: *Computational Intelligence: A Logical Approach* (1998) and *Artificial Intelligence: Foundations of Computational Agents* (2010; 2<sup>nd</sup> Ed. 2017). He served as the founding Director of the UBC Laboratory for Computational Intelligence and the President of AAAI and CAIAC. He is a Fellow of AAAI, CAIAC, CIFAR and the Royal Society of Canada.

\*\*\*\*\*

**Alison Gopnik** is a professor of psychology and affiliate professor of philosophy at the University of California at Berkeley. She received her BA from McGill University and her PhD. from Oxford University. She is an internationally recognized leader in the study of children’s learning and development and introduced the idea that probabilistic models and Bayesian inference could be applied to children’s learning. She is an elected member of the American Academy of Arts and Sciences and a fellow of the Cognitive Science Society and the American Association for the Advancement of Science.

She is the author or coauthor of over 100 journal articles and several books including “Words, thoughts and theories” MIT Press, 1997, and the bestselling and critically acclaimed popular books “The Scientist in the Crib” William Morrow, 1999, “The Philosophical Baby; What children’s minds tell us about love, truth and the meaning of life”, and “The Gardener and the Carpenter”, Farrar, Strauss and Giroux. She has also written widely about cognitive science and psychology for The New York Times, The Atlantic, and The New Yorker among others. Her TED talk on her work has been viewed more than 3 million times. And she has frequently appeared on TV and radio including “The Charlie Rose Show” and “The Colbert Report”. Since 2013 she has written the Mind and Matter column for the Wall Street Journal.

\*\*\*\*\*

**Barbara Grosz** is Higgins Professor of Natural Sciences in the School of Engineering and Applied Sciences at Harvard University and a member of the External Faculty of Santa Fe Institute. She has made many contributions to the field of Artificial Intelligence (AI) through her pioneering research in natural language processing and in theories of multi-agent collaboration and their application to human-computer interaction. Her current research explores ways to use models developed in this research to improve health care coordination and science education. A member of the National Academy of Engineering and the American Philosophical Society, she is a fellow of the American Academy of Arts and Sciences, the Association for the Advancement of Artificial Intelligence, the Association for Computing Machinery, and the American Association for the Advancement of Science, and a corresponding fellow of the Royal Society of Edinburgh. She received the 2009 ACM/AAAI Allen Newell Award, the 2015 IJCAI Award for Research Excellence, AI’s highest honor, and the 2017 Association for Computational Linguistics Lifetime Achievement Award. She was founding dean of science and then dean of Harvard’s Radcliffe Institute for Advanced Study, and she is known for her role in the establishment and leadership of interdisciplinary institutions and for her contributions to the advancement of women in science. Professor Grosz serves on the boards of several scientific, scholarly and academic institutions.

\*\*\*\*\*

**Bart Selman** is a Professor of Computer Science at Cornell University. He previously was at AT&T Bell Laboratories. His research interests include computational sustainability, efficient reasoning procedures, planning,

knowledge representation, and connections between computer science and statistical physics. He has (co-)authored over 100 publications, including six best paper awards. His papers have appeared in venues spanning Nature, Science, Proc. Natl. Acad. of Sci., and a variety of conferences and journals in AI and Computer Science. He has received the Cornell Stephen Miles Excellence in Teaching Award, the Cornell Outstanding Educator Award, an NSF Career Award, and an Alfred P. Sloan Research Fellowship. He is a Fellow of the American Association for Artificial Intelligence and a Fellow of the American Association for the Advancement of Science.

\*\*\*\*\*

**Brenden M. Lake** is an Assistant Professor of Psychology and Data Science at New York University. He received his M.S. and B.S. in Symbolic Systems from Stanford University in 2009, and his Ph.D. in Cognitive Science from MIT in 2014. He was a postdoctoral Data Science Fellow at NYU from 2014-2017. Brenden is a recipient of the Robert J. Glushko Prize for Outstanding Doctoral Dissertation in Cognitive Science, he is a MIT Technology Review Innovator Under 35, and his research was selected by Scientific American as one of the most important advances of 2016. His research focuses on computational problems that are easier for people than they are for machines, such as learning new concepts, creating new concepts, learning-to-learn, and asking questions.

\*\*\*\*\*

**Bruno Olshausen** is Professor of Neuroscience and Optometry at the University of California, Berkeley. He also serves as Director of the Redwood Center for Theoretical Neuroscience, an interdisciplinary research group focusing on mathematical and computational models of brain function. He received B.S. and M.S. degrees in Electrical Engineering from Stanford University, and a Ph.D. in Computation and Neural Systems from the California Institute of Technology. From 1996-2005 he was a faculty member of the Departments of Psychology and Neurobiology, Physiology & Behavior at UC Davis. Since 2005 he has been at UC Berkeley. Olshausen's research aims to understand the information processing strategies employed by the brain for doing tasks such as object recognition and scene analysis. The aim of this work is not only to advance our understanding of the brain, but also to discover new algorithms for scene analysis based on how brains work.

\*\*\*\*\*

**Chris Wood** received his Ph.D. from Yale University in 1973. Following a postdoctoral appointment at Walter Reed Army Institute of Research, he returned to Yale as faculty with joint appointments in Neurology, Neurosurgery, and Psychology. He remained at Yale until 1989 when he became Group Leader of the Biophysics Group at Los Alamos National Laboratory, a group responsible for a wide range of physics and biophysics research, including protein crystallography, quantum information, and human brain imaging.

Chris left Los Alamos in 2005 to become Vice President of the Santa Fe Institute, a position he held until 2016 when he retired after chemotherapy and major surgery for Stage 4 cancer. He has remained active in retirement, organizing SFI Working Groups on "The Meaning of Information" (co-organized with philosopher Daniel Dennett), and "Cybersecurity is Turing's Achilles Heel", which brought together experts in cybersecurity from leading industry and government organizations.

Chris' research interests include imaging and modeling the human brain, and the fundamental relationships between abstract computation and its physical implementations in natural and human-constructed systems. Chris began his long relationship with digital computers as a graduate student engaged in real-time data acquisition from the brains of humans undergoing neurosurgical procedures.

\*\*\*\*\*

**Christopher Mole** is a Professor of Philosophy at the University of British Columbia, Vancouver, where he chairs the Programme in Cognitive Systems. He is the author of *Attention is Cognitive Unison* (2011), and *The Unexplained Intellect: Complexity, Time, and the Metaphysics of Embodied Thought* (2016). He is also the co-author of a forthcoming book on philosophical issues in psychiatry. His research is concerned with various questions that arise from our attempts to understand the mind scientifically.

\*\*\*\*\*

**Cristopher Moore** received his B.A. in Physics, Mathematics, and Integrated Science from Northwestern University, and his Ph.D. in Physics from Cornell. From 2000 to 2012 he was a professor at the University of New Mexico, with joint appointments in Computer Science and Physics. Since 2012, Moore has been a resident professor at the Santa Fe Institute; he has also held visiting positions at École Normale Supérieure, École Polytechnique, Université Paris 7, the Niels Bohr Institute, Northeastern University, and the University of Michigan. He has published over 150 papers at the boundary between physics and computer science, ranging from quantum computing, to phase transitions in NP-complete problems, to the theory of social networks and efficient algorithms for analyzing their structure. He is an elected Fellow of the American Physical Society, the American Mathematical Society, and the American Association for the Advancement of Science. With Stephan Mertens, he is the author of *The Nature of Computation* from Oxford University Press.

\*\*\*\*\*

**Cosma Shalizi** is an associate professor in the Department of Statistics at Carnegie Mellon University, with joint appointments in the Machine Learning Department, Center for the Neural Basis of Cognition, and the H. John Heinz III College. He is also external faculty with the Santa Fe Institute. Cosma earned an A.B. in physics from the University of California, Berkeley, in 1993, and a Ph.D. in physics from the University of Wisconsin-Madison in 2001. Cosma came to Carnegie Mellon in 2005 following post-doc positions at the Santa Fe Institute and the University of Michigan Center for the Study of Complex Systems. His research interests include statistical inference for complex systems; nonparametric prediction for stochastic processes; causal inference; large deviations and ergodic theory; networks and information flow in neuroscience, economics and social sciences; heavy-tailed distributions; and self-organization.

\*\*\*\*\*

**David Krakauer** is President and William H. Miller Professor of Complex Systems at the Santa Fe Institute. David's research explores the evolution of intelligence on earth. This includes studying the evolution of genetic, neural, linguistic, social and cultural mechanisms supporting memory and information processing, and exploring their generalities. At each level David asks how information is acquired, stored, transmitted, robustly encoded, and processed. This work is undertaken through the use of empirically supported computational and mathematical models.

David served as the founding Director of the Wisconsin Institute for Discovery, the Co-Director of the Center for Complexity and Collective Computation, and was Professor of mathematical genetics at the University of Wisconsin, Madison. David has previously served as chair of the faculty and a resident professor and external professor at the Santa Fe Institute.

David has been a visiting fellow at the Genomics Frontiers Institute at the University of Pennsylvania, a Sage Fellow at the Sage Center for the Study of the Mind at the University of Santa Barbara, a long-term Fellow of the Institute for Advanced Study in Princeton, and visiting Professor of Evolution at Princeton University.

In 2012 Krakauer was included in the *Wired Magazine Smart List* as one of 50 people "who will change the World." In 2016 Krakauer was included in *Entrepreneur Magazine's* visionary leaders advancing global research and business.

\*\*\*\*\*

**David Wolpert** is a professor at the Santa Fe Institute, visiting professor at MIT, and adjunct professor at ASU. He is the author of three books (and co-editor of several more), over 200 papers, has three patents, is an associate editor at over half a dozen journals, has received numerous awards, and is a fellow of the IEEE.

He has over 19,000 citations, in fields including thermodynamics of computation, molecular biology, foundations of physics, machine learning, dynamics of languages, game theory and distributed optimization. In particular his machine learning technique of stacking was instrumental in both winning entries for the Netflix competition, and his papers on the no free lunch theorems have over 7,000 citations. (Details at <http://davidwolpert.weebly.com>).

Most of his current research involves two topics: (1) Combining recent breakthroughs in nonequilibrium statistical physics with computer science theory to lay the foundations of a complete theory of the thermodynamics of computation. (2) Modeling social organization (command and information networks within social groups) using information theory.

Before his current position he was the Ulam scholar at the Center for Nonlinear Studies, and before that he was at NASA Ames Research Center and a consulting professor at Stanford University, where he formed the Collective Intelligence group. He has worked at IBM and a data mining startup, and is external faculty at numerous international institutions. His degrees in Physics are from Princeton and the University of California.

\*\*\*\*\*

**Dawn Song** is a Professor in the Department of Electrical Engineering and Computer Science at UC Berkeley. Her research interest lies in deep learning, security, and blockchain. She has studied diverse security and privacy issues in computer systems and networks, including areas ranging from software security, networking security, distributed systems security, applied cryptography, blockchain and smart contracts, to the intersection of machine learning and security. She is the recipient of various awards including the MacArthur Fellowship, the Guggenheim Fellowship, the NSF CAREER Award, the Alfred P. Sloan Research Fellowship, the MIT Technology Review TR-35 Award, the Faculty Research Award from IBM, Google and other major tech companies, and Best Paper Awards from top conferences in Computer Security and Deep Learning. She is ranked the most cited scholar in computer security (AMiner Award). She obtained her Ph.D. degree from UC Berkeley. Prior to joining UC Berkeley as a faculty, she was a faculty at Carnegie Mellon University from 2002 to 2007. She is also a serial entrepreneur, including Ensign Security (acquired by FireEye Inc.) and Menlo Security. She is currently the founder and CEO of Oasis Labs, creating privacy-first cloud computing on blockchain.

\*\*\*\*\*

**Dileep George** is co-founder of Vicarious AI, where he leads the development of a general AI layer for robotics. Vicarious is an AI+robotics startup with backing from several well known investors that include Mark Zuckerberg, Jeff Bezos, Founders Fund, and Khosla Ventures. Vicarious has pioneered brain-inspired and data-efficient AI algorithms, and they gained world-wide attention for fundamentally breaking text-based CAPTCHAs with very little training data. Dileep has authored several influential papers, published in Science, NIPS, ICML and CVPR and his research has been featured in NYT, WSJ and NPR. Before cofounding Vicarious, Dileep was CTO of Numenta, an AI company he cofounded with Jeff Hawkins and Donna Dubinsky during his graduate studies at Stanford. Before Numenta, Dileep was a Research Fellow at the Redwood Neuroscience Institute. Dileep's research on hierarchical models of the brain earned him a PhD in Electrical Engineering from Stanford University. He has an MS in Electrical Engineering from Stanford and a B.Tech in EE from IIT Mumbai.

\*\*\*\*\*

**Douwe Kiela** is a Research Scientist at Facebook AI Research in New York. He holds a PhD in Computer Science from the University of Cambridge, where he was supervised by Stephen Clark. His thesis was on grounding semantics in perceptual modalities. He has recently given a TEDx talk about grounded language learning and a tutorial on multi-modal learning and reasoning at ACL. He has co-authored over thirty articles in the leading conferences and journals in natural language processing and machine learning. His current interests lie in grounding emergent communication in multi-agent systems and learning hierarchical representations for improved natural language understanding.

\*\*\*\*\*

**Fernando Pereira** is VP and Engineering Fellow at Google, where he leads research and development in natural language understanding and machine learning. His previous positions include chair of the Computer and Information Science department of the University of Pennsylvania, head of the Machine Learning and Information Retrieval department at AT&T Labs, and research and management positions at SRI International. He received a Ph.D. in Artificial Intelligence from the University of Edinburgh in 1982, and has over 120 research publications on computational linguistics, machine learning, bioinformatics, speech recognition, and logic programming, as well as several patents. He was elected AAAI Fellow in 1991 for contributions to computational linguistics and logic programming, ACM Fellow in 2010 for contributions to machine learning models of natural language and biological sequences, and ACL Fellow for contributions to sequence modeling, finite-state methods, and dependency and deductive parsing. He was president of the Association for Computational Linguistics in 1993.

\*\*\*\*\*

**Garrett T. Kenyon** received his Bachelor's degree in Physics from the University of California, Santa Cruz and his Ph.D. in Physics from the University of Washington, Seattle. He has worked for over thirty years on problems involving various aspects of neural computation. Dr. Kenyon is listed as the first or senior author on tens of refereed publications and regularly serves as a reviewer for the NSF, NIH and other government agencies. Dr. Kenyon has led a number of research projects funded by the NSF, DARPA and DoD as well as projects funded by the DOE's LDRD program. Recently, Dr. Kenyon led the effort to develop neuromorphic object detection and tracking algorithms for the Sparse Adaptive Local Learning for Sensing and Analytics (SALLSA) project, itself a component of the DARPA program in Unconventional Processing of Signal for Intelligent Data Exploitation (UPSIDE). Dr. Kenyon also recently led a DOE LDRD funded effort to use PetaVision, a high-performance neural simulation toolbox developed over the last decade by Dr. Kenyon's research team, to implement a Sparse Prediction Machine on LANL's Trinity supercomputer. Dr. Kenyon is currently PI of an NSF funded project to integrate neural processing mechanisms, including top-down feedback and synchronous oscillations, into convolutional neural network architectures. He has recently started a research project with Prof. Ed Kim, Villanova, funded through the Intel Neural Research Community, to investigate robustness to adversarial examples using sparse coding algorithms implemented on the Loihi neuromorphic processor.

\*\*\*\*\*

**Irene Pepperberg** (S.B, MIT, '69; Ph.D., Harvard, '76), a Research Associate at Harvard, has been a Northwestern University visiting Assistant Professor, University of Arizona tenured Associate Professor, MIT Media Lab visiting Associate Professor, Brandeis University adjunct Associate Professor and MIT Senior Lecturer. She received John Simon Guggenheim, Whitehall, Harry Frank Guggenheim, Selby, and Radcliffe Fellowships, was an alternate for the Cattell Award for Psychology, won the 2005 Frank Beach Award (best paper—comparative psychology), was nominated for Weizmann, L'Oréal, and Grawemeyer Awards, and the Animal Behavior Society's 2001 Quest and 2015 Exemplar Awards. She received St. Johns University's 2013 Clavius Award. Her book, *The Alex Studies*, on Grey parrot cognition/interspecies communication, received favorable reviews in publications as diverse as the NY Times and Science. Her memoir, *Alex & Me*, a NY Times bestseller, won a Christopher Award. She has published over 100 scholarly articles in peer reviewed journals and as book chapters, is a Fellow of the Animal Behavior Society, the American Psychological Association, the American Psychological Society, the American Ornithologists' Union, AAAS, Psychonomics Society, and the Midwestern and Eastern Psychological Associations. She is a consulting editor for several journals, the previous associate editor for *Journal of Comparative Psychology*, and member-at-large for APA Divisions 3 and 6.

\*\*\*\*\*

**Jessica Flack** is a professor at the Santa Fe Institute and director of SFI's Collective Computation Group(C4). The focus of Flack's research program is development of a theory for how adaptive systems collectively compute. Research in C4 draws on evolutionary theory, cognitive neuroscience and behavior, statistical mechanics, information theory, dynamical systems and theoretical computer science to study the roles of information processing and collective computation in the emergence of robust structure and function in biological

and social systems. Goals include identifying the computational principles that allow nature to overcome subjectivity due to information processing to produce ordered states and understanding why adaptive systems typically have multiple space and timescales. A central idea is that noisy information processors construct their macroscopic worlds through collective coarse-graining in evolutionary and/or learning time. In other words, how the appropriate aggregation of information from individuals making decisions under uncertainty can produce good collective forecasts. Flack was previously founding director of University of Wisconsin-Madison's Center for Complexity and Collective Computation in the Wisconsin Institutes for Discovery. Flack's work has been covered by scientists and science journalists in many publications and media outlets, including the *BBC*, *NPR*, *Nature*, *Science*, *The Economist*, *New Scientist*, *Current Biology*, *The Atlantic*, and *Quanta Magazine*.

\*\*\*\*\*

**Josh Bongard** is a professor in the Department of Computer Science at the University of Vermont. He was named a Microsoft New Faculty Fellow (2006), an MIT Technology Review member of the "Top 35 Innovators Under the Age of 35" (2007), and the recipient of a PECASE award (2011). His funded research covers the crowdsourcing of robotics, embodied cognition, human-robot interaction, autonomous machines that recover function after damage, soft robotics, and white box machine learning. His work has been funded by NSF, DARPA, ARO, AFRL, the Sloan Foundation, and NASA. He is the co-author of the book "How the Body Shapes the Way We Think: A New View of Intelligence".

\*\*\*\*\*

**Julia Hirshberg** is Percy K. and Vida L. W. Hudson Professor of Computer Science and was Chair of the Computer Science Department at Columbia University from 2012-2018. She received her PhD in Computer Science from the University of Pennsylvania. She worked at Bell Laboratories and AT&T Laboratories -- Research from 1985-2003 as a Member of Technical Staff and a Department Head, creating the Human-Computer Interface Research Department. She served as editor-in-chief of Computational Linguistics from 1993-2003 and co-editor-in-chief of Speech Communication from 2003-2006 and is now on the Editorial Board. She was on the Executive Board of the Association for Computational Linguistics (ACL) from 1993-2003, on the Permanent Council of International Conference on Spoken Language Processing (ICSLP) since 1996, and on the board of the International Speech Communication Association (ISCA) from 1999-2007 (as President 2005-2007, Advisory Council 2007--), on the Executive Board of the Computing Research Association from 2013-2014, on the AAAI Council from 2012-2015, and the Executive Board of the North American Chapter of the ACL, 2012-2015, and on the IEEE Speech and Language Processing Technical Committee from 2011-2018. She currently co-chairs the CRA's Committee on Women (CRA-W), on which she has served on as board member since 2009. She has been active in working for diversity at AT&T and at Columbia. She has been a AAAI fellow since 1994, an ISCA Fellow since 2008, and a (founding) ACL Fellow since 2011, an ACM Fellow since 2015, an IEEE Fellow since 2017, a member of the National Academy of Engineering since 2017, and a member of the American Academy of Arts and Sciences since 2018. She was elected to the American Philosophical Society in 2014 and as an Honorary Member of the Association for Laboratory Phonology in the same year. She received an Honorary Doctorate (Hedersdoktor) from the Royal Institute of Technology (KTH) Stockholm in 2007, a Columbia Engineering School Alumni Association (CESAA) Distinguished Faculty Teaching Award in 2009 and the Janette and Armen Avanesians Diversity Award in 2018. She received the IEEE James L. Flanagan Speech and Audio Processing Award in 2011, and the ISCA Medal for Scientific Achievement in 2011. She was selected as the Fillmore Professor at the LSA Summer Institute in 2017.

\*\*\*\*\*

**Melanie Mitchell** is Professor of Computer Science at Portland State University, and External Professor and Member of the Science Board at the Santa Fe Institute. She has held faculty or professional positions at the University of Michigan, the Santa Fe Institute, Los Alamos National Laboratory, the OGI School of Science and Engineering, and Portland State University. She is the author or editor of six books and many scholarly papers in the fields of artificial intelligence, cognitive science, and complex systems. Her book *Complexity: A Guided Tour* (Oxford, 2009) won the 2010 Phi Beta Kappa Science Book Award. It was also named by Amazon.com as one of the ten best science books of 2009. Melanie originated the Santa Fe Institute's *Complexity Explorer* project, which

offers online courses and other educational resources related to the field of complex systems. Melanie's newest book, *Artificial Intelligence: A Guide for Thinking Humans*, will be published by Farrar, Straus, and Giroux in early 2019. Melanie's current research is on integrating deep learning with probabilistic models and symbolic knowledge for the task of recognizing visual situations in images.

\*\*\*\*\*

**Melanie Moses** is a Professor of Computer Science at the University of New Mexico. She studies complex biological and information systems, the scaling properties of networks, and the general rules governing the acquisition of energy and information in complex adaptive systems. She uses computational and mathematical models to understand scaling properties and search processes in distributed biological and computational systems including ant colonies, immune systems, and robot swarms.

\*\*\*\*\*

**Michael Strevens** was born and raised in New Zealand. He moved to the US in 1991 to study for a PhD at Rutgers University; currently, he teaches philosophy of science at New York University. His academic work is principally concerned with the nature of science, covering topics such as scientific explanation, understanding, complex systems, probability of various sorts, causation, and the social structure of science; he also applies contemporary research in cognitive psychology to explain aspects of both philosophical and scientific thinking. His book *\_Thinking Off Your Feet: How Empirical Psychology Vindicates Armchair Philosophy\_* is forthcoming with Harvard University Press in Fall 2018, and he is completing a trade book that explains why science is so successful at creating knowledge and why it took so long for humans to figure out how to do it right. The most recent interest of his interests is the nature of understanding and the question whether there is any meaningful difference between the understanding provided by the humanities and that provided by the social sciences.

\*\*\*\*\*

**Mirta Galesic** is Professor and Cowan Chair in Human Social Dynamics at the Santa Fe Institute, and Associate Researcher at the Harding Center for Risk Literacy at the Max Planck Institute for Human Development in Berlin, Germany. She studies how simple cognitive mechanisms interact with our social networks to produce complex social phenomena. In one line of research, she investigates how apparent cognitive biases in social judgments emerge as a product of the interplay of well-adapted minds and the statistical structure of social environments. In another, she studies how collective performance depend on the interaction of group decision strategies and network structures. A third line of research investigates opinion dynamics in real-world societies using cognitively-enriched models from statistical physics. And, she studies how people understand and cope with uncertainty and complexity inherent in many everyday decisions about health, financial, or environmental decisions.

\*\*\*\*\*

**Percy Liang** is an Assistant Professor of Computer Science at Stanford University (B.S. from MIT, 2004; Ph.D. from UC Berkeley, 2011). His research spans machine learning and natural language processing, with the goal of developing trustworthy agents that can communicate effectively with people and improve over time through interaction. Specific topics include question answering, dialogue, program induction, interactive learning, and reliable machine learning. His awards include the IJCAI Computers and Thought Award (2016), an NSF CAREER Award (2016), a Sloan Research Fellowship (2015), and a Microsoft Research Faculty Fellowship (2014).

\*\*\*\*\*

**Rodney Brooks** is the Panasonic Professor of Robotics (emeritus) at MIT. His academic research has been in computer vision, robotics, and artificial life. From 1997 until 2007 he was director of the MIT AI Lab, and then in 2003 the founding director of the MIT Computer Science and Artificial Intelligence Lab (a reconstitution of Project MAC from 1963).

In the 1980s he developed the behavior based approach to Artificial Intelligence, and has commercialized that innovation. He was co-founder, Chairman and CTO of iRobot until 2008, and since then has been founder, Chairman and CTO of Rethink Robotics. He has delivered more than 20 million robots, mostly to clean floors in homes, but also to handle roadside IEDs in war zones, shut down and clean up Fukushima Dai-Ichi, roam about on Mars, and work in close proximity to humans in factories. In 2018 he made a move to be a full-time blogger and thinker, based in San Francisco.

He is a member of the National Academy of Engineering and the American Academy of Arts and Sciences and is a fellow of the IEEE, ACM, AAAI, and AAAS. His blog can be found at [rodnebrooks.com/blog](http://rodnebrooks.com/blog), and he tweets at @rodnebrooks.

\*\*\*\*\*

**Tom M. Mitchell** is the E. Fredkin University Professor at Carnegie Mellon University, where he founded the world's first Machine Learning Department. Mitchell's research lies in machine learning, artificial intelligence, cognitive neuroscience, and the impact of AI on society. Over the past decade he has studied knowledge representations in both the human brain (e.g., using brain imaging to study representations of words and sentences), and computers (e.g., developing computer programs to understand language). Mitchell thinks recent progress in deep network learning holds important clues to how intelligent systems might learn pragmatic knowledge representations that support practical inference in ways that general theorem provers do not.

Mitchell is a member of the U.S. National Academy of Engineering, the American Association of Arts and Sciences, and is Past President of the Association for the Advancement of Artificial Intelligence. In 2015 he was awarded an Honorary Doctor of Laws Degree from Dalhousie University for his contributions to machine learning and cognitive neuroscience.

\*\*\*\*\*

**Yarden Katz** is a departmental fellow in Systems Biology at Harvard Medical School and an affiliate of the Berkman Klein Center for Internet & Society at Harvard University. He received his PhD from MIT in Brain and Cognitive Sciences. His work explores the interface between biology and cognition.