An Introduction to Human-Agent-Robot Teamwork (HART)

Jeffrey M. Bradshaw

Florida Institute for Human and Machine Cognition (IHMC)

jbradshaw@ihmc.us

Links to papers mentioned in this presentation: www.jeffreymbradshaw.net

Outline

- ***History of HART***
- Some Basic Concepts
- Purpose of the 2015 Workshop

Brief History of HART

- Knowledge Acquisition Workshops (KAW), Banff, 1985-2000
- Distributed AI, 1980-ca. 2000
- Autonomous Agents (1997-2001)/ICMAS (1995-2001)/AAMAS, 2002-present
- Human Robot Interaction (HRI) Workshops, 2006-present
- HART Workshops and Publications, 2009-2015, ongoing
- Related Lorentz Center Workshops, 2009-2014, ongoing
- HRI 2015 Workshops (Human-Robot Teamwork, Towards a Framework for Joint Action)

Bradshaw, J.M. From knowledge science to symbiosis science. Invited paper for special issue on "Twenty-Five years of Knowledge Acquisition", *International Journal of Human-Computer Studies* 71:2 (2012) 171-176

Co-located with esaw2009

21 November 2009
Utrecht, The Netherlands

Sponsored by: SIKS

HART Workshop

Supporting Joint Activity in Human-Agent-Robot Teamwork

http://www.jeffreymbradshaw.net/HART/

Maarten Sierhuis
(Carnegie Mellon/NASA Ames)
Maarten.Sierhuis@nasa.gov

Jeffrey M. Bradshaw
(Institute for Human and Machine Cognition)
jbradshaw@ihmc.us

Jurriaan van Diggelen (TNO)
(Local organization)
jurriaan.vandiggelen@tno.nl

Human-Agent-Robot Teamwork

Jeffrey M. Bradshaw, Ph.D.
Florida Institute for Human and Machine Cognition (IHMC)
Pensacola, FL USA 32502
+1 (850) 202-4462
jbradshaw@ihmc.us

Virginia Dignum, Ph.D.
Catholijn M. Jonker, Ph.D.
Delft University of Technology
Delft, The Netherlands
M.V.Dignum@tudelft.nl
C.M.jonker@tudelft.nl

Maarten Sierhuis, Ph.D.
Agent iSolutions
Moffett Field, CA
sierhuis@ejenta.com

Contribution and Benefit

Teamwork has become a widely accepted metaphor for describing the nature of multi-robot and multi-agent cooperation. By virtue of teamwork models, team members attempt to manage general responsibilities and commitments to each other in a coherent fashion that both enhances performance and facilitates recovery when unanticipated problems arise. Whereas early research on teamwork focused mainly on interaction within groups of autonomous agents or robots, there is a growing interest in leveraging human participation effectively. Unlike autonomous systems designed primarily to take humans out of the loop, many important applications require people, agents, and robots to work together in close and relatively continuous interaction. For software agents and robots to participate in teamwork alongside people in carrying out complex real-world tasks, they must have some of the capabilities that enable natural and effective teamwork among groups of people. Just as important, developers of such systems need tools and methodologies to assure that such systems will work together reliably and safely, even when they have been designed independently.

The purpose of the HART workshop is to explore theories, methods, and tools in support of humans, agents and robots working together in teams. Position papers that combine findings from fields such as computer science, artificial intelligence, cognitive science, anthropology, social and organizational psychology, human-computer interaction to address the problem of HART are strongly encouraged. The workshop will formulate perspectives on the current state-of-the-art, identify key challenges and opportunities for future studies, and promote community-building among researchers and practitioners.

Workshop Overview

The workshop will be structured around four two-hour sessions on themes relevant to HART. Each session will consist of presentations and questions on selected position papers, followed by a whole-group discussion of the current state-of-the-art and the key challenges and research opportunities relevant to the theme. During the final hour, the workshop organizers will facilitate a discussion to determine next steps. The workshop will be deemed a success when collaborative scientific projects for the coming year are defined, and publication venues are explored. For example, results from the most recent HART workshop (Lorentz Center, Leiden, The Netherlands, December 2010) will be reflected in a special issue of *IEEE Intelligent Systems* on HART that is slated to appear in January/February 2012.

HRI'12, March 5–8, 2012, Boston, Massachusetts, USA.
ACM 978-1-4503-1063-5/12/03.

Maarten Sierhuis
(Carnegie Mellon/NASA Ames
Maarten.Sierhuis@nasa.gov

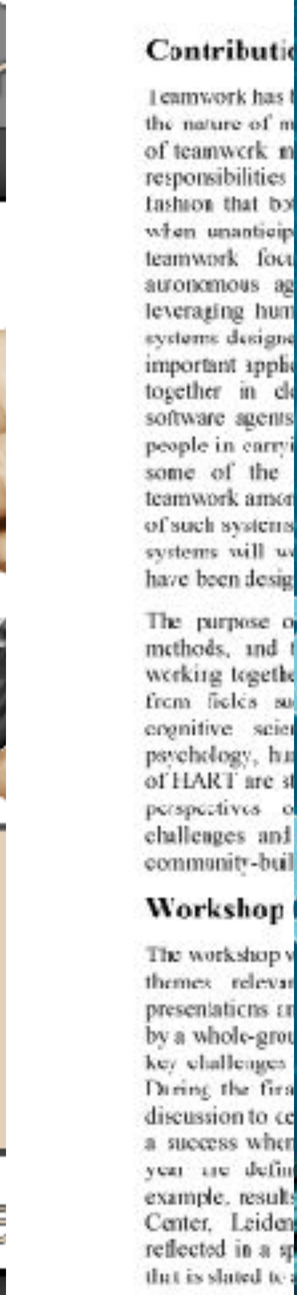

Human-Agent-Robot Teamwork

Jeffrey M. Bradshaw, Ph.D. Virginia Dignum, Ph.D. Maarten Sierhuis, Ph.D.

Copyright is held by the author/owner(s).
HRI'12, March 5–8, 2012, Boston, Massachusetts, USA.
ACM 978-1-4503-1063-5/12/03.

Collected Essays on Human-Centered Computing, 2001–2011

Edited and Coauthored by Robert R. Hoffman

With Coeditors Pat Hayes, Kenneth M. Ford, and Jeffrey M. Bradshaw

IEEE computer society

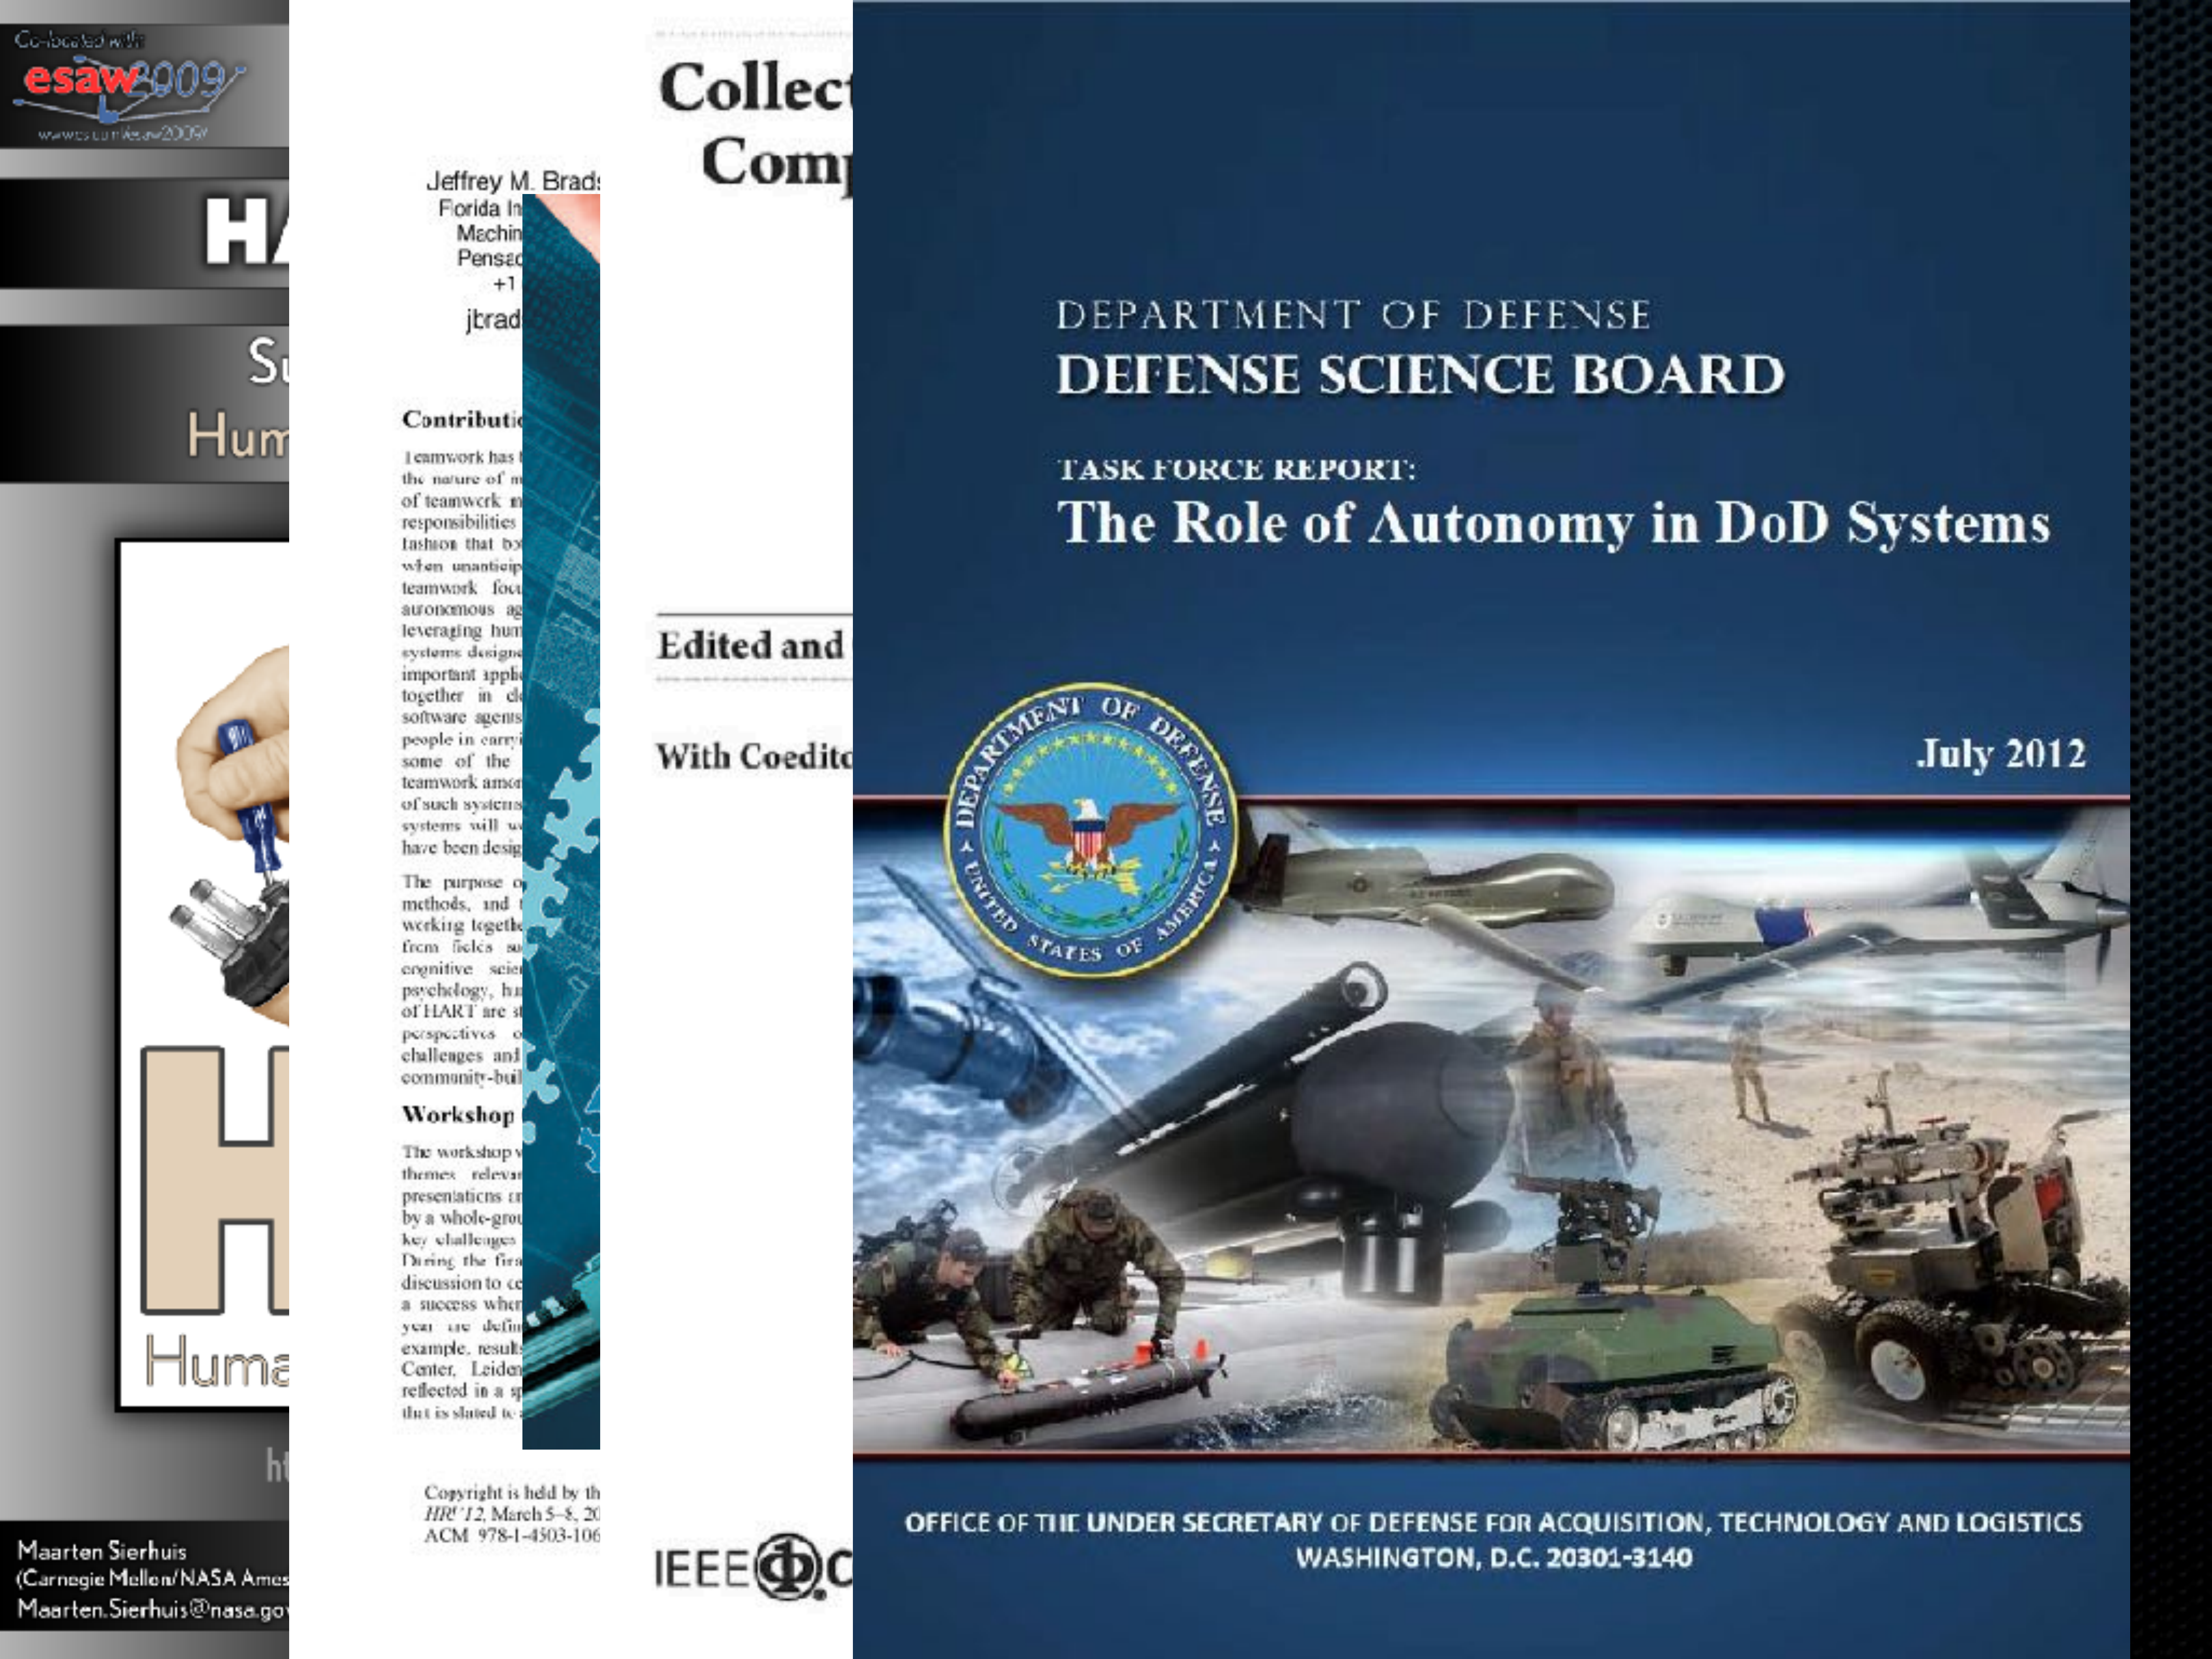
Co-located with
esaw2009
Maarten Sierhuis
DEPARTMENT OF DEFENSE
DEFENSE SCIENCE BOARD
TASK FORCE REPORT:
The Role of Autonomy in DoD Systems
July 2012
OFFICE OF THE UNDER SECRETARY OF DEFENSE FOR ACQUISITION, TECHNOLOGY AND LOGISTICS
WASHINGTON, D.C. 20301-3140

Co-located with
esaw2009
Maarten Sierhuis
Contributi
Teamwork has
Workshop
Edited and
With Coedito
OFFICE OF THE UN
INTELLIGENT
HUMAN-MACHINE COLLABORATION
Summary of a Workshop
Grounding
Decision
Making
Goals
Responsibility
Intelligent
Autonomy
Perception
Communication
Teamwork
Collaborative
Coordination
Knowledge
Adaptation
Interaction
Situation
Awareness
User
Human
Robot
Intention
Shared
Mental
Model
Joint
Learning
Team
NATIONAL RESEARCH COUNCIL
OF THE NATIONAL ACADEMIES

Lorentz center
Rich Cognitive Models
for Policy Design and Simulation
Workshop January 12 – 16 2009, Leiden, The Netherlands
Scientific Organizers
• Virginia Dignum, Utrecht
• Wander Jager, Groningen
• Catholijn M. Jonker, Delft
Keynote Speakers
• Jeff Bradshaw, Pensacola
• Rosaria Conte, Rome
• Frank Dignum, Utrecht
• Nigel Gilbert, Guildford
• Gert Jan Hofstede, Wageningen
• Barry Silverman, Philadelphia
The Lorentz Center is an international center in the Sciences.
Its aim is to organize workshops for scientists in an atmosphere which fosters collaborative work, discussions and interactions.
For registration see www.lorentzcenter.nl
Design SuperNova Studios, Den Haag
BNVKI
SIKS
Lorentz center
www.lorentzcenter.nl

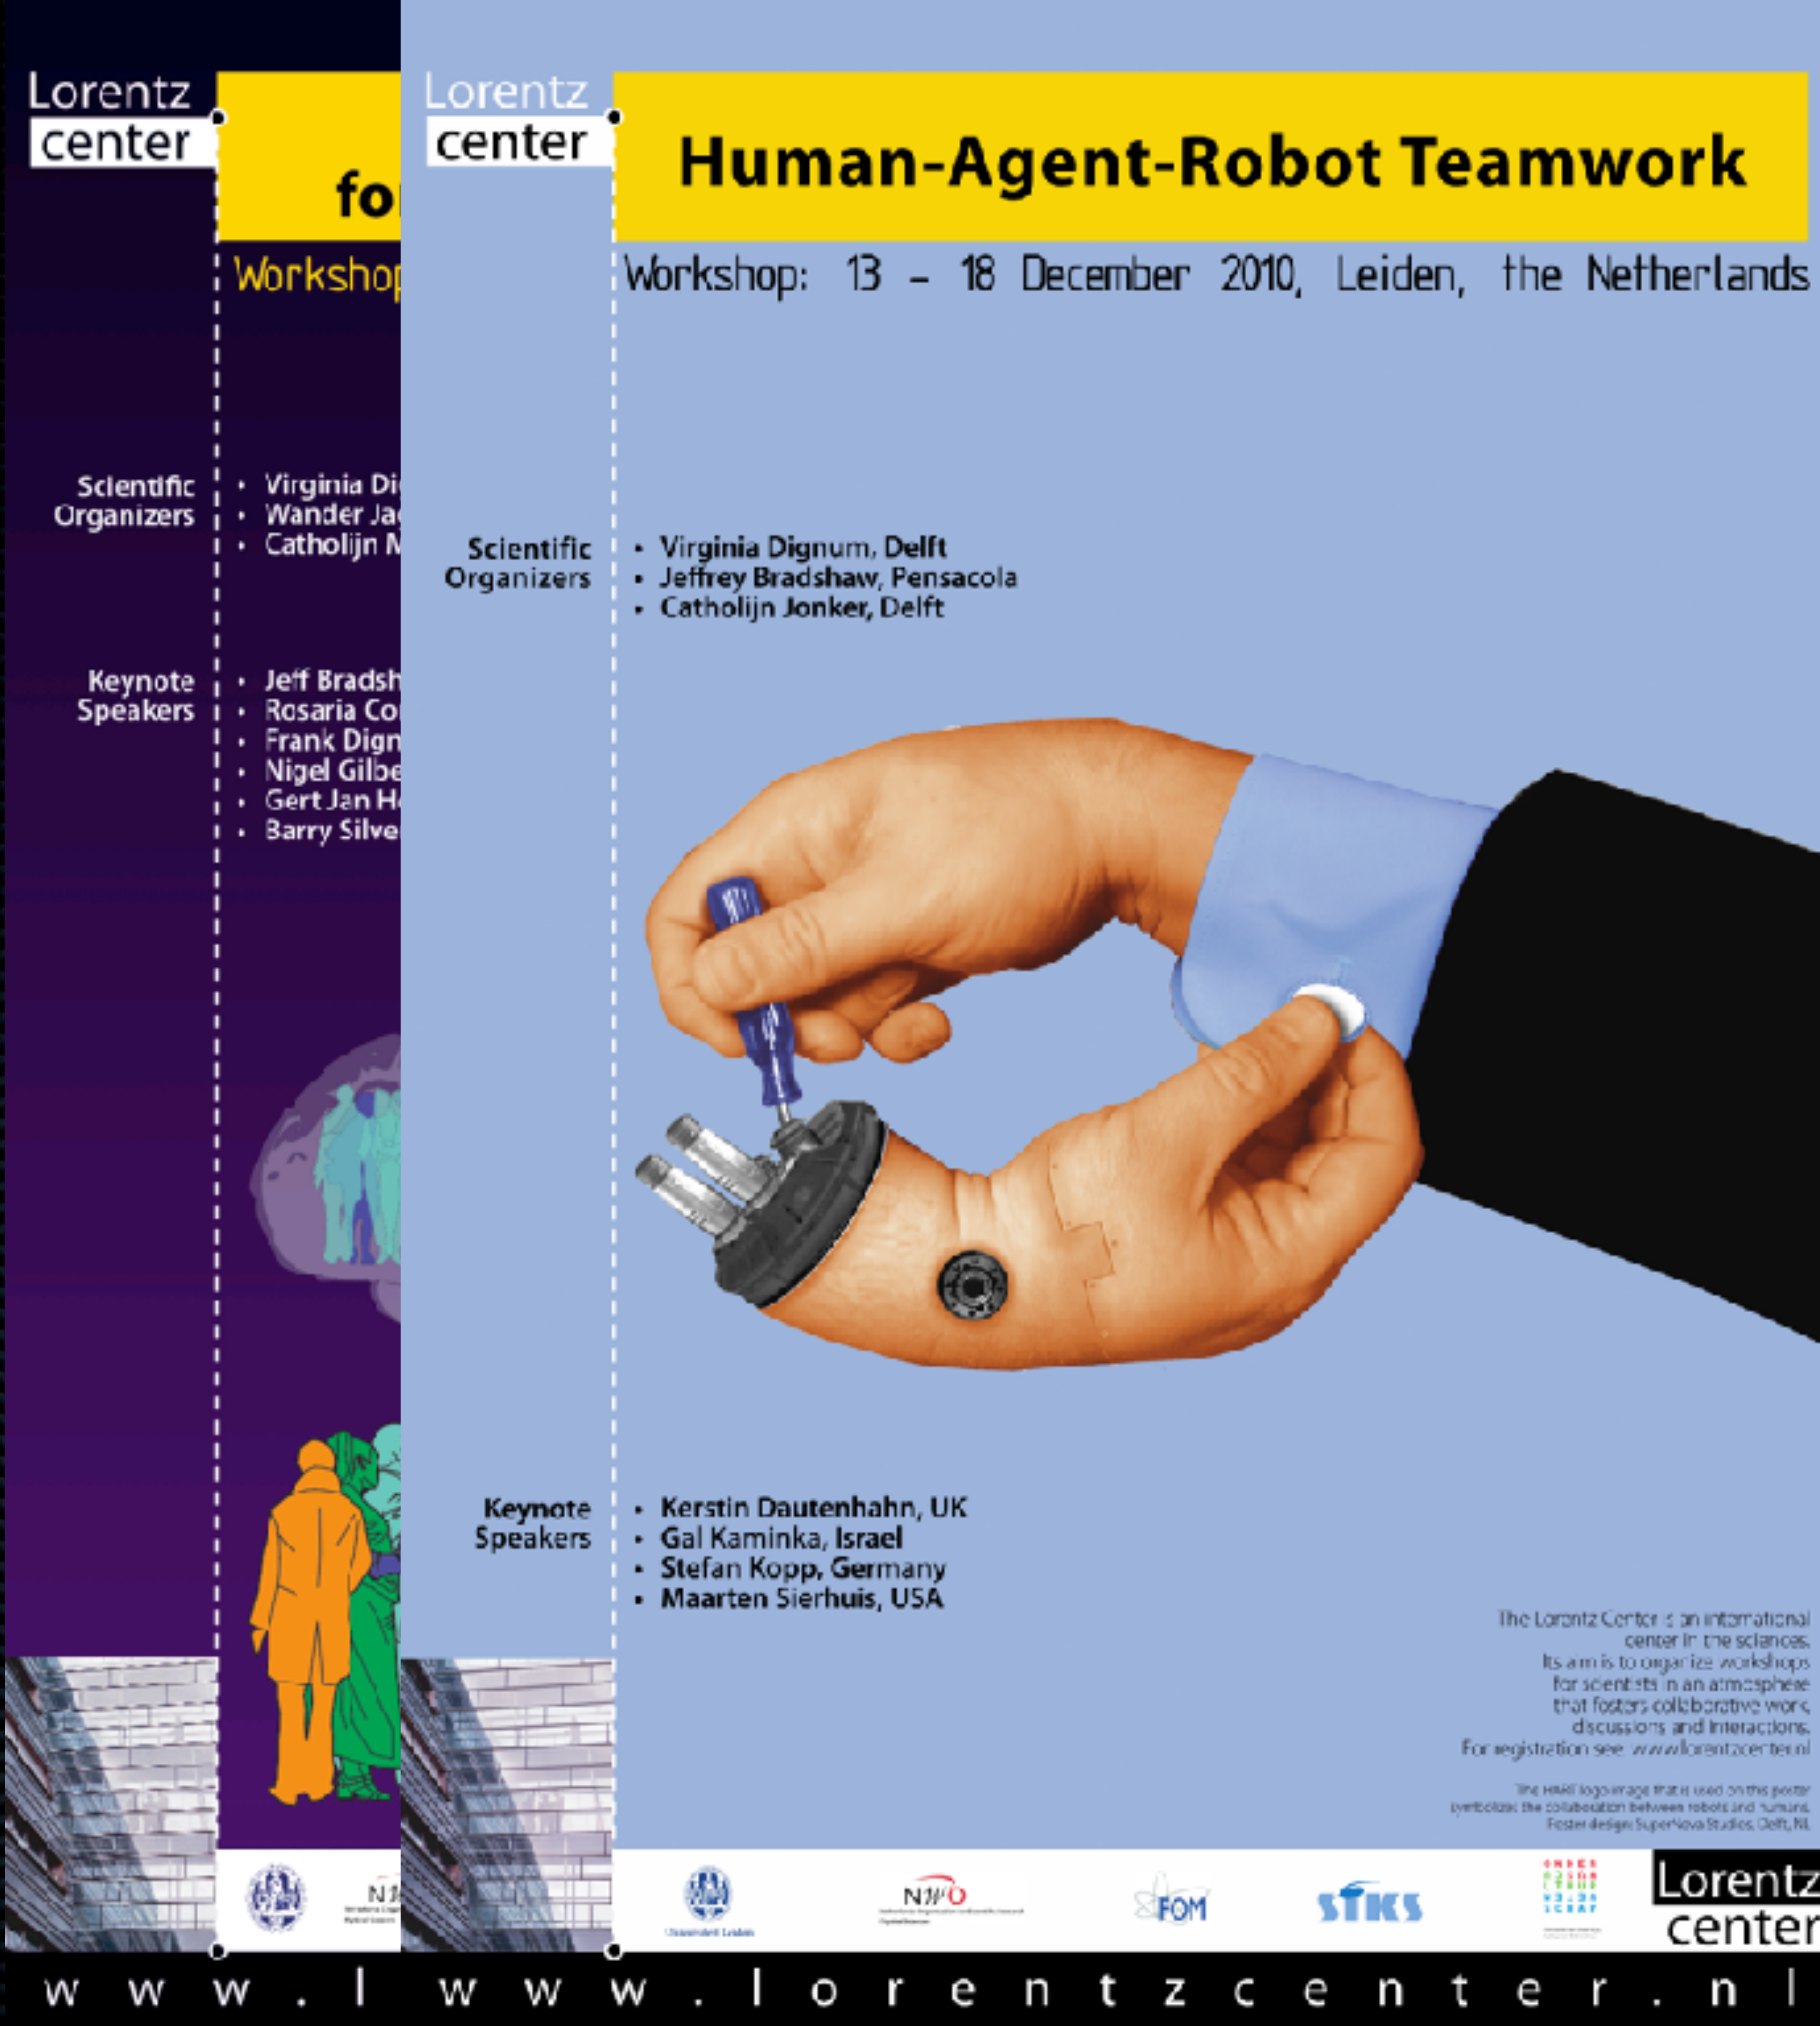
Lorentz
center
fo
Workshop
Scientific
Organizers
Virginia Di
Wander Ja
Catholijn N
Keynote
Speakers
Jeff Bradsh
Rosaria Co
Frank Dign
Nigel Gilbe
Gert Jan H
Barry Silve
Lorentz
center
Human-Agent-Robot Teamwork
Workshop: 13 – 18 December 2010, Leiden, the Netherlands
Scientific
Organizers
• Virginia Dignum, Delft
• Jeffrey Bradshaw, Pensacola
• Catholijn Jonker, Delft
Keynote
Speakers
• Kerstin Dautenhahn, UK
• Gal Kaminka, Israel
• Stefan Kopp, Germany
• Maarten Sierhuis, USA
The Lorentz Center is an international
center in the sciences.
Its aim is to organize workshops
for scientists in an atmosphere
that fosters collaborative work,
discussions and interactions.
For registration see www.lorentzcenter.nl
Lorentz
center
w w w . l o r e n t z c e n t e r . n l

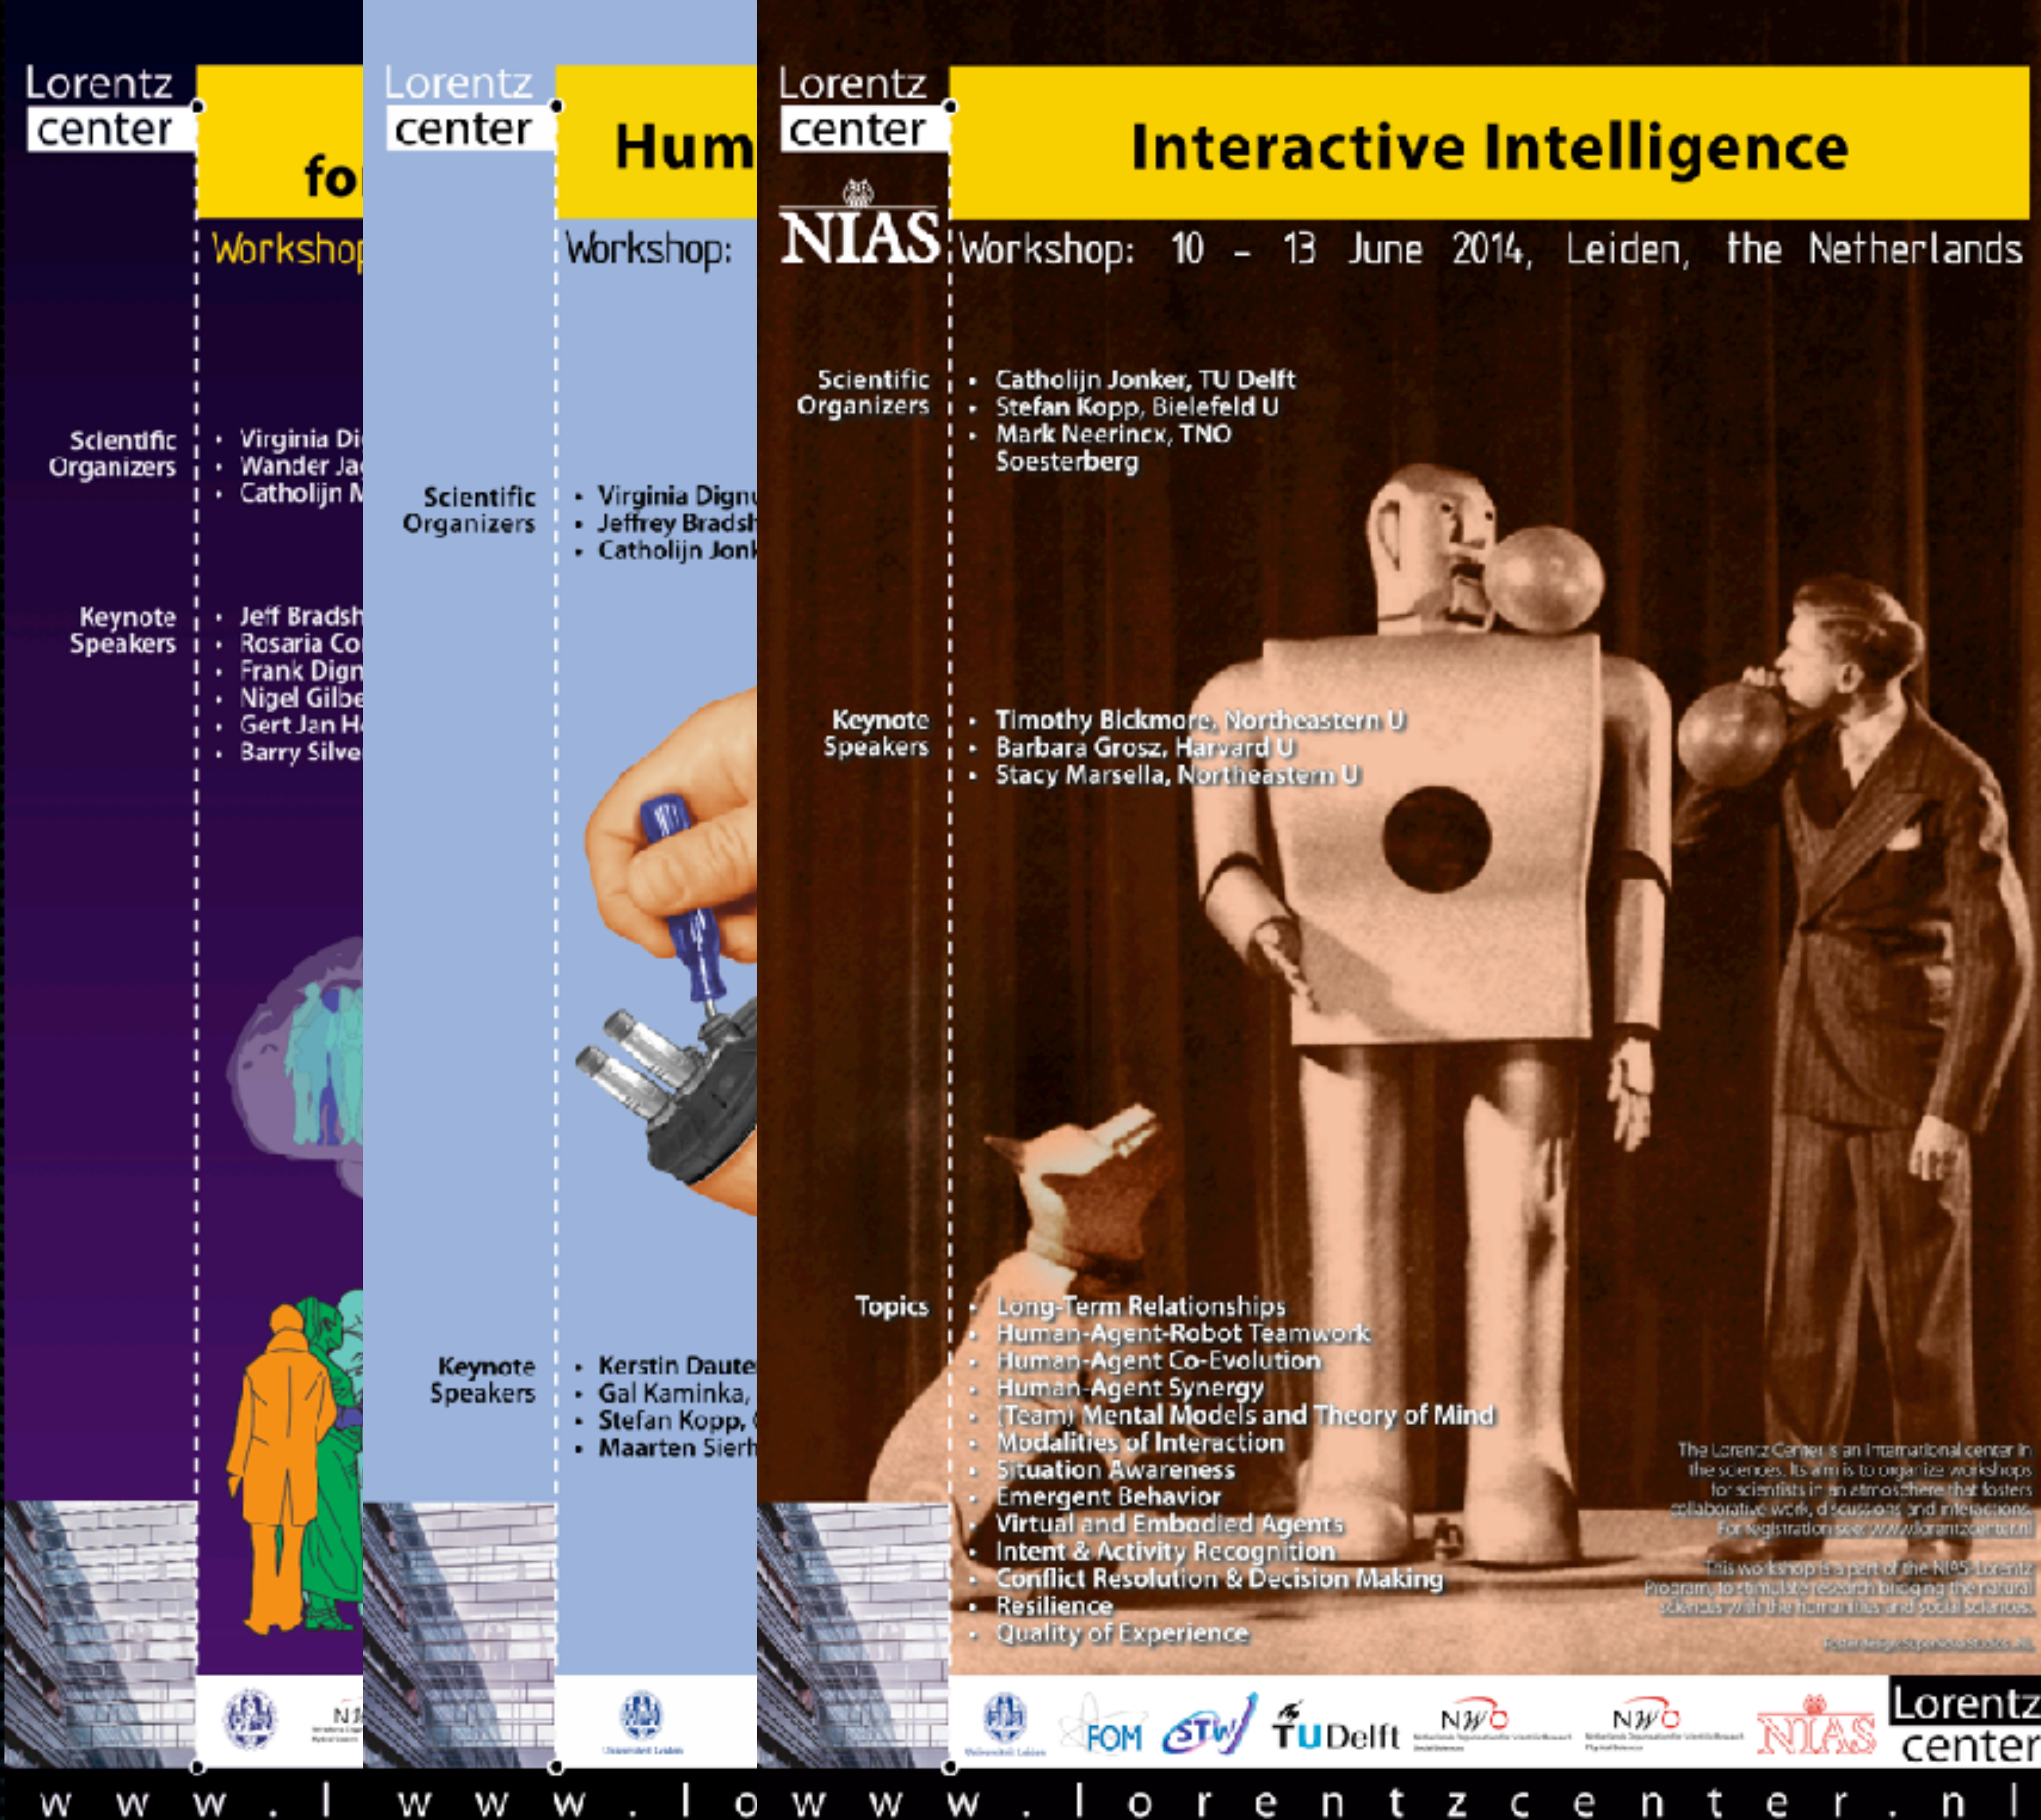
Lorentz center
Interactive Intelligence
NIAS
Workshop: 10 – 13 June 2014, Leiden, the Netherlands
Scientific Organizers
• Catholijn Jonker, TU Delft
• Stefan Kopp, Bielefeld U
• Mark Neerincx, TNO Soesterberg
Keynote Speakers
• Timothy Bickmore, Northeastern U
• Barbara Grosz, Harvard U
• Stacy Marsella, Northeastern U
Topics
• Long-Term Relationships
• Human-Agent-Robot Teamwork
• Human-Agent Co-Evolution
• Human-Agent Synergy
• (Team) Mental Models and Theory of Mind
• Modalities of Interaction
• Situation Awareness
• Emergent Behavior
• Virtual and Embodied Agents
• Intent & Activity Recognition
• Conflict Resolution & Decision Making
• Resilience
• Quality of Experience
FOM
STW
TUDelft
NWO
NIAS
Lorentz center
Lorentz center
Hum
Workshop:
Scientific Organizers
Keynote Speakers
Lorentz center
Workshop
Scientific Organizers
Keynote Speakers
w w w . l o r e n t z c e n t e r . n l

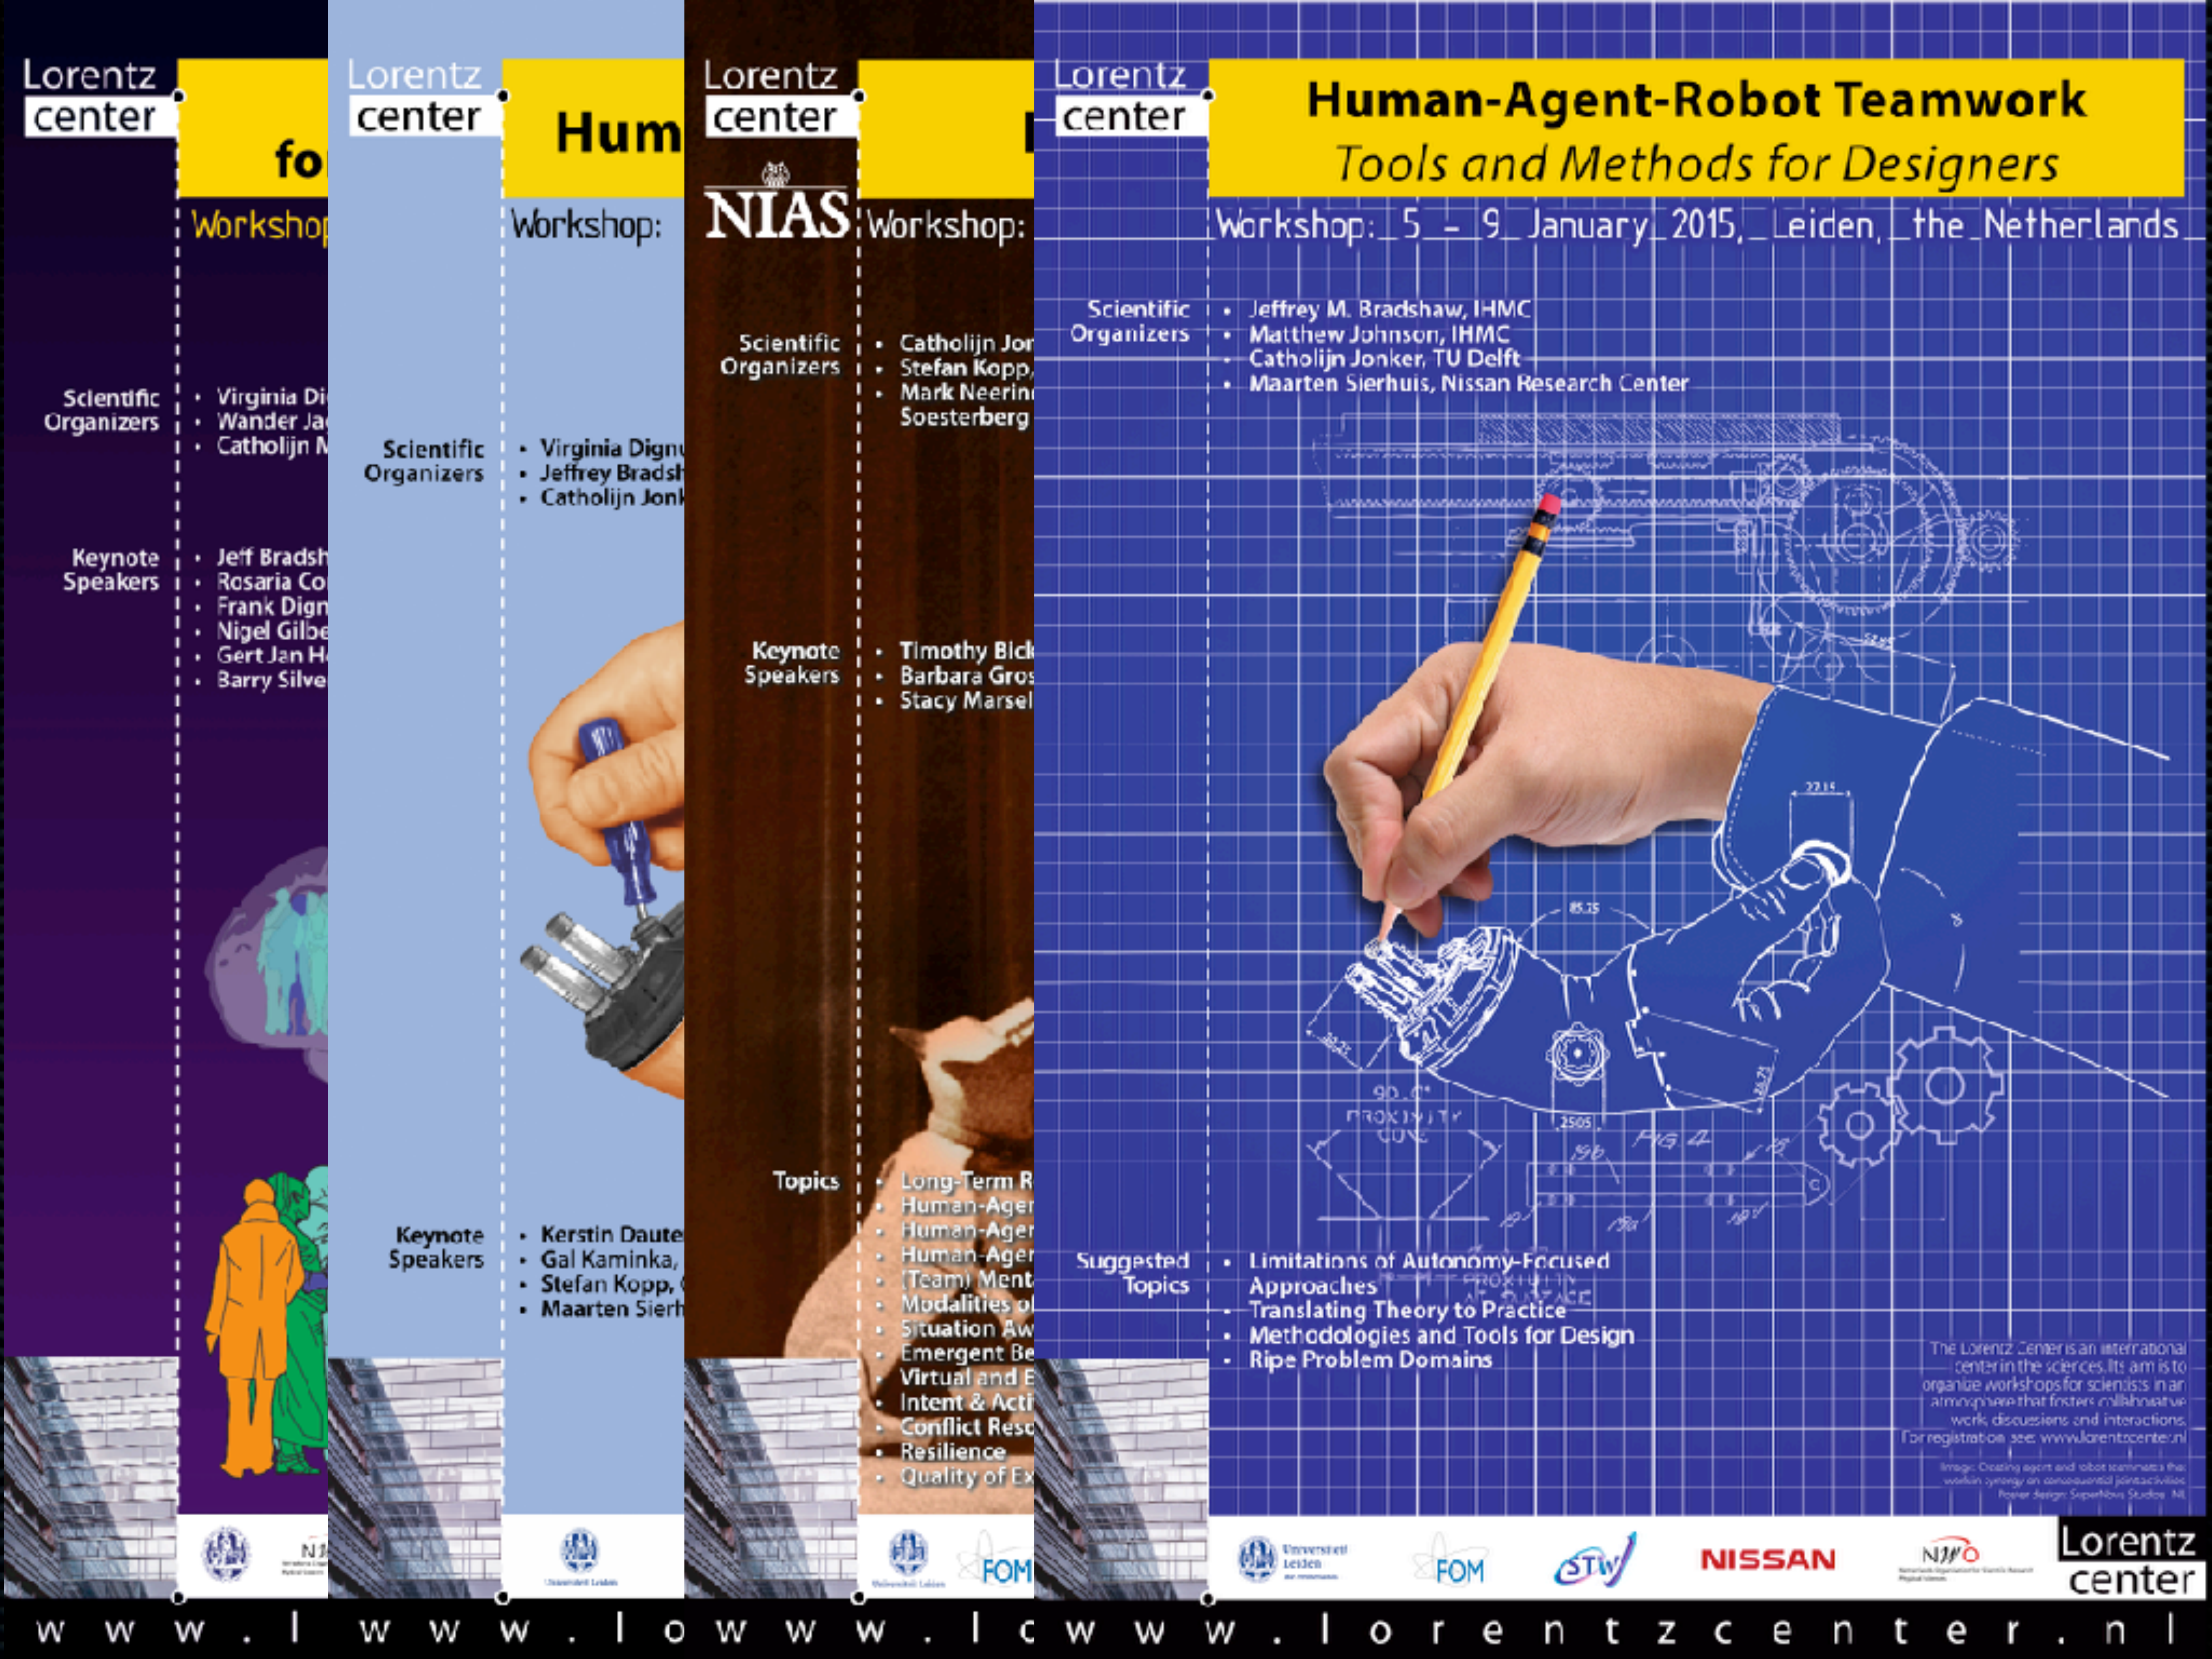
Lorentz center
Human-Agent-Robot Teamwork
Tools and Methods for Designers
Workshop: 5 – 9 January 2015, Leiden, the Netherlands
Scientific Organizers
Jeffrey M. Bradshaw, IHMC
Matthew Johnson, IHMC
Catholijn Jonker, TU Delft
Maarten Sierhuis, Nissan Research Center
Suggested Topics
Limitations of Autonomy-Focused Approaches
Translating Theory to Practice
Methodologies and Tools for Design
Ripe Problem Domains
FOM
STW
NISSAN
NWO
NIAS
Keynote Speakers
Topics
Workshop:
www.lorentzcenter.nl


Lorentz center
Astronomy Workshops 2010
Lorentz center
Life Sciences Workshops 2010
Lorentz center
Physics Workshops 2010
Lorentz center
Mathematics and Computer Science Workshops 2010

towards a Framework for Joint Action (2nd e
2 Mar 2015 Portland, USA. (United States)

Main menu
Home
Call for Papers
Important Dates
Organizers
Contact
Submit

My Space

User name

login

Password

password

Login

› Lost password ?
› Create account

HELP
@ Contact

Call for Papers

Aim

The aim of this workshop is to allow resea
joint action, roboticists but also philosoph
psychologists, to have a context for discu
progressive elaboration of a framework fo
action.

To achieve this goal, we propose to the co
common example (as it is sometimes don
planning competition) with the goal to ide
and skills needed for the successful perfor
action and to see which of these are prese
missing in any of our architectures. This s
build upon each other's experience to furt
work. We hope that this could be a first m
workshop in the next years.

Topics

We are seeking to frame joint action, inter
include:

- joint goal establishment and negotia

HRI 2015 PORTLAND/USA

Workshop on Human-Robot Teaming
March 2nd, 2015

Home	Submissions	Program	Organizers

Overview

Human-Robot Teaming is a full day workshop to be held on March 2nd, 2015 at HRI2015 in Portland, Oregon. We seek an interdisciplinary discussion covering the various facets of teamwork facilitation between humans and robots.

Developing collaborative robots that can productively and safely operate out of isolation in uninstrumented, human-populated environments is an important goal for the field of robotics. The development of such agents, those that handle the dynamics of human environments and the complexities of interpreting human interaction, is a strong focus within Human-Robot Interaction and involves underlying research questions deeply relevant to the broader robotics community.

This workshop aims to bring together peer-reviewed technical and positional contributions spanning a multitude of topics within the domain of human-robot teaming. This workshop seeks to bring together researchers from a wide array of human-robot interaction research concentrations with the common ideal of enabling humans and robots to better work together towards common goals. The morning session will be devoted to gaining insight from invited speakers and contributed papers, while the afternoon session will heavily emphasize participant interaction via poster presentations, breakout sessions, and an expert panel discussion.

Papers and extended abstracts may be submitted here:
https://easychair.org/conferences/?conf=hrihrt2015

Important Dates

- **26 Jan 2015 -- Submission Deadline**
- 12 Feb 2015 -- Notification of Acceptance
- 24 Feb 2015 -- Camera ready papers due
- 02 Mar 2015 -- Workshop

Contact

bradley.h.hayes@yale.edu

Topic Summary

- Task planning under uncertainty
- Empirical Methods for Team Evaluation
- Motion planning in multi-agent or dynamic environments
- Collaborator action and preference modeling
- Interpreting social signals for intention recognition
- Applications of sliding autonomy
- Requesting assistance or failure recovery in teams

Outline

- History of HART
- ***Some Basic Concepts***
- Purpose of the 2015 Workshop

Some Basic Concepts

- Working Separately vs. Working Together
- Joint Activity Theory
- Taskwork vs. Teamwork
- Ten Challenges for Making Automation a Team Player
- Seven Deadly Myths of Autonomous Systems
- Seven Cardinal Virtues of Effective Human-Machine Teamwork

Some Basic Concepts

- ***Working Separately vs. Working Together***
- Joint Activity Theory
- Taskwork vs. Teamwork
- Ten Challenges for Making Automation a Team Player
- Seven Deadly Myths of Autonomous Systems
- Seven Cardinal Virtues of Effective Human-Machine Teamwork

Paul Fitts: HABA-MABA Chart

HUMANS SURPASS MACHINES IN THE:

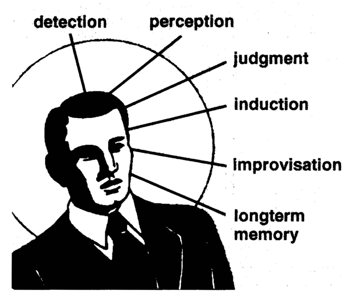

- Ability to detect small amounts of visual or acoustic energy
- Ability to perceive patterns of light or sound
- Ability to improvise and use flexible procedures
- Ability to store very large amounts of information for long periods and to recall relevant facts at the appropriate time
- Ability to reason inductively
- Ability to exercise judgment

MACHINES SURPASS HUMANS IN THE:

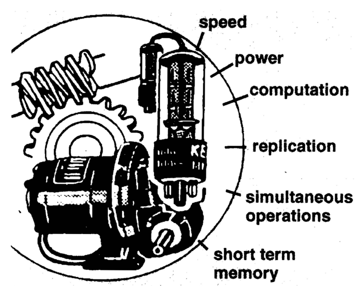

- Ability to respond quickly to control signals, and to apply great force smoothly and precisely
- Ability to perform repetitive, routine tasks
- Ability to store information briefly and then to erase it completely
- Ability to reason deductively, including computational ability
- Ability to handle highly complex operations, i.e., to do many different things at once.

Woods and Hoffman: An “Un-Fitts List”

Machines	
Are constrained in that:	**Need people to:**
Sensitivity to context is low and is ontology-limited	Keep them aligned to context
Sensitivity to change is low and recognition of anomaly is ontology-limited	Keep them stable given the variability and change inherent in the world
Adaptability to change is low and is ontology-limited	Repair their ontologies
They are not “aware” of the fact that the model of the world is itself in the world	Keep the model aligned with the world
People	
Are not limited in that:	**Yet they create machines to:**
Sensitivity to context is high and is knowledge- and attention-driven	Help them stay informed of ongoing events
Sensitivity to change is high and is driven by the recognition of anomaly	Help them align and repair their perceptions because they rely on mediated stimuli
Adaptability to change is high and is goal-driven	Effect positive change following situation change
They are aware of the fact that the model of the world is itself in the world	Computationally instantiate their models of the world

Hoffman, Robert, Paul Feltovich, Kenneth M. Ford, David D. Woods, Gary Klein, and Anne Feltovich. "A rose by any other name… would probably be given an acronym." *IEEE Intelligent Systems*, July-August 2002, 72-80.

Why do we need a new approach?

- *Function Allocation (Fitts)*
  - characterize the general strengths and weaknesses of humans and machines
- *Supervisory Control (Sheridan)*
  - a human oversees autonomous systems, statically allocating tasks to them.
- Adjustable Autonomy (Dorais)
  - autonomous systems operate with dynamically varying levels of independence
- Sliding Autonomy (Dias)
  - Same as adjustable autonomy
- Adaptive Automation (Sheridan)
  - the system must decide at runtime which functions to automate
- Flexible autonomy (Technology horizons)
  - the system can vary the degree of autonomy from essentially none to full
- Mixed-initiative interaction (Allen)
  - An interaction strategy, where each agent can contribute what it does best
- Collaborative Control (Fong)
  - Allows the human to close the perceptual or cognitive loops
- Cognitive Task Analysis, Human Factors and others
  - Provides an understanding of human needs, usability, etc.

Task Allocation

Dynamic
Task Allocation

Both parties

Not just task
allocation

Human side

Why do we need a new approach?

Focusing solely on autonomy ignores issues that have plagued systems from delivering the promised improvements in performance

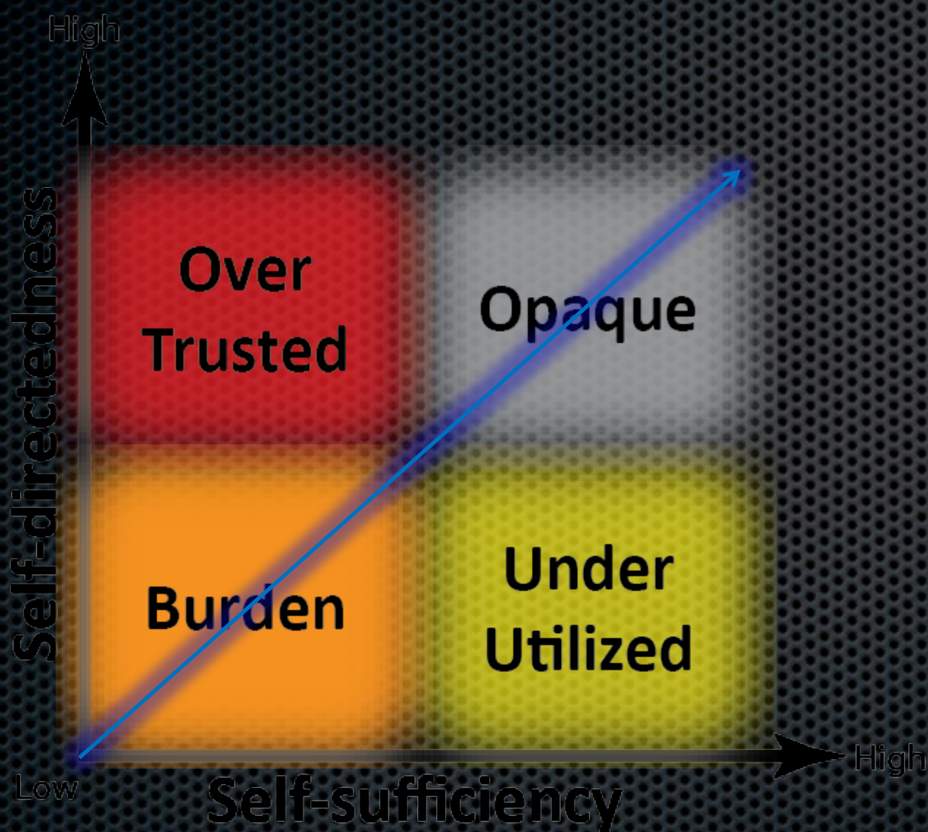

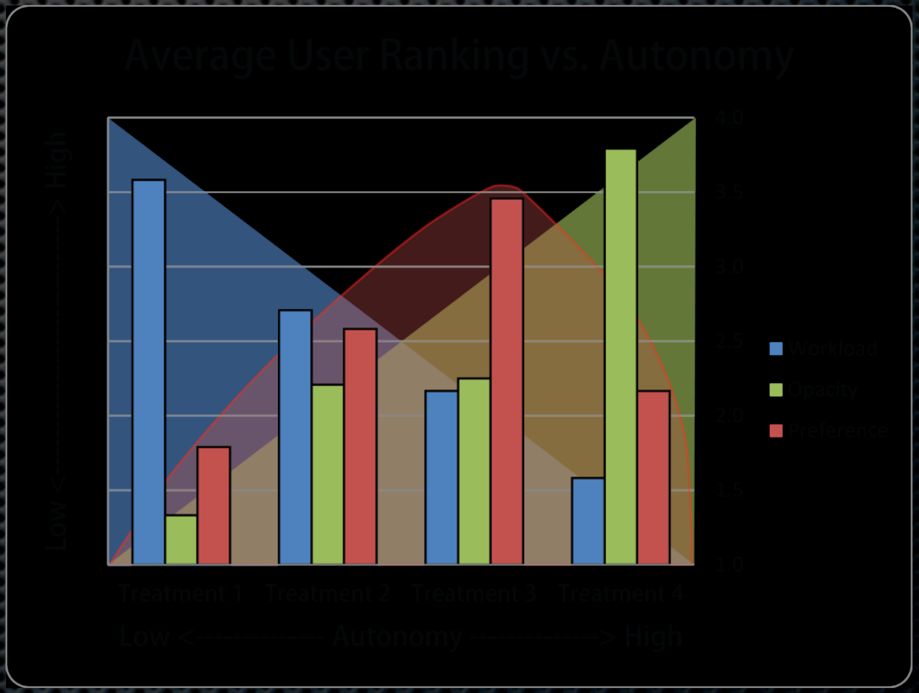

Why do we need a new approach?

Few of these approaches provide a method or a comprehensive approach to determining requirements and most are based on LOA.

- Functional Differences Matter
- Levels Are Neither Ordinal nor Representative of Value
- Autonomy is Relative to the Context of the Activity
- Levels of Autonomy Encourage Reductive Thinking
- The Levels of Autonomy Concept Is Insufficient to Meet Future Challenges
- Levels Provide Insufficient Guidance to the Designer

Working *for* People vs. Working *with* People

- There are situations where the goal of minimizing human involvement with autonomous systems can be argued effectively
- However, many of the most challenging deployments of autonomous systems in the future will continue to involve people in complementary roles (not just as supervisors of autonomy), with the autonomous systems working as part of a world filled with people
  - E.g., DARPA Robotic Challenge

Some Basic Concepts

- Working Separately vs. Working Together
- ***Joint Activity Theory***
- Taskwork vs. Teamwork
- Ten Challenges for Making Automation a Team Player
- Seven Deadly Myths of Autonomous Systems
- Seven Cardinal Virtues of Effective Human-Machine Teamwork

Aspects of Joint Activity

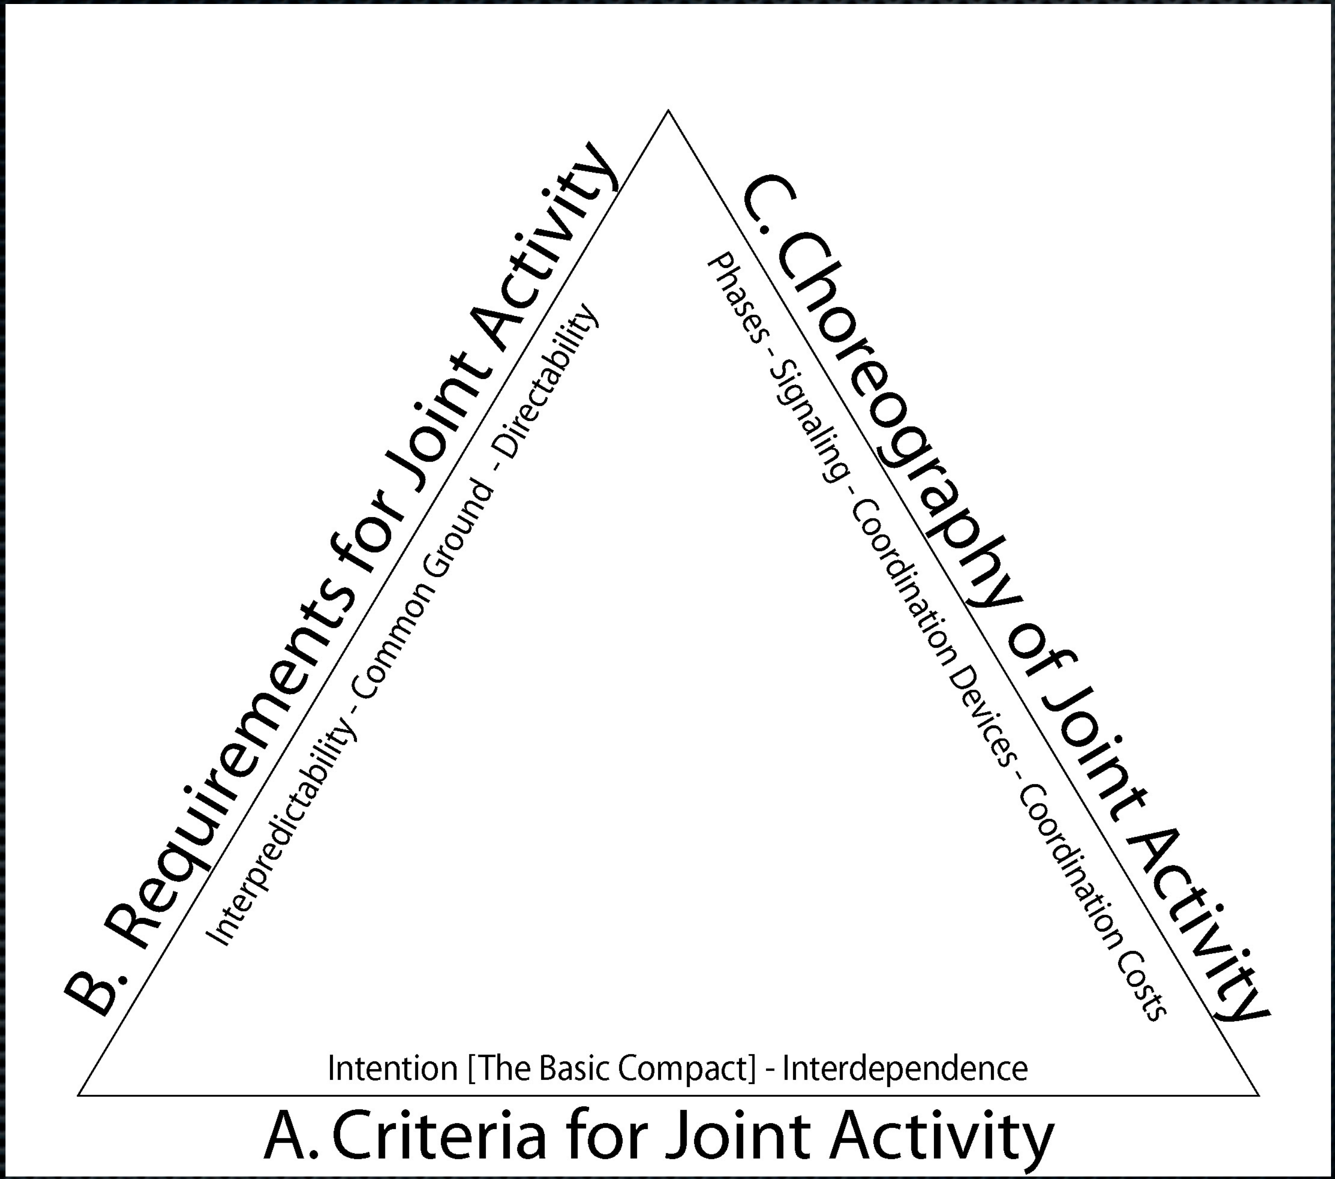

Klein, G., Feltovich, P., Bradshaw, J. M., & Woods, D. D. (2005). Common ground and coordination in joint activity. *Organizational Simulation.* W. B. Rouse and K. R. Boff. New York City, NY, John Wiley.

Some Basic Concepts

- Working Separately vs. Working Together
- Joint Activity Theory
- ***Taskwork vs. Teamwork***
- Ten Challenges for Making Automation a Team Player
- Seven Deadly Myths of Autonomous Systems
- Seven Cardinal Virtues of Effective Human-Machine Teamwork

Cohen and Levesque: Joint Intentions

- Basic concepts:
  - Agents form teams by adopting joint persistent goals (JPG's) to achieve a team action
  - JPG's hold if and only if all team members mutually believe:
    - the goal is not yet achieved
    - they want the goal to be achieved
    - until the goal is known to be achieved, unachievable, or no longer relevant, they should persist in holding the goal
  - If a team member discovers the goal to be achieved, unachievable, or no longer relevant, it will tell its teammates
- Key points
  - Teamwork involves more than simple coordination
  - Teamwork knowledge should be explicitly modeled as a separate domain

Cohen, P. R. and H. J. Levesque (1991). Teamwork, Menlo Park, CA: SRI International.

Teamwork and Taskwork are Separable

Soccer Taskwork:

- Kicking to a target
- Dribbling, tackling
- Tracking the ball, goal …

Soccer Teamwork:

- Allocating players to roles
- Synchronizing tactics
- Sharing relevant information
- …….

Slide from Gal A. Kaminka, *Robots are Agents, Too!*

Some Basic Concepts

- Working Separately vs. Working Together
- Joint Activity Theory
- Taskwork vs. Teamwork
- ***Ten Challenges for Making Automation a Team Player***
- Seven Deadly Myths of Autonomous Systems
- Seven Cardinal Virtues of Effective Human-Machine Teamwork

Ten Research Challenges

1. Forming and maintaining the Basic Contract
2. Forming and maintaining adequate models of others' intentions and actions
3. Maintaining predictability without hobbling adaptivity
4. Maintaining adequate directability
5. Effective signaling of pertinent aspects of status and intentions
6. Observing and interpreting signals of status and intentions
7. Engagement in goal negotiation
8. Autonomy and planning technologies that are incremental and collaborative
9. Attention management
10. Controlling the costs of coordinated activity

Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., & Feltovich, P. (2004). "Ten challenges for making automation a "team player" in joint human-agent activity." *IEEE Intelligent Systems* 19(6): 91-95.

Some Basic Concepts

- Working Separately vs. Working Together
- Joint Activity Theory
- Taskwork vs. Teamwork
- Ten Challenges for Making Automation a Team Player
- ***Seven Deadly Myths of Autonomous Systems***
- Seven Cardinal Virtues of Effective Human-Machine Teamwork

Seven Deadly Myths of Autonomous Systems

- 1. Autonomy is unidimensional
- 2. The conceptualization of “levels of autonomy” is a useful scientific grounding for the development of autonomous system roadmaps
- 3. Autonomy is a widget
- 4. “Autonomous systems” are autonomous
- 5. Once “achieved,” “full autonomy” obviates the need for human-machine collaboration
- 6. As machines acquire more “autonomy,” they work as simple multipliers of human capability
- 7. “Full autonomy” is not only possible, but always desirable

Bradshaw, J.M, Robert R. Hoffman, Matthew Johnson, and David D. Woods. The Seven Deadly Myths of "Autonomous Systems.” *IEEE Intelligent Systems*, May/June 2013 (vol. 28 iss. 3), pp. 54-61.

Levels of Autonomy and Supervisory Control

Level	Description
High	10. The computer decides everything, acts autonomously, ignoring the human.
	9. The computer informs the human only if it, the computer, decides to.
	8. The computer informs the human only if asked, or
	7. The computer executes automatically, then necessarily informs the human, and
	6. The computer allows the human a restricted time to veto before automatic execution, or
	5. The computer executes that suggestion if the human approves, or
	4. The computer suggests one alternative
	3. The computer narrows the selection down to a few, or
	2. The computer offers a complete set of decision/action alternatives, or
Low	1. The computer offers no assistance; the human must take all decisions and actions.

Adapted from Parasuraman, Sheridan, and Wickens, 2000

Things Are Not That Simple…

- The notion of “levels” of autonomy can be deceptive
  - Autonomy is not an independent property of a system, but must be described in terms of particular tasks and situations
    - No system—and, for that matter, no person—is capable enough to be able to perform “autonomously” in every task and situation
    - On the other hand, even the simplest machine can function autonomously if the task and context is sufficiently constrained.
  - Autonomy is multi-dimensional

Dynamics of Trust Calibration

This simplified diagram is meant to convey an intuition about how degrees of appropriately calibrated trust (or mistrust) vary over time and changing context. The green zone indicates acceptable bounds on trust calibration. Above the green zone is a zone of over-trust. Below it is a zone of under-reliance. Active trust management requires developing effective ways of revealing context-sensitive human and machine trust signatures, allowing human and machines to accurately calibrate degree of trust in others' capabilities in a given situation. It also requires developing means for humans and machines to actively probe others' capabilities in order to understand whether others are operating within their competence envelopes.

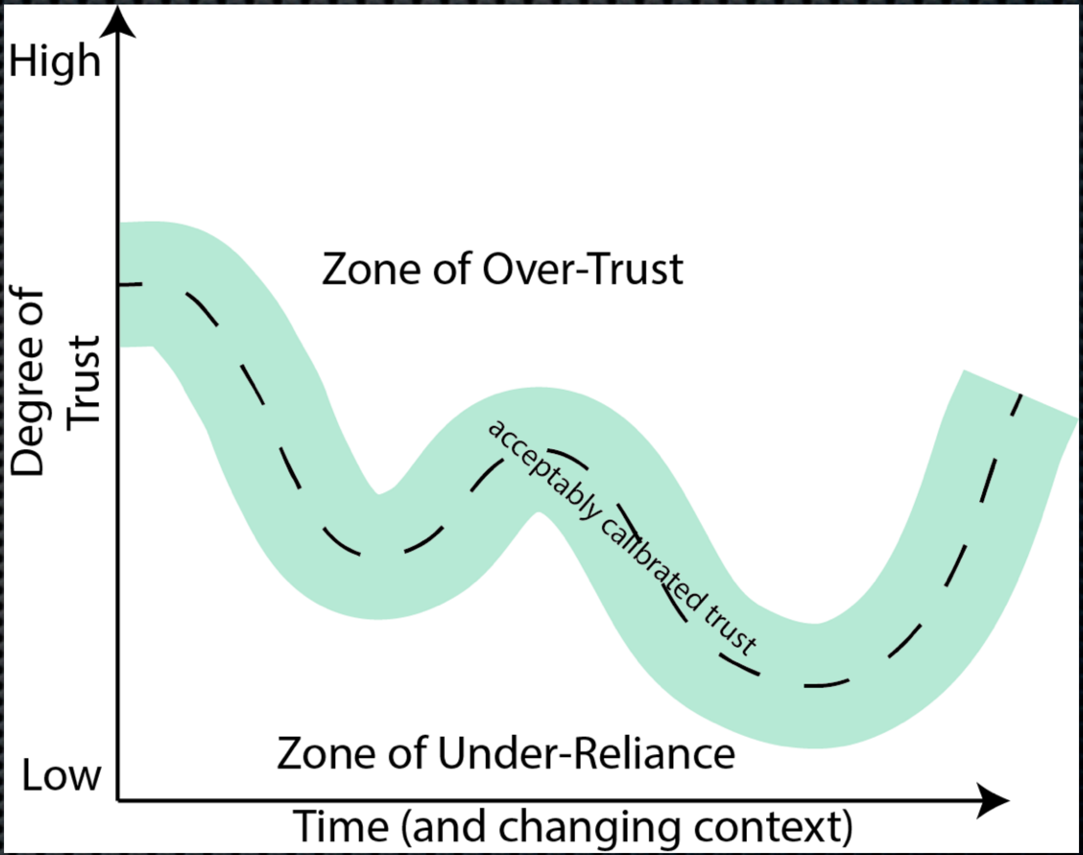

Johnson, Matthew, J.M. Bradshaw, R. R. Hoffman. Trust Calibration as applied to Micro Air Vehicles. *AAAI Spring Symposium on Trust*, Stanford, CA, March 2013

Problems with the “Levels of Autonomy” Approach

- Problem 1: Functional Differences Matter (e.g., making decision vs. performing action, teamwork vs. taskwork)
- Problem 2: Levels Are Neither Ordinal nor Representative of Value
- Problem 3: Autonomy is Relative to the Context of the Activity
- Problem 4: Levels of Autonomy Encourage Reductive Thinking (e.g., viewing parallel activities as sequential)
- Problem 5: The Concept of Levels of Autonomy Is Insufficient to Meet Future Challenges
- Problem 6: Levels Provide Insufficient Guidance to the Designer

Johnson, Matthew, Jeffrey M. Bradshaw, Paul J. Feltovich, Robert R. Hoffman, Catholijn Jonker, Birna van Riemsdijk, and Maarten Sierhuis. Beyond Cooperative Robotics: The Central Role of Interdependence in Coactive Design. *IEEE Intelligent Systems*, May/June 2011 (vol. 26 iss. 3), pp. 81-88.

Erroneous Notions about Adjustable Autonomy and Adaptive Function Allocation

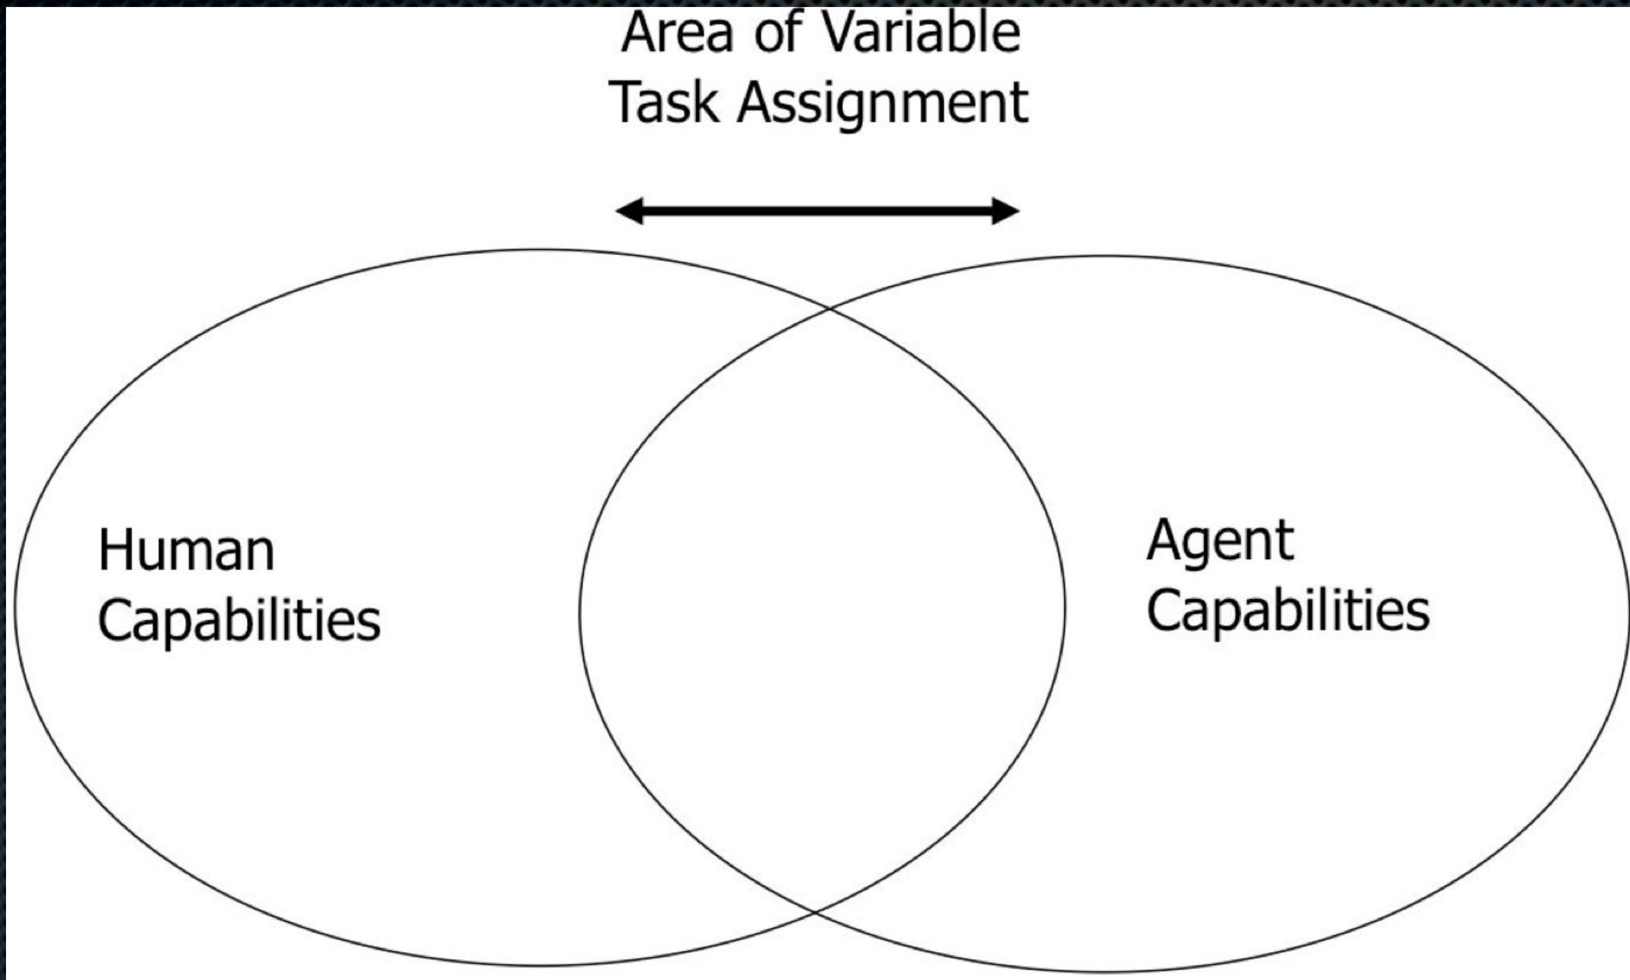

Bradshaw, J.M., Paul Feltovich, Hyuckchul Jung, Shri Kulkarni, William Taysom, and Andrzej Uszok. Dimensions of adjustable autonomy and mixed-initiative interaction. In Agents and Computational Autonomy: Potential, Risks, and Solutions. Lecture Notes in Computer Science 2969, edited by Matthias Nickles, Michael Rovatsos and Gerhard Weiss, 17-39. Berlin, Germany: Springer-Verlag, 2004.

Things Are Not That Simple…

- Many functions in complex systems are shared by humans and machines
- Automated assistance of whatever kind does not simply enhance our ability to perform the task: it changes the nature of the task itself—usually adding new kinds of work that must be executed concurrently (Don Norman)
  - Substitution Myth (David Woods)
- Overly simple approaches fail to exploit opportunities for human-machine synergy

Norman, D.A. "Cognitive artifacts." In *Designing Interaction: Psychology at the Human-Computer Interface,* edited by J.M. Carroll, 17-38. Cambridge: Cambridge University Press, 1992.
Christofferson, K., and David D. Woods. "How to make automated systems team players." In *Advances in Human Performance and Cognitive Engineering Research, Vol. 2,* edited by E. Salas. JAI Press, Elsevier, 2002.

Some Basic Concepts

- Working Separately vs. Working Together
- Joint Activity Theory
- Taskwork vs. Teamwork
- Ten Challenges for Making Automation a Team Player
- Seven Deadly Myths of Autonomous Systems
- ***Seven Cardinal Virtues of Effective Human-Machine Teamwork***

Seven Cardinal Virtues of Human-Machine Teamwork

- 1. Clarity: Focus on improving mission performance of the work system, not on maximizing autonomous capabilities
- 2. Humility: Assess the sweet spot in development effort payoff
- 3. Resilience: If you don't plan to fail, you fail to plan
- 4. Helpfulness: Think combine and succeed, not divide and conquer
- 5. Cohesiveness: Design for teamwork in addition to taskwork
- 6. Integrity: Designing for human-machine teamwork goes deeper than the user interface
- 7. Thrift: Don't simply downsize human involvement, rightsize it

Johnson, M., Bradshaw, J.M., Hoffman, R. R., Feltovich, P. J., and Woods, D. D. Seven Cardinal Virtues for Human-Machine Teamwork: Examples from the DARPA Robotic Challenge. IEEE Intelligent Systems, November/December 2014 (vol. 29 iss. 6), pp. 74-80.

Coactive Design

In sophisticated human-agent systems,
the underlying ***interdependence*** of joint activity
is the critical design feature.

Dependent — Independent — Interdependent

1997

2002

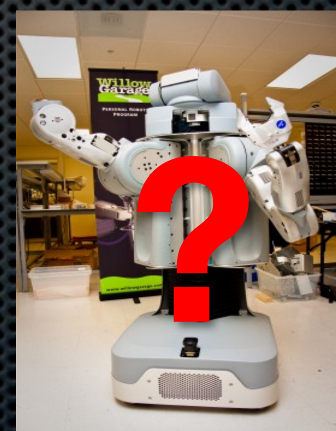

Future

Johnson, M., J.M. Bradshaw, P. J. Feltovich, C. M. Jonker, M. B. van Riemsdijk, and M. Sierhuis. Coactive design: Designing support for interdependence in joint activity. Journal of Human-Robot Interaction, Vol. 3, No. 1, 2014, pp. 43-69.

Supporting Interdependence

Human needs	Issue	Robot needs
What is the robot doing?	Mutual Transparency	What is the intent of the human?
Why did the robot do that?	Mutual Explainability	What is the task context?
What is the robot going to do next?	Mutual Predictability	What does the human need from me?
Can we make the robot do what we need?	Mutual Directability	Can the human provide help?
Does use of autonomy add value?	Mutual Cost Benefit Managment	Will my actions provide value to the human?

Coactive Emergence

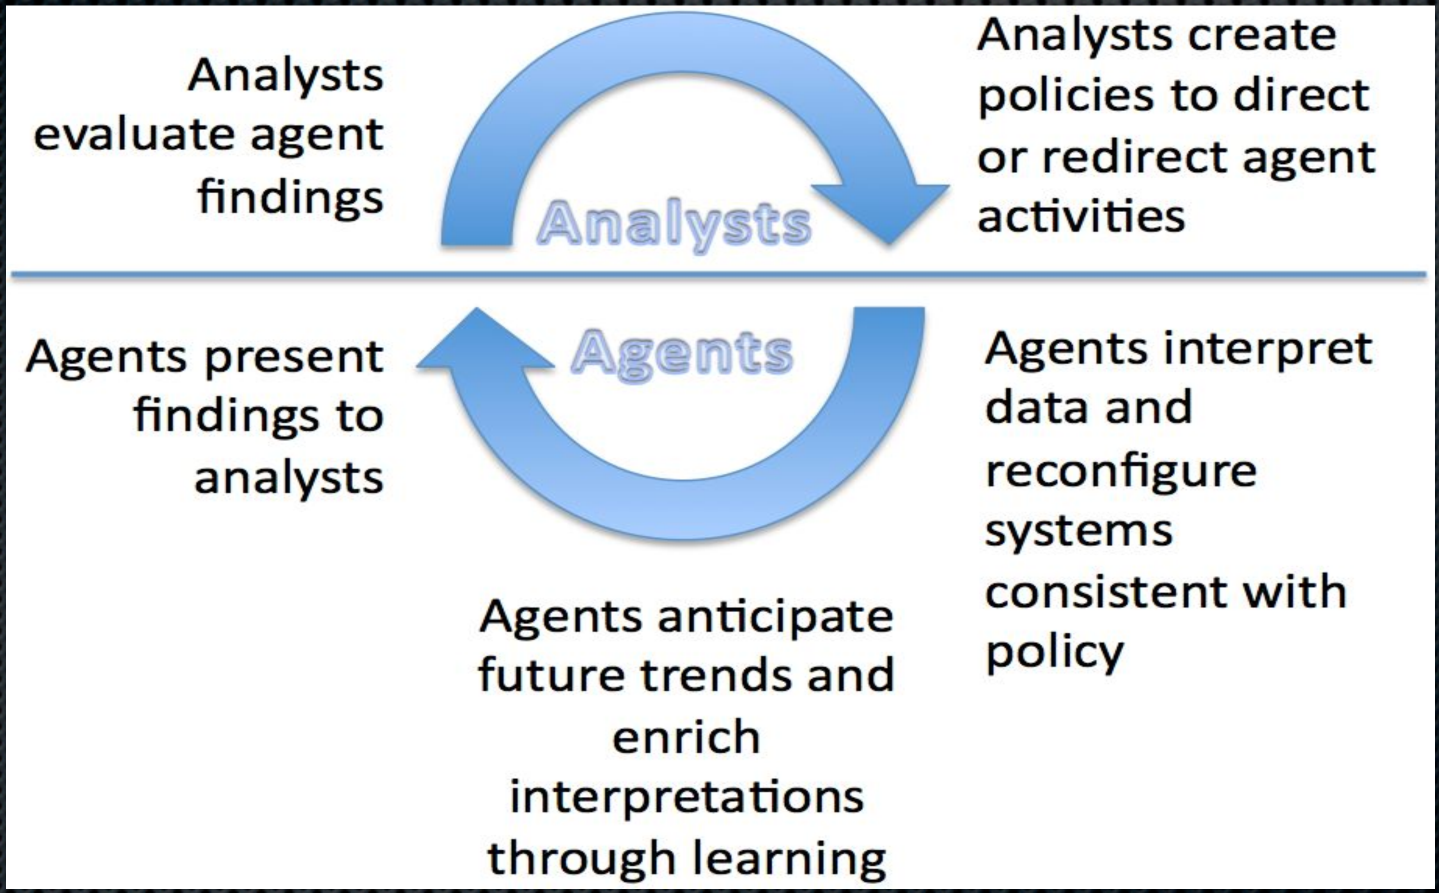

Bunch, L., J. M. Bradshaw, T. Eskridge, R. Hoffman, and M. Johnson. Principles for Human-Centered Interaction Design, Part 2: Can Machines and Humans Think *Together*? *IEEE Intelligent Systems,* in press.

Outline

- History of HART
- Some Basic Concepts
- ***Purpose of the 2015 Workshop***

Barriers to HART in Practice

- Dispositional barriers: Some agent and robot researchers get into the field specifically because they want to do research on autonomous capabilities
- Hollywood glamor: Movies and media glamorize fully autonomous systems
- Research sponsor misconceptions: Some research sponsors think that autonomous capabilities are the holy grail for the best and cheapest agent/robot performance
- Engineering and design barriers: Methods, tools, and good examples lacking to inspire useful implementations

Purpose of the 2015 Workshop

- Reducing Barriers to the Adoption of HART Approaches by Developing Usable Tools and Methods for Designers and Engineers