

2010 IEEE 9th International Conference on Development and Learning

August 18-21, 2010
University of Michigan - Ann Arbor



Conference Program

Table of Contents

Welcome to 2010 ICDL!	1
2010 ICDL Committee	2
2010 Reviewers	3
2010 ICDL Technical Program	4
Keynote Speakers: Biographies & Abstracts	9
Oral Presentations: Full Length Papers	14
Poster Presentations: Full Length Papers	21
Two Page Abstracts	29

2010 IEEE 9th ICDL Sponsors



Welcome to 2010 ICDL!

Welcome to the Ninth IEEE International Conference on Development and Learning, meeting at the University of Michigan in Ann Arbor.

As you know, we are brought together by a shared interest in the foundations of knowledge, in both humans and robots, and how that knowledge can be acquired from the foundations up, through experience and interaction with the world.

For some of us, the focus is on the nature of human knowledge, learning and development, and computational models including robots that serve as tools for creating hypotheses adequate for these complex phenomena. For others of us, the focus is on robot learning and performance, and the human example shows how the most intelligent species on the planet meets the challenges of embedded intelligence. And of course, these are not binary choices, or even ends of a spectrum, but landmarks in a complex landscape with many diverse positions.

Accordingly, ICDL-2010 is co-sponsored by the IEEE Computational Intelligence Society and by the Cognitive Science Society. Furthermore, we have received generous gifts from the Toyota Research Center in Ann Arbor, Michigan, and from Microsoft Research in Redmond, Washington. We are very grateful for these contributions.

ICDL is a small interdisciplinary conference with a single presentation track, designed to encourage cross-disciplinary interaction and discussion. This year, ICDL-2010 features four keynote addresses from prominent experts with different specialties and perspectives, 23 full papers accepted for plenary presentation, and 28 full papers accepted for poster presentation. The reviewing process selected these papers from 81 submitted full papers. In addition, 16 short poster abstracts were accepted to encourage the dissemination of late-breaking results.

In the spirit of the interdisciplinary nature of ICDL, we ensured that each paper was assigned to program committee members and auxiliary reviewers from both natural intelligence and computational intelligence backgrounds. Final decisions were made by the Program Chairs, based on the reviews. The reviews were careful, extensive, and helpful to the authors, whatever the outcome of the review process for each paper.

In addition to the organizing and reviewing, ICDL has depended on the efforts of a large number of other people. Publicity Co-Chairs Ian Fasel (North America), Jun Tani (Asia), and Jochen Triesch (Europe) helped publicize the conference. Publication Chair Nicholas Butko assembled the program and the Proceedings DVD, and worked with IEEE to publish and archive your excellent work. Webmasters Vindhya Baddela, Seth Levine, and Jingen Liu created the website and kept it up to date. At the University of Michigan, Karen Alexa and Francine Dolins worked tirelessly to make the conference work. They organized a great group of student volunteers, including Chris Cutcher, Jay Hack, Bangxin Hu, Jong Jin Park, and Grace Tsai. Any attempt to list individual people inevitably omits someone who contributed mightily and deserves recognition (I mean you!), so thank you to that person especially.

The real source of success for a conference starts with the work that has been done at many different home institutions, preparing for good papers, posters, and presentations. The heart of it is the discussions, in the sessions, in the hallways, and over meals and drinks, that guide the work done over the next year(s), and carry the field forward.

Welcome to Ann Arbor, and best wishes for a great conference.

Benjamin Kuipers, General Co-Chair
Thomas Shultz, General Co-Chair
Alex Stoytchev, Program Co-Chair
Chen Yu, Program Co-Chair

2010 ICDL Committee

General Chairs

Benjamin Kuipers, University of Michigan
Thomas Shultz, McGill University

Program Chairs

Alexander Stoytchev, Iowa State University
Chen Yu, Indiana University, Bloomington

Program Committee

Andrea Thomaz, Georgia Institute of Technology
Angelo Cangelosi, University of Plymouth, UK
Anna Fish, Carnegie Mellon University
Aude Billard, EPFL, Switzerland
Brian Scassellati, Yale University
Christopher Conway, Saint Louis University
Cota Nabeshima, Cyberdyne, Inc, Japan
Eliana Colunga, University of Colorado
Erol Sahin, Middle East Technical University, Turkey
Freek Stulp, University of South California
Gedeon Deak, University of California at San Diego
Giorgio Metta, University of Genoa
Hanako Yoshida, University of Houston
Hui Zhang, Indiana University at Bloomington
Jeff Siskind, Purdue University
Jerry Zhu, University of Wisconsin, Madison
Jivko Sinapov, Iowa State University
Jessica Horst, University of Sussex, UK
Karl F. MacDorman, Indiana University
Katharina Rohlfing, University of Bielefeld, Germany
Kerstin Dautenhahn, The University of Hertfordshire, UK
Kevin Gold, Wellesley College
Kris Hauser, Indiana University at Bloomington

Lakshmi Gogate, Florida Gulf Coast University
Lorenzo Natale, Italian Institute of Technology
Luc Berthouze, University of Sussex, UK
Luis Montesano, Instituto Superior Técnico, Lisboa, Portugal
Matthew Schlesinger, Southern Illinois University
Nathan Sprague, Kalamazoo College
Pierre-Yves Oudeyer, INRIA, France
Rick Dale, University of Memphis
Shohei Hidaka, Japan Advanced Institute of Science and Technology
Stephen W. Hart, Italian Institute of Technology, Genoa
Susan Jones, Indiana University, Bloomington
Su Li, Chinese Academy of Sciences
Tim Oates, University of Maryland Baltimore County
Yukie Nagai, Osaka University, Japan
Zhengyou Zhang, Microsoft Research

Publicity Co-Chairs

Asia: Jun Tani, Riken Brain Science Institute
Europe: Jochen Triesch, J.W. Goethe University, Frankfurt
North America: Ian Fasel, University of Arizona

Publication Chair

Nicholas Butko, University of California, San Diego

Website

Jingen Liu, University of Michigan
Vindhya Baddela, University of Michigan
Seth Levine, University of Michigan

2010 Reviewers

The ICDL Program Chairs would like to acknowledge and thank all of the reviewers. Without you, it would be impossible to maintain such a high quality conference year after year. Thank you.

Adam Sheya
Ales Ude
Alessandra Sciutti
Alex Doumas
Alexandra Kirsch
Alfredo Pereira
Andrei Barbu
Andrew Barto
Andrew Olney
Andrew Stout
Aris Alissandrakis
Arjan Gijsberts
Ben Goertzel
Brenna Argall
Brian MacWhinney
C. Anton Rytting
Catherine Havasi
Chris Buckley
Chris Robinson
Chris Vigorito
Christian Balkenius
Christopher Chatham
Christopher Crick
Cindy Bethel
Colin Bannard
Cordula Vesper
Cornelia Caragea
Dan Bush
David Andrzejewski
David Feil-Seifer
David Rakison

Davide Marocco
Dirk Kraft
Dov Katz
Emre Ugur
Erhan Oztop
Eric Meisner
Erik Thiessen
Eris Chinellato
Francesco Nori
Francesco Rea
Frederick Shic
Gary Lupyan
George Konidaris
Heidi Kloos
Hiroki Mori
Ian Fasel
James Kole
Jennifer Richler
Jeremy Loebach
Joe Saunders
John Purdy
Jon Willits
Jonas Ruesch
Jonathan Mugan
Joseph Modayil
Justin Hart
Kathryn Merrick
Katsunari Shibata
Kerr Wesley
Kristy Snyder-Tapp
Louis ten Bosch

Ludovic Righetti
Manuel Lopes
Marco Mirolli
Mark Lee
Martha Palmer
Matt Miller
Maya Cakmak
Mehmet R. Dogar
Michele Brumley
Naomi Feldman
Nicholas Butko
Norbert Krüger
Peter Ford Dominey
Rebecca Brand
Remi van Trijp
Ruben Martinez-Cantin
Ryo Saegusa
Sarah Osentoski
Scott Robinson
Shane Griffith
Shiraj Sen
Si Wu
Siddharth Narayanaswamy
Stefanie Tellex
Steven Cadavid
Tom Erez
Ulrik Beierholm
Vadim Tikhanoff
Vishwanathan Mohan
Vladimir Sloutsky
Vladimir Sukhoy

2010 ICDL Technical Program

Wednesday, August 18

1:00 - 6:00	Registration
6:00 - 7:00	Welcome Reception
7:00 - 8:00	Keynote Presentation @ The Campus Inn <i>David Vernon, Italian Institute of Technology</i>
8:00 -	Social Hour

Thursday, August 19

8:30 - 9:30	Keynote Presentation Rod Grupen, University of Massachusetts, Amherst
9:30 - 10:00	Break
10:00 - 11:30	Paper Session: Multimodal and Sensorimotor Representations Discovering Sensor Space: Constructing spatial embeddings that explain sensor correlations <i>Joseph Modayil</i> The Odd One Out Task: Toward an Intelligence Test for Robots <i>Jivko Sinapov, Alexander Stoytchev</i> Towards the Object Semantic Hierarchy <i>Changhai Xu, Benjamin Kuipers</i> Using the Head to Stabilize Action: Reaching by Young Children <i>Hongwei Shen, Thomas Baker, T. Rowan Candy, Chen Yu, Linda Smith</i>
11:30 - 1:30	Lunch break + set up posters
11:30 - 1:30	ICDL Governing Board meeting and IEEE TAMM editorial meeting
1:30 - 3:00	Paper Session: Language Learning Distribution of Object types of "light" and "heavy" early-learned English verbs <i>Josita Maouene, Aarre Laakso, Mounir Maouene, Linda Smith</i> A Computational Model for Grounding Words in the Perception of Agents <i>Claudius Gläser, Frank Joublin</i> Bootstrapping Syntax from Morpho-Phonology <i>Thomas Shultz, Vincent Berthiaume, Frédéric Dandurand</i> Using Information Gain to Build Meaningful Decision Forests for Multilabel Classification <i>Kevin Gold, Allison Petrosin</i>

3:00 - 3:30	Break
3:30 - 5:00	Poster spotlight (6-page papers) [28 x 3 min each]
	Bootstrapping inverse Kinematics with Goal Babbling <i>Matthias Rolf, Jochen Steil, Michael Gienger</i>
	Effect of Neuromodulation on Performance in Game Playing: A Modeling Study <i>Derrick Asher, Andrew Zaldivar, Jeffrey Krichmar</i>
	Becoming a Gaze User: Joint Attention and Cooperation in One- and Two-year-olds <i>Zhen Wu, Jingtong Pan, Yanjie Su, Julie Gros-Louis</i>
	Developmental Learning of Integrating Visual Attention Shifts and Bimanual Behavior in Object Grasping and Manipulation Tasks <i>Sungmoon Jeong, Minho Lee, Hiroaki Arie, Jun Taniguchi</i>
	An Integrated Three-Stage Model Towards Grammar Acquisition <i>Yo Sato, Joe Saunders, Frank Broz, Caroline Lyon, Chrystopher Nehaniv</i>
	Investigating word learning processes in an artificial agent <i>Michele Gubian, Christina Bergmann, Lou Boves</i>
	Prenatal to Postnatal Transfer of Motor Skills Through Motor-Compatible Sensory Representations <i>Timothy Mann, Yoonsuck Choe</i>
	Real Time Targeted Exploration in Large Domains <i>Todd Hester, Peter Stone</i>
	Maturationally-Constrained Competence-Based Intrinsically Motivated Learning <i>Adrien Baranes, Pierre-Yves Oudeyer</i>
	A developmental approach to the emergence of communication in socially situated embodied agents <i>Michael Sheldon, Mark Lee</i>
	A Reinforcement Learning Model of Reaching Integrating Kinematic and Dynamic Control in a Simulated Arm Robot <i>Daniele Caligiore, Eugenio Guglielmelli, Anna Borghi, Domenico Parisi, Gianluca Baldassarre</i>
	Batch versus Interactive Learning by Demonstration <i>Peng Zang, Runhe Tian, Andrea Thomaz, Charles Isbell</i>
	Toddlers' Transitions on Non-verbal False-belief Tasks Involving a Novel Location: A Constructivist Connectionist Model <i>Vanessa Evans, Vincent Berthiaume, Thomas Shultz</i>
	How can bottom-up information shape learning of top-down attention control skills? <i>Dimitri Ognibene, Giovanni Pezzulo, Gianluca Baldassarre</i>
	On the Influence of Sensor Morphology on Eye Motion Coordination <i>Harold Martinez, Max Lungarella, Rolf Pfeifer</i>
	Scenes and Tracking with Dynamic Neural Fields: How to Update a Robotic Scene Representation <i>Stephan Zibner, Christian Faubel, Ioannis Iossifidis, Gregor Schöner, John Spencer</i>
	Serial order in an acting system: a multidimensional dynamic neural fields implementation <i>Yulia Sandamirskaya, Gregor Schöner</i>
	Competence Progress Intrinsic Motivation <i>Andrew Stout, Andrew Barto</i>

Imitation Learning with Hierarchical Actions <i>Abram Friesen, Rajesh Rao</i>
Autism as an impairment in detecting invariants <i>Norbert Michael Mayer, Ian Fasel</i>
Robot Motivator: Increasing User Enjoyment and Performance on a Physical/Cognitive Task <i>Juan Fasola, Maja Matarić</i>
WWN-Text: Cortex-Like Language Acquisition with "What" and "Where" <i>Kajal Miyan, Juyang Weng</i>
The Similarity-Attraction Effect in Human-Robot Interaction <i>Emily Bernier, Brian Scassellati</i>
Mutually constrained multimodal mapping for simultaneous development: modeling vocal imitation and lexicon acquisition <i>Yuki Sasamoto, Yuichiro Yoshikawa, Minoru Asada</i>
Towards an Extended Dynamic Field Theoretic Model of Iowa Gambling Task Performance <i>Robert Lowe, Boris Duran, Tom Ziemke</i>
Discovering the concept of anaphora from grounded verb models <i>Kruti Neema, Amitabha Mukerjee</i>
Where-What Network-4: The Effect of Multiple Internal Areas <i>Matthew Luciw, Juyang Weng</i>
External Rotation as Morphological Bootstrapping for Emergence of Biped Walking <i>Koh Hosoda, Yoichiro Ishii</i>

5:00 - 7:00	Poster session (6-page papers)
-------------	---------------------------------------

7:00 -	Dinner on the town
--------	---------------------------

Friday, August 20

8:30 - 9:30	Keynote Presentation Felix Warneken, Harvard University
-------------	--

9:30 - 10:00	Break
--------------	--------------

10:00 - 11:30	Paper Session: Social Learning
	Joint Attention through the Hands: Investigating the Timing of Object Labeling in Dyadic Social Interaction <i>Martin Rickert, Chen Yu, Amanda Favata</i>
	Optimality of Human Teachers for Robot Learners <i>Maya Cakmak, Andrea Thomaz</i>
	Developing feedback: how children of different age contribute to a tutoring interaction with adults <i>Anna-Lisa Vollmer, Karola Pitsch, Katrin Lohan, Jannik Fritsch, Katharina Rohlfing, Britta Wrede</i>
	Evolution of Social Learning Strategies <i>Marcel Montrey, Thomas Shultz</i>

11:30 - 12:00	IM_CLeVeR Project Update Gianluca Baldassarre
12:00 - 1:30	Lunch break + set up posters
1:30 - 3:00	Paper Session: Learning from Tactile Feedback
	Learning to recognize shapes with a sensory substitution system: a longitudinal study with 4 non-sighted adolescents <i>Katia Rovira, Olivier Gapenne, Amal Ali Ammar</i>
	Tactile Guidance for Policy Refinement and Reuse <i>Brenna Argall, Eric Sauser, Aude Billard</i>
	A human fetus development simulation: Self-organization of behaviors through tactile sensation. <i>Hiroki Mori, Yasuo Kuniyoshi</i>
	Learning to Press Doorbell Buttons <i>Vladimir Sukhoy, Jivko Sinapov, Liping Wu, Alexander Stoytchev</i>
3:00 - 3:30	Break
3:30 - 4:40	Paper Session: Vision and Visual Development
	Learning to Look <i>Nicholas Butko, Javier Movellan</i>
	A Model of the Emergence of Early Imitation Development based on Predictability Preference <i>Takashi Minato, Dale Thomas, Yuichiro Yoshikawa, Hiroshi Ishiguro</i>
	Intrinsically Motivated Information Foraging <i>Ian Fasel, Andrew Wilt, Nassim Mafi, Clayton Morrison</i>
4:40 - 5:00	CogSci 2010 Invited Guest Paper
5:00 - 5:10	Break
5:10 - 5:45	Poster spotlight (2-page abstracts) [16 x 2 min each]
	An experiment with human-robot interaction to study of intentional agency in joint visual attention <i>Takeshi Konno, Takashi Hashimoto</i>
	Does embodiment affect tutoring behavior? <i>Katrin Solveig Lohan, Sebastian Gieselmann, Anna-Lisa Vollmer, Katharina Rohlfing, Britta Wrede</i>
	Student computers: Six-year-olds believe that computers can learn. <i>Maria Vázquez, Megan Saylor, Daniel Levin</i>
	Variability Facilitates Learning of Relations Between Words and Actions <i>Heidi Baumgartner, Courtney Millhoff, Katharine Graf Estes, Lisa Oakes</i>
	Development of Perception and Representation of Rhythmic Information: Towards a Computational Model <i>Inês Salselas, Perfecto Herrera</i>
	Modeling Real-time Multimodal Interaction with Virtual Humans <i>Hui Zhang, Damian Fricker, Chen Yu</i>

Development of Semantic Network: Evidence from a Forced Choice Association <i>Shohei Hidaka</i>
Infants' learning about human and non-human faces: An eye-tracking study <i>Emily Spring, Karinna Hurley, Lisa Oakes</i>
How to Build a Schema-Representation <i>Nikolas Hemion, Frank Joublin, Katharina Rohlfing</i>
A computational model of learning intuitive object dynamics <i>Lewis Fishgold, Dana Ballard, Benjamin Kuipers</i>
Learning-to-Grasp: from an Infant to a Troublemaker <i>Baro Hyun, Anouck Girard, Benjamin Kuipers</i>
Learning representations and detecting changes in complex environments <i>Katherine Snyder, Dmitry Kit, Brian Sullivan, Mary Hayhoe, Dana Ballard</i>
A Boosting Approach to Understanding Referential Intent <i>Daniel Yurovsky, Martin Rickert, Chen Yu, Linda Smith</i>
Preliminary Findings on Characteristics of Handwriting in Children With Autism and Asperger's Disorder <i>Beth Johnson, Nicole Rinehart, James Phillips, Joanne Fielding, Kerryn Saunders</i>
A Comparative Study of Multimodal Human-Robot and Human-Human Interactions <i>Damian Fricker, Chen Yu, Matthias Scheutz, Paul Schermerhorn</i>
Discovering complex categories through active sampling <i>Doug Markant, Todd Gureckis</i>

5:45 - 7:00 **Poster session (2-page abstracts)**

7:00 - **Banquet: Campus Inn**

Saturday, August 21

8:30 - 9:30 **Keynote Presentation**
Susan Gelman, University of Michigan, Ann Arbor

9:30 - 10:00 **Break**

10:00 - 11:30	Paper Session: Models and Algorithms
	kx-trees: An unsupervised learning method for use in developmental agents <i>Brandon Rohrer</i>
	Dealing with uncertain input in word learning <i>Maarten Versteegh, Louis ten Bosch, Lou Boves</i>
	Recognizing behaviors and the Internal State of the Participants <i>Wesley Kerr, Paul Cohen</i>
	A Spiking Neural Model for the Spatial Coding of Cognitive Response Sequences <i>Suresh Vasa, Tao Ma, Kiran Byadarhaly, Mithun Perdoor, Ali Minaei</i>

11:30 - 1:30 **Autonomous Mental Development Technical Committee meeting. Buffet lunch provided. All ICDL participants invited. Conference ends.**

Keynote Speakers: Biographies & Abstracts

Cognitive Development and the iCub Humanoid Robot



David Vernon
Italian Institute of Technology

Time: 7:00pm-8:00pm, Wednesday, August 18
Location: The Dahlmann Campus Inn

Abstract:

This talk addresses the central role played by development in cognition. The focus is on applying our knowledge of development in natural cognitive systems, i.e. human infants, to the problem of creating artificial cognitive systems in the guise of humanoid robots.

The work described in this talk is founded on the premise that (a) cognition is the process by which an autonomous self-governing agent acts effectively in the world in which it is embedded, that (b) the dual purpose of cognition is to increase the agent's repertoire of effective actions and its power to anticipate the need for and outcome of future actions, and that (c) development plays an essential role in the realization of these cognitive capabilities.

Our goal is to identify the key design principles for cognitive development. We do this by bringing together insights from four areas: enactive cognitive science, developmental psychology, neurophysiology, and computational modelling. We then discuss progress in applying these principles to the implementation of a cognitive architecture for the iCub, a open-systems humanoid robot which has been designed specifically as a common platform for research on embodied cognitive systems.

Bio Sketch:

David Vernon works as a freelance research scientist in the broad area of computer vision, robotics, and cognition. Recently, he has served as the coordinator of euCognition: the European Network for the Advancement of Artificial Cognitive Systems and as a member of the project team working on the creation of the iCub, an open-source cognitive humanoid robot. Over the past 32 years, he has held positions at Westinghouse Electric, Trinity College Dublin, the European Commission, the National University of Ireland, Maynooth, Science Foundation Ireland, University of Genoa, among others. He has authored five books on computer vision and has published over ninety papers in the fields of computer vision, robotics, and cognitive systems. His current research focus is on enactive approaches to cognition. He is a Senior Member of the IEEE, a Chartered Engineer of the Institution of Engineers of Ireland, and a past Fellow of Trinity College Dublin. He is an editor of the Springer series of Cognitive Systems Monographs (COSMOS).

The Developmental Discovery and Organization of Environmental Affordances



Rod Grupen
University of Massachusetts, Amherst

Time: 8:30am-9:30am, Thursday, August 19
Location: The Horace H. Rackham Building

Abstract:

Robots are clumsy and inflexible mostly because they are programmed by humans with little insight into what the world is like for a robot. Drawing on theories of how animals develop, I will argue for a different kind of machine that invents new states and actions in order to acquire skills for the real world, forms categories about controllable contexts, and learns about us and from us. Using UMass robots, I'll describe recent work that uses intrinsically motivated learning to discover controllable interactions with open, unstructured worlds.

Bio Sketch:

Professor Grupen has degrees in Physics, Mechanical Engineering, and received a Ph.D. in Computer Science from the University of Utah in 1988. He conducts research that integrates signal processing, control, dynamical systems, learning, and development as a means of designing controllers for intelligent systems. Grupen and his students develop hierarchical representations for sensorimotor policies that facilitate learning and transfer. This approach guides experimental research using dexterous mobility and manipulation platforms and has yielded techniques for motion control and collision avoidance, methods for coordinating multiple distributed robots, policies for dexterous grasping and walking gaits, and new concepts in the control of whole-body mobile manipulators. These tools have been used as a computational account of sensorimotor development in animals, including early childhood development, and new developmental programming techniques for intelligent machines that interact with open environments.

Professor Grupen is a Co-Editor-in-Chief of Robotics and Autonomous Systems, Editor of AIEDAM (AI in Engineering Design and Manufacture), and serves on several international program committees. He has been the Principal Investigator on grants from the NSF, DARPA, NASA, ONR, ARO, AFOSR, Microsoft, and iRobot. He co-founded the Embedded Systems instructional laboratory in which undergraduate and graduate students from many disciplines learn about building integrated sensorimotor systems. He received an Outstanding Teacher Award and recently received the Chancellor's medal and was designated a Distinguished Faculty Lecturer at the University of Massachusetts Amherst.

Core Capacities for Cooperation: Examples from Human Children and Chimpanzees



**Felix Warneken
Harvard University**

Time: 8:30am-9:30am, Friday, August 20
Location: The Horace H. Rackham Building

Abstract:

Cooperation is the hallmark of human social life, spanning from simple acts of helping another person to large-scale collaborative practices in which multiple people pool their efforts to achieve goals that lie beyond the means of any one individual. These behaviors raise interesting questions for cognitive science because they require specific social-cognitive capacities such as the ability to represent goal-directed action for helping and the formation of shared plans of actions for collaborative activities. In my talk, I will review studies from developmental and comparative psychology which aim at identifying the core capacities that (a) enable the most basic forms of cooperation in humans and our evolutionary relatives and (b) constitute the building blocks for later-emerging and more complex forms of cooperation in human development. These insights from developmental psychology might help to specify and test the cognitive architecture required for artificial agents to engage in diverse cooperative interactions with humans.

Bio Sketch:

Dr. Felix Warneken is Assistant Professor and Director of the Social Cognitive Development Group in the Psychology Department at Harvard University. Trained as a developmental and comparative psychologist, he conducts research with a focus on cooperation and social-cognitive development in children and great apes. He studied in Germany and the United States, receiving his doctoral degree from the Universität Leipzig while working at the Max Planck Institute for Evolutionary Anthropology (MPI EVA). He continued as a postdoctoral researcher at MPI EVA and joined the Harvard Psychology Department in 2009. He received several awards and fellowships, including an Outstanding Dissertation Award by the Society for Research in Child Development in 2009 and a Novartis Fellowship in 2006. A study demonstrating altruistic helping in children and chimpanzees was named one of the 100 most important science stories in 2007 by Discover Magazine.

**Sharks attack humans, but most sharks don't attack humans:
Learning to express generalizations in language**



**Susan Gelman
University of Michigan**

Time: 8:30am-9:30am, Saturday, August 21
Location: The Horace H. Rackham Building

Abstract:

My talk will examine the challenges that learners face when learning to express general categories (e.g., "sharks" in "Sharks attack humans"). These expressions, known as "generics", are both conceptually and linguistically challenging. Although we experience the world in terms of individual objects and events, we must form abstractions that extend beyond these individual entities. Moreover, the same forms of language that are used to express generics are also used to refer to particular individuals (compare "The dog is a 4-legged animal" to "The dog is sleeping"). I discuss different kinds of learning models, and suggest that generics are a default mode of generalization for human learners.

Bio Sketch:

Susan A. Gelman earned her Ph.D. in Psychology from Stanford University in 1984. She then joined the faculty of the University of Michigan, where she is currently the Frederick G. L. Huetwell Professor of Psychology. Her research interests focus on concept and language development in young children. She is interested in how children organize their experiences into categories, children's early theories of the world around them (particularly their emerging biological theories), children's "essentialist" beliefs, children's understanding of causality, and the role of language in expressing and conveying children's concepts. Dr. Gelman has published 6 books and monographs and over 100 articles and chapters, including *The Essential Child* (Oxford, 2003), which won the Eleanor E. Maccoby book prize of APA Division 7 and the 2005 book prize of the Cognitive Development Society. Professor Gelman received the APA Distinguished Scientific Award for Early Career Contribution to Psychology in the Developmental Area, the American Psychological Foundation Robert L. Fantz Award, the Chase Memorial Award, the Boyd McCandless Young Scientist Award from APA Division 7, and a J. S. Guggenheim Fellowship. She has served on the editorial boards of several journals, and is a fellow of the American Academy of Arts and Sciences.

Oral Presentations: Full Length Papers

Proceedings: Pages 120-125

Discovering Sensor Space: Constructing Spatial Embeddings That Explain Sensor Correlations

Joseph Modayil

Abstract— A fundamental task for a developing agent is to build models that explain its uninterpreted sensory-motor experience. This paper describes an algorithm that constructs a sensor space from sensor correlations, namely the algorithm generates a spatial embedding of sensors where strongly correlated sensors will be neighbors in the embedding. The algorithm first infers a sensor correlation distance and then applies the fast maximum variance unfolding algorithm to generate a distance preserving embedding. Although previous work has shown how sensor embeddings can be constructed, this paper provides a framework for understanding sensor embedding, introduces a sensor correlation distance, and demonstrates embeddings for thousands of sensors on intrinsically curved manifolds.

Proceedings: Pages 126-131

The Odd One Out Task: Toward an Intelligence Test for Robots

Jivko Sinapov and Alexander Stoytchev

Abstract— Detecting the outlier in a set of objects is a fundamental task used in a wide variety of intelligence tests. This paper proposes a theoretical model that allows a robot to interactively estimate the pairwise similarity between everyday objects and use this knowledge to solve the odd one out task. That is, given a set of objects, the robot's task is to select the one object that does not belong in the group. In our experiments, the robot interacted with fifty different household objects (by applying five different exploratory behaviors on them) and perceived auditory and proprioceptive sensory feedback. Pairwise object similarity was estimated for different behavior and modality contexts. In a series of subsequent tests, three objects from a given category (e.g., three cups or three pop cans) along with one object from outside that category were selected and the robot's internal models were queried to pick the object that does not belong in the group. The object similarity relations learned by the robot were used to pick the most dissimilar object, with success rates varying from 45% to 100%, depending on the category. The results show that the learned similarity measures were sufficient to capture some of the common properties of human-defined object categories, such as cups, bottles, and pop cans.

Proceedings: Pages 39-45

Towards the Object Semantic Hierarchy

Changhai Xu and Benjamin Kuipers

Abstract— An intelligent agent, embedded in the physical world, will receive a high-dimensional ongoing stream of low-level sensory input. In order to understand and manipulate the world, the agent must be capable of learning high-level concepts. Object is one such concept. We are developing the Object Semantic Hierarchy (OSH), which consists of multiple representations with different ontologies. The OSH factors the problems of object perception so that intermediate states of knowledge about an object have natural representations, with relatively easy transitions from less structured to more structured representations. Each layer in the hierarchy builds an explanation of the sensory input stream, in terms of a stochastic model consisting of a deterministic model and an unexplained "noise" term. Each layer is constructed by identifying new invariants from the previous layer. In the final model, the scene is explained in terms of constant background and object models, and low-dimensional pose trajectories of the observer and the foreground objects.

The object representations in the OSH range from 2D views, to 2D planar components with 3D poses, to structured 3D models of objects. This paper describes the framework of the Object Semantic Hierarchy, and presents the current implementation and experimental results.

Proceedings: Pages 108-113

Using the Head to Stabilize Action: Reaching by Young Children

Hongwei Shen, Thomas J. Baker,
T. Rowan Candy, Chen Yu, and Linda B. Smith

Abstract— Even seemingly simple reaching task requires complex integration or remapping of reference frames with respect to eye, head and hand. Head-centered reference frame and head movement may play important roles in natural tasks and in the development of reaching. To understand the underlying control mechanism that support smooth reaching in particular and the seamless coordination of head and hand movements more generally, a semi-naturalistic experiment is designed for 1 1/2 to 5 years old children, in which children are free to reach or not reach for the ball coming out from the slots on the puppet show. The fine-grained head and hand motion data were recorded. The analysis results show that (1) All the children participants stabilize their heads right before the reach; (2) During the reach, the head and hands move synergistically in the similar directions. (3) Large and variable head movements co-occur with more variable and jerky reaches. The coupling of the head and hands may suggest a dynamical integration of multiple reference frames.

Proceedings: Pages 88-94

Distribution of Object Types of “Light” and “Heavy” Early-Learned English Verbs

Josita Maouene, Aarre Laakso, Mounir Maouene, and Linda B. Smith

Abstract— *In the developmental psycholinguistic literature, it is common to distinguish verbs that are semantically light from those that are not. One important reason is that the light verbs (take, get, make, do, go, etc.) – excellent substitutes for specific verbs and very frequent in adult speech to children – are thought to help children learn the verb system. Although quantitative and qualitative criteria (e.g., frequency, grammaticalization, semantic generality, high transitivity) have been proposed for distinguishing light and heavy verbs, some puzzling questions remain: how good are criteria that define heavy verbs as non-light ones? Are verbs bimodally distributed? Do children’s light and heavy verbs align with adult ones? This paper proposes a new candidate -- using the number of objects (free associations and co-occurrences) a verb has as an indicator of its semantic generality -- and applies it to 80 early-learned English verbs. The results suggest that early-learned light and heavy verbs differ in the breadth of the objects they are associated with: light verbs have weak associations with specific objects, whereas heavy verbs are strongly associated with specific objects. There is also a hint that some verbs have narrower associations with objects in speech from and to children.*

Proceedings: Pages 52-57

Bootstrapping Syntax from Morpho-Phonology

Thomas R. Shultz, Vincent G. Berthiaume, and Frédéric Dandurand

Abstract— *It has been a puzzle how the syntax of natural language could be learned from positive evidence alone. Here we present a hybrid neural-network model in which artificial syntactic categories are acquired through unsupervised competitive learning due to grouping together lexical words with consistent phonological endings. These relatively large syntactic categories then become target signals for a feed-forward error-reducing network that learns to pair these lexical items with smaller numbers of function words to form phrases. This hybrid model learns phrasal syntax from positive evidence alone, while covering the essential findings in recent experiments on adult humans learning an artificial language. The model further predicts generalization to novel lexical words (exceptions) from knowledge of function words.*

Proceedings: Pages 26-32

A Computational Model for Grounding Words in the Perception of Agents

Claudius Gläser and Frank Joublin

Abstract— *In this paper we present a computational model for incremental word meaning acquisition. It is designed to rapidly build category representations which correspond to the meaning of words. In contrast to existing approaches, our model further extracts word meaning-relevant features using a statistical learning technique. Both category learning and feature extraction are performed simultaneously. To achieve the contradictory needs of rapid as well as statistical learning, we employ mechanisms inspired by Complementary Learning Systems theory. Therefore, our framework is composed of two recurrently coupled components: (1) An adaptive Normalized Gaussian network performs a one-shot memorization of new word-scene associations and uses the acquired knowledge to categorize novel situations. The network further reactivates memorized associations based on its internal representation. (2) Based on the reactivated patterns an additional component subsequently extracts features which facilitate the categorization task. An iterative application of the learning mechanism results in a gradual memory consolidation which let the internal representation of a word meaning become more efficient and robust. We present simulation results for a scenario in which words for object relations concerning position, size, and color have been trained. The results demonstrate that the model learns from few training exemplars and correctly extracts word meaning-relevant features.*

Proceedings: Pages 58-63

Using Information Gain to Build Meaningful Decision Forests for Multilabel Classification

Kevin Gold and Allison Petrosino

Abstract— *“Gain-Based Separation” is a novel heuristic that modifies the standard multiclass decision tree learning algorithm to produce forests that can describe an example or object with multiple classifications. When the information gain at a node would be higher if all examples of a particular classification were removed, those examples are reserved for another tree. In this way, the algorithm performs some automated separation of classes into categories; classes are mutually exclusive within trees but not across trees. The algorithm was tested on naive subjects’ descriptions of objects to a robot, using YUV color space and basic size and distance features. The new method outperforms the common strategy of separating multilabel problems into L binary outcome decision trees, and also outperforms RAKEL, a recent method for producing random multilabel forests.*

Proceedings: Pages 114-119

Joint Attention through the Hands: Investigating the Timing of Object Labeling in Dyadic Social Interaction

Martin Rickert, Chen Yu, and Amanda Favata

Abstract— *Previous studies of joint attention and its' role in language learning have focused on eye-gaze cues. The goal of the present study is to discover fine-grained patterns of joint hand activities in child-parent social interaction that facilitate successful word learning. To this end, we address the following three topics: 1) quantifying joint manual actions between the parent and the child and in particular how the child follows the parent's bid of attention through manual actions; 2) discovering the timings between joint manual actions and object naming events; and 3) linking those timings with language learning results. Multiple high-resolution data streams were examined for episodes involving object-labeling events that either preceded or followed joint attentional focus as established through the hand actions of the dyad. Our findings suggest that the success of word learning through social interaction depends on the specific timing between follow-in joint hand activities and naming events.*

Proceedings: Pages 64-69

Optimality of Human Teachers for Robot Learners

Maya Cakmak and Andrea L. Thomaz

Abstract— *In this paper we address the question of how closely everyday human teachers match a theoretically optimal teacher. We present two experiments in which subjects teach a concept to our robot in a supervised fashion. In the first experiment we give subjects no instructions on teaching and observe how they teach naturally as compared to an optimal strategy. We find that people are suboptimal in several dimensions. In the second experiment we try to elicit the optimal teaching strategy. People can teach much faster using the optimal teaching strategy, however certain parts of the strategy are more intuitive than others.*

Proceedings: Pages 76-81

Developing feedback: how children of different age contribute to a tutoring interaction with adults

*Anna-Lisa Vollmer, Karola Pitsch,
Katrin S. Lohan, Jannik Fritsch,
Katharina J. Rohlfing, and Britta Wrede*

Abstract— *Learning is a social and interactional endeavor, in which the learner generally receives support from his/her social environment. In this process, the learner's feedback is important as it provides information about the learner's current understanding which, in turn, enables the tutor to adjust his/her presentation accordingly. Thus, through their feedback learners can actively shape the tutor's presentation - a resource which is highly valuable, if we aim at enabling robot systems to learn from a tutor in social interaction. But what kind of feedback should a robot produce and at which time? In this paper, we analyze the interaction between parents and their infants (8 to 30 months) in a tutoring scenario with regard to the feedback provided by the learner in three different age groups. Our combined qualitative and quantitative analysis reveals which features of the feedback change with the infant's progressing age and cognitive capabilities.*

Proceedings: Pages 95-100

Evolution of Social Learning Strategies

Marcel R. Montrey and Thomas R. Shultz

Abstract— *We study three types of learning with Bayesian agent-based modeling. First, we show that previous results obtained from learning chains can be generalized to a more realistic lattice world involving multiple social interactions. Learning based on the passing of posterior probabilities converges to the truth more quickly and reliably than does learning based on imitation and sampling from the environment; and the latter method gets closer to the truth than does pure imitation. The passing of posterior probability distributions can be viewed as teaching by explanation, and as an implementation of the cultural ratchet, which allows rapid progress without backsliding. We also show that evolution selects these learning strategies in proportion to their success. However, if the environment changes very rapidly, evolution favors the imitation-plus-reinforcement strategy over the more sophisticated posterior passing. Implications for developmental robotics, human uniqueness, and interactions between learning and evolution are discussed.*

Proceedings: Pages 1-6

Learning to recognize shapes with a sensory substitution system: a longitudinal study with 4 non-sighted adolescents

*Katia Rovira, Olivier Gapenne,
and Amal Ali Ammar*

Abstract— Inspired by the work on sensory substitution which shows the technical and practical possibilities for perceptual learning, we have carried out a global longitudinal study of the learning capacities for reading digital graphical objects by four blind high-school students. We report here the results obtained with the Tactos system during certain phases of this study, in particular the acquisition of reading strategies and the possibility of deploying these strategies in the case of a geometry exercise.

Proceedings: Pages 7-12

Tactile Guidance for Policy Refinement and Reuse

*Brenna D. Argall, Eric L. Sauser,
and Aude G. Billard*

Abstract— Demonstration learning is a powerful and practical technique to develop robot behaviors. Even so, development remains a challenge and possible demonstration limitations can degrade policy performance. This work presents an approach for policy improvement and adaptation through a tactile interface located on the body of a robot. We introduce the Tactile Policy Correction (TPC) algorithm, that employs tactile feedback for the refinement of a demonstrated policy, as well as its reuse for the development of other policies. We validate TPC on a humanoid robot performing grasp-positioning tasks. The performance of the demonstrated policy is found to improve with tactile corrections. Tactile guidance also is shown to enable the development of policies able to successfully execute novel, undemonstrated, tasks.

Proceedings: Pages 82-87

A human fetus development simulation: Self-organization of behaviors through tactile sensation

Hiroki Mori and Yasuo Kuniyoshi

Abstract— Recent progresses of ultrasound imaging technology have led observations of fetal intrauterine behavior and a perspective of intrauterine learning. Understanding fetal behavior in uterus is important for medical cares for prenatal infants, because the intervention like “nesting” or “swaddling” in NICU (Neonatal Intensive Care Unit) is based on a perspective of intrauterine learning. However, fetal behavior is not explained sufficiently by the perspective. In this study, we have proposed a hypothesis in which two fetal behaviors, Isolated leg/arm movements and hand and face contact, emerge within self-organization of interaction among an uterine environment, a fetal body, and a nervous system. through tactile sensation in uterus. We have conducted computer experiments with a simple musculoskeletal model in uterus and a whole body fetal musculoskeletal model with tactile for the hypothesis. We confirmed that tactile sensation induces motions in the experiments of the simple model, and the fetal model with human like tactile distribution have behaved with the two motions similar to real fetal behaviors. Our experiments indicated that fetal intrauterine learning is possibly core concept for the fetal motor development.

Proceedings: Pages 132-139

Learning to Press Doorbell Buttons

*Vladimir Sukhoy, Jivko Sinapov, Liping Wu,
and Alexander Stoytchev*

Abstract— This paper describes an approach that a robot can use to learn to press doorbell buttons. This approach combines exploratory behaviors with an active learning strategy to enable the robot to learn faster how and where it should press a button in order to trigger the buzzer. The framework was tested with an upper-torso humanoid robot on seven different doorbell buttons. Three different active learning exploration strategies were evaluated: random, stimulus-driven, and uncertainty-driven. The results show that an active learning strategy can significantly speedup the robot’s learning progress. Among the three strategies that were evaluated, the uncertainty-driven strategy was the most effective.

Proceedings: Pages 70-75

Learning To Look

Nicholas J. Butko and Javier R. Movellan

Abstract— How can autonomous agents with access to only their own sensory-motor experiences learn to look at visual targets? We explore this seemingly simple question, and find that naïve approaches are surprisingly brittle. Digging deeper, we show that learning to look at visual targets contains a deep, rich problem structure, relating sensory experience, motor experience, and development. By capturing this problem structure in a generative model, we show how a Bayesian observer should trade off different sources of uncertainty in order to discover how their sensors and actuators relate. We implement our approach on two very different robots, and show that both of them can quickly learn reliable intentional looking behavior without access to anything beyond their own experiences.

Proceedings: Pages 19-25

A Model of the Emergence of Early Imitation Development based on Predictability Preference

Takashi Minato, Dale Thomas, Yuichiro Yoshikawa, and Hiroshi Ishiguro

Abstract— This article presents a mechanism for the early development of imitation through a simulation of infant-caregiver interaction. A model was created to acquire a body mapping (a mapping from observed body motions to motor commands), which is necessary for imitation, while discriminating self-motion from the motion of the other. The simulation results show that the development of a body mapping depends on a predictability preference (a function of how an agent decides regarding its options of ‘what to imitate’). The simulated infants are able to develop the components of a healthy body mapping in order, that is, relating self motion first, followed by an understanding of others’ motions, which is supported by psychological studies. This order of development emerges spontaneously without the need for any explicit mechanism or any partitioning of the interaction. These results suggest that this predictability preference is an important factor in infant imitation development.

Proceedings: Pages 101-107

Intrinsically Motivated Information Foraging

Ian Fasel, Andrew Wilt, Nassim Mafi, and Clayton T. Morrison

Abstract— We treat information gathering as a POMDP in which the goal is to maximize an accumulated intrinsic reward at each time step based on the negative entropy of the agent’s beliefs about the world state. We show that such information foraging agents can discover intelligent exploration policies that take into account the long-term effects of sensor and motor actions, and can automatically adapt to variations in sensor noise, different amounts of prior information, and limited memory conditions.

Proceedings: Pages 13-18

kx-trees: An unsupervised learning method for use in developmental agents

Brandon Rohrer

Abstract— Acquiring concepts from experience is a key aspect of development and one that is commonly neglected in learning agents. In this work, concept acquisition is formulated as an unsupervised learning problem and is addressed with a novel algorithm: kx-trees. kx-trees differ from prior approaches to unsupervised learning in that they require very little information; four user selected parameters determine all aspects of kx-trees' performance. Notably, and in contrast with most other unsupervised learning approaches, they do not require that the input state space be well-scaled. kx-trees' operation is described in detail and illustrated with two simulations. The second simulation shows some similarities between kx-trees and feature construction in the human visual processing system.

Proceedings: Pages 46-51

Dealing with uncertain input in word learning

Maarten Versteegh, Louis ten Bosch,
and Lou Boves

Abstract— In this paper we investigate a computational model of word learning, that is embedded in a cognitively and ecologically plausible framework. Multi-modal stimuli from four different speakers form a varied source of experience. The model incorporates active learning, attention to a communicative setting and clarity of the visual scene. The model's ability to learn associations between speech utterances and visual concepts is evaluated during training to investigate the influence of active learning under conditions of uncertain input. The results show the importance of shared attention in word learning and the model's robustness against noise.

Proceedings: Pages 33-38

Recognizing Behaviors and the Internal State of the Participants

Wesley Kerr and Paul Cohen

Abstract— Psychological research has demonstrated that subjects shown animations consisting of nothing more than simple geometric shapes perceive the shapes as being alive, having goals and intentions, and even engaging in social activities such as chasing and evading one another. While the subjects could not directly perceive affective state, motor commands, or the beliefs and intentions of the actors in the animations, they still used intentional language to describe the moving shapes. We present representations and algorithms that enable an artificial agent to correctly recognize other agents' activities by observing their behavior. In addition, we demonstrate that if the artificial agent learns about the activities through participation, where it has access to its own internal affective state, motor commands, etc., it can then infer the unobservable internal state of other agents.

Proceedings: Pages 140-146

A Spiking Neural Model for the Spatial Coding of Cognitive Response Sequences

Suresh Vasa, Tao Ma, Kiran V. Byadarhaly,
Mithun Perdoor, and Ali A. Minai

Abstract— The generation of sequential responses is a fundamental aspect of cognitive function, encompassing processes such as motor control, linguistic expression, memory recall and thought itself. There is considerable evidence that complex cognitive responses (such as voluntary actions) are constructed as chunked sequences of more elementary response primitives or synergies, which can themselves be seen often as sequences of simpler primitives. Almost all neural models of sequence representation are based on the principle of recurrence, where each successive item is generated by preceding items. However, it is also interesting to consider the possibility of purely spatial neural representations that result in sequential readout of pre-existing response elements. Such representations offer several potential benefits, including parsimony, efficiency, flexibility and generalization. In particular, they can allow response sequences to be stored in memory as chunks encoded by fixed point attractors. In this paper, we present a simple spiking neuron model for the flexible encoding and replay of response sequences through the impulsive triggering of coding patterns represented as fixed point attractors. While not intended as a detailed description of a specific brain region, the model seeks to capture fundamental control mechanisms that may apply in many parts of the nervous system.

Poster Presentations: Full Length Papers

Poster Spotlights: 3:30PM-5:00PM, Thursday, 19 August 2010
Poster Session: 5:00PM-7:00PM, Thursday, 19 August 2010

Proceedings: Pages 147-154

Bootstrapping inverse Kinematics with Goal Babbling

Matthias Rolf, Jochen J. Steil,
and Michael Gienger

Abstract— We present an approach to learn inverse kinematics of redundant systems without prior- or expert-knowledge. The method allows for an iterative bootstrapping and refinement of the inverse kinematics estimate. We show that the information structure induced by goal-directed exploration enables an efficient resolution of inconsistent samples solely from observable data. The bootstrapped solutions are aligned for a maximum of movement efficiency, i.e. realizing an effector movement with a minimum of joint motion. We derive and illustrate the exploration and learning process with a low-dimensional kinematic example and show that the same procedure scales for high dimensional problems, such as hyperredundant planar arms with up to 50 degrees of freedom.

Proceedings: Pages 155-160

Effect of Neuromodulation on Performance in Game Playing: A Modeling Study

Derrik E. Asher, Andrew Zaldivar,
and Jeffrey L. Krichmar

Abstract— Neuromodulators can have a strong effect on how organisms learn and compete for resources. Neuromodulators, such as dopamine (DA) and serotonin (5-HT), are known to be important in predicting rewards, costs, and punishments. To better understand the effect of neuromodulation on decision-making, a computational model of the dopaminergic and serotonergic systems was constructed and tested in games of conflict. This neural model was based on the assumptions that dopaminergic activity increases as expected reward increases, and serotonergic activity increases as the expected cost of an action increases. Specifically, the neural model guided the learning of an agent that played a series of Hawk-Dove games against an opponent. The model responded appropriately to changes in environmental conditions or to changes in its opponent's strategy. The neural agent became Dove-like in its behavior when its dopaminergic system was compromised, and became Hawk-like in its behavior when its serotonergic system was compromised. Our model suggests how neuromodulatory systems can shape decision-making and adaptive learning in competitive situations.

Proceedings: Pages 161-164

Becoming a Gaze User: Joint Attention and Cooperation in One- and Two-year-olds

Zhen Wu, Jingtong Pan, Yanjie Su,
and Julie Gros-Louis

Abstract— The goal of the current research was to investigate the potential relationship between joint attention and cooperation. Sixty-six children at 14-30 months engaged in a cooperative activity with an adult partner, who stopped participating at a specific point during the tasks. We explored whether children would continue the cooperative action when the cooperative activity was interrupted, and how they would initiate it. In addition, their joint attention abilities were assessed with Early Social Communication Scales. Results showed that children's successful cooperation with adults increased with age, and was predicted by Initiative Joint Attention (IJA). In addition, gaze alternation predicted infants' reengagement behavior during the interruption. However, no significant relations were found between Responding to Joint Attention (RJA) and cooperation. The research suggests a complex relationship between joint attention and coordinated activity when considering RJA and IJA separately.

Proceedings: Pages 165-170

Developmental Learning of Integrating Visual Attention Shifts and Bimanual Object Grasping and Manipulation Tasks

Sungmoon Jeong, Minho Lee, Hiroaki Arie,
and Jun Tani

Abstract— In order to achieve visual-guided object manipulation tasks via learning by example, the current neuro-robotics study considers integration of two essential mechanisms of visual attention and arm/hand movement and their adaptive coordination. The present study proposes a new dynamic neural network model in which visual attention and motor behavior are associated with task specific manners by learning with self-organizing functional hierarchy required for the cognitive tasks. The top-down visual attention provides a goal-directed shift sequence in a visual scan path and it can guide a generation of a motor plan for hand movement during action by reinforcement and inhibition learning. The proposed model can automatically generate the corresponding goal-directed actions with regards to the current sensory states including visual stimuli and body postures. The experiments show that developmental learning from basic actions to combinational ones can achieve certain generalizations in learning by which some novel behaviors without prior learning can be successfully generated.

An Integrated Three-Stage Model Towards Grammar Acquisition

Yo Sato, Joe Saunders, Frank Broz, Caroline Lyon, and Chrystopher L. Nehaniv

Abstract— *This paper presents a three-stage model of language acquisition that integrates phonological, semantic and syntactic aspects of language learning. With the assumption that these three stages arise roughly in sequence, we test the model using the experimental methodology of cognitive robotics, where an emphasis is placed on situating the robot in a realistic, interactive environment. The first, phonological stage consists in learning sound patterns that are likely to correspond to words. The second stage concerns word-denotation association, which relies not only on sensory input but also on the learner's speech output in 'dialogue'. The data thus gathered allows us to invoke semantic bootstrapping in the third, grammar induction stage, where sets of words are mapped with simple logical types. We have started implementing the model and report here on the initial results of the human-robot interaction experiments we conducted.*

Investigating word learning processes in an artificial agent

Michele Gubian, Christina Bergmann, and Lou Boves

Abstract— *Researchers in human language processing and acquisition are making an increasing use of computational models. Computer simulations provide a valuable platform to reproduce hypothesised learning mechanisms that are otherwise very difficult, if not impossible, to verify on human subjects. However, computational models come with problems and risks. It is difficult to (automatically) extract essential information about the developing internal representations from a set of simulation runs, and often researchers limit themselves to analysing learning curves based on empirical recognition accuracy through time. The associated risk is to erroneously deem a specific learning behaviour as generalisable to human learners, while it could also be a mere consequence (artifact) of the implementation of the artificial learner or of the input coding scheme.*

In this paper a set of simulation runs taken from the ACORNS project is investigated. First a look 'inside the box' of the learner is provided by employing novel quantitative methods for analysing changing structures in large data sets. Then, the obtained findings are discussed in the perspective of their ecological validity in the field of child language acquisition.

Prenatal to Postnatal Transfer of Motor Skills Through Motor-Compatible Sensory Representations

Timothy A. Mann and Yoonsuck Choe

Abstract— *How can sensory-motor skills developed as a fetus transfer to postnatal life? We investigate a simulated reaching task by training controllers under prenatal conditions (i.e. confined space) and evaluating them based on postnatal conditions (i.e. targets outside of the confined training space). One possible solution is to identify a sensory representation that is easy to extrapolate over. We compared two kinds of sensory representations: world-centered sensory representation based on Cartesian coordinates and agent-centered sensory representation based on polar coordinates. Despite similar performance under prenatal conditions, controllers using agent-centered sensory representation had significantly better performance than controllers using world-centered sensory representation under postnatal conditions. It turns out that the success of the agent-centered sensory representation is (in part) due to being complementary to the action encodings. Further analysis shows that the action encodings (i.e. changes in joint angles) were highly predictive of the change in state when agent-centered sensory representation was used (but not world-centered). This suggests that a powerful strategy for transferring sensory-motor skills to postnatal life involves selecting a sensory representation that complements the action encodings used by an agent.*

Real Time Targeted Exploration in Large Domains

Todd Hester and Peter Stone

Abstract— *A developing agent needs to explore to learn about the world and learn good behaviors. In many real world tasks, this exploration can take far too long, and the agent must make decisions about which states to explore, and which states not to explore. Bayesian methods attempt to address this problem, but take too much computation time to run in reasonably sized domains. In this paper, we present *TEXPLORE*, the first algorithm to perform targeted exploration in real time in large domains. The algorithm learns multiple possible models of the domain that generalize action effects across states. We experiment with possible ways of adding intrinsic motivation to the agent to drive exploration. *TEXPLORE* is fully implemented and tested in a novel domain called *Fuel World* that is designed to reflect the type of targeted exploration needed in the real world. We show that our algorithm significantly outperforms representative examples of both model-free and model-based RL algorithms from the literature and is able to quickly learn to perform well in a large world in real-time.*

Maturationally-Constrained Competence-Based Intrinsically Motivated Learning

Adrien Baranes and Pierre-Yves Oudeyer

Abstract— *This paper studies the coupling of intrinsic motivation and physiological maturational constraints, and argues that both mechanisms may have complex bidirectional interactions allowing to actively control the growth of complexity in motor development. First, we introduce the self-adaptive goal generation algorithm (SAGG), instantiating an intrinsically motivated goal exploration mechanism for motor learning of inverse models. Then, we introduce a functional model of maturational constraints inspired by the myelination process in humans, and show how it can be coupled with the SAGG algorithm, forming a new system called McSAGG. We then present experiments to evaluate qualitative properties of these systems when applied to learning a reaching skill with an arm with initially unknown kinematics.*

A developmental approach to the emergence of communication in socially situated embodied agents

Michael Sheldon and Mark Lee

Abstract— *This paper reports on a developmental approach to the learning of communication in embodied agents, taking inspiration from child development and recent advances in the understanding of the mirror neuron system within the brain. We describe a part of the ROSSI project which focuses upon gestural communication in the form of pointing. We are examining the idea that pointing may be a key step towards simple spoken communication and exploring the internal representations that may be formed during this process.*

The possible developmental stages leading to proto-imperative pointing actions in a robotic system are outlined, and how this may be built upon to result in an understanding of two word speech is discussed. The learning mechanism is based around Piagetian schema learning whilst the developmental path follows a mixture of Piagetian and Vygotskian theories.

A Reinforcement Learning Model of Reaching Integrating Kinematic and Dynamic Control in a Simulated Arm Robot

*Daniele Caligiore, Eugenio Guglielmelli,
Anna M. Borghi, Domenico Parisi,
and Gianluca Baldassarre*

Abstract— *Models proposed within the literature of motor control have polarised around two classes of controllers which differ in terms of controlled variables: the Force-Control Models (FCMs), based on dynamic control, and the Equilibrium-Point Models (EPMs), based on kinematic control. This paper proposes a bioinspired model which aims to exploit the strengths of the two classes of models. The model is tested with a 3D physical simulator of a 2DOF-controlled arm robot engaged in a reaching task which requires the production of curved trajectories to be solved. The model is based on an actor-critic reinforcement-learning algorithm which uses neural maps to represent both percepts and actions encoded as joint-angle desired equilibrium points (EPs), and a noise generator suitable for fine tuning the exploration/exploitation ratio. The tests of the model show how it is capable of exploiting the simplicity and speed of learning of EPMs as well as the flexibility of FCMs in generating curved trajectories. Overall, the model represents a first step towards the generation of models which exploit the strengths of both EPMs and FCMs and has the potential of being used as a new tool for investigating phenomena related to the organisation and learning of motor behaviour in organisms.*

Batch versus Interactive Learning by Demonstration

*Peng Zang, Runhe Tian, Andrea L. Thomaz,
and Charles L. Isbell*

Abstract— *Agents that operate in human environments will need to be able to learn new skills from everyday people. Learning from demonstration (LfD) is a popular paradigm for this. Drawing from our interest in Socially Guided Machine Learning, we explore the impact of interactivity on learning from demonstration. We present findings from a study with human subjects showing people who are able to interact with the learning agent provide better demonstrations (in part) by adapting based on learner performance which results in improved learning performance. We also find that interactivity increases a sense of engagement and may encourage players to participate longer. Our exploration of interactivity sheds light on how best to obtain demonstrations for LfD applications.*

Proceedings: Pages 225-230

Toddlers' Transitions on Non-verbal False-belief Tasks Involving a Novel Location: A Constructivist Connectionist Model

Vanessa C. Evans, Vincent G. Berthiaume,
and Thomas R. Shultz

Abstract— Some argue that children learn a Theory of Mind (ToM), the understanding that others have mental states, at around 3.5 years. This is evidenced by their transition from failure to success on verbal false-belief tasks, when they begin to verbally predict an actress will search for a toy where she falsely believes it to be, rather than in its actual location. However, non-verbal measures have recently been used to show that children in their second year of life may already have some understanding of others' false beliefs. We present a Sibling-Descendant Cascade-Correlation neural-network model of one study that found 25-month-old toddlers correctly anticipated an actress would search according to her false belief. Networks were trained on true- and false-belief search patterns, simulating toddlers' everyday experience with true and false beliefs, and then tested on non-verbal true- and false-belief tasks involving a novel location. Networks transitioned from incorrectly predicting true-belief searches in both true- and false-belief tasks to making correct predictions in both tasks. Our model thus (1) reproduced the transition that has been observed in older children and (2) generalized its learning to a novel location. The model can be used to refine our understanding of the transitions while again demonstrating the usefulness of SDCC as an algorithm for modeling cognitive development.

Proceedings: Pages 231-237

How can bottom-up information shape learning of top-down attention-control skills?

Dimitri Ognibene, Giovanni Pezzulo,
and Gianluca Baldassarre

Abstract— How does bottom-up information affect the development of top-down attention control skills during the learning of visuomotor tasks? Why is the eye fovea so small? Strong evidence supports the idea that in humans foveation is mainly guided by task-specific skills, but how these are learned is still an important open problem. We designed and implemented a simulated neural eye-arm coordination model to study the development of attention control in a search-and-reach task involving simple coloured stimuli. The model is endowed with a hard-wired bottom-up attention saliency map and a top-down attention component which acquires task-specific knowledge about spatial relationships and retains information about potential gaze targets. This architecture achieves high performance very fast. To explain this result, we argue that: (a) the interaction between bottom-up and top-down

mechanisms supports the development of task-specific attention control skills by allowing an efficient exploration of potentially useful gaze targets; (b) bottom-up mechanisms permits the exploitation of the initial limited task-specific knowledge by actively selecting areas where it can be suitably applied; (c) bottom-up processes shape objects representation, their value, and their roles (these can change during learning, e.g. distractors can become useful attentional cues); (d) increasing the size of the fovea alleviates perceptual aliasing, but at the same time increases input processing costs and the number of trials required to learn. Overall, the results indicate that bottom-up attention mechanisms can play a relevant role in attention control, especially during the acquisition of new task-specific skills.

Proceedings: Pages 238-243

On the Influence of Sensor Morphology on Eye Motion Coordination

H. Martinez, M. Lungarella, and R. Pfeifer

Abstract— Developmental robotics focuses on how to endow robots with adaptive capabilities. Even though embodiment has been recognized as an essential factor for understanding development, there is yet not much work that investigates how the morphology of sensors and actuators shapes adaptivity and learning processes. Moreover, these studies are largely at an intuitive and qualitative level. In this paper, we address the issue by studying how in an active vision system sensor morphology and bodily features affect a behavior such as vergence. Specifically, we present an information-theoretic analysis of two experiments showing how adequate sensor morphology influences statistical dependencies in the sensorimotor loop. The results show that an appropriate morphology reduces the amount of input without disrupting the information structure in the sensorimotor loop. The second result shows how the later morphology under the vergence behavior increases the information structure among the motor actions and the pixels. We also speculate on the implications of our results for attention, reaching and grasping.

Proceedings: Pages 244-250

Scenes and Tracking with Dynamic Neural Fields: How to Update a Robotic Scene Representation

Stephan K. U. Zibner, Christian Faubel,
Ioannis Iossifidis, Gregor Schöner,
and John P. Spencer

Abstract— We present an architecture based on the Dynamic Field Theory for the problem of scene representation. At the core of this architecture are three-dimensional neural fields linking feature to spatial information. These three-dimensional fields are coupled to lower-dimensional fields that provide both

a close link to the sensory surface and a close link to motor behavior. We highlight the updating mechanism of this architecture, both when a single object is selected and followed by the robot's head in smooth pursuit and in multi-item tracking when several items move simultaneously.

Proceedings: Pages 251-256

Serial order in an acting system: a multidimensional dynamic neural fields implementation.

Yulia Sandamirskaya and Gregor Schöner

Abstract— Learning and generating serially ordered sequential behavior in a real, embodied agent that is situated in a partially unknown environment requires that noisy sensory information is used both to control appropriate motor actions and to determine that a particular action has been successfully terminated. While most current models do not address these conditions of embodied sequence generation, we have earlier proposed a neurally inspired model based on Dynamic Field Theory that enables sequences in which each action may take unpredictable amounts of time. Here we extend this earlier work to accommodate heterogeneous sets of actions. We show that a set of matching conditions- of-satisfaction can be used to stably represent the terminal condition of each action and trigger the cascade of instabilities that switches the system from one stable state to the next. A robotic implementation on a vehicle with a camera and a simple robot arm demonstrates the stability of the resulting scheme.

Proceedings: Pages 257-262

Competence Progress Intrinsic Motivation

Andrew Stout and Andrew G. Barto

Abstract— One important role of an agent's motivational system is to choose, at any given moment, which of a number of skills the agent should attempt to improve. Many researchers have suggested "intrinsically motivated" systems that receive internal reward for model learning progress, but for the most part this notion has not been applied with respect to skill competence, or to choose between skills. In this paper we propose an agent motivated to gain competence in its environment by learning a number of skills, addressing head-on the mechanism of competence progress motivation for the purpose of governing the efficient learning of skills. We demonstrate this new approach in a simple illustrative domain and show that it outperforms a naïve agent, achieving higher competence faster by focusing attention and learning effort on skills for which progress can be made while ignoring those skills that are already learned or are at the moment too difficult.

Proceedings: Pages 263-268

Imitation Learning with Hierarchical Actions

Abram L. Friesen and Rajesh P. N. Rao

Abstract— Imitation is a powerful mechanism for rapidly learning new skills through observation of a mentor. Developmental studies indicate that children often perform goal-based imitation rather than mimicking a mentor's actual action trajectories. Further, imitation, and human behavior in general, appear to be based on a hierarchy of actions, with higher-level actions composed of sequences of lower-level actions. In this paper, we propose a new model for goal-based imitation that exploits action hierarchies for fast learning of new skills. As in human imitation, learning relies only on sample trajectories of mentor states. Unlike apprenticeship or inverse reinforcement learning, the model does not require that mentor actions be given. We present results from a large-scale grid world task that is modeled after a puzzle box task used in developmental studies for investigating hierarchical imitation in children. We show that the proposed model rapidly learns to combine a given set of hierarchical actions to achieve the subgoals necessary to reach a desired goal state. Our results demonstrate that hierarchical imitation can yield significant speed-up in learning, especially in large state spaces, compared to learning without a mentor or without an action hierarchy.

Proceedings: Pages 269-273

Autism as an impairment in detecting invariants

Norbert Michael Mayer and Ian Fasel

Abstract— Individuals with Autism Spectrum Disorders (ASD) show a variety of atypical behavioral and perceptual patterns which seem to lack a common underlying cause. We hypothesize that there may be a general impairment of the autistic brain characterized by a restriction on the class of connectivity patterns, i.e. features, that can be utilized for learning perceptual and cognitive tasks. In particular, we suggest that the autistic brain may not make proper use of features which pool information over larger areas of the input space, which would otherwise allow them to make use of symmetries and develop invariants to permutations.

Proceedings: Pages 274-279

Robot Motivator: Increasing User Enjoyment and Performance on a Physical/Cognitive Task

Juan Fasola and Maja J. Mataric

Abstract— *We describe the design and implementation of a socially assistive robot that is able to monitor the performance of a user during a combined cognitive and physical task, with the purpose of providing motivation to the user to complete the task and to improve task performance. The work presented aims to study the effects of verbal praise, encouragement, and motivation on a user's enjoyment of the task, and to analyze the effectiveness of online adaptation of the task in response to user performance, with the goal of reducing user frustration and increasing user intrinsic motivation. A three-condition study was constructed for evaluation of the robot; the results of the robot's interaction with human users are presented.*

Proceedings: Pages 280-285

WWN-Text: Cortex-Like Language Acquisition with "What" and "Where"

Kajal Miyan and Juyang Weng

Abstract— *Based on some recent advances in understanding and modeling cortical processing for space and time, we propose a developmental, general purpose model for language acquisition using multiple motor areas. Through the ventral pathway, the "what" motor learns, abstracts and projects (as recurrent top-down context) information that is related to the meaning of the text. Via the dorsal pathway, the "where/how" motor learns, abstracts and projects (as top-down context) information that relates to the spatial information of text, e.g., where is the text on a page. This is a major departure from the traditional symbolic and connectionist approaches to natural language processing (NLP) — the nature of the motor areas, i.e., actions, of the developmental agent play the role of "hubs" in language acquisition and understanding. As any human communicable concept can be either verbally stated (what) or demonstrated through actions (how), this model seems to be the first general purpose developmental model for general language acquisition, although the size of experiments is still very small. Furthermore, unlike traditional NLP approaches, we do not use hand-crafted language structure but allow primary and secondary associations as seen in animal learning, as a general scheme for language acquisition.*

Proceedings: Pages 286-290

The Similarity-Attraction Effect in Human-Robot Interaction

Emily P. Bernier and Brian Scassellati

Abstract— *Constructing valid robotic models of social development requires that we accurately characterize the social learning and interaction that can take place between a robotic agent and a human adult. To that end, this study examined the effect of perceived attitudinal similarity on human-robot interaction. 28 participants rated toys by order of preference and then interacted with a small, socially-expressive robot to determine the robot's preferences for the same toys. The robot displayed either the same preferences as the participant or exactly the opposite preferences. Participants in the Similar-Preferences condition rated the robot as significantly friendlier than did participants in the Dissimilar-Preferences condition. However, there was no difference between conditions in how participants rated their enjoyment of the interaction. These findings have interesting implications for human-robot interaction studies in general, and for work in robotic models of developmental social cognition specifically.*

Proceedings: Pages 291-296

Mutually constrained multimodal mapping for simultaneous development: modeling vocal imitation and lexicon acquisition

Yuki Sasamoto, Yuichiro Yoshikawa, and Minoru Asada

Abstract— *This paper presents a method of simultaneous development of vocal imitation and lexicon acquisition with a mutually constrained multimodal mapping. A caregiver is basically assumed to give matched pairs for mappings, for example by imitating the learner's voice or labelling an object that it is looking at. However, the tendency cannot be always expected to be reliable. Subjective consistency is introduced to judge whether to believe the observed experiences (external input) as reliable signal for learning. It estimates the value of one layer by combining the values from other layers and external input. Based on the proposed method, a simulated infant robot learns mappings among the representations of its caregiver's phonemes, those of its own phonemes, and those of objects. The proposed mechanism enables correct mappings even when caregivers do not always give correct examples, as real caregivers do not for their infants.*

A Dynamic Field Theoretic Model of Iowa Gambling Task Performance

Robert Lowe, Boris Duran, and Tom Ziemke

Abstract— Choice behaviour where outcome-contingencies vary or are probabilistic has been the focus of many benchmark tasks of infant to adult development in the psychology literature. Dynamic field theoretic (DFT) investigations of cognitive and behavioural competencies have been used in order to identify parameters critical to infant development. In this paper we report the findings of a DFT model that is able to replicate normal functioning adult performance on the Iowa gambling task (IGT). The model offers a simple demonstration proof of the parsimonious reversal learning alternative to Damasio's somatic marker explanation of IGT performance. Our simple model demonstrates a potentially important role for reinforcement/reward learning to generating behaviour that allows for advantageous performance. We compare our DFT modeling approach to one used on the A-not-B infant paradigm and suggest that a critical aspect of development lies in the ability to flexibly trade off perseverative versus exploratory behaviour in order to capture statistical choice-outcome contingencies. Finally, we discuss the importance of an investigation of the IGT in an embodied setting where reward prediction learning may provide critical means by which adaptive behavioural reversals can be enacted.

Discovering the concept of anaphora from grounded verb models

Kruti Neema and Amitabha Mukerjee

Abstract— A number of computational models simulate the grounded learning of units of language in the early learner. But can this initial lexical knowledge be used to acquire complex grammatical notions such as anaphora? We build on earlier work, where we simulate a language learner with perceptual attention and learn, in an unsupervised manner, a set of action models along with the participating agents, and then the corresponding linguistic units (verbs). We consider how this knowledge may be used to bootstrap the learning of anaphora. Given an input video with moving shapes, the system considers human narratives that refer to this scene. The acquired perceptual schemas and their arguments are mapped to the appropriate verbs and nouns in the discourse. We first detect the synonyms of the arguments as the repeated labels used in the constructions referring to a known action scene. After ruling out the synonyms, we find that the anaphora remain as units that are referring to more than one grounded object. We show that both third-person singular and plural anaphors and even a common reciprocal anaphor ("each other") can be discovered. Finally, we show that in situations where the referent is missing altogether (zero anaphora), certain correlations may also be inferred from the regularities in the mapping between perceptual schemas and language.

Where-What Network-4: The Effect of Multiple Internal Areas

Matthew Luciw and Juyang Weng

Abstract— The general visual attention-recognition (AR) problem remains open. Given a set of images, each with a single target foreground over some complex background, it requires output of both location and type of this single foreground. First, many approaches cannot deal with the richness of the class of possible backgrounds, which has a huge number of variations and also could include distractor-like patterns. This potentially leads to an explosion of resources required within the model. Second, all current approaches break down as the number of locations, types, and variations (within each type) increases towards human-level. This paper is concerned with model selection for networks dealing with the general AR problem. The major challenge is ensuring the model remains as simple as possible as the complexity of the data increases. In developmental general AR, the model must be adapted on the fly. We discuss these issues in context of the latest version of the biologically-inspired developmental Where-What Network. We show how local detectors reduce the number of neurons exponentially and deal with the complex background problem. The purpose of multiple layers seems to be to allow combinatorial patterns to emerge. Top-down connections cause more discriminative features to develop, but since complex data requires a bank of shared features, top-down connections are probably not beneficial for the early layer(s). When a layer's features are class-specific and there is no combinatorial structure to exploit on top of this layer, it is not useful to add another layer but it is useful to utilize top-down connections to develop more discriminative features.

External Rotation as Morphological Bootstrapping for Emergence of Biped Walking

Koh Hosoda and Yoichiro Ishii

Abstract— Many researchers are interested in the onset and learning of bipedal walking, but still not much is known how a human (and even a robot) can acquire the ability. In this paper, we hypothesize that external rotation of the hip joint plays an essential role for emergence of bipedal walking in human infancy. We build an infant robot "Pneu-born 13" driven by pneumatic artificial muscles so that we would get biologically plausible results. The robot can walk without any intention to move forward by utilizing the external rotation of the hip. We also analyze a quasi-static model of the robot and find that the joint compliance plays an important role in this scheme.

Two Page Abstracts

Poster Spotlights: 5:10PM-5:45PM, Friday, 20 August 2010

Poster Session: 5:45PM-7:00PM, Friday, 20 August 2010

An experiment with human-robot interaction to study intentional agency in joint visual attention

Takeshi Konno and Takashi Hashimoto

Does embodiment affect tutoring behavior?

Katrin Solveig Lohan, Sebastian Gieselmann, Anna-Lisa Vollmer, Katharina J. Rohlfing, and Britta Wrede

Student computers: Six-year-olds believe that computers can learn

Maria D. Vázquez, Megan M. Saylor, and Daniel T. Levin

Variability Facilitates Learning of Relations Between Words and Actions

Heidi A. Baumgartner, Courtney Millhoff, Katharine Graf Estes, and Lisa M. Oakes

Development of Perception and Representation of Rhythmic Information: Towards a Computational Model

Inês Salselas and Perfecto Herrera

Modeling Real-time Multimodal Interaction with Virtual Humans

Hui Zhang, Damian Fricker, and Chen Yu

Development of Semantic Network: Evidence from a Forced Choice Association

Shohei Hidaka

Infants' learning about human and non-human faces: An eye-tracking study

Emily A. Spring, Karinna B. Hurley, and Lisa M. Oakes

How to Build a Schema-Representation

Nikolas J. Hemion, Frank Joublin, and Katharina J. Rohlfing

A computational model of learning intuitive object dynamics

Lewis Fishgold, Dana H. Ballard, and Benjamin Kuipers

Learning-to-Grasp: from an Infant to a Troublemaker

Baro Hyun, Anouck Girard, and Benjamin Kuipers

Learning representations and detecting changes in complex environments

Kat Snyder, Dmitry Kit, Brian Sullivan, Mary Hayhoe, Dana Ballard

A Boosting Approach to Understanding Referential Intent

Daniel Yurovsky, Martin Rickert, Chen Yu, and Linda B. Smith

Preliminary Findings on Characteristics of Handwriting in Children With Autism and Asperger's Disorder

Beth Johnson, Nicole Rinehart, James Phillips, Joanne Fielding, and Kerry Saunders

A Comparative Study of Multimodal Human- Robot and Human-Human Interactions

Damian Fricker, Chen Yu, Matthias Scheutz, and Paul Schermerhorn

Discovering complex categories through active sampling

Doug Markant and Todd Gureckis

