

Agents Learning Interactively
from Human Teachers

ALIHT 2011

W. Bradley Knox
Jake Beal

Brenna Argall
Sonia Chernova
Peter Stone
Matt Taylor
Andrea Thomaz

These slides are posted on
the ALIHT website's Program page.

Welcome!

Quick stats

- **14 papers**
- **5 invited talks**
 - **Joanna Bryson** (University of Bath)
 - **Thomas G. Dietterich** (Oregon State)
 - **Ian Fasel** (University of Arizona)
 - **Jan Peters** (Max-Planck Institute)
 - **Dan Roth** (University of Illinois - Urbana Champaign)

Best Presentation Award



[Home](#) [Submissions](#) [Program](#) [Accepted Papers](#) [Organizers](#)

Voting for Best Presentation

Email your vote to Brenna Argall at brennadee.argall@epfl.ch between the end of the last paper presentation and midnight local time that night (Sunday, 7/17). Include paper title and authors. We will post the winner shortly thereafter.

Full schedule:

Day 1 (Saturday, 7/16)

8:30 - 9:00 : Registration

9:00 - 11:00: Session 1A

- Introductory talk (30 min.)
- Invited talk - Joanna Bryson (1 hr)
- Towards Task Understanding through Multi-State Visuo-Spatial Perspective Taking for Human-Robot Interaction

11:00 - 11:30: Coffee break

11:30 - 13:00: Session 1B - Learning the meaning of human communication / Multiple teaching modalities (Session Chair - Joanna Bryson)

- Game-based Language Tutoring
- Simultaneous Acquisition of Task and Feedback Models
- A Hidden Domain for Human and Electronic Students

13:00 - 14:30: Lunch break

Agents Learning

Interactively

human sees an effect of learning before teaching finishes (teach -> observe learning -> teach)

from Human Teachers

implies the human considers the student and communicates intentionally

Outline

- Why?
- Taxonomy
- Discussion points/questions

Why? (grounded answers)

- Programming for non-programmers
- Customization/extension by the end-user
- Faster and/or less costly learning
- “You don’t know something until you teach it.”
- To study how people teach

Why? (speculative answers)

- Interaction may build trust and human understanding of the agent
- Learning creates social connection
- The thrill of teaching
- Human-centered AI

**From many contributions,
sorting it out**

Purpose of teaching

- Autonomous task completion
 - Teaching new tasks
 - Customizing existing task solutions
- Improving communication
- Learning through teaching

Human-to-agent communication modalities

- Demonstration
- Reward/punishment
- Verbal advice/directions
- Curriculum design / Environment shaping
- Gestures

- Unconstrained interaction
- Unintentional signals (e.g., facial expressions)

Agent-to-human communication modalities

- Observable behavior
- Asking (for help, information, guidance, etc.)
- Belief/prediction statements
- Emotional expression

Interaction scheme

- Iterations between teacher and student
- Teacher and student act concurrently

Knowledge representation

- Behavior parameters
- Value functions
- Probabilistic/predictive models
- Logical formulas

Learning from multiple sources

- Multiple teaching modalities (demonstration and feedback)
- Combining with non-teaching information (e.g., MDP reward for reinforcement learning)

Evaluation metrics

- Effectiveness - learned performance
- Efficiency
 - Human time
 - Training cost by performance
- User satisfaction

Taxonomy

- Purpose of teaching
- Human-to-agent communication
- Agent-to-human communication
- Interaction scheme
- Knowledge representation
- Learning from multiple sources
- Evaluation metrics

Let's discuss

(over the next two days)

Comparative evaluation

Interactive algorithms often aren't compared.

But we must evaluate relative strengths to move forward.

Standardized challenge task?

- room for robots?

Theory

What should we try to prove?

What assumptions must be made?

At what cost to applicability?

Perhaps one of our goals should be to provide the correct assumptions.

Gathering/reusing data

Ease:

Supervised learning > reinforcement learning
> learning interactively from a human

In what situations can data be reused?

Strategies for reducing cost of human data?

Experimental logistics

Experiments with authors or colleagues as subjects yield narrower results.

But technical academic departments often lack infrastructure for facilitating human studies.

Tap our collective experience in creating such infrastructure.

Publishing venues

General AI - IJCAI, AAAI

Machine learning – ICML, ECML, NIPS,

Agents-focused – AAMAS, GECCO, IVA

Robots/Interaction – HRI, ICRA, IROS, ROMAN, RSS(?)

HCI/Interfaces – IUI, UMAP, CHI, SIGGRAPH(?)

Developmental learning – ICDL

NLP - ACL, CoNLL, EMNLP, NAACL

Journals - TAMM (and *many* others)

Reviewers

ALIHT straddles several areas, and reviewers often come from narrower backgrounds.

Strategies for addressing reviewer's biases?

(e.g., from the RL community, arguably misplaced standards for theory and extensiveness of experiments and too much lenience on number and source of subjects)

At community and individual levels

Fundamentals of ALIHT

Is our task to integrate developments from machine learning, psychology, etc.?

Or are there fundamental contributions that generalize across the ALIHT subfield?

- Biggest bottlenecks?
- What can we offer our larger communities?

And what can we take from each other?

Proposed discussion topics

- Comparative evaluation
- Theory
- Gathering/reusing data
- Experimental logistics
- Publishing venues
- Reviewers
- Fundamentals of ALIHT

Enjoy! (And discuss!)