The Human Experience in Interactive Machine Learning

Karen M. Feigh
Samantha Krening

Cognitive Engineering Center
Georgia Institute of Technology

GOAL

To enable people
to naturally and intuitively
teach agents to perform tasks.

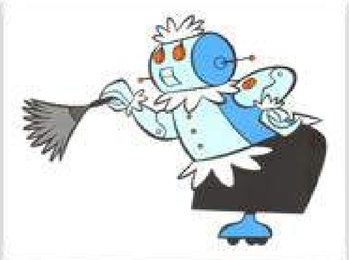

Widespread integration of robotics requires ML agents that are
- more accessible,
- easily customizable,
- more intuitive for people to understand.

How do people teach?

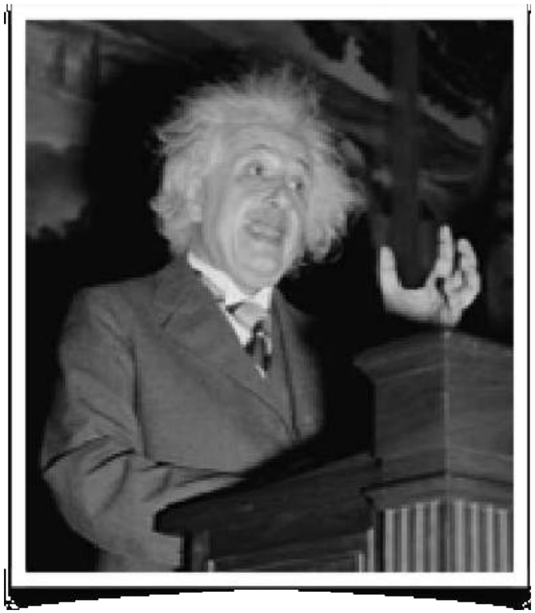

Explanation

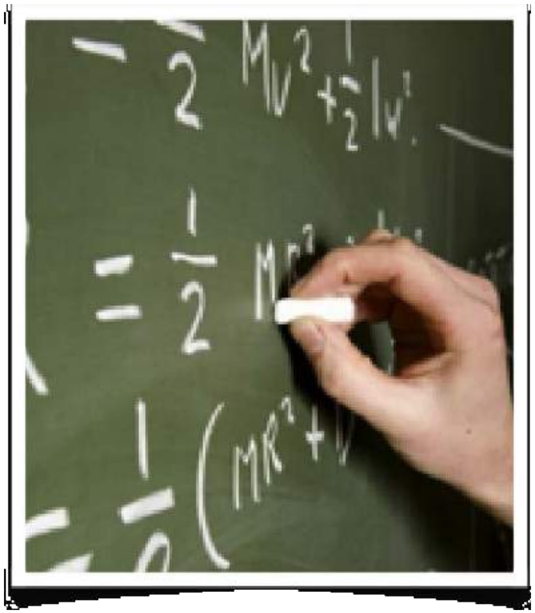

Demonstration

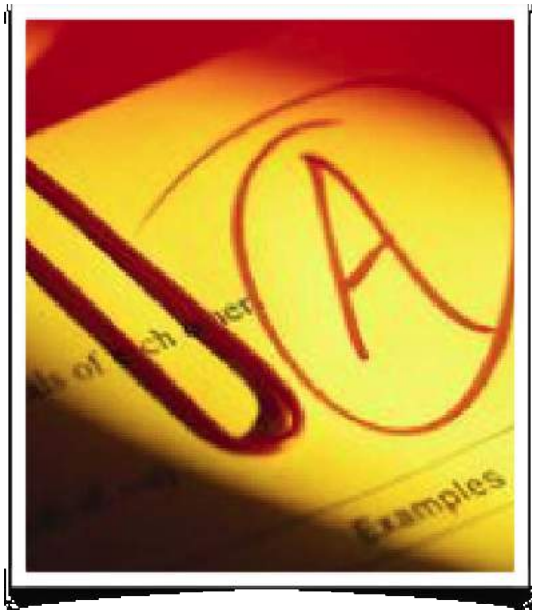

Critique

Design Algorithm

- Design with the human in mind!
  - Expected behavior/teaching template
  - Design interaction to improve human factors
- Don’t tack on HF analysis as an afterthought. Instead, use it to direct design.

ML Testing

- Oracles
- Simulations
- Traditional ML measures
  - Learning curve
  - Training time
  - # inputs required
- Compare measures to other algorithms and say it is better based on quantitative ML measures

Human-Subject Experiment

- Traditional ML measures
- Human Factors
  - Frustration
  - Perceived performance and intelligence
  - Immediacy
  - Clarity
  - Expected behavior

Human Factors should be used to direct the design process

A lot of research stops with ML testing

Many Human-Subject experiments are proof of concept and do not measure human factors

Reinforcement Learning with Human Verbal Input

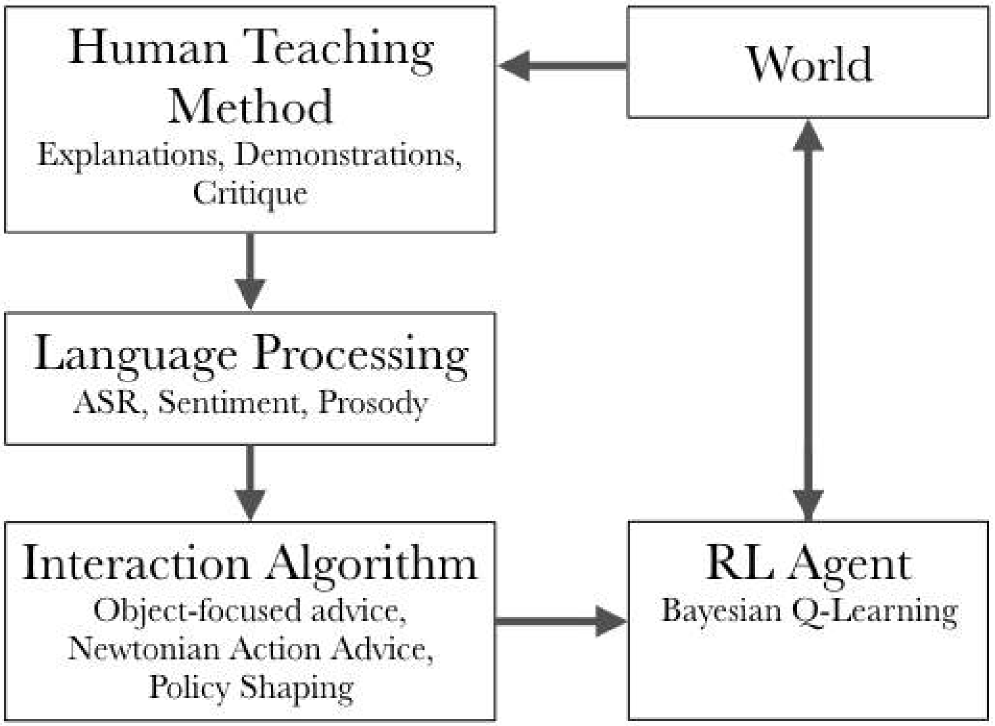

Initial Study:

ADVICE VS CRITIQUE

Research Questions

- How does the interaction method affect
  - The experience of the human teacher?
  - Perceived intelligence of the agent?

Created two different IML agents

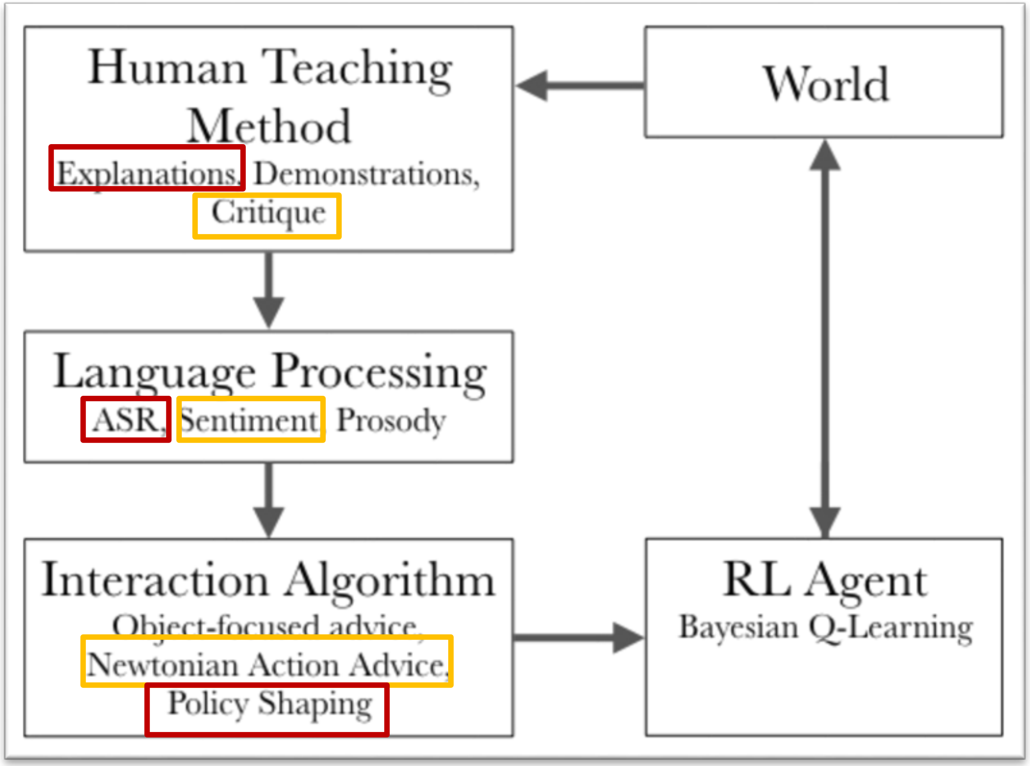

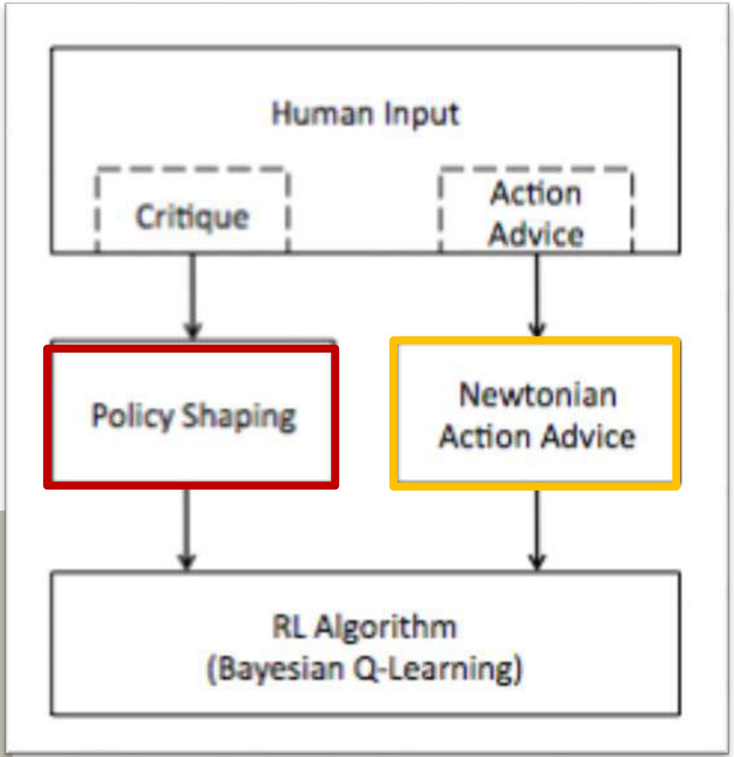

Development

- Create the Newtonian Action Advice algorithm
- Create method for filtering critique using sentiment analysis
- Human factors design and analysis

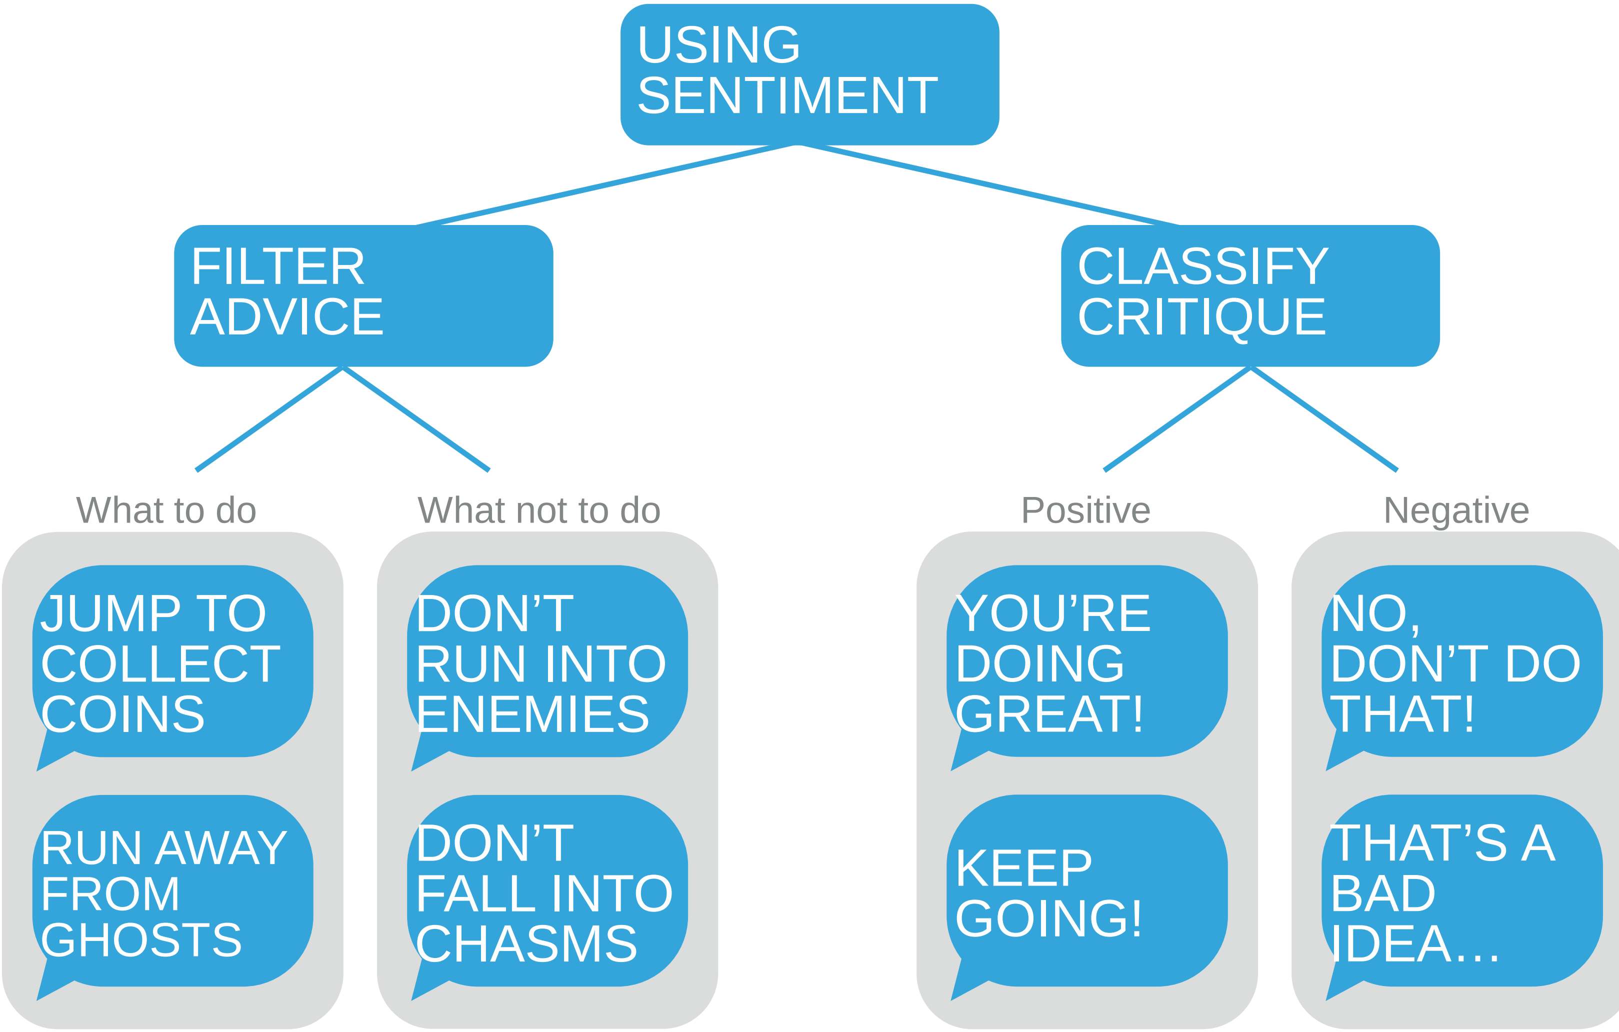

Published in IEEE Transaction on Cognitive and Developmental Systems. Special Issue on Cognitive Agents and Robotics for Human-Centred Systems. Publication: December.

Newtonian Action Advice

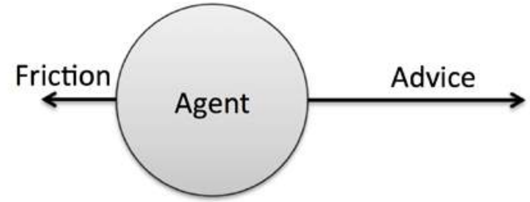

NAA is an IML algorithm that connects action advice ("move left") to an RL agent.

- The advice is a 'force' that causes an initial push.
- Afterward, 'friction' works to stop the agent from following the advice after a time
- Then, the agent reverts to normal exploration vs. exploitation

More Research Questions

- How does the interaction method affect
  - The experience of the human teacher?
  - Perceived intelligence of the agent?

- Can sentiment analysis filter natural language critique?
- Can prosody be used as an objective metric for frustration?
- Is NAA intuitive to train?

Task/Game domain

Simple Grid-world game, Radiation World

World is static & fully observable.
Humans usually know correct & optimal solution

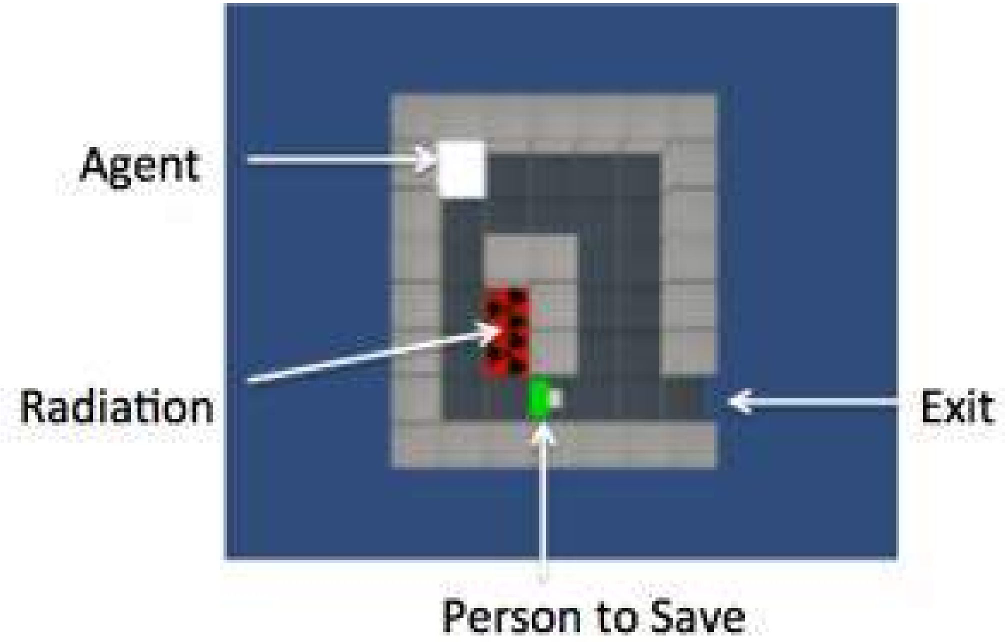

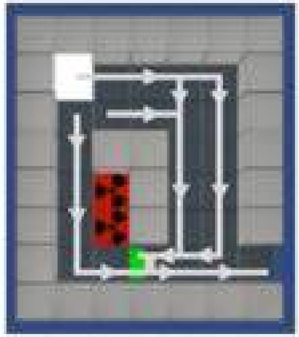

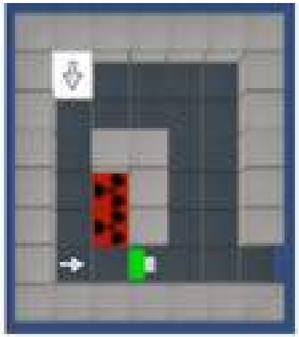

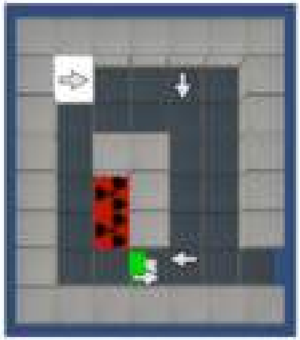

Human-in-the-loop Experiment

Procedure

For each agent:

- Participants given instructions about how to train the agent and allowed to practice.
- Participants asked to train an agent for as many training episodes as they felt necessary or until they decided to give up
- Participants completed questionnaire about their experience

- After training both agents:
    - Questionnaire comparing the experiences of training both agents

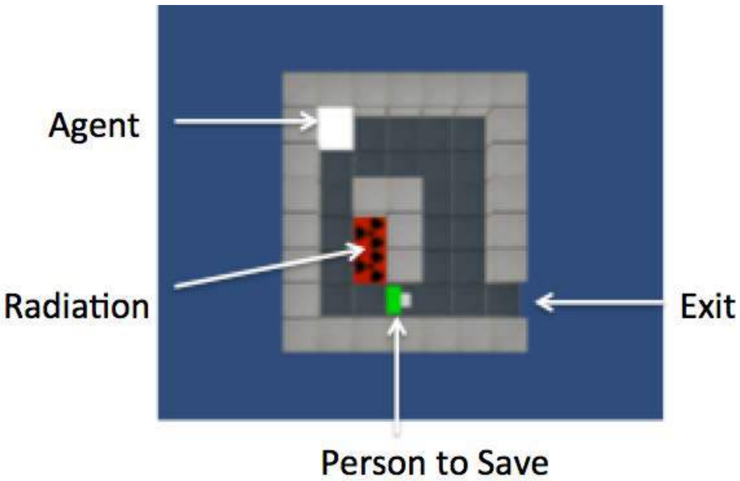

Domain: Radiation World (unity)

24 Participants with little to no experience with ML participated
Training order was balanced.

Metrics

- Wanted to understand the human teacher's experience training the agent.
- Wanted to understand how the human teacher's perceived the intelligence of the ML agent.
- *We modified a common workload scale to rate qualities that had been found in the literature to impact experience and intelligence.*
- *We also asked for free-form explanation of responses*

Human Factors Metrics

- **Perceived Intelligence**
  - How smart the participants felt the algorithm was
- **Frustration**
  - Degree of frustration participant felt training the agent.
- **Perceived Performance**
  - How well the participants felt the algorithm learned
- **Transparency**
  - How well the participants feel they understood what the agent was doing
- **Immediacy**
  - Degree to which the agent follow advice as fast as desired

Traditional ML Metrics

- Performance metrics
  - Cumulative reward
- Efficiency metrics
  - Training time
  - Human input
  - Number of actions to complete episode

Traditional Metrics

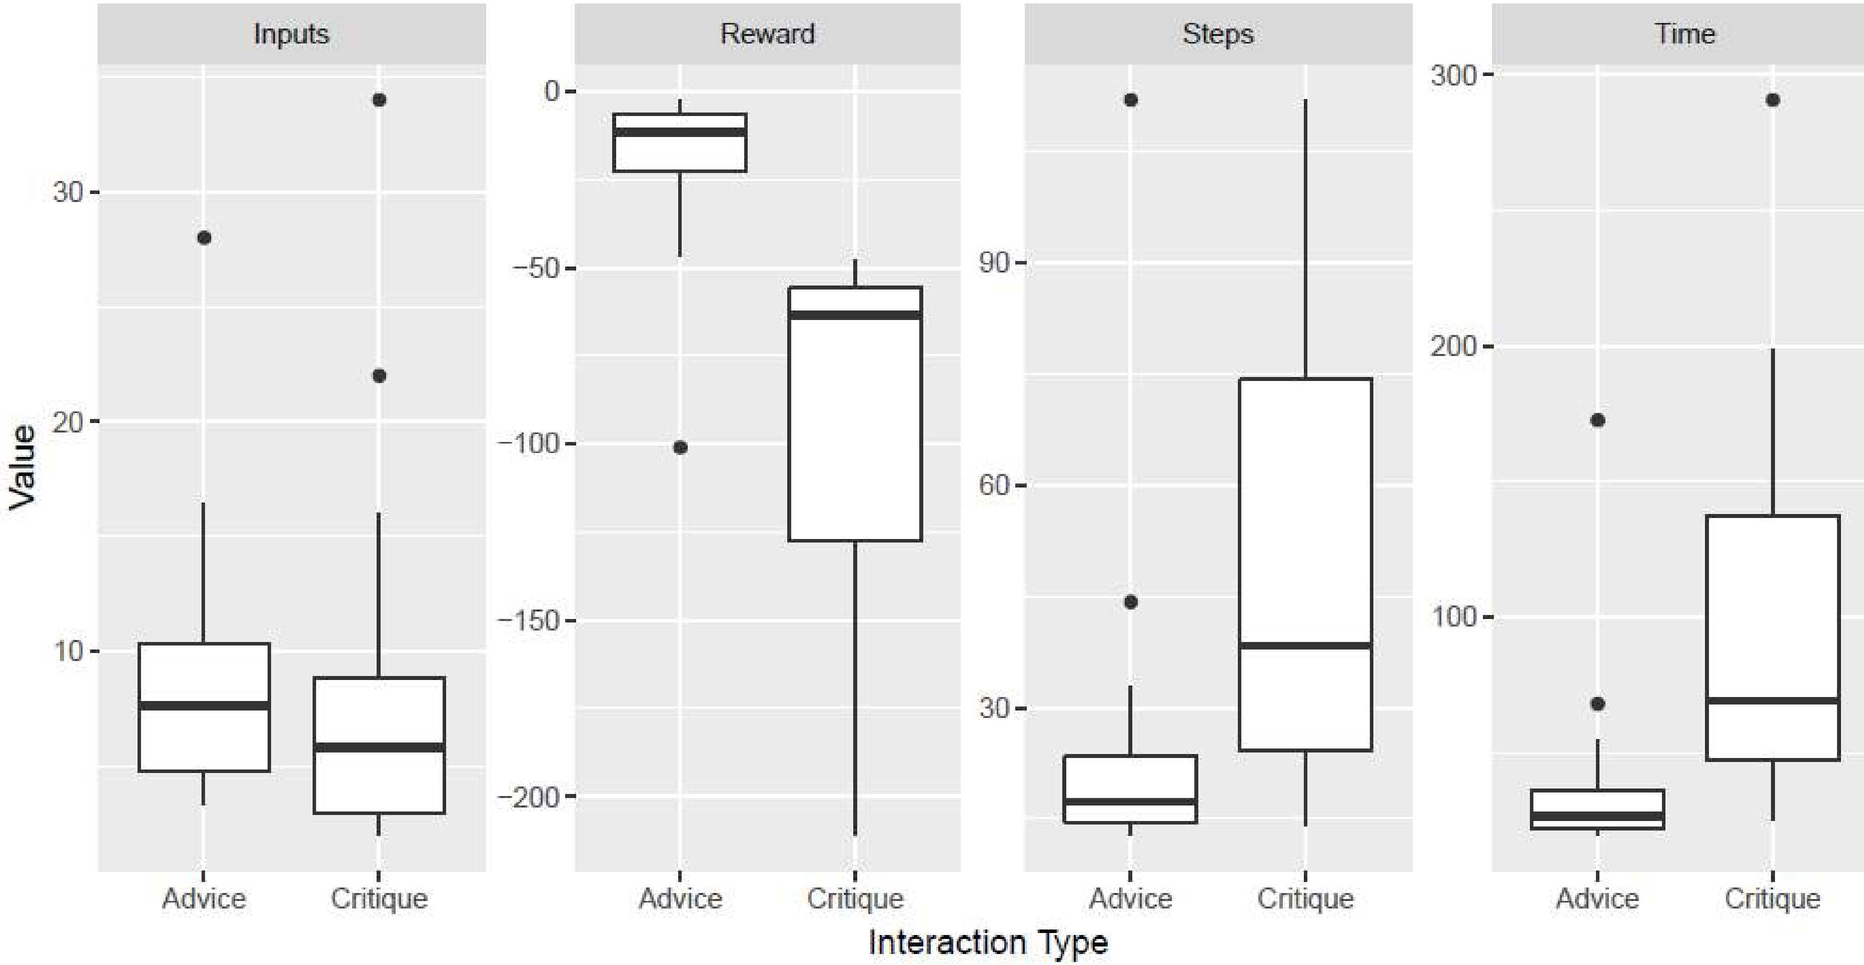

Human Factors Metrics

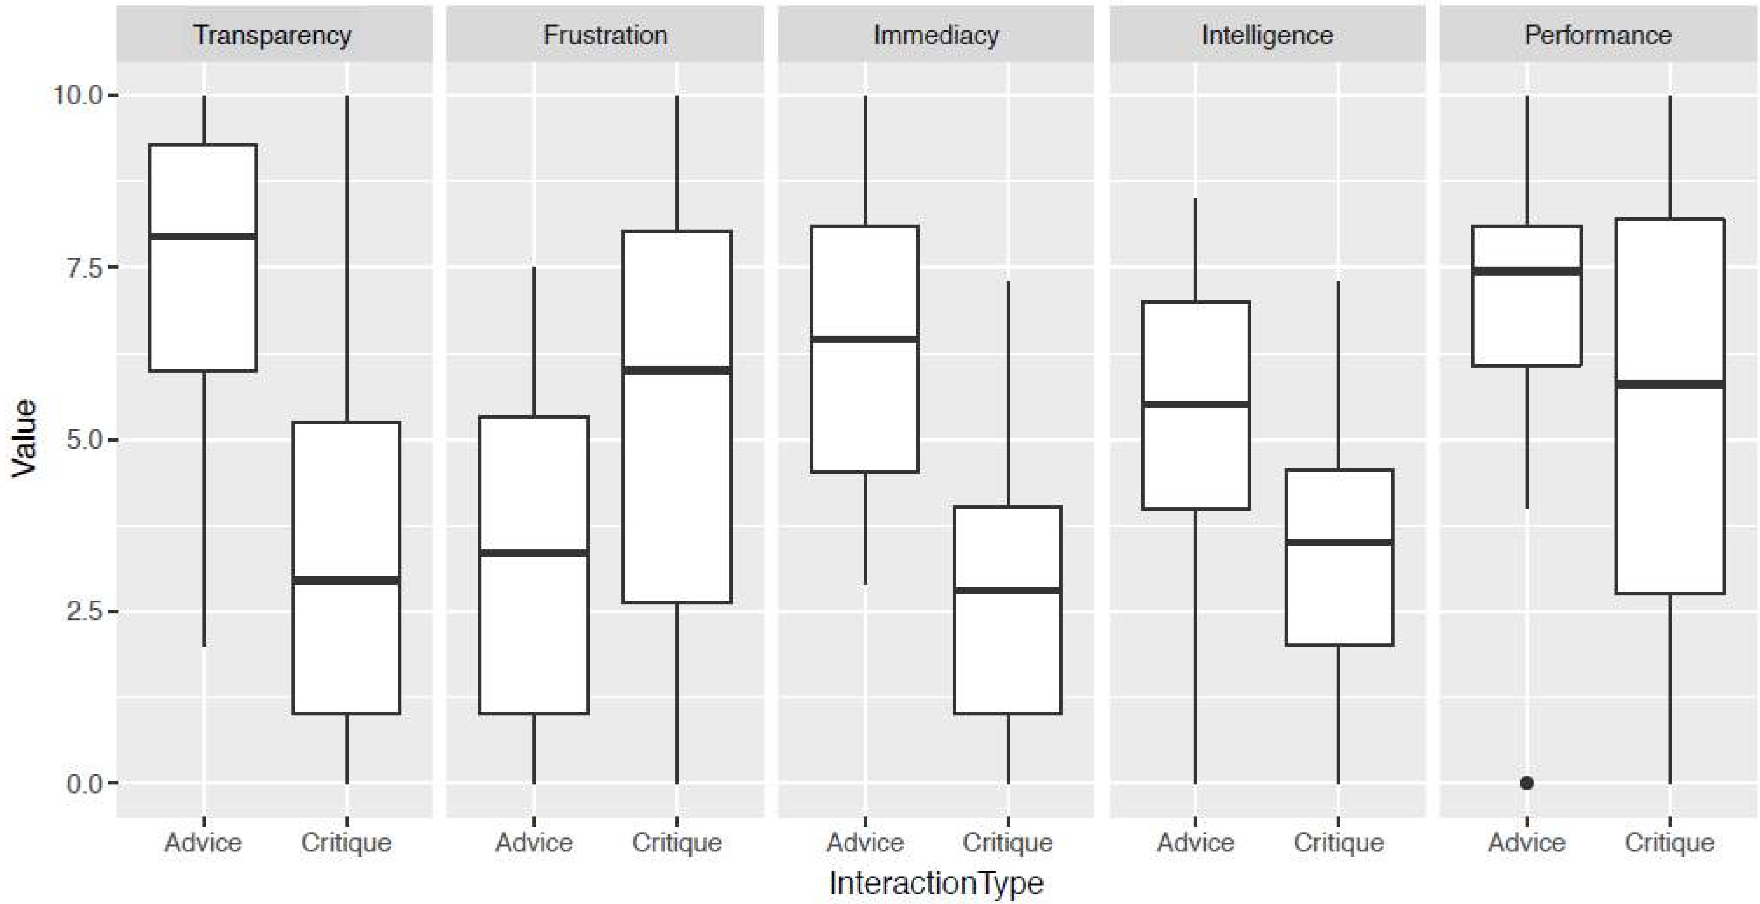

PERCEIVED INTELLIGENCE

Overall, the Action Advice agent was considered more intelligent than Critique 54% scored 3+

Main factors:

- *Compliance* with input: whether the agent did what it was told
- *Immediacy*: how quickly the agent learned
- *Effort*: the amount of input needed to train the agent

Explanations:

P22 *"The Action Advice was significantly more intelligent then the Critique. It followed my comments and completed the task multiple times."*
P11 *"I felt that the action advice agent was more intelligent because it seemed to learn faster and recover from mistakes faster."*
P3 *"The Advice agent responded with the correct results and was able to perform the tasks with minimal effort."*

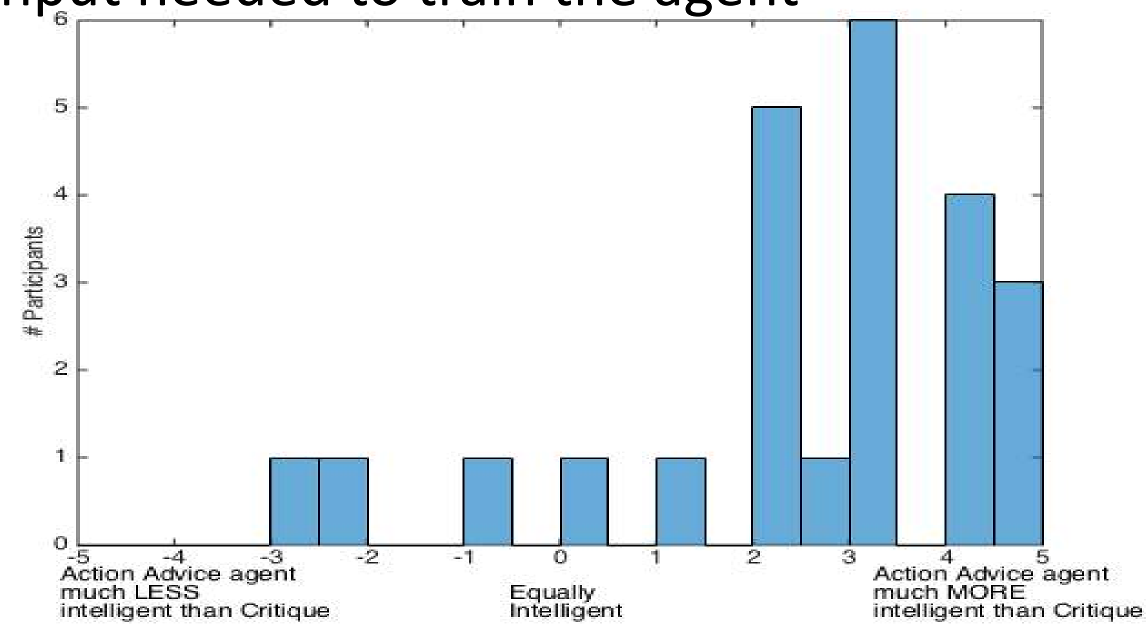

FRUSTRATION

Overall, the Action Advice agent was considered less frustrating than Critique

Main factors:

- *Powerlessness:* whether the agent's behavior made the human operator feel powerless
- *Transparency:* whether the human understands why the agent made its choices
- *Complexity:* the complexity of allowed human instruction

Explanations:

P14 *"In the critique case, I felt powerless to direct future actions, especially to avoid the agent jumping into the radioactive pit."*

P15 *"I did not understand how the critique would use my inputs."*

P12 *"I wanted to give more complex advice to 'help' the Critique Agent."*

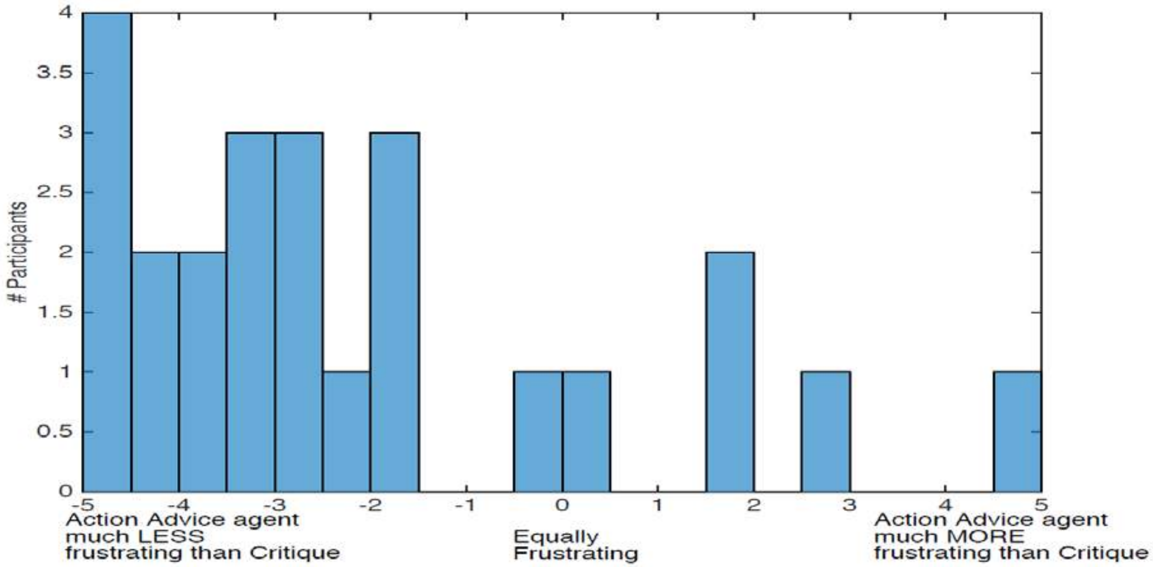

WHAT IMPACTED METRICS

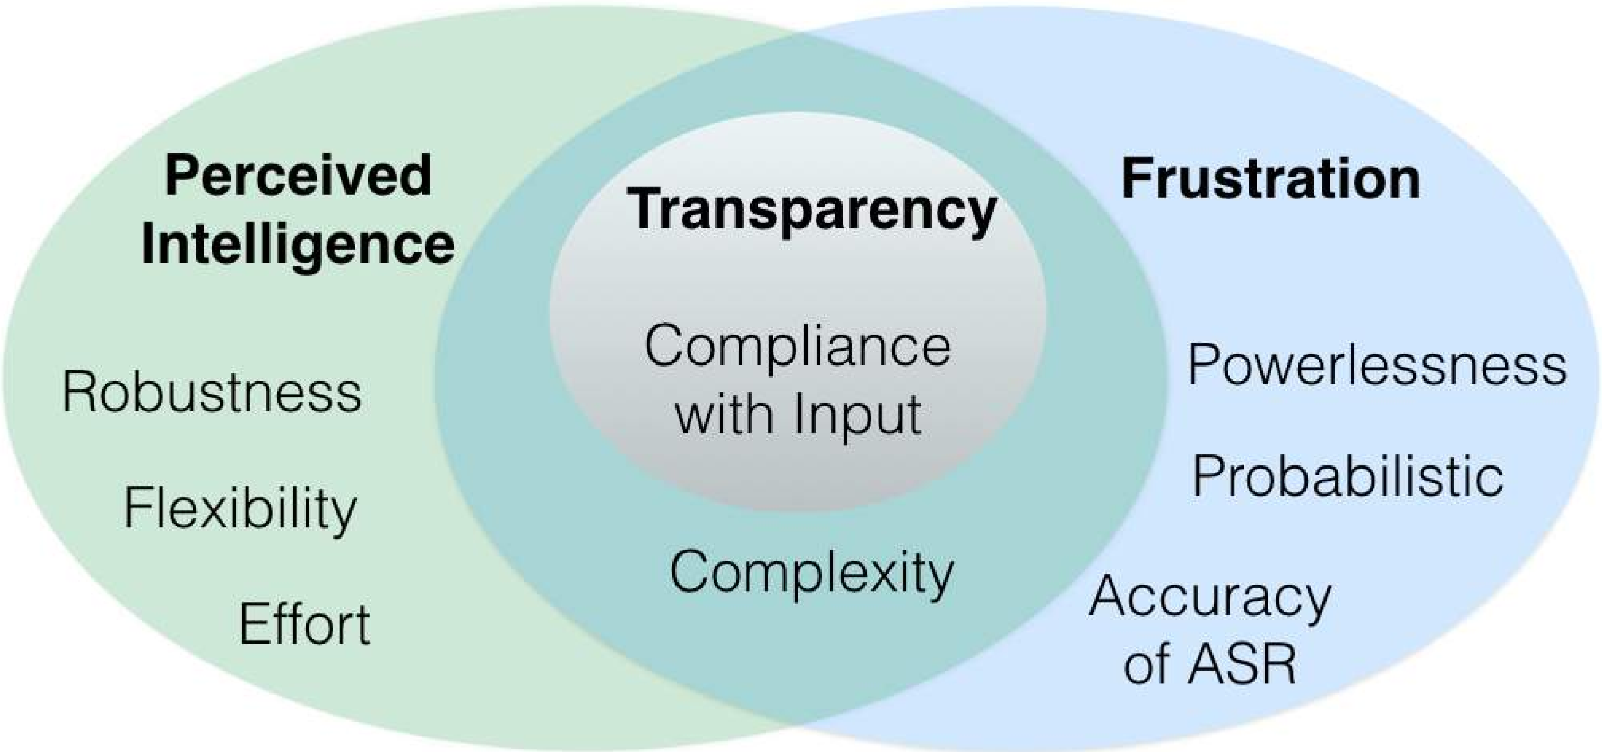

Second Study:

ADVICE VS CRITIQUE

What impacts human perception of ML algorithms?

- Our initial study indicated that a few specific characteristics of ML algorithms might impact human perception.

- We conducted an additional study to try to understand what elements of the algorithm impacted this perception and what specific elements.

Design Considerations

Design Consideration	Reason
Instructions about future, not past	Increases perceived control, transparency, immediacy, rhetoric (action advice, not critique)
Compliance with Input	Decreases frustration and increases perceived intelligence and performance
Empowerment	Clearly, immediately, and consistently follow the human's instructions. Decreases frustration.
Transparency	Immediately comply with instructions. Decreases frustration, increases perceived intelligence.
Immediacy	fication.
Deterministic Interaction	, manner. Increase
Complexity	will decrease
ASR accuracy	l processing time to
Robustness & Flexibility	es improves
Generalization through time	

In a follow-up experiment, we tested how 3 of these design considerations impact the user experience.

FOUR TYPES OF ALGORITHMS:

STANDARD – **SINGLE STEP**

Advice was followed for one time step. Similar to learning from demonstration collecting state-action pairs.

VARIATION: **GENERALIZATION OVER TIME**

When a human provided advice, the agent follows advice for 5 time steps.

VARIATION: **TIME DELAY**

This variation introduced a delay of 2 seconds between when advice was given and executed. Advice was followed for 5 time steps.

VARIATION: **PROBABILISTIC**

When a human provided advice, the agent chose whether to follow advice based on a probability for 5 time steps. Similar to policy shaping.

All algorithms were variants of Q learning.

Procedure

- Participants trained four agents that have the same underlying ML algorithm (Q learning) but small differences in the design of the interaction.
- For each agent
  - Participant is given instructions
  - Trains the agent until satisfied or decide to quit
    - Often ~ 4 minutes and 2-10 episodes
  - Answers questions about their experience
- Training is based on verbal instructions
  - left, right, up, down
- 24 participants with no prior ML experience
- Order of agents trained was balanced

Metrics

- **Frustration**
  - Degree of frustration participant felt training the agent.
- **Immediacy**
  - Degree to which the agent follow advice as fast as desired
- **Perceived Intelligence**
  - How smart the participants felt the algorithm was
- **Perceived Performance**
  - How well the participants felt the algorithm learned
- **Transparency**
  - How well the participants feel they understood what the agent was doing

HUMAN EXPERIENCE RATINGS

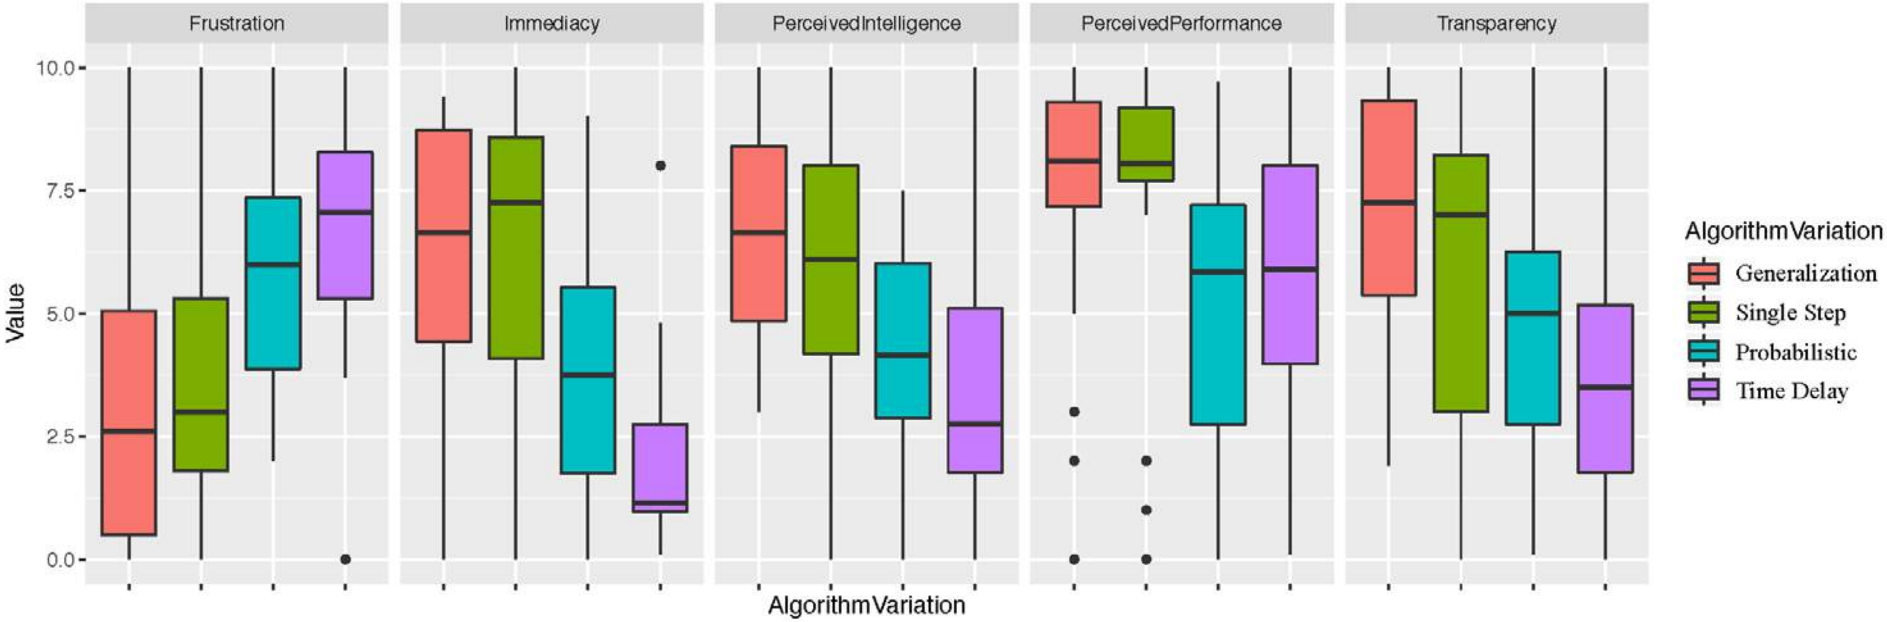

Overall, the baseline Generalization agent created the best human experience.

The Time Delay variation was the worst in terms of immediacy, transparency, and perceived intelligence.

CLOSER LOOK AT FRUSTRATION

Participants found the Generalization agent to be less frustrating than any of the variations

The variation with a time delay between when advice is given and used was the most frustrating

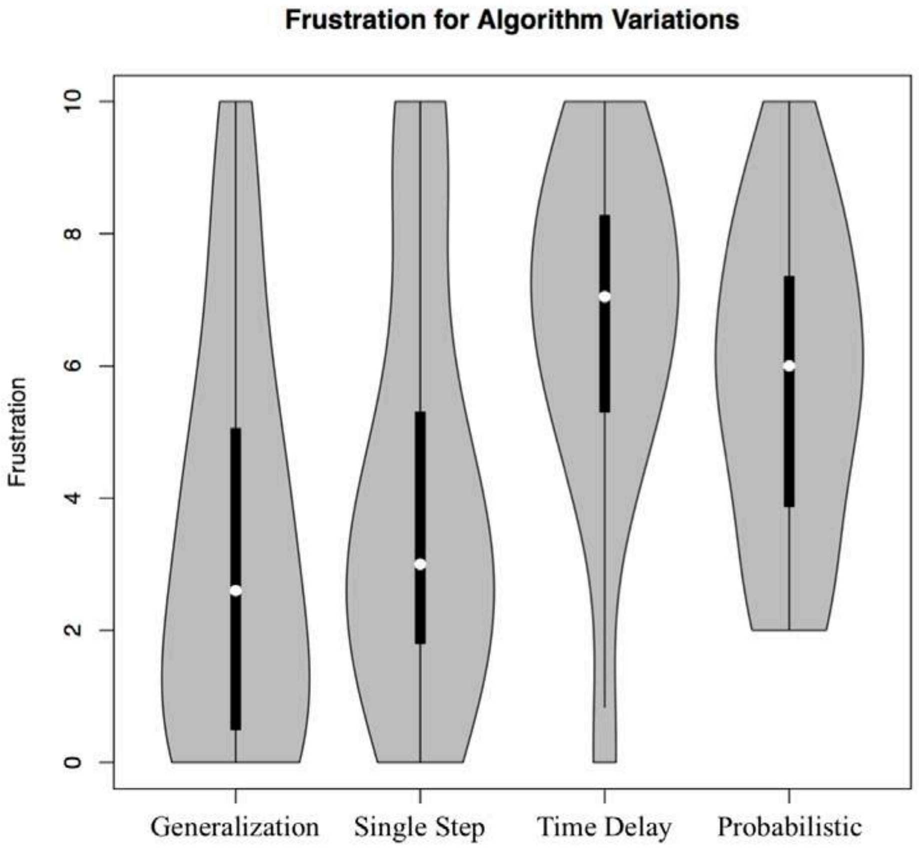

PERFORMANCE

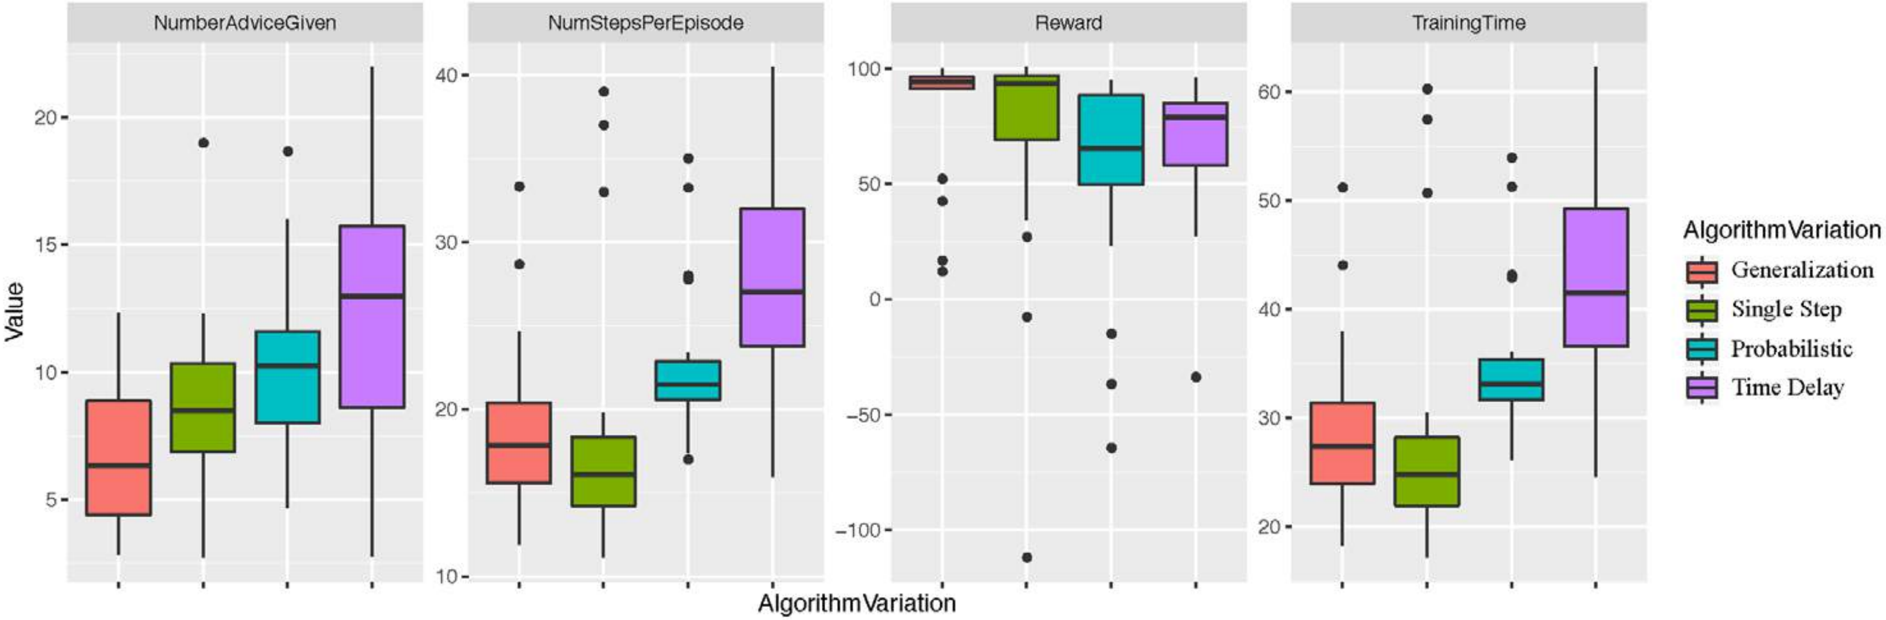

The Generalization agent was able to earn a higher reward in less time while using less information from participants than the Probabilistic or Time Delay variations.

People perform worse using the algorithm variation they like the least.

Take Aways

What makes human teachers like ML agents:

- **Compliance with input:** whether the agent did what it was told
- **Responsiveness:** how quickly the agent learned
- **Effort:** the amount of input needed to train the agent

- **Complexity:** the complexity of allowed human instruction
- **Transparency:** whether the human understands why the agent made its choices
- **Robustness and Flexibility:** the agent's ability to correct mistakes and learn alternate policies

Future Directions

- See how these results generalize to more complex domain spaces
  - Where teacher does not have access to total state
  - Where the state changes over time

- Expand detailed investigation to alternate teaching methods
  - Critique – what aspects are well received which ones are poorly received
  - Others

RECENT PUBLICATIONS

1. **Characteristics that Influence Perceived Intelligence in AI Design.** Proceedings of the Human Factors and Ergonomics Society (HFES) Annual Meeting (To appear) 2018

2. **Interaction Algorithm Effect on Human Experience.** *ACM Transactions on Human-Robot Interaction (THRI). Special Issue Journal on Artificial Intelligence for Human-Robot Interaction*. (Accepted) 2018

3. **Newtonian Action Advice: Integrating Human Verbal Instruction with Reinforcement Learning.** On ArXiv. (In preparation for submission to AAAI.)

4. **Shifting Role for Human Factors in an "Unmanned" Era.** *Journal of Theoretical Issues in Ergonomics Science (TTIE) 2017*

QUESTIONS?