WORKSHOP ON INTELLIGENT HUMAN MACHINE COLLABORATION

June 12-14, 2012 • Washington, DC

A Meeting Organized by The Board on Global Science and Technology of The National Research Council

A MEETING SPONSORED BY: The Board on Global Science and Technology THE NATIONAL ACADEMIES *Advisers to the Nation on Science, Engineering, and Medicine*

Table of Contents

The Board on Global Science and Technology (BGST)

The Board on Global Science and Technology (BGST), a joint activity of Policy and Global Affairs Division and the Division on Engineering and Physical Sciences, is a Board of the National Research Council. It is charged with overseeing activities to assess the implications for U.S. policy makers of global advances in science and technology. To that end, BGST has established a program of workshops and other convening activities, both within the United States and overseas, to build and sustain an international, interactive community of scientists, engineers, medical and health researchers, and entrepreneurs who are engaged in the research and development of emerging technologies.

In addition, BGST has developed related web-based activities for rapid public dissemination of information resulting from the workshops, and for sustaining the interactive engagement within the scientific and technical communities, thereby enhancing global transparency with regard to scientific and technological advances.

The U.S. National Academies

The National Academies are made up of four organizations— the National Academy of Sciences, the National Academy of Engineering, the Institute of Medicine and the National Research Council. These four organizations are part of a private, nonprofit institution that provides science, technology and health policy advice under a congressional charter signed by President Abraham Lincoln that was originally granted to the National Academy of Sciences in 1863. The operating arm of the National Academies is the National Research Council (NRC). NRC's mission is to improve government decision making and public policy, increase public education and understanding, and promote the acquisition and dissemination of knowledge in matters involving science, engineering, technology, and health. The NRC does not receive direct federal appropriations for its work. Individual projects are funded by federal agencies, foundations, other governmental and private sources, and the institution's endowment. The work is made possible by 6,000 of the world's top scientists, engineers, and other professionals who volunteer their time without compensation to serve on committees and participate in activities. The NRC is administered jointly by the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine.

Board on Global Science and Technology
THE NATIONAL ACADEMIES
Advisers to the Nation on Science, Engineering, and Medicine

Workshop on Intelligent Human-Machine Collaboration

This meeting will bring together an international community of researchers from science and engineering disciplines (e.g., robotics, human-robotic interaction, software agents/multi-agents, social/cognitive sciences, and human-machine teamwork) to identify current technological research challenges that have yet to be overcome in intelligent human-machine collaboration in dynamic, unstructured environments. This workshop will also address anticipated applications of human and machine teamwork and the impact this could have on future generations of intelligent human-machine collaboration.

This meeting was organized by The Workshop on Intelligent Human-Machine Collaboration Planning Committee, an ad hoc committee of the National Research Council. The members of this committee include:

Dr. Jeffrey Bradshaw	Florida Institute for Human and Machine Cognition
Dr. Dianne Chong	The Boeing Company
Dr. Gal Kaminka	Bar Ilan University
Dr. Geert-Jan Kruijff	German Research Centre for Artificial Intelligence (DFKI)
Dr. Brian Williams	Massachusetts Institute of Technology

Principal Project Staff
Dr. William (Bill) Berry, BGST Director
Dr. Patricia Wrightson, BGST Associate Director
Dr. Ethan Chiang, BGST Program Officer
Neeraj P. Gorkhaly, BGST Research Associate

Workshop on Intelligent Human-Machine Collaboration

Planning Committee Biographical Sketches

Jeffrey M. Bradshaw, Ph.D., is a Senior Research Scientist at the Florida Institute for Human and Machine Cognition, IHMC, (Pensacola, FL) where he leads the research group developing the KAoS policy and domain services framework. Though his earliest publications were focused on memory and language, Dr. Bradshaw's research focus soon turned to a wide variety of topics relating human and machine intelligence. With Ken Ford, he edited the seminal volume Knowledge Acquisition as a Modeling Activity, and became well-known for his role in helping develop a suite of successful methodologies and tools for automated knowledge acquisition (ETS, Aquinas, Axotl, Canard, DDUCKS, eQuality). While at Boeing, he also led groundbreaking industry-wide efforts in aviation safety and training technologies, founding the emerging technologies group of the Aviation Industry Computer-Based Training Committee (AICC). He also provided technical leadership for a suite of projects to improve long-term follow-up care delivery for bone-marrow transplantation at the Fred Hutchinson Cancer Research Center. Dr. Bradshaw has helped pioneer the research area of multi-agent systems, and his first book on the topic, Software Agents, became a classic in the field and a best-seller for The MIT Press. At IHMC, he has further broadened his research interests and is currently involved in research on topics such as policy-based coordination of joint activity in humans and machines, Semantic Web technologies, adjustable autonomy and mixed-initiative interaction, cognitive systems, biologically-inspired security, visualization and performance support for complex analysis problems, network science, and augmented cognition.

Dianne Chong, Ph.D., is the Vice President of Materials and Process Planning in the Boeing Engineering, Operations &Technology organization (Bellevue, WA). In this position she leads the organization responsible for development and support of materials & manufacturing processes for the Boeing Enterprise. Prior to this she was the Director of Materials & Process Technology for Boeing Commercial Airplanes. Dr. Chong was also the Director of Strategic Operations and Business for IDS Engineering. In this capacity, she was the lead director defining and implementing a solid strategy for all Boeing Engineering. She has also been the Department Head / team leader of MSE, liaison, and process control groups in Phantom Works and Integrated Defense Systems. Dr. Chong received Bachelors degrees in biology and psychology from the University of Illinois. She also earned Masters degrees in physiology and metallurgical engineering. In 1986, Dr. Chong received her Ph.D. in Metallurgical Engineering from the University of Illinois. She also completed an Executive Master of Manufacturing Management at Washington University. Dr. Chong has served as the St. Louis representative to Military Handbook 5 where she has chaired the Aerospace Users' Group and the titanium casting group. Dr. Chong is a member of TMS, AIAA, ASM International, SME, SWE, Beta Gamma Sigma, and Tau Beta Pi. She has been recognized for managerial achievements and as a diversity change agent. She was also recognized as an outstanding alumna of The University of Illinois in 2006. Dr. Chong is a member of the National Materials Advisory Board. She has served on the Board of Trustees and is a Fellow of the ASM International. In 2007-08, she served as the President of ASM International.

Gal Kaminka, Ph.D., is an associate professor at the computer science department, and the brain sciences research center, at Bar Ilan University (Ramat Gan, Israel). His research expertise includes multi-agent and multi-robot systems, teamwork and coordination, behavior and plan recognition, and modeling social behavior. He received his Ph.D. from the University of Southern California (2000), and spent two years as a post-doctorate fellow at Carnegie Mellon University. Today, Prof. Kaminka leads the MAVERICK research group at Bar Ilan, supervising over a dozen MSc. and Ph.D. students. He was awarded an IBM faculty award and top places at international robotics competitions. He served as the program chair of the 2008 Israeli Conference on Robotics, and the program co-chair of the 2010 Int'l Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS). He has served on the international executive bodies of IFAAMAS (International Foundation of Autonomous Agents and Multi-Agent Systems) and AAAI (Association for Advancement of Artificial Intelligence). Currently, he is spending his sabbatical as a Radcliffe Fellow, at Harvard University's Radcliffe Institute for Advanced Study.

Geert-Jan Kruijff, Ph.D., graduated as an engineer from the University of Twente, combining Computer Science and Analytical Philosophy. Because of his interest in computational linguistics and dialogue systems, he ended up in Prague doing a Ph.D. in mathematical linguistics, finishing in 2001. While at Prague, he spent two years in Edinburgh, doing a lot of work on categorical grammar and description logic-like semantics. From 2001 until 2004 Dr. Kruijff worked at the Department of Computational Linguistics at Saarland University, working on grammar formalisms and building LEGO robots. In late 2004 he turned building talking robots into a full-time job at DFKI GmbH (Kaiserslautern, Germany), leading several international projects. Their Talking Robots group is now about 14 people strong, plus MSc students. Since early 2010 he has been the coordinator for a large international project on HRI and urban search & rescue.

Brian C. Williams, Ph.D., is a Professor of Aeronautics and Astronautics at the Massachusetts Institute of Technology (Cambridge, MA), and a member of the Computer Science and Artificial Intelligence Laboratory. He received his S.B., S.M. and Ph.D. from MIT in CS and AI in 1989. His research focuses on autonomous explorers, multi-robot coordination and human-robot teamwork, enabled through the of use model-based programming methods and highly deductive, reactive execution kernels. Prof. Williams is a pioneer in the fields of model-based autonomy, model-based diagnosis and qualitative reasoning, with extensive experience developing automated reasoning methods for model-based planning, execution, diagnosis and repair, methods for hybrid estimation diagnosis and control, and for intent recognition. At Xerox PARC from 1989 to 1994, Prof. Williams co-developed the GDE and Sherlock model-based diagnosis systems, which have served as the foundation for many practical systems in model-based diagnosis. He pioneered model-based autonomy in the 90's through the Livingstone model-based health management and Burton model-based execution systems at NASA and the Titan, Kirk, Sulu and Kongming model-based execution systems at MIT. At NASA Ames from '94 to '99 he formed the Autonomous Systems area, co-invented the Remote Agent autonomous control system, and co-led the flight demonstration of the Livingstone fault management system on the NASA Deep Space One probe. Prof. Williams is an AAAI fellow and AIAA associate fellow. He was a member of the Caltech JPL Advisory Council and JPL Technical Division Advisory Board, and was a member of the Young Panel, which assessed future Mars missions in light of the Mars Climate Orbiter and Polar Lander incidents. Prof. Williams is co-chair of ICAPS-12 and AAAI Compsust-12, as well as having been co-chair of Compsust-10 and AAAI Compsust-11. He has been a guest editor for Artificial Intelligence, and on the editorial boards of the Journal of Artificial Intelligence Research, the Journal of Field Robotics and AI Press.

Workshop on Intelligent Human-Machine Collaboration

Staff Biographical Sketches

William O. (Bill) Berry, Ph.D.**,** is Director of the Board on Global Science and Technology at the National Academies' National Research Council. The board's charge is to oversee activities to assess the implications for U.S. policy makers of global advances in science and technology. Previously Dr. Berry was a Senior Research Scientist in the Department of Biology at the University of Maryland and a Distinguished Research Fellow at the Center for Technology and National Security Policy, National Defense University where he taught and conducted research at the interface of science, technology and national security policy. At the Department of Defense, he served as Deputy Under Secretary of Defense for Laboratories and Basic Sciences and Director for Basic Research where he was responsible for providing scientific leadership, management oversight, policy guidance and coordination of the basic research programs of the military and defense research agencies. In addition, he was responsible for science, technology, engineering and mathematics education and workforce issues, policy for grants, policy for defense laboratories and international science and technology programs. Dr. Berry served as the U.S. Principal Member on the NATO Research and Technology Board and as the U.S. Washington Deputy for The Technical Cooperation Program (TTCP) with Australia, Canada, New Zealand and the United Kingdom. He served as the Associate Deputy Assistant Secretary of the Air Force for Science Technology and Engineering and Director of the Washington Office of the Air Force Research Laboratory following a long career at the Air Force Office of Scientific Research where he directed the Air Force's basic research programs in chemistry, life, and environmental sciences. Dr. Berry's research publications are in the fields of environmental toxicology, neuroscience, international collaboration, science and technology policy and national security policy. Dr. Berry earned a BS in Biology from Lock Haven University, a MAT in Zoology from Miami University, and a Ph.D. in Zoology from the University of Vermont.

Patricia Wrightson, Ph.D., is Associate Director of the Board on Global Science and Technology, a board of the National Academies that explores the global impact of emerging technologies. She came to the National Academies in 2000 to organize the Workshop on Scientific Communication and National Security and in 2003 became the co-staff director of the Roundtable on Scientific Communication and National Security, a joint project of the National Research Council and the Center for Strategic and International Studies, and subsequently became the director of the NRC's Committee on Scientific Communication and National Security. Dr. Wrightson was the study director of *Beyond "Fortress America": National Security Controls on Science and Technology in a Globalized World (2009)* and of *Challenges to Securing the Homeland: The Impact of Export Controls on the Department of Homeland Security (2012)*. She also directed Global Dialogues on Emerging Science and Technology, a series of four international workshops on fundamental research in diverse emerging areas of science. At the National Academies, Dr. Wrightson has also worked on South Asia-related S&T issues and on arms control. Before coming to the National Academies she taught at Georgetown University in the Department of Government and the School of Foreign Service. She has also taught at the Elliot School of International Affairs and at the U.S. Naval Academy. Currently, Dr. Wrightson is teaching Ethics in a Globalized World at Georgetown University's Public Policy Institute.

Ethan Chiang, Ph.D., is a Program Officer with the Board on Global Science and Technology. Dr. Chiang is currently the lead staffer for an upcoming report sponsored by the Department of Defense that assesses the global S&T landscape for responding to the challenges of sustaining historical trends in computing performance and to the challenge presented by the shift to multicore processors. Prior to this, he served as a staff officer for the 2012 NRC study, *Export Control Challenges Associated with Homeland Security*, and for the 2010-2011 BGST Workshop Series on Data Intensive Sciences. Dr. Chiang first came to the National Academies in 2009 as a Christine Mirzayan Science and Technology Policy Fellow, where he worked with the Committee on Scientific Communication and National Security, a standing committee of the National Research Council. Dr. Chiang's research publications are in the field of Chemistry, Chemical Biology, Nanobiotechnology, Immunology, Materials Science, and Polymer Chemistry. Dr. Chiang holds a Ph.D. and M.S. in Chemistry and Chemical Biology from Cornell University and a B.A. in Chemistry from Whitman College. During his graduate studies, he received Molecular Biophysics and Nanobiotechnology training grants from the National Institute of Health and the National Science Foundation. He was also the recipient of a Cornell University – Rockefeller University – Sloan Kettering Cancer Institute Tri-Institutional Chemical Biology Fellowship.

Neeraj P. Gorkhaly is a Research Associate with the Committee on Science, Engineering, and Public Policy and BGST. He grew up in Kathmandu, Nepal and completed his B.A. in Economics and International Studies from The Ohio State University. He was a John Glenn Policy Fellow at the World Bank and Ohio State representative for the NAFTA Study Abroad Program at El Colegio De Postgraduados En Ciencias Agrarias, Texecoco, Mexico. He currently chairs the board of directors for VENT publication, an e-magazine promoting citizen journalism and also serves on the board for The SaTTya program, both in Nepal. He occasionally mentors John Glenn Policy Fellows in DC. Neeraj has been at The National Academies since 2004.

Workshop Participant Roster

Michael Beetz
Technische Universitat Muenchen, Germany

Jeffrey Bradshaw
Inst. for Human & Machine Cognition, USA

Xiaoping Chen
University of Science and Technology of China

Frank Dignum
Utrecht University, The Netherlands

Terry Fong
NASA Ames, USA

Michael Freed
SRI International, USA

Tal Oron-Gilad
Ben-Gurion University, Israel

Michael Goodrich
Brigham Young University, USA

Robert Hoffman
Inst. for Human & Machine Cognition, USA

Andreas Hofmann
Vecna Technologies, USA

Geert-Jan Kruijff
German Research Center for Artificial Intelligence (DFKI), Germany

Paul Maglio
IBM Research, USA
University of California, Merced, USA

Alex Morison
The Ohio State University, USA

Don Mottaz
The Boeing Company, USA

Yukie Nagai
Osaka University, Japan

Daniele Nardi
University of Rome, Italy

Mark Neerincx
Delft University of Technology, The Netherlands

Lin Padgham
Royal Melbourne Institute of Technology (RMIT), Australia

Sarvapali Ramchurn
Southampton University, UK

Matthias Scheutz
Tufts University, USA

Jean Scholtz
Pacific Northwest National Laboratory, USA

Dirk Schulz
Fraunhofer Institute for Communications, Information Processing and Ergonomics, Germany

Lakmal Seneviratne
Khalifa University, UAE
King's College London, UK

Candy Sidner
Worcester Polytechnic University, USA

Liz Sonenberg
University of Melbourne, Australia

Satoshi Tadokoro
Tohoku University, Japan

Manuela Veloso
Carnegie Melon University, USA

Rong Xiong
Zhejiang University, China

Tom Wagner
iRobot, USA

Brian Williams
Massachusetts Institute of Technology, USA

Holly Yanco
University of Massachusetts Lowell, USA

Workshop on Intelligent Human-Machine Collaboration

Biographical Sketches of Workshop Participants

Michael Beetz, Ph.D., is a professor of Computer Science in the Department of Informatics at the Technische Universitat Muenchen (Munich, Germany) and heads the intelligent Autonomous Systems Group. Between 2006 to 2011, he was Vice Coordinator of the German national cluster of excellence, CoTeSys (Cognition for Technical Systems), where he coordinated the research area, "Knowledge and Learning". Dr. Beetz received his diploma degree in Computer Science with distinction from the University of Kaiserslautern. He received his MSc, MPhil, and Ph.D. degrees from Yale University in 1993, 1994 and 1996, and his Venia Legendi from the University of Bonn in 2000. Dr. Beetz was a member of the steering committee of the European network of excellence in AI planning (PLANET) and coordinated the research area "Robot Planning". He is associate editor of the AI Journal. His research interests include: plan-based control of robotic agents, knowledge processing and representation for robots, integrated robot learning, and cognitive perception. Email: Michael.beetz@in.tum.de.

Jeffrey M. Bradshaw, Ph.D.**,** is a Senior Research Scientist at the Florida Institute for Human and Machine Cognition, IHMC, (Pensacola, FL) where he leads the research group developing the KAoS policy and domain services framework. Though his earliest publications were focused on memory and language, Dr. Bradshaw's research focus soon turned to a wide variety of topics relating human and machine intelligence. With Ken Ford, he edited the seminal volume Knowledge Acquisition as a Modeling Activity, and became well-known for his role in helping develop a suite of successful methodologies and tools for automated knowledge acquisition (ETS, Aquinas, Axotl, Canard, DDUCKS, eQuality). While at Boeing, he also led groundbreaking industry-wide efforts in aviation safety and training technologies, founding the emerging technologies group of the Aviation Industry Computer-Based Training Committee (AICC). He also provided technical leadership for a suite of projects to improve long-term follow-up care delivery for bone-marrow transplantation at the Fred Hutchinson Cancer Research Center. Dr. Bradshaw has helped pioneer the research area of multi-agent systems, and his first book on the topic, Software Agents, became a classic in the field and a best-seller for The MIT Press. At IHMC, he has further broadened his research interests and is currently involved in research on topics such as policy-based coordination of joint activity in humans and machines, Semantic Web technologies, adjustable autonomy and mixed-initiative interaction, cognitive systems, biologically-inspired security, visualization and performance support for complex analysis problems, network science, and augmented cognition. Email: jbradshaw@ihmc.us.

Xiao-Ping Chen, Ph.D., is a full professor in the School of Computer Science and Technology and the Director of the Center for Artificial Intelligence Research at the University of Science and Technology of China (Hefei, China). He also serves as a trustee of the International RoboCup Federation, a member of the Editorial Board of the Journal of Artificial Intelligence Research, a member of the Editorial Board of the Knowledge Engineering Review, and the Chair of the Chinese National RoboCup Committee. He was a Senior Program Committee member of IJCAI'11, AAAI'11 and AAMAS'12, Co-chair of RoboCup 2008. Prof. Chen has been working in the fields of Artificial Intelligence and Autonomous Robotics. He established and has led the USTC Multi-Agent Systems Lab and robot team, WrightEagle, which won 5 champions and 9 runners-up in RoboCup world championships. Prof. Chen found and has led the KeJia Project, a long-term Cluster of Excellence at USTC which aims at developing human-level intelligent robots to serve people under real-world environments. Prof. Chen won the USTC President Award for Research Excellence in 2010. Email: xpchen@ustc.edu.cn.

Frank Dignum, Ph.D., is associate professor of Computer Science at Utrecht University (The Netherlands). He is currently also honorary senior research fellow at the University of Melbourne. He leads research in adaptive interactive systems. He supervises three Ph.D. students and one postdoc. The projects concern the use of agent technology for games and simulations, with the aim of making the games and simulations more flexible and adaptive to the user, In particular, where agents incorporate cultural and personality aspects. He is the technical coordinator of the EU projects. He has obtained numerous national and international grants leading to around 30 researchers working (or having worked) on his grants. He is the initiator and organizer of the international workshop on Agents for Games and Simulations. Besides this workshop he has organized other agent-related workshops and conferences and was local organizer of the main agent conference (AAMAS) in 2005. He was area chair of IJCAI 2011 and workshops chair of AAMAS 2011. He is on the (senior) program committee of many workshops and conferences. He is a keynote and invited speaker at conferences all over the world. He has written over 250 papers in international journals and conferences and has an H-index of 41. Email: f.p.m.dignum@uu.nl.

Terry Fong, Ph.D., is the Director of the Intelligent Robotics Group at the NASA Ames Research Center (Moffett Field, CA). From 2002 to 2004, Dr. Fong was the deputy leader of the Virtual Reality and Active Interfaces Group at the Swiss Federal Institute of Technology (EPFL). From 1997 to 2000, he was Vice President of Development for Fourth Planet, Inc., a developer of real-time visualization software. Dr. Fong has published more than one hundred papers in field robotics, human-robot interaction, virtual reality user interfaces, and parallel processing. Dr. Fong received his B.S. and M.S. in Aeronautics and Astronautics from the Massachusetts Institute of Technology and his Ph.D. in Robotics from Carnegie Mellon University. Email: terry.fong@nasa.gov.

Michael Freed, Ph.D., is a research program director for artificial intelligence at SRI International (Menlo Park, CA) where he leads mainly DARPA-funded projects focused on problems such as email overload and allowing end users to modify software by teaching it as they would a human worker. Before SRI, he spent 12 years at NASA running the Apex Lab, which pioneered new techniques in computational human factors, unmanned aerial surveillance and open-source toolkits for developing intelligent agents. Dr. Freed earned his Ph.D. for Northwestern University in 1998. Email: freed@ai.sri.com

Tal Oron-Gilad, Ph.D., is the Head of the Human Factors Engineering Graduate program at the Department of Industrial Engineering Management at Ben-Gurion University (Beer-Sheva, Israel). She has a MSc. degree in Industrial Engineering from the Technion, and a Ph.D. from Ben-Gurion University. Previously, she was a research associate for several years at the University of Central Florida. Email: orontal@bgu.ac.il.

Michael A. Goodrich, Ph.D., is a Professor of Computer Science at Brigham Young University (Provo, UT). Prior to this, he spent one year as a research assistant professor in the Computer Science Department at BYU. Dr. Goodrich received his Ph.D. at Brigham Young University in 1996 from the Electrical and Computer Engineering Department under the direction of Dr. Wynn Stirling. Following graduation, he completed a two year postdoctoral research associate position at Nissan CBR in Cambridge, Massachusetts. His doctoral work was in intelligent control, and his post-doctoral research was in computational models of intelligent human behavior. Email: mike@cs.byu.edu.

Andreas Hofmann, Ph.D., works in the area of reactive robotic motion control systems, a field at the nexus of artificial intelligence, machine learning, expert systems and robotics. His expertise focuses on advanced planning and motion control tasks for humanoid robots, with an emphasis on balance control in the presence of disturbances. He is also focusing on the development of robots that can do significant mechanical work in unstructured environments, particularly for logistics and manufacturing applications. Dr. Hofmann is Vice-president of Autonomous Systems at Vecna Technologies, a robotics company in Cambridge, MA. He is also a research scientist in the MIT MERS lab, led by Prof. Brian Williams. Previously, Dr. Hofmann co-founded Gensym Corp., a company specializing in automatic, supervisory-level control software for the chemical process and manufacturing industries. Among other duties, he was responsible for the development of key components of Gensym's flagship G2 real-time expert system. Dr. Hofmann received a bachelor's degree in Electrical Engineering at MIT, a masters in Electrical Engineering from Rensselaer Polytechnic Institute, and a Ph.D. in computer science from MIT. Email: ahofmann@vecna.com.

Robert Hoffman, Ph.D., is recognized as one of the world leaders in the fields of Cognitive Systems Engineering and Human-Centered Computing. He is a Fellow of the Association for Psychological Science and a Fulbright Scholar. His Ph.D. is in experimental psychology from the University of Cincinnati, where he received McMicken Scholar, Psi Chi, and Delta Tau Kappa Honors. Following a Postdoctoral Associateship at the Center for Research on Human Learning at the University of Minnesota, Hoffman joined the faculty of the Institute for Advanced Psychological Studies at Adelphi University. He began his career as a psycholinguist, and founded the journal, *Metaphor and Symbol*. His subsequent research leveraged the psycholinguistics background in the study of methods for eliciting the knowledge of domain experts. Hoffman has been recognized internationally in disciplines including psychology, remote sensing, weather forecasting, and artificial intelligence; for his research on human factors in remote sensing; for his work in the psychology of expertise and the methodology of cognitive task analysis; and for his work on HCC issues in intelligent systems technology and the design of macrocognitive work systems. Hoffman is a Co-Editor for the Department of Human-

Centered Computing in *IEEE: Intelligent Systems*. He is Editor for the book Series, "Expertise: Research and Applications." He has also co-founded the *Journal of Cognitive Engineering and Decision Making.* His major current projects involve evaluating the effectiveness of knowledge management, and performance measurement for macrocognitive work systems. A full vita and all of his publications are available for download at [www.ihmc.us/users/rhoffman/main]. Email: rhoffman@ihmc.us.

Geert-Jan Kruijff, Ph.D., graduated as an engineer from the University of Twente combining Computer Science and analytical Philosophy. Because of his interest in computational linguistics and dialogue systems, he ended up in Prague doing a Ph.D. in mathematical linguistics, finishing in 2001. While in Prague, he also spent two years in Edinburgh, working on categorial grammar and description logic-like semantics. From 2001 until 2004 Dr. Kruijff worked in the Department of Computational Linguistics at Saarland University on grammar formalisms and building LEGO robots. In late 2004 he turned building talking robots into a full-time job at DFKI GmbH (Kaiserslautern, Germany), leading several international projects. Their Talking Robots group is now about 14 people strong, plus MSc. students. Since early 2010 he has been the coordinator for a large international project on HRI and urban search and rescue. Email: gj@dfki.de.

Paul P. Maglio, Ph.D., is a research scientist at IBM Research – Almaden (San Jose, CA), and a Professor of Technology Management at the University of California, Merced. He has published more than 100 refereed papers in various areas of cognitive science, computer science, human-computer interaction, and management. In the area of HCI, Dr. Maglio has worked on attentive interfaces, peripheral information, and systems management. His book on computer systems management, co-authored with colleagues at IBM Research, "Taming information technology: Lessons from studies of computer systems administrators", will be published by Oxford University Press this summer. A founder of the field of service science, Dr. Maglio serves on editorial boards for the Journal of Service Research (Sage) and Service Science (INFORMS), is lead editor of the "Handbook of Service Science" (Springer), and has co-chaired many related conferences, including Art and Science of Service (2011), International Conference on Service Oriented Computing (2010), and Frontiers in Service (2007). Dr. Maglio holds an S.B. in computer science and engineering from MIT and a Ph.D. in cognitive science from UCSD. Email: pmaglio@ucmerced.edu.

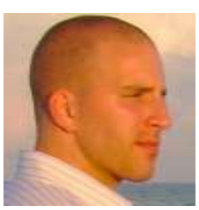

Alexander Morison, Ph.D., is a Research Scientist in the Integrated System Engineering Department at The Ohio State University (Columbus, OH). Currently, Dr. Morison is also acting as a technical adviser and Research Scientist with the 711th Human Performance Wing, Air Force Research Laboratory under an Intergovernmental Personnel Act Mobility Program (IPA). Dr. Morison completed his Ph.D. in Cognitive Systems Engineering in 2010. Prior to completing his doctorate, Dr. Morison completed a B.S. in Electrical Engineering and Applied Physics from Case Western Research University and an M.S. in Computer Science and Engineering from The Ohio State University. Dr. Morison studies the growing challenge of coupling human observers to remote sensor systems (Morison et. al., 2009; Morison and Woods, in progress). Specifically, he is using fundamentals from human visual perception to define a new paradigm in human-robot interaction called extending perception (Morison, 2010).

Inspired by models of human perception and attention, he has invented solutions to the image overload, keyhole effect (Morison, Woods, and Davis, 2009), and multiple feeds problems associated with layered sensing systems and mobile sensor platforms. He is currently developing new devices (Morison, Woods, and Roesler, 2009) and algorithms (Davis, Morison, and Woods, 2007) that implement the basic principles of perspective control for surveillance systems, human-robot interaction, and layered sensing systems to help human decision makers use of the large volume and flux of data these systems generate. Email: morison.6@osu.edu.

Don Mottaz is the Director for Assembly and Integration for Boeing Research and Technology (BR&T) (Bellevue, WA). In his current assignment he leads a team of 300 engineers and technicians with responsibility for the development and implementation of technologies used for the assembly of Boeing products. This includes robotic/automated assembly systems for component buildup, and wings and fuselages assembly, the application of smart tools and information systems to the factory floor to improve assembly methods and situational awareness in manufacturing, new drilling and fastener systems, augmented reality technologies and factory infrastructure projects to improve part logistics and delivery throughout the supply chain. Prior to his current assignment, Don served as the Director for Materials and Process Technology (M&PT) support to BCA programs and M&PT Chief Engineer for BCA Product Development. Don has held assignments supporting all areas of commercial aircraft manufacturing and design, as well as leading the M&PT Technology Development activities. Don's 33 plus years experience covers all areas of manufacturing operations from detail part manufacturing, chemical finishing and processing, sub structure assembly and major aircraft structure join, systems installation, functional test and final paint and delivery. Don has extensive experience in the development and introduction of new materials and process into the manufacturing of commercial airplanes. He has a Bachelor of Science in Chemical Engineering and an MBA from the University of Washington in Seattle. Email: Donald.a.mottaz@boeing.com.

Yukie Nagai, Ph.D., is a Specially Appointed Associate Professor of Osaka University (Japan). She received a Ph.D. in Engineering from Osaka University in 2004 and worked as a Research Associate at Osaka University from 2002 to 2004. From 2004 to 2006, she was a Postdoc Researcher at the National Institute of Information and Communications Technology in Kyoto, Japan. From 2006 to 2009, she was a Postdoc Researcher of the research group for Applied Computer Science and of Research Institute for Cognition and Robotics at Bielefeld University, Germany. Email: yukie@ams.eng.osaka-u.ac.jp.

Daniele Nardi, Ph.D., is Full Professor at Facoltà Ingegneria, Sapienza Universita di Roma (Rome, Italy) in the Dipartimento Informatica e Sistemistica, since 2000. His current research interests are in Artificial Intelligence, Cognitive Robotics, Multi-Agent/Multi-Robot Systems and Search and Rescue Robotics. He is author of more than 100 scientific publications, recipient of "IJCAI-91 Publisher's Prize" and of Prize "Intelligenza Artificiale 1993" and ECCAI Fellow. He is currently Vice President of RoboCup Federation, Coordinator of the Curricula in Computer Engineering at Sapienza Univ. Roma and Director of the research laboratory "Cognitive Robot Teams". Email: nardi@dis.uniroma1.it.

Mark Neerincx, Ph.D., is professor in Human-Machine Interaction (HMI) at The Delft University of Technology (The Netherlands), and senior researcher at TNO Perceptual and Cognitive Systems. He obtained a MSc. in Cognitive Psychology from Leiden University and a Ph.D. in Human-Computer Interaction from the University of Groningen, both in The Netherlands. Dr. Neerincx has extensive experience in fundamental and applied research on human-computer interaction, among other things in the domains of health, security, defense, and space. Important results are (1) cognitive task load and emotion models for performance and health support, (2) models of human-machine partnership for attuning assistance to the individual user and momentary usage context, (3) prototypes of electronic partners for self-efficacy, (4) cognitive engineering methods and tools, and (5) a diverse set of usability "best practices". Mark Neerincx acquired numerous—national and international—Research & Development projects and has been involved in diverse scientific and educational activities. Email: mark.neerincx@tno.nl.

Lin Padgham, Ph.D., is Professor of Artificial Intelligence in the School of Computer Science and I.T. at RMIT University (Melbourne, Australia). She obtained a Ph.D. from the University of Linkoping in 1989 (Sweden). Dr. Padgham's research interests are in various aspects of commonsense reasoning with an emphasis on formal methods for knowledge representation which are coupled with computationally realizable algorithms. She has spent more than 15 years researching intelligent multi-agent systems and has developed (with others in her group) tools and a methodology for building agent systems. She co-authored the first detailed book (published 2004) on a practical approach to building multi-agent systems. In 2005, the supporting tool for this methodology, the Prometheus Design Tool, won the award for the best demonstration at AAMAS'05. The RMIT Intelligent Agents group is internationally recognized in the area of Agent Based Software Engineering. A recent focus is in the use of agent based modeling and simulation for complex applications in policy and planning or training, and trans-disciplinary collaboration to facilitate use of these technologies to address complex social issues. Much of Dr. Padgham's current research focuses on modeling, building and understanding intelligent agents for complex application areas requiring a balance between goal directed long-term behavior and reactive response to a dynamic environment. She has investigated various extensions to standard 'Belief Desire Intention' (BDI) reasoning, including planning, goal conflicts, and learning. She has recently been investigating use of the BDI paradigm in modeling human agents in simulations. Dr. Padgham has worked with a range of Industry partners on research projects funded by national competitive grants, and has developed long standing collaborations with some of these partners. Dr. Padgham also has a successful history of obtaining national competitive grants for less applied research. She has been Program Chair for AAMAS 2008, General Chair for AAMAS 2012, is on the Editorial Board of the Journal for Autonomous Multi Agent Systems, and regularly serves on the (Senior) Program Committees for major international conferences such as KR, IJCAI and ECAI. Email: lin.padgham@rmit.edu.au.

Sarvapali D. Ramchurn, Ph.D., is a lecturer at the School of Electronics and Computer Science, University of Southampton (UK). He obtained his Ph.D. in Multi-Agent Systems under the supervision of Prof. Nick Jennings (a pioneer of multi-agent systems). Prior to his appointment as lecturer, he was a lead research fellow on the ALADDIN and IDEAS projects. Dr. Ramchurn has published over 35 peer-reviewed articles in prestigious journals (AIJ, JAIR) and conferences (AAMAS, IJCAI, AAAI) on the topics of multi-agent coordination, coalition formation, and more recently, on agent-based energy management applications. With his colleagues, he won the best paper award at AAMAS 2010 (the premier international conference) for their work on agent-based micro-storage in the smart grid. His work on agent-based demand-side management was also nominated for the best innovative agents application award at AAMAS 2011. Previously Dr. Ramchurn also led teams that won the Prisoner's dilemma competition (2004&2005), and the RobocupRescue Infrastructure competition (2007). Email: sdr@ecs.soton.ac.uk.

Matthias Scheutz, Ph.D., is an associate professor in the Department of Computer Science at Tufts University (Medford, MA). He received degrees in Philosophy (M.A. 1989, Ph.D. 1995) and Formal Logic (M.S. 1993) from the University of Vienna and in Computer Engineering (M.S. 1993) from the Vienna University of Technology (1993) in Austria. He also received the joint Ph.D. in Cognitive Science and Computer Science from Indiana University in 1999. Before moving to Tufts, Matthias was an Associate Professor of Informatics and Computer Science in the School of Informatics, Associate Professor of Cognitive Science in the Cognitive Science Program, and Adjunct Associate Professor in the Department of Psychological and Brain Sciences at Indiana University, Bloomington. He has over one hundred peer-reviewed publications in artificial intelligence, artificial life, agent-based computing, natural language processing, cognitive modeling, robotics, human-robot interaction, and foundations of cognitive science. His current research and teaching interests include multi-scale agent-based models of social behavior and complex cognitive affective robots with natural language capabilities for natural human-robot interaction. He serves as Co-Director for Computer Science in the new interdisciplinary Tufts program in cognitive and brain science. Email: mscheutz@cs.tufts.edu.

Jean Scholtz, Ph.D., has worked in the area of user-centered evaluation of technology for 20 years. She is semi-retired and currently works part-time for the Pacific Northwest National Laboratory (Richland, WA), primarily in user-centered evaluation of visual analytic systems. As such she focuses on users in law enforcement and intelligence analysts. Previously, she worked for the National Institute of Standards and Technology (NIST) where she developed metrics and methodologies for user centered evaluations of human-robot interaction and intelligence analysis programs for the Defense Department and the Intelligence Community. She is a founder and co-chair of the Visual Analytics Science and Technology (VAST) Challenge which offers participants an opportunity to use their visual analytics systems to analyze synthetic data with embedded ground truth. Dr. Scholtz's work history includes positions at Bell Telephone Laboratories and Intel Corporation. She was also on the Computer Science faculty at Portland State University, Portland, OR. Dr. Scholtz holds a Ph.D. in Computer Science from the University of Nebraska. She is a member of ACM and IEEE and has over 60 publications in user-centered evaluation. Email: jean.scholtz@pnnl.gov.

Dirk Schulz, Ph.D., is head of the research group on unmanned systems at the Fraunhofer Institute for Communications, Information Processing and Ergonomics, FKIE (Munich, Germany). His main research interest lies in the development of effective and robust state estimation, motion control, and human-robot interaction techniques for mobile ground robots carrying out service and surveillance tasks in challenging outdoor environments. Application domains of interest are, e.g., autonomous inspection of industrial plants, CBRNE reconnaissance, and multi-robot systems for surveillance. Dr. Schulz studied Computer Science at the University of Bonn. There, he also received his doctorial degree in 2002. After a one year visit to the University of Washington (Seattle, WA) he returned as a post-doctorial researcher to the University of Bonn, and then joined FKIE in 2008. Email: dirk.schulz@fkie.fraunhofer.de.

Lakmal Seneviratne, Ph.D., is the Director of the Robotics Institute at Khalifa University (Abu Dhabi, UAE). He is also Professor of Mechatronics at King's College London (UK). He was the founding Director of the Centre for Mechatronics at KCL, from 1994 till 2005. He was the Head of the Division of Engineering at KCL from 2004-2009. He is currently the Associate Provost for research and graduate studies at Khalifa University. His main research interests are centered on robotics and automation, with particular emphasis on increasing the autonomy of robotic systems interacting with complex dynamic environments. He has published over 250 peer reviewed articles on these topics. Email: Lakmal.seneviratne@kustar.ac.ae.

Candy Sidner, Ph.D., is a research professor in the Computer Science dept. at Worcester Polytechnic Institute (Worcester, MA). Dr. Sidner is interested in multi-modal computational models of verbal and non-verbal interaction, especially dialogue and discourse, with social and collaborative models, both for embedded conversational (on screen) agents and for human-robot interaction. Her current research involves creating models of interaction for agents, both intelligent virtual agents and humanoid robots, and includes modelling relationships between people and agents, when the agents are "always-on" in people's personal environments. Dr. Sidner's work includes the study of phenomena in human-human interaction, both dialogue and gestures, developing computational models and implementing and testing those models with everyday people. This work is joint with Prof. Chuck Rich at WPI and Prof. Tim Bickmore at Northeastern University. Email: sidner@spi.edu.

Liz Sonenberg, Ph.D., is a professor in the Department of Computing and Information Systems at Melbourne University (Australia), and since August 2009 has also had the part-time role of Pro Vice-Chancellor (Research Collaboration) in Melbourne Research. The integrating theme of her research is the conceptualization and construction of more adaptive, distributed, and intelligent information systems. Much of the work focuses on agent technology, which views a distributed system in terms of interacting autonomous software entities. Using the agent metaphor can allow system developers to adopt a level of abstraction in design that is useful for modeling complex tasks and environments, and in building software systems that are robust in the face of change and unexpected events. An important aspect of the research is the requirement of the human-machine interface and consequent implications for the development of computational mechanisms to support decision-making in complex settings. Her specialized interests are: Multi-agent systems - especially collaboration and teamwork, Automated negotiation and decision support, Context-aware computing and technologies for personalization, and Computational modeling of human problem solving. Email: l.sonenberg@unimelb.edu.au.

Satoshi Tadokoro, Ph.D., received a B. E. degree in precision machinery engineering in 1982, a M.E. degree in 1984 from the University of Tokyo, and a D.E. degree in 1991. He was an associate professor of Kobe University between 1993 and 2005, and since 2005 has been a professor of the Graduate School of Information Sciences (GSIS) at Tohoku University (Sendai, Japan). He was also the Deputy Dean of GSIS between 2012 and 2013. He was a project leader of the MEXT DDT Project on rescue robotics between 2002 and 2007, which includes more than 100 professors nationwide, and the NEDO Project that developed the rescue robot Quince, which has been used at the Fukushima-Daiichi Nuclear Power Plant Accident since June 2011. He established RoboCupRescue in 1999, TC on Rescue Engineering of SICE in 2000 (the first chair), IEEE Robotics and Automation Society (RAS) TC on Safety, Security and Rescue Robotics in 2001 (the first co-chair), and International Rescue System Institute (IRS) in 2002. He was IEEE RAS Japan Chapter Chair in 2003-2005, trustee of The RoboCup Federation in 2005-2010, Chair of JSME Robotics Mechatronics Division (RMD) in 2009, IEEE Robotics and Automation Society (RAS) AdCom Member in 2008-2010. He is at present President of IRS, and IEEE RAS Vice President for Technical Activities in 2012-2013. He received IEEE Fellow in 2009, JSME Fellow in 2005, SICE Fellow in 2011, The Robot Award 2008, FDMA Commissioner Highest Award in 2008, JSME Funai Award in 2007, Best Book Author Award from AEM Society in 2006, JSME Robotics and Mechatronics Award in 2011, JSME Robotics and Mechatronics Academic Achievement Award in 2005, etc. He published Rescue Robotics from Springer, RoboCupRescue from Kyoritsu Publ., etc. His research interest is in rescue robotics, virtual reality and new actuators. Email: tadokoro@rm.is.tohoku.ac.jp.

Manuela Veloso, Ph.D., is the Herbert A. Simon Professor of Computer Science at Carnegie Mellon University (Pittsburgh, PA). Her long-term research goal is the effective construction of autonomous agents where cognition, perception, and action are combined to address planning, execution, and learning tasks. Her vision is that multiple intelligent robots with different sets of complementary capabilities will provide a seamless synergy of intelligence. Dr. Veloso's research focuses on the continuous integration of reactive, deliberative planning, and control learning for teams of multiple agents acting in adversarial, dynamic, and uncertain environments. Her multi-agent and multi-robot research interests have been motivated by and experimented in the domain of robot soccer. Since 2009, she has been investigating indoor mobile, service, companion robots, CoBots, such that robots and humans interact in a symbiotic relationship building upon individual strengths and limitations. Veloso created and directs her CORAL overarching research lab for the research on intelligent agents that Collaborate, Observe, Reason, Act, and Learn. As of 2010, she has ten Ph.D. students and has graduated another twenty one Ph.D. students, whose theses are available at her website www.cs.cmu.edu/~mmv. She received her Ph.D in Computer Science from Carnegie Mellon University in 1992, M.A. in Computer Science from Boston University in 1986, and M.SC in Electrical and Computer Engineering from Instituto Superior Tecnico in 1984. Email: veloso@cmu.edu.

Tom Wagner, Ph.D., is CTO of iRobot (Bedford, MA). In this position, his areas of responsibility include technology, software and research. Previously, he served as vice president and technical director of iRobot's Government and Industrial Robots division. Prior to joining iRobot, Wagner was a program manager at the Defense Advanced Research Projects Agency (DARPA), the research and development agency of the U.S. Department of Defense, where he managed programs in robotics, communications, command and control, tele-health and artificial intelligence. Earlier in his career, Wagner served as a professor at the University of Maine, as a principal lead at Honeywell, and in various advisory and leadership roles in small companies. He holds a Ph.D. in computer science from the University of Massachusetts. Email: twagner@irobot.com.

Brian C. Williams, Ph.D., is a Professor of Aeronautics and Astronautics at the Massachusetts Institute of Technology, and a member of the Computer Science and Artificial Intelligence Laboratory. He received his S.B., S.M. and Ph.D. from MIT (Cambridge, MA) in Computer Science and Artificial Intelligence in 1989. His research focuses on autonomous explorers, multi-robot coordination and human-robot teamwork, enabled through the of use model-based programming methods and highly deductive, reactive execution kernels. Prof. Williams is a pioneer in the fields of model-based autonomy, model-based diagnosis and qualitative reasoning, with extensive experience developing automated reasoning methods for model-based planning, execution, diagnosis and repair, methods for hybrid estimation diagnosis and control, and for intent recognition. At Xerox PARC from 1989 to 1994, Prof. Williams co-developed the GDE and Sherlock model-based diagnosis systems, which have served as the foundation for many practical systems in model-based diagnosis. He pioneered model-based autonomy in the 90's through the Livingstone model-based health management and Burton model-based execution systems at NASA and the Titan, Kirk, Sulu and Kongming model-based execution systems at MIT. At NASA Ames from '94 to '99 he formed the Autonomous Systems area, co-invented the Remote Agent autonomous control system, and co-led the flight demonstration of the Livingstone fault management system on the NASA Deep Space One probe. Prof. Williams is a AAAI fellow and AIAA associate fellow. He was a member of the Caltech JPL Advisory Council and JPL Technical Division Advisory Board, and was a member of the Young Panel, which assessed future Mars missions in light of the Mars Climate Orbiter and Polar Lander incidents. Prof. Williams is co-chair of ICAPS-12 and AAAI Compsust-12, as well as having been co-chair of Compsust-10 and AAAI Compsust-11. He has been a guest editor for Artificial Intelligence, and on the editorial boards of the Journal of Artificial Intelligence Research, the Journal of Field Robotics and AI Press. Email: Williams@csail.mit.edu.

Rong Xiong, Ph.D., is the associate professor at the Institute of Cyber Systems and Control of Zhejiang University (Hangzhou, China), and the Director of the Robotics Laboratory. Her research interests include visual recognition, simultaneous localization and mapping, motion planning and control for humanoid robots. Dr. Xiong received her M.S. degree in Computer Science and Engineering in 1997, and her Ph.D. degree in Control Science and Engineering in 2009, from Zhejiang University, China. Email: rxiong@iipc.zju.edu.cn.

Holly Yanco, Ph.D., is a Professor and Associate Chair of the Computer Science Department at the University of Massachusetts Lowell. Her research interests include human-robot interaction, multi-touch computing, interface design, robot autonomy, fostering trust of autonomous systems, evaluation methods for human-robot interaction, and the use of robots in K-12 education to broaden participation in computer science. Dr. Yanco's research has been funded by the National Science Foundation, including a Career Award, the Army Research Office, Microsoft, and the National Institute of Standards and Technology. Dr. Yanco was the General Chair of the 2012 ACM/IEEE International Conference on Human-Robot Interaction. She served on the Executive Council of the Association for the Advancement of Artificial Intelligence (AAAI) from 2006-2009, was the Symposium Chair for AAAI from 2002-2005, and was the Exhibition Co-Chair of the ACM/IEEE Conference on Human-Robot Interaction from 2007-2009. She is a senior member of AAAI. Dr. Yanco is a co-developer of the Artbotics program, which combines art and robotics in programs for middle school, high school, and undergraduate students. She was the PI of the NSF-funded development of Pyro, a Python-based robotics curriculum, which was selected as the Premier Courseware of 2005 by NEEDS. Dr. Yanco has a Ph.D. and MS in Computer Science from the Massachusetts Institute of Technology (MIT) and a BA in Computer Science and Philosophy from Wellesley College. Email: holly@cs.uml.edu.

Board on Global Science and Technology
THE NATIONAL ACADEMIES
Advisers to the Nation on Science, Engineering, and Medicine

Workshop on Intelligent Human-Machine Collaboration

Agenda

Monday, June 11th

6:00 – 7:30 PM	Welcome Reception (State Plaza Hotel, Ambassador Room)

Tuesday, June 12th

Closed Session (Invited Workshop Participants Only)

8:00 – 8:30 AM	Breakfast (Room 120)
8:30 – 8:45 AM	Welcome and setting the stage
8:45 – 9:45 AM	Participant introductions
9:45 – 10: 00 AM	Introduction to scenario exercise Moderator: Brian Williams
10:00 - 12:30 PM	Breakout Groups: Real-world Applications of Intelligent Human-Machine Collaboration Scenario A: Disaster Management Moderator: Michael Goodrich Scenario B: Small Lot Agile Manufacturing Moderator: Matthias Scheutz Scenario C: Hospital Service Robotics Moderator: Candy Sidner Scenario D: Virtual Team Training Moderator: Mark Neerincx Scenario E: Personal Satellite Assistants Moderator: Terry Fong

Open Session

12:30 – 1:30 PM	Lunch
1:30 – 3:00 PM	Group discussion: Moderator: Jean Scholtz

	Breakout groups A, B, and C report back on findings from earlier scenario exercise (30 min each)
3:00 – 3:15 PM	Break (refreshments available)
3:15 – 4:15 PM	Group discussion: Moderator: Tal Oron-Gilad
	Breakout groups D and E report back on findings from earlier scenario exercise (30 min each)
4:15 – 4:30 PM	Break
4:30 - 5:30 PM	Group discussion: Moderator: Lin Padgham

- What international, global, or cross-cultural considerations were raised during your scenario discussions?
- What are the benefits of intelligent human-machine collaboration vs. traditional autonomy?
- What are some of the commonalities in human-machine issues that were raised across the scenarios?
- What are the issues that were not raised?
- What are the biggest overall research challenges? Which of these challenges would require significant breakthroughs? Which of these breakthroughs are unlikely to occur in the next ten years? In twenty years?

Based on this discussion, workshop participants will select and self-organize into 5 topics for the next day's Collaboration Panels.

5:30 PM — End of Day One. Participants should plan on making their own dinner arrangements.

Wednesday, June 13th

Closed Session (Invited Workshop Participants Only)

8:00 – 8:30 AM	Breakfast
8:30 – 8:45 AM	Reflections from Day One
8:45 – 9:45 AM	Panel Breakout Groups
	Participants will meet with their respective panels (organized the previous afternoon) to organize a 30-minute discussion. Each panel should create a PowerPoint presentation for the discussion.

Open Session

9:45 – 10:15 AM	Panel I: (30 minutes)

10:15 – 10:45 AM	Panel II: (30 minutes)
10:45 – 11:00 AM	Break (refreshments available)
11:00 – 11:30 AM	Panel III: (30 minutes)
11:30 – 12:00 AM	Panel IV: (30 minutes)
12:00 – 12:30 PM	Panel V: (30 minutes)
12:30 – 1:30 PM	Lunch
1:30 – 2:45 PM	Group discussion: Moderator: Liz Sonenberg What are the global (or trans-national) challenges that intelligent human-machine collaboration can help to solve?
2:45 – 3:00 PM	Break (refreshments available)
3:00 – 4:15 PM	Group discussion: Moderator: Manuela Veloso What kinds of breakthroughs would be game-changers for significantly improved intelligent human-machine collaboration? What are the implications of these breakthroughs for national and global security, competitiveness, and human well-being?
4:15 – 5:00 PM	Group discussion: Moderator: Geert-Jan Kruijff Summary of challenges and solutions discussed throughout workshop
5:00 PM	Group Picture at Albert Einstein Statue End of Day Two. Participants should plan on making their own dinner arrangements.

Thursday, June 14th

Selected Research Topics in Intelligent Human-Machine Collaboration (Room 125)

8:00 – 8:30 AM	Breakfast
8:30 – 10:30 AM	Session 1: Socio-Cognitive Issues Moderator: Jeff Bradshaw **Yukie Nagai**, Osaka University *Robots that Learn to Communicate with Humans* **Alex Morison**, Ohio State University *Expanding Human Perception and Attention to New Spatial-Temporal Scale Through Networks of Sensor Systems*

	Candy Sidner, Worcester Polytechnic University *Agents for Long Term Relationships with Isolated Older Adults*
	Frank Dignum, Utrecht University *Interaction in Context*
10:30 – 10:45 AM	Break (refreshments available)
10:45 – 12:15 PM	Session 2: Learning and Adaptation in Dynamic Settings Moderator: Geert-Jan Kruijff
	Michael Freed, SRI International *A Virtual Assistant for Email Overload*
	Satoshi Tadokoro, Tadokoro University *Disaster Response Robot Quince and Lessons at Fukushima-Daiichi Nuclear Power Plant Accident*
	Michael Goodrich, Brigham Young University *A New Perspective on Levels of Autonomy*
12:15 – 1:15 PM	Lunch
1:15 – 2:45 PM	Session 3: Challenging Applications Moderator: Jeff Bradshaw
	Xiao-Ping Chen, University of Science and Technology of China *An Approach to Enabling Robots to Fulfill Open-ended Requests*
	Lakmal Seneviratne, Khalifa University & King's College London *Force Feedback and Haptic Interfaces During Robot Assisted Surgical Interventions*
	Rong Xiong, Zhejiang University *A Study on Humanoid Robots Playing Table Tennis*
2:45 – 3:00 PM	Break (refreshments available)
3:00 – 4:00 PM	Session 4: Human-Machine Interaction and Teaming Moderator: Brian Williams
	Holly Yanco, University of Massachusetts Lowell *Human-in-the-Loop Control of Robot Systems*
	Tom Wagner, iRobot Corporation *Importance of Human Robot Teams For Fielding Product Grade Systems*
4:00 – 4:30 PM	Final Discussion
4:30 PM	Meeting adjourned

Scenario Breakout Groups

Scenario A. <u>Preparing for and Managing a Major Natural Disaster</u> (Room 119)

- Michael Goodrich (Moderator)
- Lin Padgham
- Satoshi Tadokoro
- Alex Morison
- Daniele Nardi
- Geert-Jan Kruijff

Scenario B. <u>Small-Lot Agile Manufacturing</u> (Room 280)

- Matthias Scheutz (Moderator)
- Xiao-Ping Chen
- Lakmal Seneviratne
- Tal Oron-Gilad
- Sarvapali Ramchurn
- Don Mottaz
- Brian Williams

Scenario C. <u>Hospital Service Robotics</u> (Room 120)

- Candy Sidner (Moderator)
- Rong Xiong
- Paul Maglio
- Holly Yanco
- Liz Sonenberg
- Tom Wagner
- Ethan Chiang

Scenario D. <u>Virtual Team Training</u> (Room 227)

- Mark Neerincx (Moderator)
- Frank Dignum
- Yukie Nagai
- Michael Freed
- Michael Beetz
- Jeff Bradshaw
- Bill Berry

Scenario E. <u>Personal Satellite Assistants</u> (Member's Room)

- Terry Fong (Moderator)
- Robert Hoffman
- Dirk Shulz
- Andreas Hofmann
- Manuela Veloso
- Jean Scholtz
- Patricia Wrightson

Description of Scenarios and Report-back Questions

A. Preparing for and Managing a Major Natural Disaster Scenario

Background:
Located in central Mexico, Mount Popocatepetl is an active volcano that has had intermittent eruptions since the mid-1990s. Its most violent eruption in 1,200 years occurred in December, 2000. Situated 60 km southeast of Mexico City (population 18 million) thirty million people overall live within its sight. In late April, the National Disaster Prevention Center (Centro Nacional de Prevencion de Desastres or CENAPRED) moved the alert status to the highest level of the yellow stage; the next stage is a red alert, which would prompt evacuations to begin. Should the volcano experience a major eruption, millions of people could die, hundreds of square miles of arable land could be destroyed, and the nation's capital itself could experience major destruction and disruption.

Having a plan ready could mean the difference between life and death for hundreds of thousands of people, and could also prevent the mass destruction of public and private property. However, no single plan can be sufficient to cover all possible scenarios—there is no way to know before the fact, for example, what the weather will be like on the day of the eruption, the direction and force of the wind, or the strength of the eruption itself. The extent of the evacuation, search and rescue and ultimately, rebuilding, will all depend on the size of the eruption and the concomitant weather patterns. Yet these are not the only factors that matter: the social, economic and political infrastructure—at both the local and national levels—will also determine how smoothly each of these stages will proceed.

Recent advances in computation and software agents can help officials make sense of complex, uncertain, high-tempo events, such as the one described here. For complex evacuation and search and rescue mission, for example, robots and autonomous vehicles, including autonomous aerial vehicles (small or big) and possibly satellites can help. There are useful technologies all around, but not a lot of technology to tie them together.

Goal:
Your company, "007 And Beyond," has decided to respond to an RFP from CENAPRED to develop the prototype for a system capable of providing complex decision support and disaster management before, during, and after a major eruption of Mount Popocatepetl. The initial prototype is to be developed in two years and fielded in five. In responding to this RFP, you must describe the kinds of agents and robots you will require, and how they will be integrated and made use of along the time-line of the entire event (from predicting the strength and impact of the volcano; to evacuation plans; post-event logistics, i.e., search and rescue, and rebuilding). Concerns about safety, usability, and trust should be addressed. In your design, you should also take into account that experts from different disciplines, languages, and cultures around the world may be called in as part of the effort should a disaster occur. However, the most important consideration for the purposes of this design exercise will not be on the capabilities of specific robots and agents per se, but on how to handle the physical and social interdependencies among groups of robots, agents, and people working together in real-world multi-cultural settings.

Constraints:
You may assume that you have resources for (including access to the top engineers) any software or hardware that you want to develop – but you need to be realistic about how much work they can carry out; some things just cannot be done in a day! If you require technology that is not available off the shelf (e.g., a new type of robot, or a specific GIS service) you may add it to the RFP, but its development time and cost must be included within the schedule limits.

At the end of this exercise, your group will present a set of slides that includes answers to – but are not necessarily limited to –the following questions:

- What design did your group develop to address the problems raised by your scenario?
- What aspects of intelligent human-machine collaboration research were relevant/irrelevant to your design approach for this scenario?
- What advantages/disadvantages did you find in an intelligent human-machine collaboration approach that would not be found in a human-only or traditional autonomy (e.g., master-slave) approach to the problem?
- What issues are raised by combining both software agents and robots as team members (in addition to people) in your project?
- What kind of progress can be demonstrated by the end of the second year? (In other words, what can your system do and what can't it do?) What could be done in years three through five?
- In what circumstances is this system likely to fail when deployed in a real-world environment?
- What international, global, or cross-cultural considerations will your project raise?
- What individuals or research groups could you draw on to help address these challenges (or their precursors)?
- What were the biggest challenges faced by your group in working through your scenario.

B. **Human-Machine Collaboration for Small-Lot Agile Manufacturing Scenario**

Background:
George is the owner of a small furniture manufacturing company. The company makes custom furniture for clients to their specifications. While some of George's products are similar, no two are exactly the same due to differences in client preferences in color, dimension, and shape. Due to variations in material costs, as well the low volume of each specific design eliminates economies of scale, George is finding it difficult to make ends meet financially. Moreover, typical manufacturing robots require significant expertise in robot programming to make needed adaptations.

Goal:
Your company has just put you and your group in charge of designing a new robot called the Pengo9000. Without checking the feasibility with your R&D department first, sales and marketing agents have already promised that customers like George will be able to triple their productivity without added manpower or significant "retooling" cost, due to the Pengo9000's envisioned adaptability and end-user programming-by-demonstration features. Pengo9000 should be able to generalize its knowledge to new situations, and to accept commonsense task descriptions from non-specialists expressed through speech, gestures, and diagrams. The Pengo9000 will then formulate plans and perform George's tasks robustly by monitoring plan execution and continuously modifying the plan reactively as changes in the environment and ongoing input from users dictates. The marketing video shows the following scene:

"Hello, how are you today?" says the Pengo9000.

"I'm fine. Would you please help me drill some holes in this chair?" asks George.

"Sure, but I do not know how to drill holes. Could you please show me?" politely asks the Pengo9000.

"Absolutely, here's how to do it" says George, as he picks up the robot's dexterous hand, places in it a cordless drill, and drags the hand to drill a hole in a piece of furniture. "You see, you need to push a little harder in this part otherwise the hole won't come out good. When it's done, it should look very clean and have no leftover debris inside, like this."

"Thanks I got it," chimes the robot.

"Great. Now please go ahead and drill 50 holes in each of the chairs over in the next room. Can you do that in 20 minutes?" says George.

"I may not be able to do that in 20 minutes, but I am confident that I'll have it done in 40 minutes." answers the robot.

"No, that's too long." says George.

"Then can you help me drill half of the chairs?" asks the Pengo9000.

"That's fine." says George. The Pengo9000 then flies out into the next room to perform its task.

"Pengo9000, please come here and help me lift up this heavy bed and move to that room." asks George.

"Okay, I will pause my current task and come help you." Again, as though an occult hand had reached down and set in motion a flurry of wheels and cogs, the Pengo9000 is at George's side helping him move the bed frame.

After the task is completed, the Pengo9000 returns to its original plan of drilling all of the chairs. Watching the robot move again into the other room, George smiles, confident that the Pengo9000 will greatly improve the efficiency of his business.

Concerns about safety, usability, and trust should be addressed. However, the most important consideration for the purposes of this design exercise will not be on the capabilities of specific robots and agents per se, but on how to handle the physical and social interdependencies among groups of robots, agents, and people working together in real-world settings. Based on the success of your two-year program, there may be additional funding for years three through five.

Constraints:
You have five years of funding, but you must show significant progress in concepts and implementation by the end of the second year. You may assume that you have resources for (including access to the top engineers) any software or hardware that you want to develop – but you need to be realistic about how much work they can carry out; some things just cannot be done in a day!

At the end of this exercise, your group will present a set of slides that include answers to – but are not necessarily limited to –the following questions:

- What design did your group develop to address the problems raised by your scenario?
- What aspects of intelligent human-machine collaboration research were relevant/irrelevant to your design approach for this scenario?
- What advantages/disadvantages did you find in an intelligent human-machine collaboration approach that would not be found in a human-only or traditional autonomy (e.g., master-slave) approach to the problem?
- What issues are raised by combining both software agents and robots as team members (in addition to people) in your project?
- What kind of progress can be demonstrated by the end of the second year? (In other words, what can your system do and what can't it do?) What could be done if you were funded for years three through five?
- In what circumstances is this system likely to fail when deployed in a real-world environment?
- What international, global, or cross-cultural considerations will your project raise?
- What individuals or research groups could you draw on to help address these challenges (or their precursors)?
- What were the biggest challenges faced by your group in proposing your solutions to the scenario?

C. <u>Human-Machine Collaboration in Hospital Service Robotics Scenario</u>

Background:

Congratulations! One of the world's largest (and richest) worldwide healthcare organizations has decided to use your expertise to change the way hospitals work, by using robots for delivery, logistics, and moving patients around. Many hours of work are currently wasted on moving equipment, supplies, and trash from the main hospital storage areas to the different departments, from the laundromat to the nurses to the patients and back to the laundromat, from the kitchen to patients and back, from patients to labs, and so forth. Additional hours of work are spent pushing patient beds and wheelchairs around, moving them from their beds to the x-ray or other labs, and back. It's true that we might always want a nurse or a doctor to accompany the patient, but the goal is to leave most of the heavy lifting and pushing to the robots.

Goal:

As a group, you are given two years to develop a commercial-grade prototype system which would be demonstrated at hospitals in three different regions of the world (involving different languages, cultures, medical practices, healthcare policies, and so forth). The system will be composed of multiple robots, software agents, and whatever portable or fixed hardware devices (e.g., workstations, tablets) or local features (e.g., speech understanding) are necessary to operate, call, coordinate, monitor, task, and retask or stop the robots when things go wrong. Working together, the robots, agents, and people should be able to carry out the tasks described above within existing hospital facilities and under normal hospital conditions – lots of patients, staff, and visitors moving about in the vicinity of the patients, multiple staff scheduling requests for pickup and delivery, and the possibility of making mistakes (e.g., kitchen staff placing the wrong order on the tray). Concerns about safety, usability, and trust should be addressed. However, the most important consideration for the purposes of this design exercise will not be on the capabilities of specific robots and agents per se, but on how to handle the physical and social interdependencies among groups of robots, agents, and people working together in real-world multi-cultural settings. Based on the success of your two-year program, there may be additional funding for years three through five.

Constraints:

You may assume that you have resources for (including access to the top engineers) any software or hardware that you want to develop – but you need to be realistic about how much work they can carry out; some things just cannot be done in a day! The emphasis should be on technology that does not require significant changes to the environment. Thus the use of place markers on walls—as an example of a small change to the environment—is allowed, but not encouraged. In contrast, assuming that staff and patients have a way of calling an agent or robot remotely is definitely okay, as long as you make it explicit in your design. Hospitals have multiple floors, and the system should find a way to address these and other physical considerations.

At the end of this exercise, your group will present a set of slides that include answers to – but are not necessarily limited to –the following questions:

- What design did your group develop to address the problems raised by your scenario?
- What aspects of intelligent human-machine collaboration research were relevant/irrelevant to your design approach for this scenario?
- What advantages/disadvantages did you find in an intelligent human-machine collaboration approach that would not be found in a human-only or traditional autonomy (e.g., master-slave) approach to the problem?
- What issues are raised by combining both software agents and robots as team members (in addition to people) in your project?
- What kind of progress can be demonstrated by the end of the second year? (In other words, what can your system do and what can't it do?) What could be done if you were funded for years three through five?
- In what circumstances is this system likely to fail when deployed in a real-world environment?
- What international, global, or cross-cultural considerations will your project raise?
- What individuals or research groups could you draw on to help address these challenges (or their precursors)?
- What were the biggest challenges faced by your group in proposing your solutions to the scenario?

D. Human-Machine Collaboration in Virtual Team Training Scenario

Background:
Maintaining common ground and dynamically coordinating individual and joint tasks are of central importance to any team faced with complex assignments. This is all the more so in ad hoc, multi-disciplinary teams with members having different cultural values, organizational backgrounds, and competence levels. Professional organizations tend to administer exercises to train their personnel. Teams in training often rely on their own judgment when interpreting the underlying processes that determine success and failure in team performance. They also must decide how to make improvements within the team. The benefits from expensive exercises can be increased significantly through an agent-based learning system that monitors relevant team processes, interprets the performance of individuals from a team perspective, and that suggests interventions to adapt the exercise (e.g., by introducing events, or delivering feedback) to optimize its learning value. Such a system would be particularly useful to improve the quality of multi-human role-playing exercises.

Goal:
Your group has been assigned to develop a new agent-based system for virtual team training. The system will be used to teach students team skills by staging training in an environment mixing humans and software agents, where the agents deliberately bring about situations that challenge the team skills of the trainees. Of course, this requires combining direction from both humans and autonomous assistants. The humans and agents work together to consider which agent behavior produces the best learning situation for the trainee, and then instructs the agents to act accordingly. Such a system imposes policy constraints on the autonomy of agents to the benefit of maintaining control over training. As an example, consider the following scenario:

Command Center officer Harry S. Jones walked down the stairs of his frigate to enter the hub of operations. Intel says that during the ship's passage through the Strait of Hormuz, the risk of terrorist attack are "high." The Rules of Engagement have been adjusted accordingly, perhaps with some degree of intelligent automation. "As of tomorrow, the mission will be critical," Jones thought. "I better use this final quiet day to prepare myself for things that may happen." He switched his workstation to TATOE (Tactical Adaptive Training Onboard Equipment) in training mode. On the screen Aedes (Roman god of the wind) appeared, his virtual instructor. "Welcome," said Aedes, "how may I help you?" "I want to practice, in preparation for my mission," replied officer Jones. "Fine," Aedes said, "please enter the critical components of your mission." Jones contemplated for a minute and then clicked on "tactical picture compilation in multi-threat environments"; "Anti Submarine Warfare and RoE"; "Joint and Combined SAG operations." "Okay," confirmed Aedes, "in previous training sessions you performed well, but make sure that you direct your team properly: give concise and univocal commands; monitor the workload and output of your team, and take action if needed." With a short "Okay, I'm ready," officer Jones concluded the dialogue with Aedes. TATOE started the scenario. Because the members of officer Jones' team were involved in other tasks, their roles are filled in by software agents. TATOE has also initialized realistic opponent agents. Aedes has instructed the team agents to select a low-probability response should officer Jones issue an ambiguous order. This is done so as to confront officer Jones with the consequences of equivocal commands. The opponent agents show proactive *behavior and are likely to demonstrate unexpected behavior or refraining from expected behavior. During the exercise, officer Jones makes a serious error, and TATOE freezes the scenario. Aedes appears and prompts officer Jones to think critically about what happened. Officer Jones suddenly acknowledges that a critical decision of him was based on an implicit and dubious assumption. Fortunately the agents in TATOE cannot only perform realistically, they can also explain their behavior in terms of beliefs, goals, and argumentation. Officer Jones rewinds the scenario to the point he considers critical and interrogates the opponent agent by asking questions like "what is your current goal?"; "what do you think is my goal?"; etc. This helps Jones to gain insight into how his decision error may have come about. Aedes proposes to rewind the scenario and to replay it from there. Officer Jones agrees. But then the operational system awakens with an alarm signal: "fast incoming track!" TATOE immediately ends the training and switches to operational mode.*

Concerns about usability and trust should be addressed. However, the most important consideration for the purposes of this design exercise will not be on the capabilities of specific agents per se, but on how to handle the physical and social interdependencies among groups of agents and people working together in real-world multi-cultural settings.

Constraints:
You have five years of funding, but you must show significant progress in concepts and implementation by the end of the second year. You may assume that you have resources for (including access to the top engineers) any software or hardware that you want to develop – but you need to be realistic about how much work they can carry out; some things just cannot be done in a day!

▪▪

At the end of this exercise, your group will present a set of slides that include answers to – but are not necessarily limited to –the following questions:

- What design did your group develop to address the problems raised by your scenario?
- What aspects of intelligent human-machine collaboration research were relevant/irrelevant to your design approach for this scenario?
- What advantages/disadvantages did you find in an intelligent human-machine collaboration approach that would not be found in a human-only or traditional autonomy (e.g., master-slave) approach to the problem?
- What issues are raised by combining both software agents and robots as team members (in addition to people) in your project?
- What kind of progress can be demonstrated by the end of the second year? (In other words, what can your system do and what can't it do?) What could be done if you were funded for years three through five?
- In what circumstances is this system likely to fail when deployed in a real-world environment?
- What international, global, or cross-cultural considerations will your project raise?
- What individuals or research groups could you draw on to help address these challenges (or their precursors)?
- What were the biggest challenges faced by your group in proposing your solutions to the scenario?

E. Personal Satellite Assistant Scenario

Background:
Enhancing the crew's ability to perform their duties is critical for successful, productive, and safe space operations during future long-duration space exploration missions. Crew time on such missions is a precious resource and may cost hundreds of dollars per minute per astronaut. The limited number of crew members are required to maintain complex systems, assist with life-critical environmental health monitoring and regulation, perform dozens of major simultaneous payload experiments, and perform general housekeeping. Safety considerations and size constraints are also important issues for many manned mission activities. Even if it were physically possible for an astronaut to enter congested spacecraft areas, protruding debris or other environmental hazards of one kind or another could pose serious safety risks.

Goal:
Your group is assigned to help develop the Portable Satellite Assistant (PSA), a spherical flying robot (approximately 3 in or 8 cm in diameter) that will operate onboard manned spacecraft in micro-gravity, pressurized environments. Environmental sensors for gas, temperature, and fire detection provide the ability for the PSA to monitor spacecraft, payload and crew conditions. In the microgravity environment, ducted fans provide propulsion and batteries will provide portable power. NASA is leading the development of PSA hardware and low-level flight, navigation, and sensor integration systems; the responsibility of your team is to develop the software necessary to allow it to do useful work in close collaboration with the crew, with other PSAs, with software agents performing critical tasks onboard and offboard the spacecraft, and with support personnel on the ground.

The following scenario emphasizes selected aspects of intelligent human-machine collaboration required for the PSA:

A crewmember is awakened by a PSA at the requested time. The astronaut asks for a video briefing on the latest events, schedule changes, and priorities while she washes, and eats breakfast. The PSA follows the crewmember through her routine while giving the updates and then checks the inventory database to ensure that the necessary resources are available for the astronaut's first scheduled task. The crewmember logs into her homepage and sets several notifications to be programmed into the PSA to remind her of important activities and times for today's tasks. As the crewmember works at a payload rack the PSA tracks her movements and provides a remote data terminal capability to allow her to check on procedures and training instructions, and to support remote videoconferencing and email exchanges with remote colleagues. Later the crewmember conducts a delicate investigation in the glove-box. She requests support from the Principal Investigator (PI) on earth to help her walk through the procedure. The PI calls up a second PSA and maneuvers about the astronaut and glove-box to obtain an optimum view of the operation and to provide real-time feedback to the crewmember. Since the crewmember and the remote PI are absorbed in performing their tasks, the PSAs coordinate the details of their flight and their participation in joint and individual activities themselves, without requiring constant attention from their human partners. Moreover, the PSAs are not just passively waiting to be told what to do. They are actively looking for ways to be helpful to the humans in their *current task as well as in ongoing responsibilities that have been delegated. For example, as the crewmember uses up supplies the PSA tracks the inventory tags and updates the inventory database. During a video inspection, a PSA notices that specimens in the habitat holding units need food. That evening a pair of PSAs use special integrated payload interfaces and cargo packages to inject supplies such as food into experimental units. One PSA injects the supplies and another collaborating PSA acts as a supply cargo carrier. In many of these tasks, the PSA interacts with software agents that are onboard and offboard the spacecraft. Software agents with specific capabilities for a given task can migrate on and off the PSA when appropriate.*

Concerns about safety, usability, and trust should be addressed. However, the most important consideration for the purposes of this design exercise will not be on the capabilities of specific robots and agents per se, but on how to handle the physical and social interdependencies among groups of robots, agents, and people working together in real-world multi-cultural settings.

Constraints:
You have five years to deliver a system, but you must show significant progress in concepts and implementation by the end of the second year. You may assume that you have resources for (including access to the top engineers) any software or hardware that you want to develop – but you need to be realistic about how much work they can carry out; some things just cannot be done in a day!

At the end of this exercise, your group will present a set of slides that include answers to – but are not necessarily limited to –the following questions:

- What design did your group develop to address the problems raised by your scenario?
- What aspects of intelligent human-machine collaboration research were relevant/irrelevant to your design approach for this scenario?
- What advantages/disadvantages did you find in an intelligent human-machine collaboration approach that would not be found in a human-only or traditional autonomy (e.g., master-slave) approach to the problem?
- What issues are raised by combining both software agents and robots as team members (in addition to people) in your project?
- What kind of progress can be demonstrated by the end of the second year? (In other words, what can your system do and what can’t it do?) What could be done in years three through five?
- In what circumstances is this system likely to fail when deployed in a real-world environment?
- What international, global, or cross-cultural considerations will your project raise?
- What individuals or research groups could you draw on to help address these challenges (or their precursors)?
- What were the biggest challenges faced by your group in proposing your solutions to the scenario?

Pre-Meeting Questionnaire

Intelligent Human-Machine Collaboration in Keywords:

Keywords describing YOUR research:

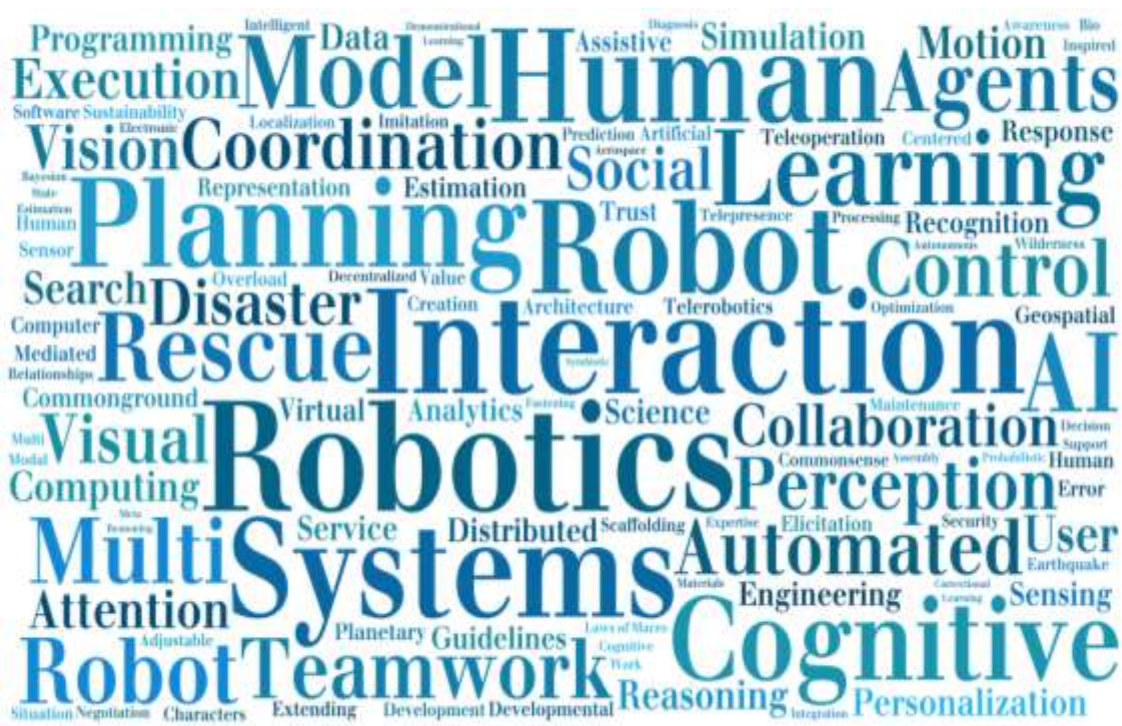

How do YOU define Intelligent Human-Machine Collaboration?

…intellectual interaction (e.g., understanding of common-sense and human desires) between humans and robots

…reasonable expectations by man and machine about what the other will do or how one will react to the other's actions

…man and machine anticipate what and when to communicate in order to support the joint activity, without communicating too much.

…collaborators each have some stake in outcomes or the performance of the system

…shared responsibility, authority, and goals

…human and machine mutually adapt their behaviors to one another to maximize the information exchange between the partners or the performance of the task

...machines reason and adapt to a situation (and its specifics) and to participating human(s)

…machines capable of taking the lead in terms of initiating actions (while being responsive to human input and override) and interacting with humans to explain its behavior if requested.

…humans and machines that complement each other

...cooperation that mimics interactions between two people.

....the undertaking of a shared goal and the development a shared plan which consists of intentions on the part of each agent, mutual beliefs about who will do what, when and where, a recipe for accomplishing the shared goal, and intentions that the other agents will do their part to accomplish the shared goal.

...technology that amplifies and extends the human abilities to know, perceive & collaborate.

...humans AND machines jointly perform tasks they would not be able to perform on their own, (or with just humans or just machines).

...intelligent machines would understand human behavior, intentions, and communications, and act appropriately against that background to collaborate as a proper team member.

...people and technology learn and act as partners to accomplish a joint goal.

…humans and machines mutual adapt their behavior, share resources and knowledge, learn and teach each other, and empower each other's capabilities.

...partnerships driven and constrained by explicit working agreements and ethical policies.

…man and machine share information and shift control by coming together or disbanding to accomplish a task or provide a service.

....joint and coordinated action involving people and computationally-based systems.

... machines are no longer a tool of the human, but a partner.

...participants work together to accomplish tasks with each participant doing a portion of the task most suited to their particular skills at the current time

…recognition by participants when others are having difficulties and helping if feasible.

...machines can request help from humans when it feels uncertain on its perception, cognition, and decision-making and take feedback from humans who believe the machine needs help.

...behaviors that exploit the respective strengths (and support the weaknesses) of the human and machine participants in a joint task.

...coordinated actions and (on the fly) decision making by both humans and machines.

...communication not requiring specialized inputs

…humans and machines react to situational changes and modify work sharing in real-time

...integration of AI into the machines.

...machines and humans combine each others' strengths and fill-in for each others' limitations.

...better performance in the mission, independently of how it is achieved.

…neither human not machine treats the other as a disturbance to be minimized

…overall problem situation awareness is shared and well understood by human and machine

…represented by the naturalness of the observed/experienced interaction.

How “intelligent” are current applications of human-machine collaboration?

…not very intelligent at all

…robots are intelligent to the extent that they can fulfill some tasks that require intelligence but intelligence is a characteristic of humans, not of robots.

…Most applications are (usually) geared towards a very specific domain and application where the environment is either stable or very restricted.

…When (impressive, complex) tasks are narrowly defined, the word "intelligent" seems wrong (e.g. Mathematica solving a complex integral or a photo album app recognizing a face). In the rare cases where it is tempting to use the word, the credit goes as much to good user interaction design than to any kind of reasoning mechanism (e.g. Siri0. So overall, there aren't even a lot of near-miss examples of intelligent machine collaboration.

What are the biggest impediments that prevent successful real-world applications of intelligent human-machine collaboration?

...limitations are analogous to social skills in humans

...Insufficient decision-making algorithms for reasoning and learning on the fly.

...Insufficient perception, modeling, and security

...the procurement process

...machine awareness of user intentions occurs under very strict or unnatural conditions

...collaboration is not robust in dynamic environmental conditions.

...applications are not well designed for usability by target clients.

...need appropriate real-time mental models that are aligned across people and systems.

...appropriately (rather than arbitrarily) combining domain knowledge with an understanding of people and with computational, robotic, and engineering knowledge.

...limited understanding of how we make machines participants in the flow of human perception and attention.

...machines are limited in their abilities to make use of open source knowledge and must rely on internal built-in knowledge.

...lacking communication and understanding of both the real world situations/problems and the technological possibilities.

...researchers have difficulty developing systems which approach real world complexity

...deep trans-disciplinary work is difficult

...learning to understand and work together with other disciplines to develop shared language and conceptual apparatus.

...human trust in intelligent systems and adoption of disruptive technologies

...insufficient sensing, adaptation, intelligent interactions

...humans and machines cannot evaluate exactly each other's abilities/inabilities in different situations or communicate in an efficient, robust way, making tacit collaboration difficult

...systems cannot explain their decisions.

...insufficient machine vision, machine learning, and natural language processing.

...insufficient knowledge of the world, speech understanding, real-world capable agents, and understanding and production of gesture.

...robots are still pretty dumb and developing a sophisticated machine capable of high levels of collaboration is prohibitively expensive in terms of both time and money.

...designing algorithms and interfaces that take into account the variability in human perception and understand human interaction with the machine.

...there is a need to move away from data-intensive modeling, and more towards the use of psychology to analyze the features of human-machine systems to determine the key interaction principles needed to make such interactions efficient and effective.

...systems must better adapt to changing situations to provide better situation awareness.

...humans change and machines typically don't, which affects trust, cohesion, etc.

....Ineffective production/managerial concepts.

…insufficient tactile sensing for teleoperation (mainly in surgery).

...insufficient sensor fusion and real-time image processing.

...systems need to be developed and tested in real-world environments.

...current hardware and software platforms are too fragile.

...fickle research funding systems that lead to fragmented development efforts.

...machine automation and processing must interact with users in novel situations.

...intelligent applications follow scripts that are not natural for the users either because of mismatched goals or unexpected ordering.

...intelligent applications do not work on the same time scale as people (too slow/fast).

...robots have dificulties reading intentions of/adapting behavior to human partners.

...lack of provably robust robotic sensing and no widely accepted methodology for bridging the gap between low-level perception and high-level cognition.

...privacy (for internet-based systems) and legal (responsibility) issues for autonomous machines.

...we don't accept that intelligent robots have (or will have for a long time) perceptual, cognitive, and actuation limitations.

...symbiotic approaches where machines ask for help, rather than just being supervised.

…social aspects of teamwork are greatly simplified and constrained.

…humans and machines make very different types of errors.

…current systems are often scripted and are narrowly vectored.

Examples of successful intelligent human-machine collaboration that exist outside the lab today:

Top 5 Responses:

1) **Robotic Surgery** (e.g. Da Vinci surgical robot) (7)
2) **Google Search/Search Engines** (7)
3) **Siri** (6)
4) **Production systems where human and robots work together at the same place** (e.g., Kiva Systems Warehouse Robots) (5)
5) **Flight Management/Navigation Systems on Commercial Aircraft** (4)

   Intelligent Vehicles (e.g. Google unmanned vehicle) (4)

   I have seen no successful examples (4)

Others:

- Disaster Response. (e.g., Fukushima-Daiichi robot operation) (3)
- Mobile apps (3)
- Driver Assistance Systems (e.g., collision/run-off-road prevention) derived from robotics search (2)
- Androids/Entertainment Robots (2)
- Mars Exploration Rover, which moves with obstacle avoidance ability and returns images captured on Mars to the Earth, but is mainly teleoperated by human on the Earth. (2)
- Amazon/Recommender systems (2)
- Home robotics systems that carry out typical house chores (roomba) (2)
- Networked, multi-player video games (including computer based teammates)
- Ebay
- 3D environments for Human Factors Analysis
- Augmented reality for satellite buildup
- Critical Systems Monitoring
- GP patient management, linking pharmacists, paramedical staff, specialists, etc. (Precedence Healthcare)
- Text analysis systems
- Q&A systems
- Hybrid assistive limb, which can detect and identify user’s motion intent, and then improve user’s physical movement.

- Segway, which is an efficient two wheeled personal transport. It detects and identifies user’s intent on the motion direction, and then carries out the movement with a self-balance control technology.
- software agents for simple tasks like calendaring
- UAVs
- Crowdsourcing platforms such as amazon turk that operate workflows on a large scale with human workers
- Traffic monitoring devices that take human reports of traffic jams and use other sensor networks to improve predictions of travel time
- Entertainment activities (emotional robots) that adjust to the human's preferences and needs.
- Email, though it is arguably not "collaborative.
- intelligent transport systems (ITS) - optimal use of road, traffic and travel data, traffic and freight management, safety and security on the road, linking the vehicle with the transport system
- Syndromic surveillance systems
- Military C2
- Agriculture - e.g., milking robots
- Mobile context-aware notification system for police officers
- Game-based (VR and simulation) training to learn to cope with demanding situations
- Virtual robotic health coach
- Adaptive automation for object identification in Naval Command Center
- Distributed Intelligent Networked Control Systems' (DINCS) to manage naval platform systems
- Phone Trees
- WII
- Online grammar correction in word processors
- Speech-to-text systems
- thermostats
- Coffee-maker
- Personal Computers
- Heating/cooling system

What have been the biggest breakthrough(s) in your field in the last 5 years?

- Robots that integrate techniques from multiple domains.
- Open knowledge sources and development of techniques that make use of these sources.
- Incorporation of more social aspects in agent systems.
- Robot Operating System (ROS) is encouraging robotics developers & researchers to build upon the work of others (instead of just "reinventing the wheel").
- Smaller and cheaper 3D sensors, such as Kinect and Velodyne, that provide very high quality data for perception systems.
- Smart phones with diverse sensors, good displays and data network access
- Within AI, machine learning and improved user interaction design for intelligent systems.
- Access to enormous amounts of human behavioral data that can be analyzed statistically.
- User modeling on the Web enabled both by tracking behavior, by pooling data across people, by access to vast amounts of data , and by innovations in machine learning.
- The most critical breakthrough in recent years is the access to vast amounts of data.
- Computation models of attention.
- Expansion of sensor systems.
- Android developments.
- Speed and effectiveness of search algorithms, such as Google
- Quality of speech recognition.
- Fast-planning algorithms
- Inexpensive sensors that provide computers with better situation awareness.
- Tactile sensing, haptic interfaces and soft interactions.
- The BigDog, i.e., advances in actuators and adaptive motion control.
- Rapidly growing numbers of robots being used by non-researchers.
- Recognition that collaboration is key to human-robot interaction, and to interaction between people and any intelligent (non-human) entity.
- UAVs
- NLP
- Multi-agent systems for solving large coordination problems in both cooperative and competitive settings with clear empirical evidence from real-world datasets and with theoretical guarantees on performance.

- The general theory of macrocognitive work systems and the fundamental trade-off functions it includes.
- Dialogue modeling.
- High degree of robot capabilities.
- Multimodal communications – integrated speech, gestures or direct digital commands
- Adaptive automation.
- Personal health coaches.
- Game-based training to cope with demanding situations.
- Work has been incremental, and cannot be characterized as 'breakthrough' work.
- Systems that can plan, reason and diagnose on the fly, based on commonsense, and mixed discrete/continuous models.
- Simplified right sized interfaces with adaptive feedback, allowing user customization and cognitive feedback systems built for human users, not based on machine limitations.
- Low-cost platforms allowing greater access and experimentation.
- Fully autonomous robot soccer teams
- 3D perception processing algorithms
- Data mining and learning
- Autonomous cars
- DaVinci surgical support
- Space and underwater robotics
- Crowdsourcing
- IBM's Watson

The stovepiped nature of academia is preventing successful applications of intelligent human-machine collaboration.

52% said TRUE

48% said FALSE

Other Responses:

- Appropriate allocation of tasks and trust in systems by users is an issue
- Stovepiped government limits progress
- Overheads/risk associated with interdisciplinary work
- I don't know.

How often do you collaborate with researchers...

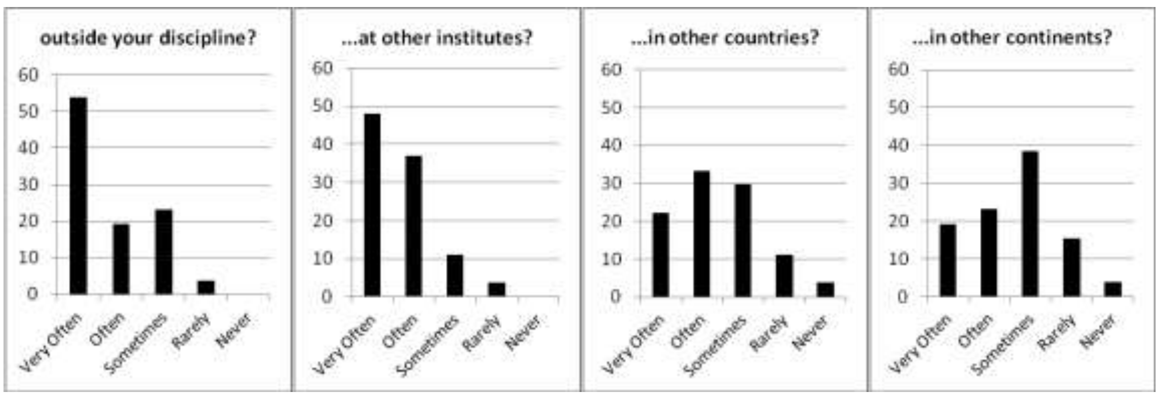

In which fields do you believe the next big breakthroughs will come from that will enable successful intelligent human-machine collaboration?

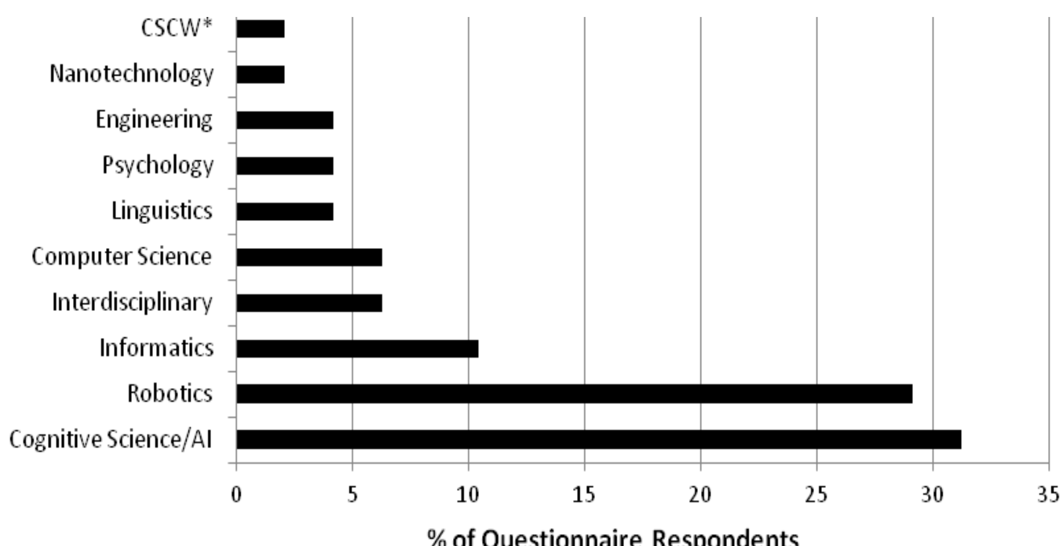

** Computer-Supported Coopperative Work*

Logistics Information

Workshop Location
Venue: National Academy of Sciences
Address: 2100 C Street, N.W., Washington, D.C. 20418
Date: 11-14 June, 2012
Phone: 202.334.2000

Welcome Reception
Date: 11 June, 2012
Time: 6:00 – 7:30 PM
Venue: The State Plaza Hotel; Ambassador Room
Address: 2117 E Street, N.W., Washington, D.C. 20006
Telephone: 202.861.8200

The National Academy of Sciences building is located 2 blocks from the State Plaza Hotel.

The Welcome Reception will provide an opportunity for workshop participants to meet one another in advance of the meeting. Light fare and refreshments will be available. Each participant has been provided with one drink coupon (redeemable for beer or wine). Workshop participants should plan on making their own dinner arrangements.

Hotel Information
Name: State Plaza Hotel
Address: 2117 E Street, N.W., Washington, D.C. 20006
Telephone: 202.861.8200

Your hotel accommodations have already been directly billed to The National Academies. You may be asked to present a credit card for any additional incidental charges you might incur, such as room service or in-room movies.

Reimbursement Information
The National Academies will reimburse (1) your transportation between the hotel and airport and (2) all meals not provided during the meeting (i.e., dinners). We will reimburse up to the government dinner per diem of $36.

For all reimbursements, you must have a corresponding receipt. Receipts for dinner expenses must be itemized as we cannot reimburse alcoholic beverages.

Included in your folder are instructions for completing your online expense reimbursement form. If you experience any difficulties with the online system, please do not hesitate to contact Mr. Neeraj Gorkhaly at ngorkhaly@nas.edu.

Nearby Restaurant Recommendations

Roti Mediterranean Grill
2221 I Street NW
Washington, DC 20037
202.499.2095

Notti Bianche
824 New Hampshire Avenue NW
Washington, DC
202.298.8085

Thai Place
2134 Pennsylvania Avenue NW
Washington, DC 20037
202.298.8204

District Commons
2200 Pennsylvania Avenue NW
Washington, DC 29937
202.587.8277

Blue Duck Tavern
1201 24th Street NW
Washington, DC 20037
202.419.6755

Bayou on Penn
2519 Pennsylvania Avenue, NW
Washington, DC 20006
202.223.6941

Circle Bistro
1 Washington Circle, NW
Washington, DC 20037
202.293.5390

DISH + drinks
924 25th Street, NW
Washington, DC 20037
202.338.8707

Café La Ruche
1039 31st Street, NW
Washington, DC 20007
202.965.2684

Clyde's Restaurant
3236 M Street, NW
Washington, DC 20007
202.333.9180

Rivers at the Watergate
600 New Hampshire Avenue, NW
Washington, DC 20037
202.333.1600

Additional restaurant suggestions can be viewed at www.opentable.com. The closest neighborhoods are Foggy Bottom, West End, and Georgetown.

Internet Connectivity

Free wireless internet is available at the workshop venue.

Key Contact Information

Workshop:
Neeraj Gorkhaly (ngorkhaly@nas.edu; 202.731.0186)
Hotel:
Concierge 202.861.8200

If you have any additional questions, please do not hesitate to contact Mr. Neeraj Gorkhaly.

Previous BGST Workshops

The Board on Global Science and Technology (BGST) conducted an international, multidisciplinary workshop on “Realizing the Value from Big Data” from February 28 to March 3, 2011, in Singapore. The workshop was jointly organized and hosted by the Institute for Infocomm Research (I2R) of Singapore’s Agency for Science, Technology and Research (A*STAR) at Fusionopolis.

On August 23-24, 2010, BGST hosted an Experts Meeting on Data Analytics and the Smart Energy Grid" that brought together scientists and engineers from major research universities, private industry representatives, and government officials to discuss the impact of large, complex, and distributed datasets and associated computational techniques on the future smart energy grid. The meeting was co-hosted by Microsoft in Redmond, Washington.