Active Comparison Based Learning Incorporating User Uncertainty and Noise

Rachel Holladay
Shervin Javdani
Anca Dragan
Siddhartha Srinivasa

Carnegie Mellon University

Berkeley
UNIVERSITY OF CALIFORNIA

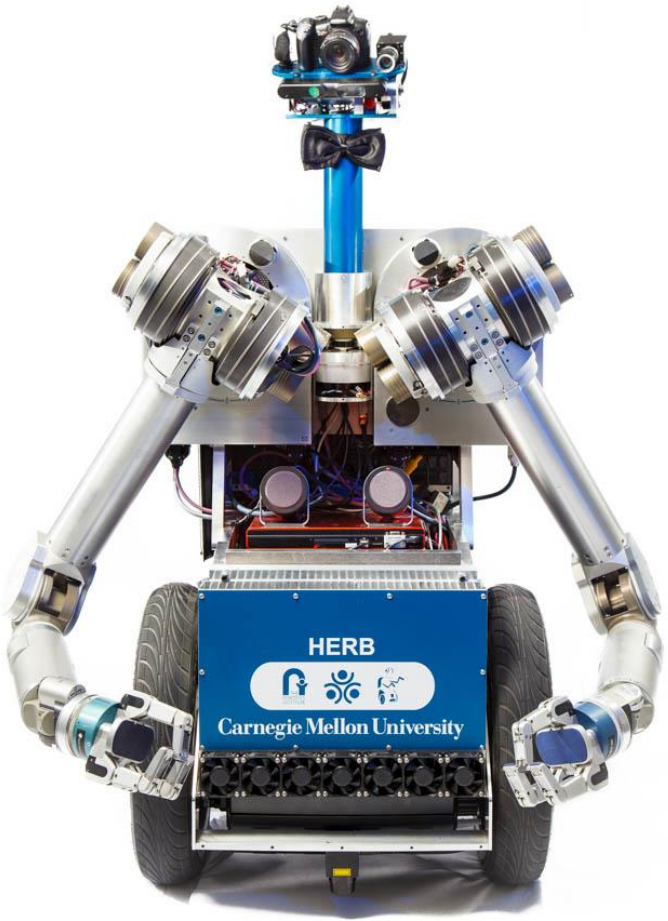

Personal Robotics Laboratory

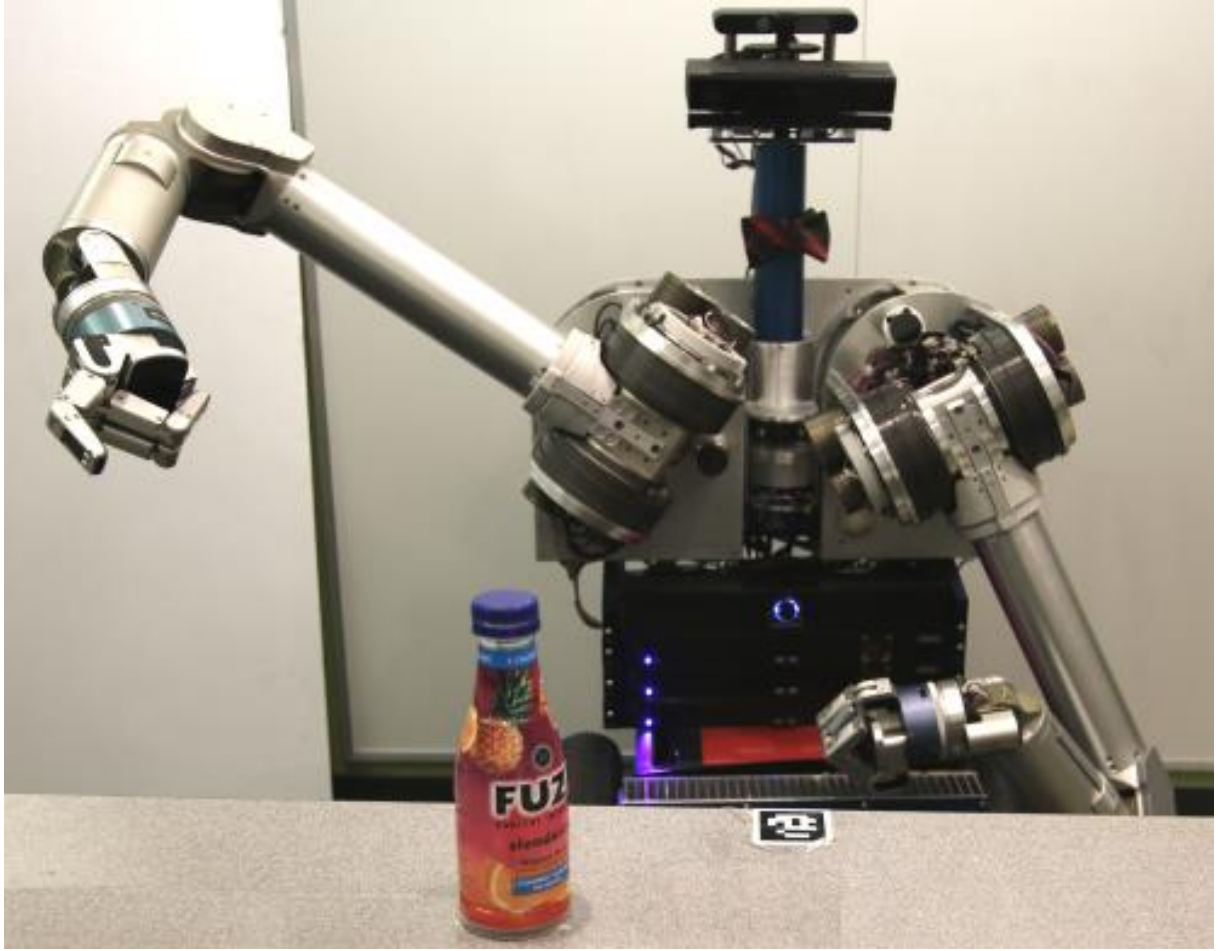

Holladay, Rachel M., Siddhartha S. Srinivasa. "RoGuE: Robot Gesture Engine." *AAAI Spring Symposium Series*. 2016.
Holladay, Rachel M., Anca D. Dragan, Siddhartha S. Srinivasa. "Legible robot pointing." *RO-MAN*. IEEE, 2014.

$$cost_{task}(\Theta)$$

$$\min_{\Theta} cost_{task}(\Theta)$$

$$cost_{task}(\Theta)$$

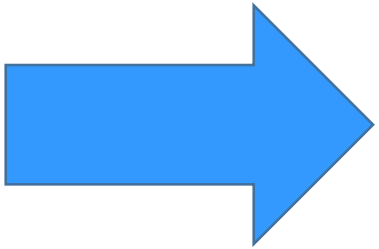

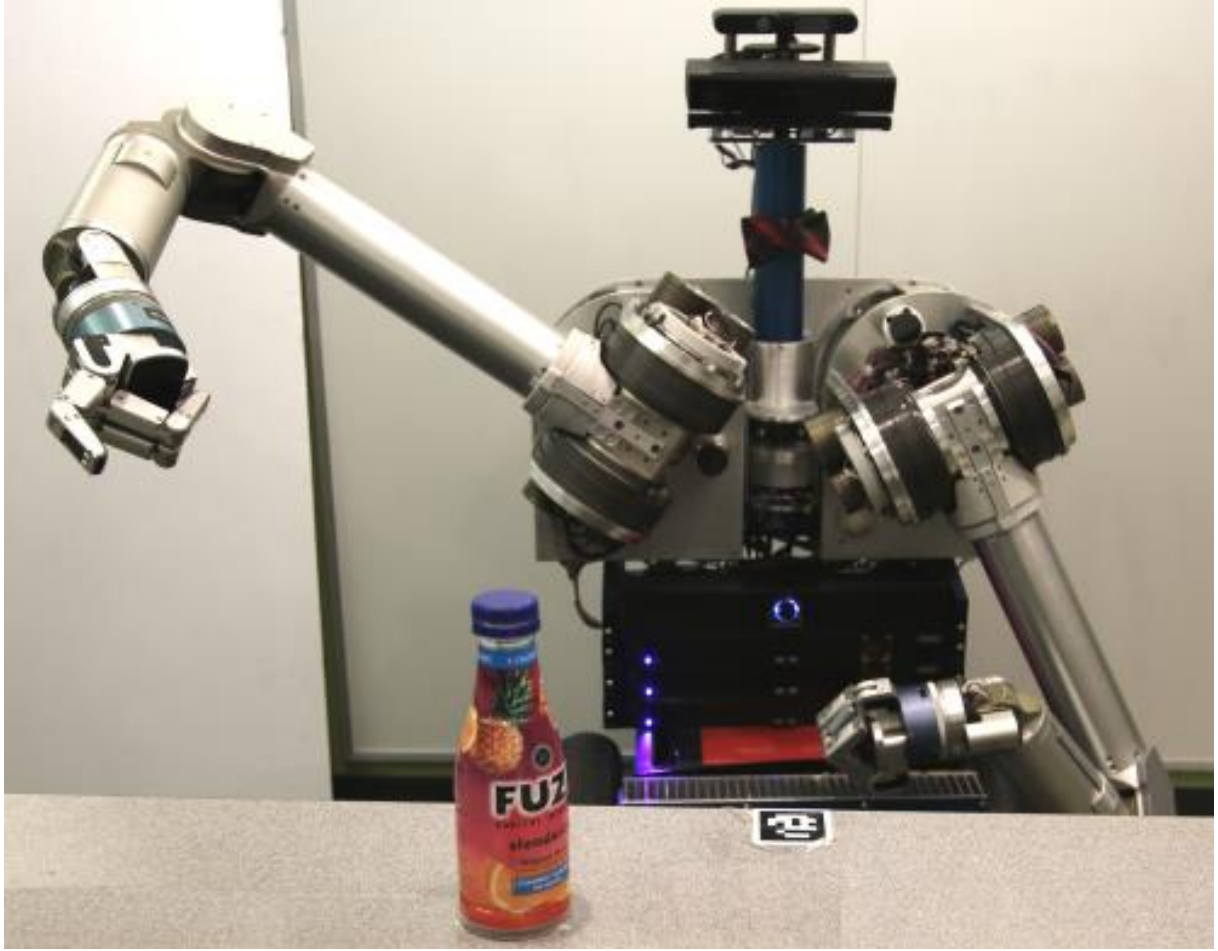

Holladay, Rachel M., Siddhartha S. Srinivasa. "RoGuE: Robot Gesture Engine." *AAAI Spring Symposium Series*. 2016.
Holladay, Rachel M., Anca D. Dragan, Siddhartha S. Srinivasa. "Legible robot pointing." *RO-MAN*. IEEE, 2014.

$cost_{task}(\Theta)$?

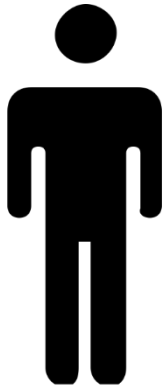

$cost_{task}(\Theta)$?

?

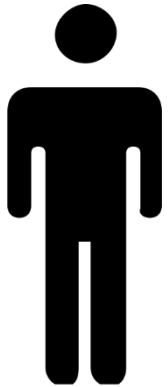

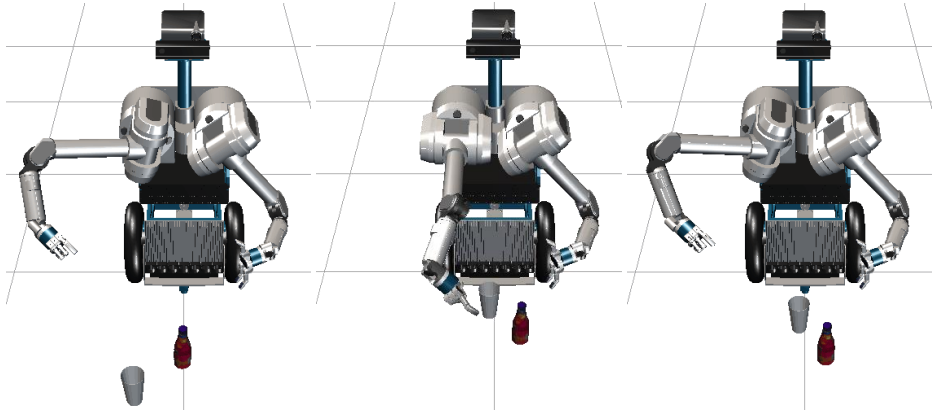

A B C

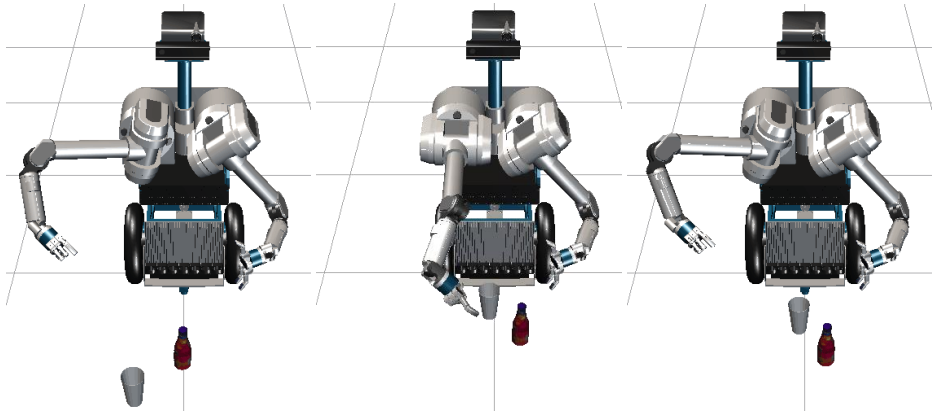

A B C

cost(A) cost(B) cost(C)

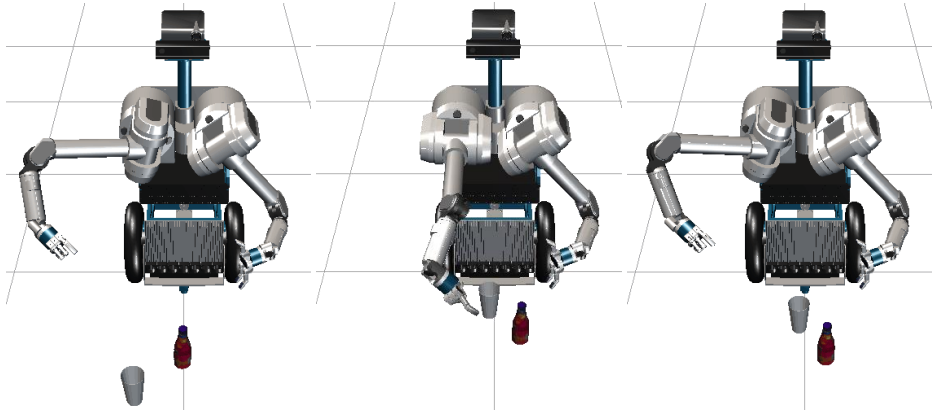

A B C

cost(A) cost(B) cost(C)

$$\text{cost}(C) < \text{cost}(A) < \text{cost}(B)$$

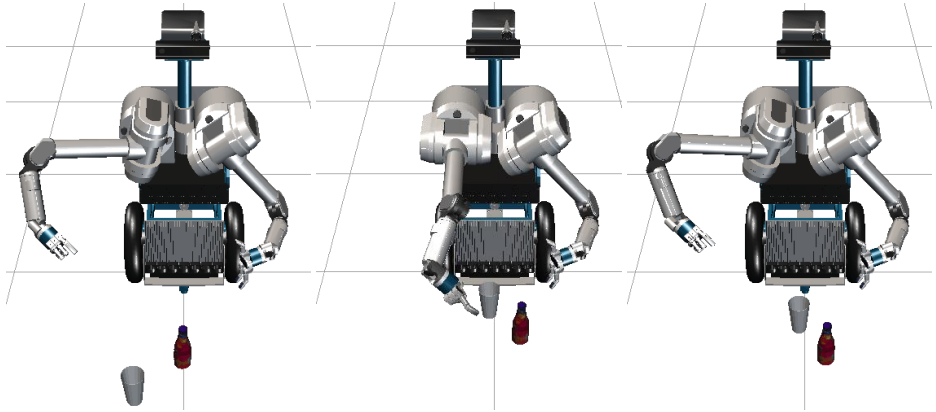

A B C

cost(A) cost(B) cost(C)

$$cost(C) < cost(A) < cost(B)$$

$$C < A < B$$

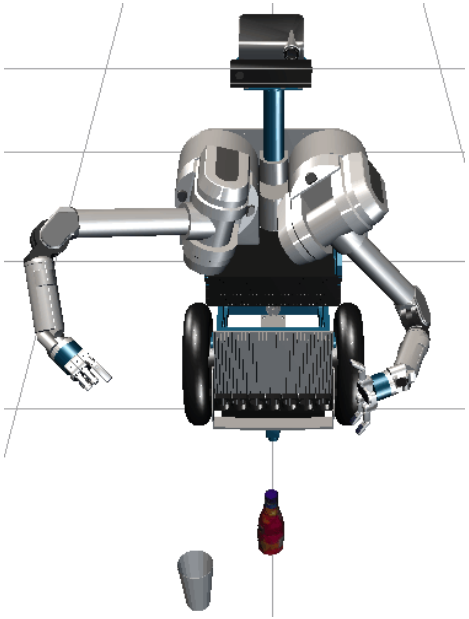

Left

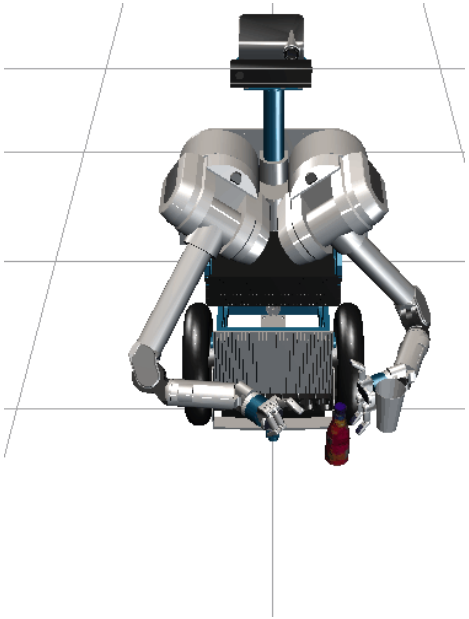

Right

Which do you prefer?

Comparison Based Learning

Ben Carterette, Paul N. Bennett, David Maxwell Chickering, and Susan T. Dumais. “Here or there: Preference judgments for relevance.” *ECIR*, 2008.

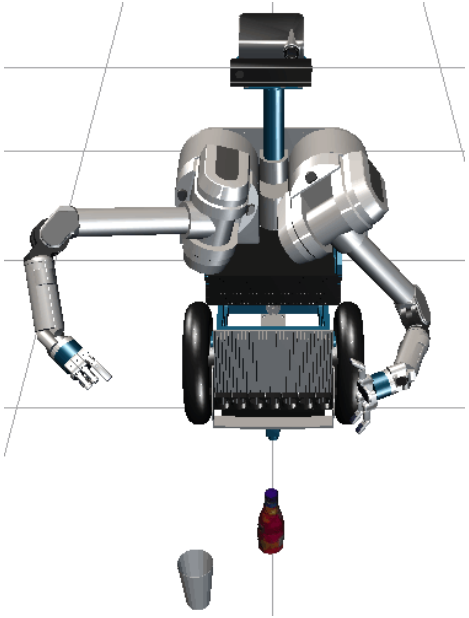

Left

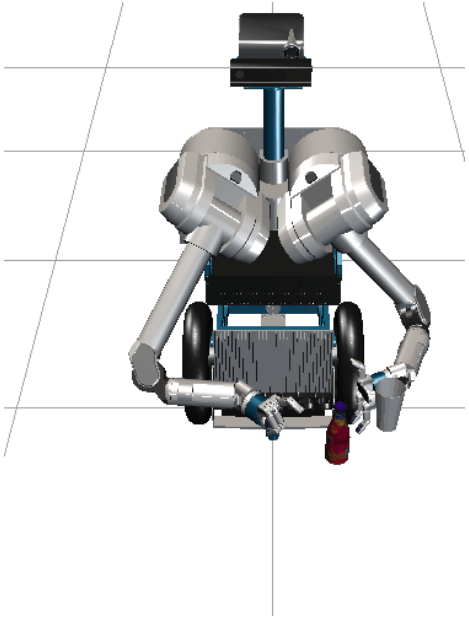

Right

Test

(Left, Right)

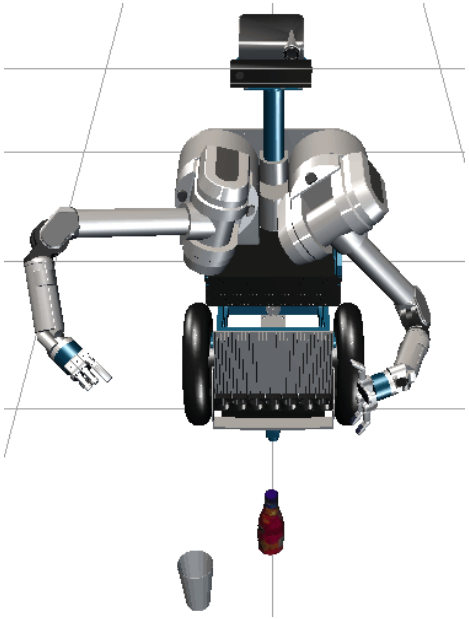

Left

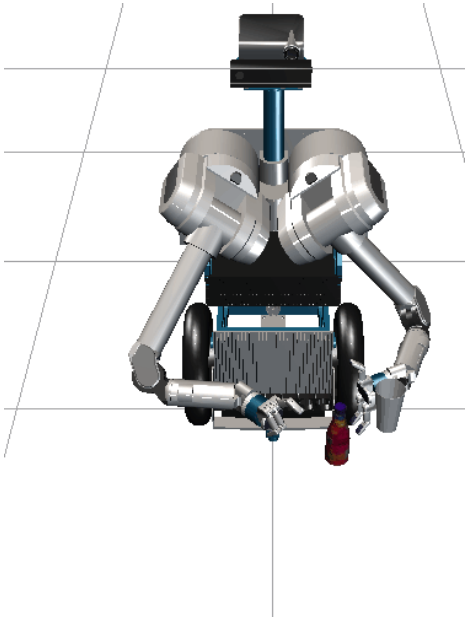

Right

Observation

Left > Right

Objective:
Learn their cost function with the
fewest number of tests to minimize user burden

Active Learning

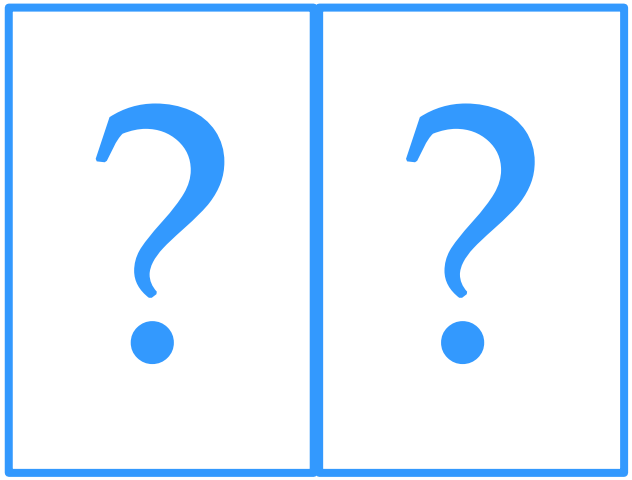

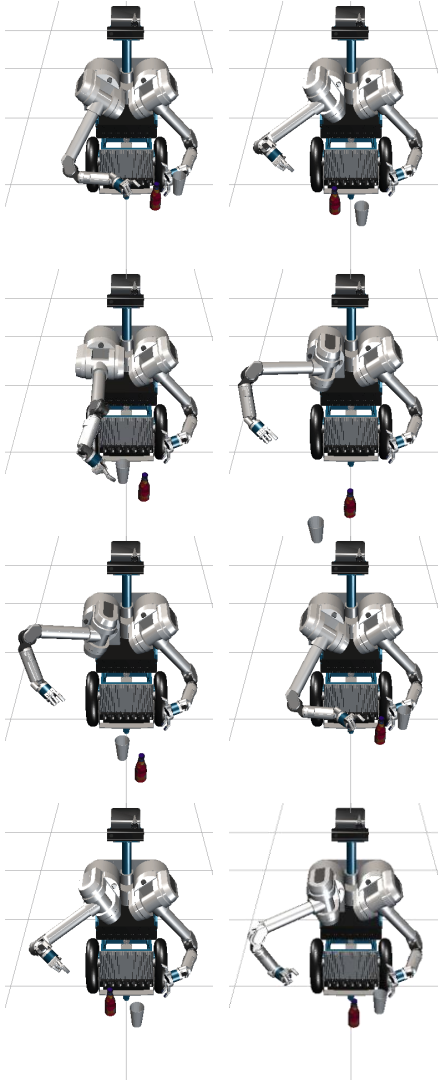

. . .

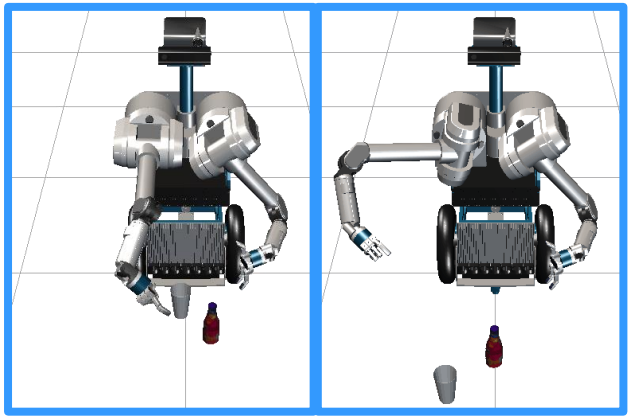

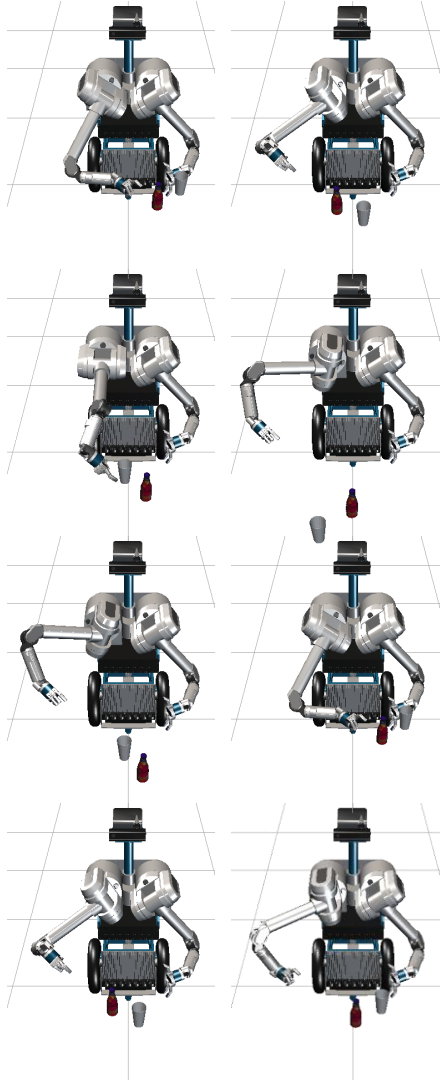

. . .

Select test that *minimizes* the *expected* number of *tests* needed

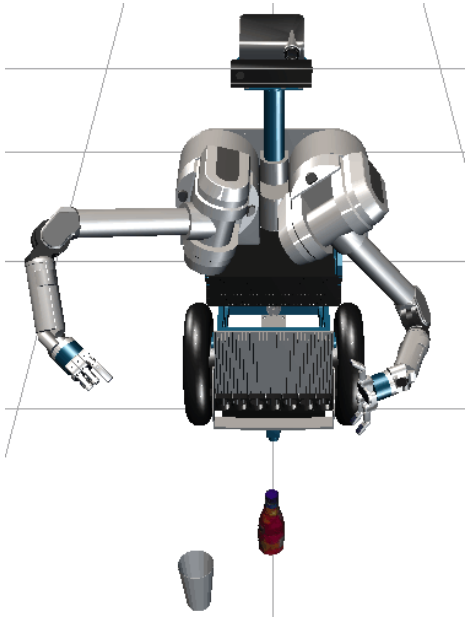

Left

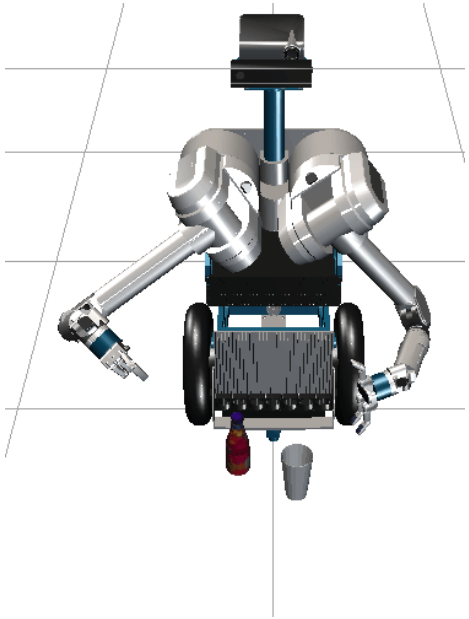

Right

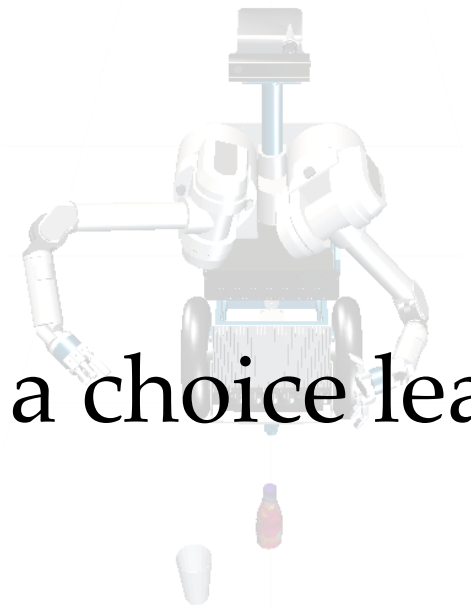

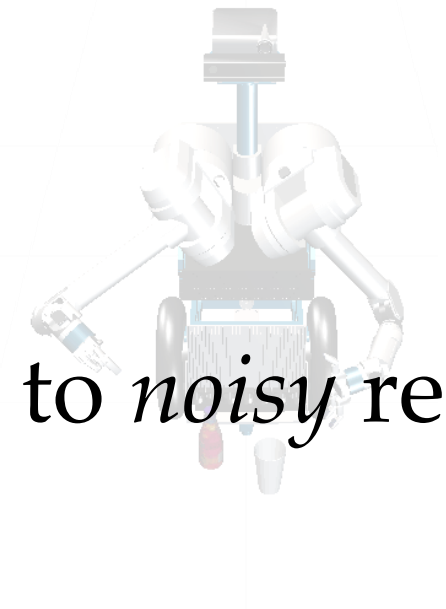

Forcing a choice leads to *noisy* responses

Left

Right

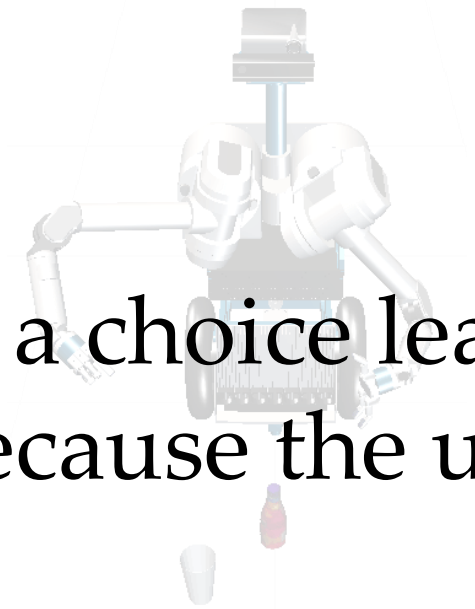

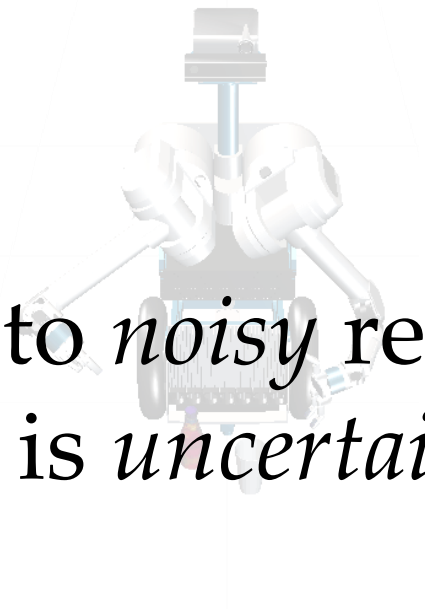

Forcing a choice leads to *noisy* responses because the user is *uncertain*.

Left

Right

Users *want* to express uncertainty.

Guillory, Andrew, Jeff A. Bilmes. "Simultaneous learning and covering with adversarial noise." *ICML-11*. 2011.

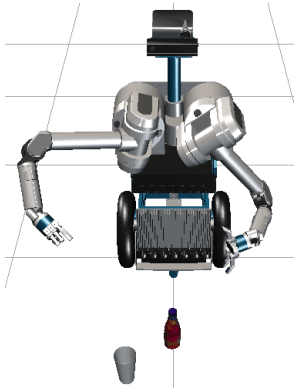

Left

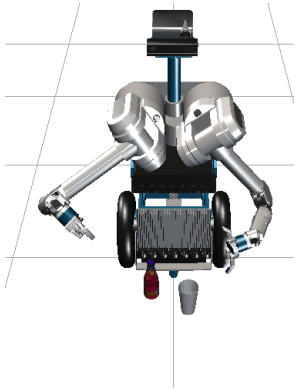

Right

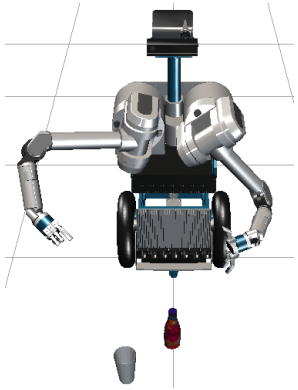

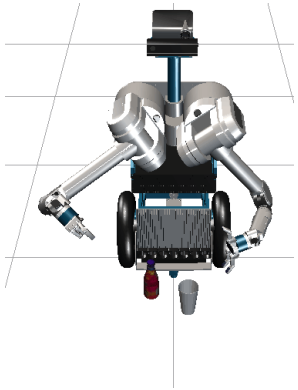

Left

Right

Forced Preference

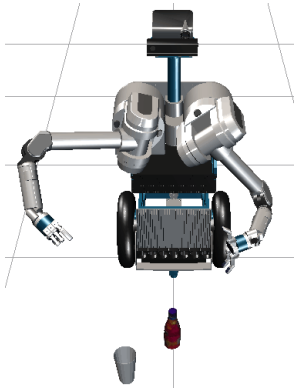

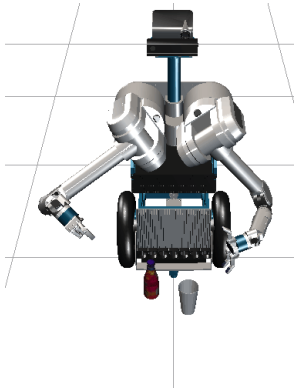

Left

Right

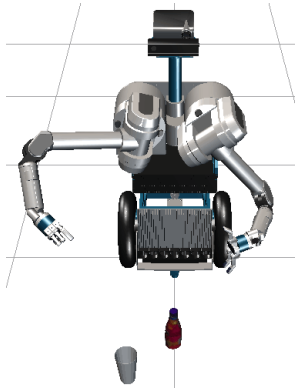

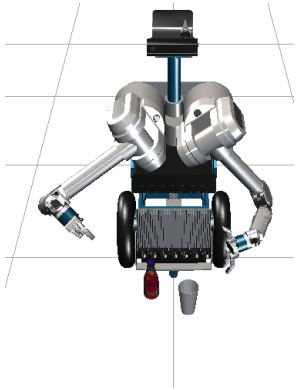

Left

?

Right

Forced Preference

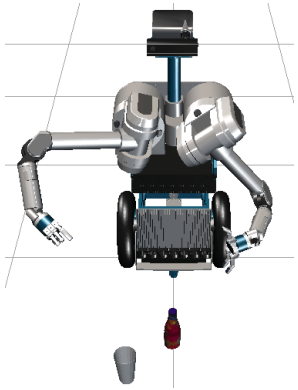

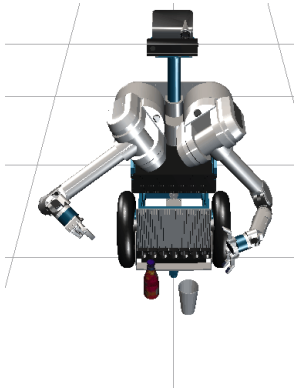

Left

Right

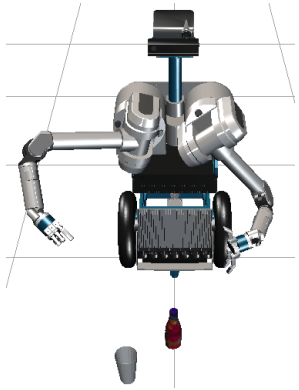

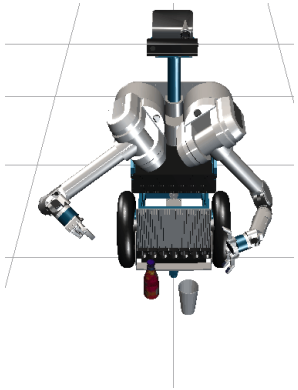

Left

?

Right

Forced Preference

Uncertain Preference

Key Insight

The user's uncertainty is informative about their cost function

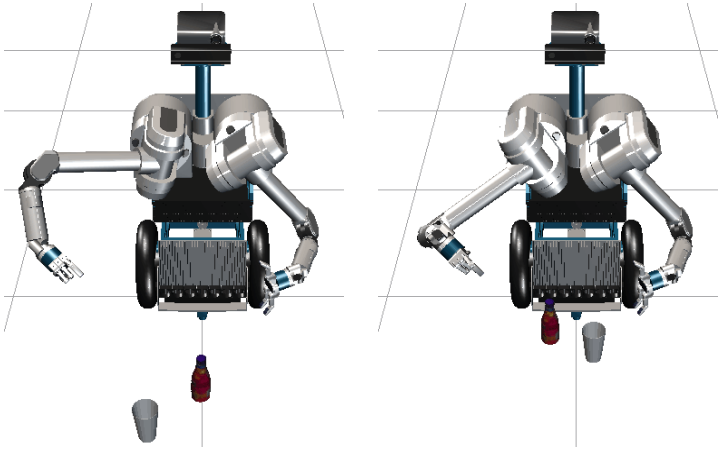

Left	?	Right

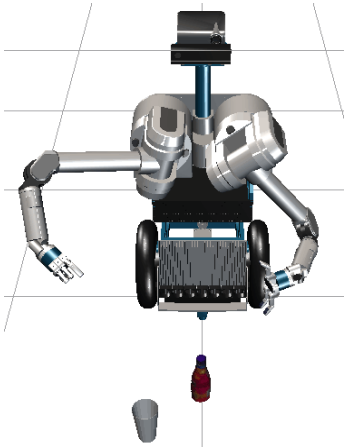

Cost(L)

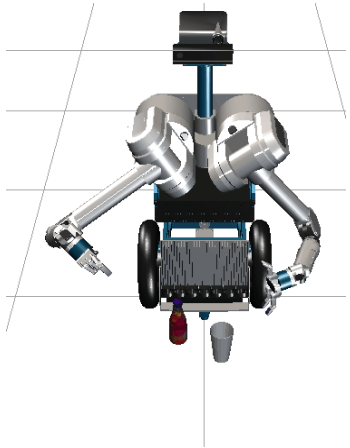

Cost(R)

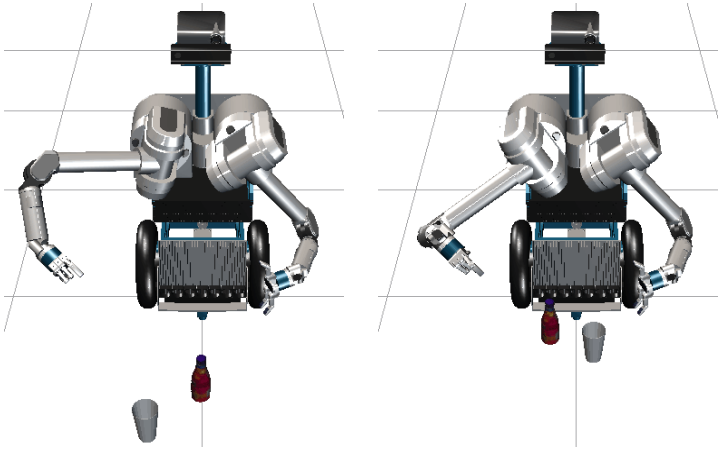

$$\text{Cost}(L) \approx \text{Cost}(R)$$

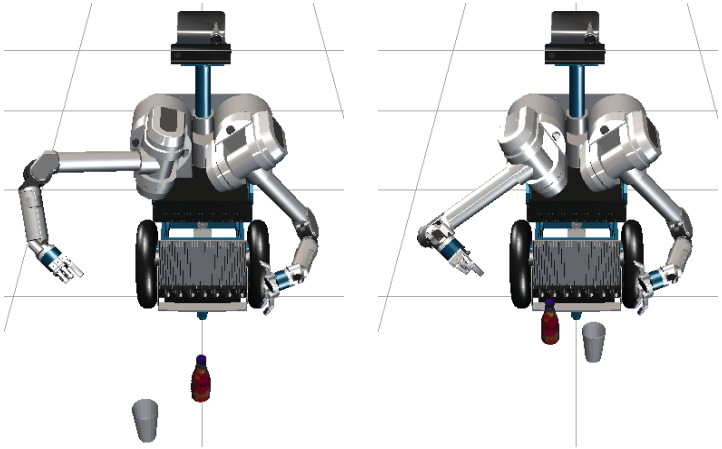

$$\text{Cost}(L) \approx \text{Cost}(R)$$

$$|\text{Cost}(L) - \text{Cost}(R)| < \varepsilon$$

Do not learn the user's ε.

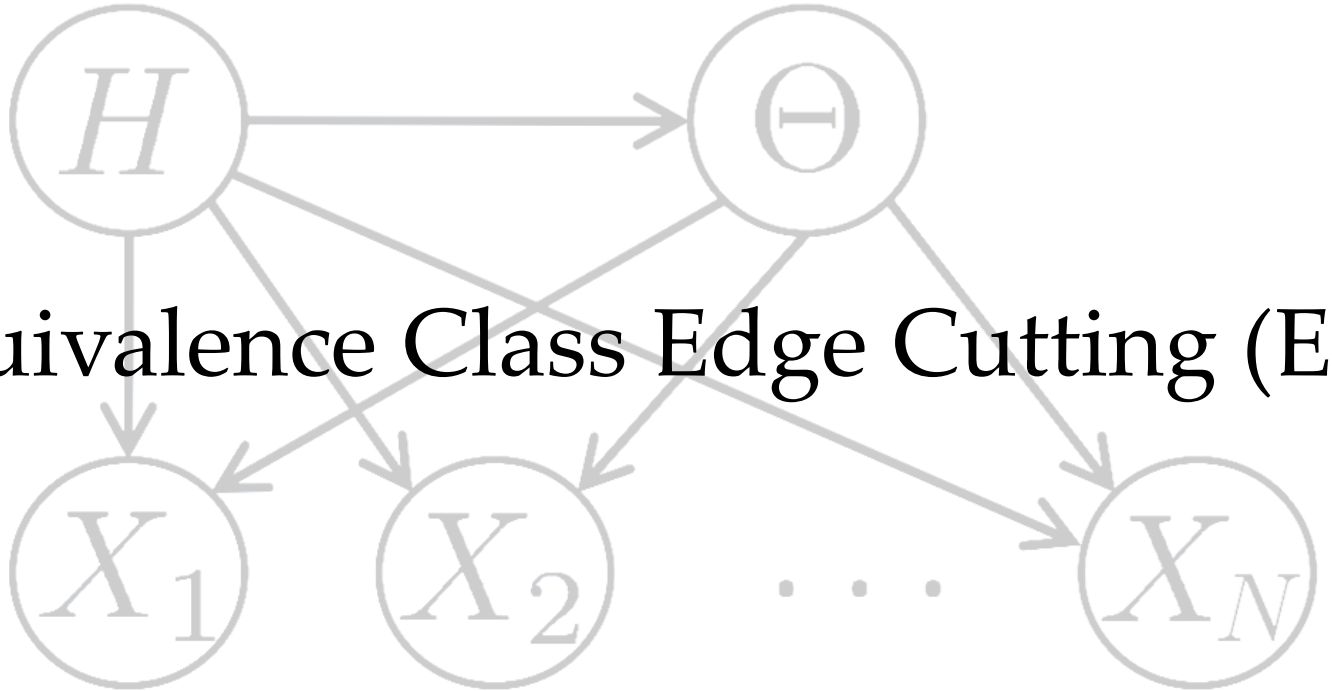

Equivalence Class Edge Cutting (EC^2)

Golovin, Daniel, Andreas Krause, and Debajyoti Ray. "Near-optimal bayesian active learning with noisy observations." *NIPS*. 2010.

Users can be *noisy.*

Saleema Amershi, Maya Cakmak, W. Bradley Knox, and Todd Kulesza. Power to the people: The role of humans in interactive machine learning. *AI Magazine*, 2014.

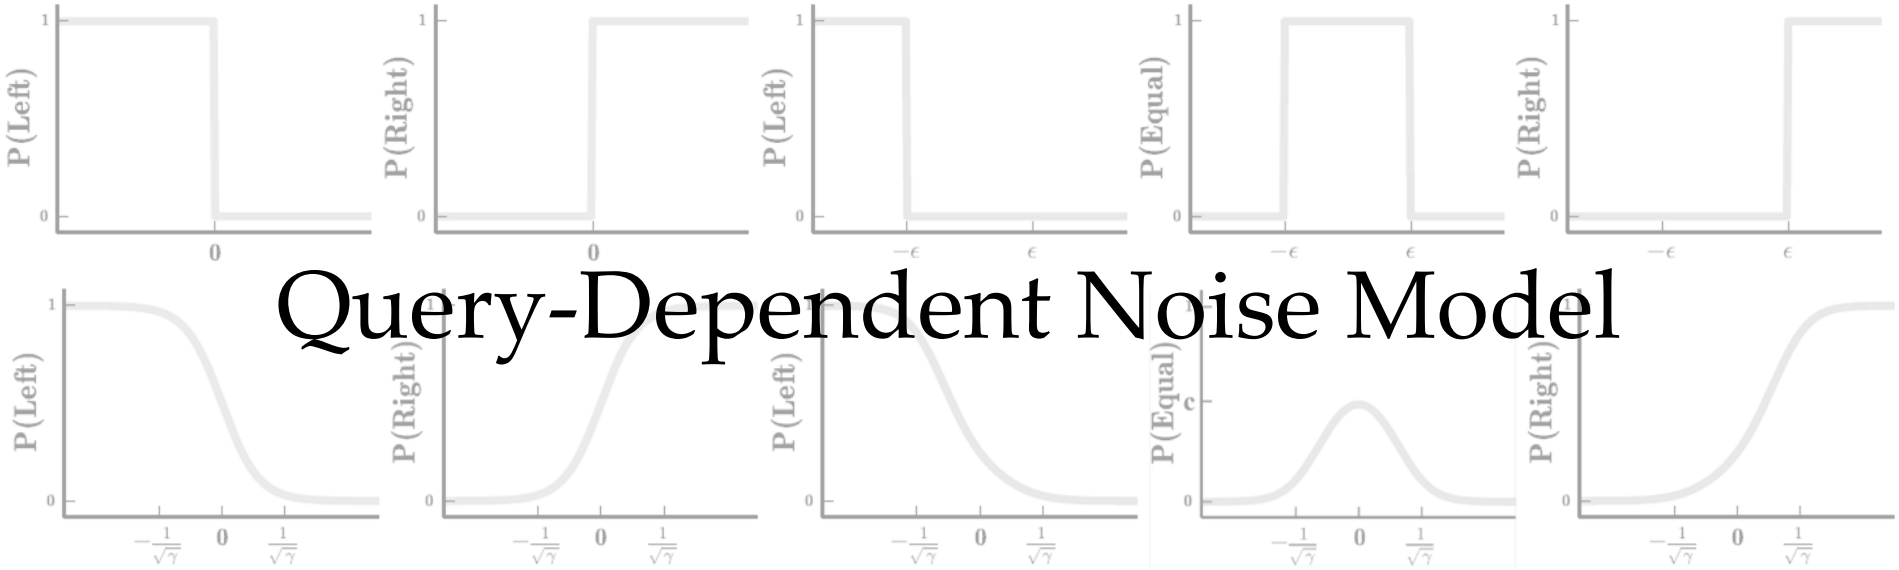

Query-Dependent Noise Model

Du, Jun, and Charles X. Ling. "Active learning with human-like noisy oracle." *ICDM*. IEEE, 2010.

CLAUS:
Comparison Learning
Algorithm for Uncertain Situations

User Evaluation

Noisy User Evaluation

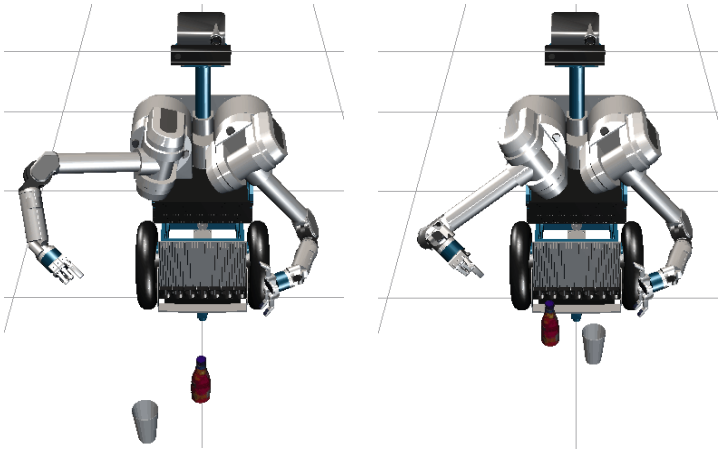

Known Cost Function

Which line is longer?

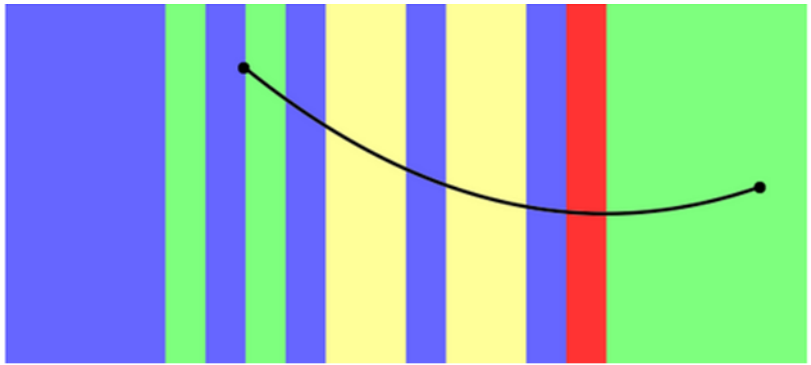

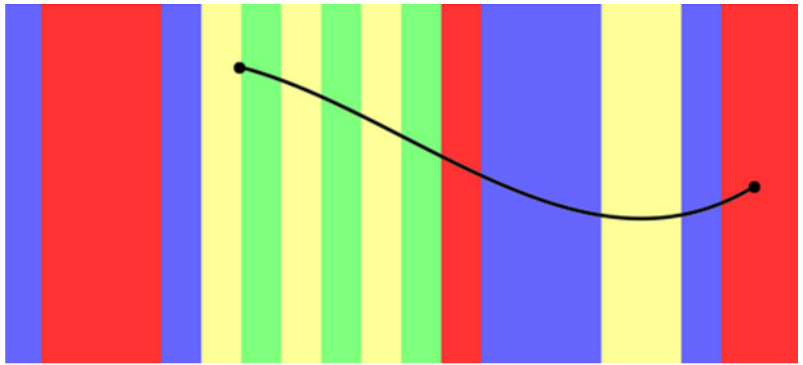

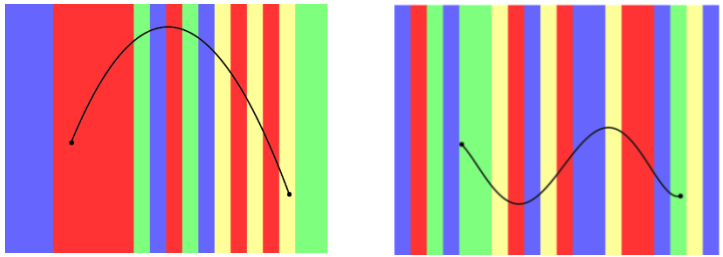

vs.

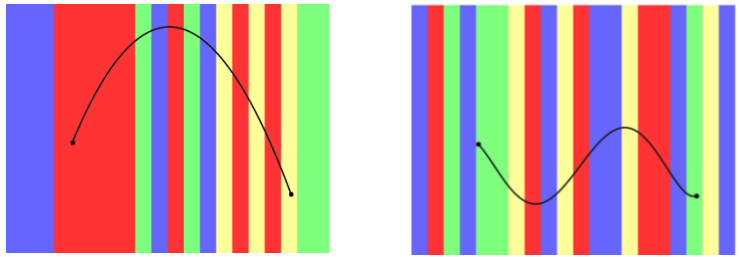

Left | Right

Left | ? | Right

Forced Preference

Uncertain Preference

How do you *label* uncertainty?

I am Not Sure

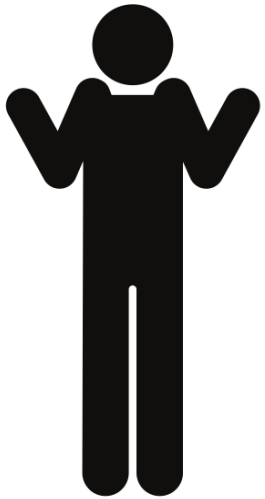

About Equal

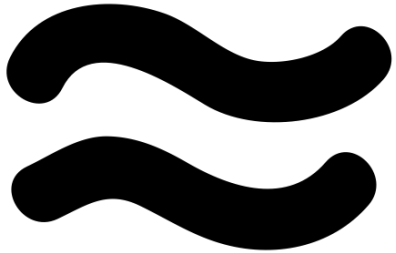

Source: Noun Project

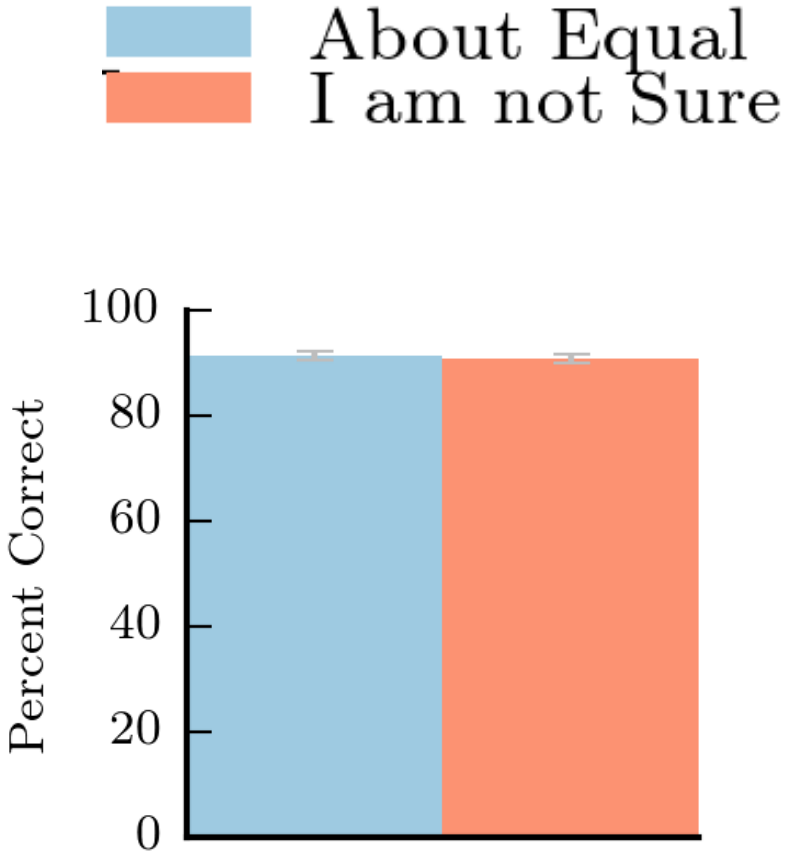
About Equal
I am not Sure
Percent Correct
100
80
60
40
20
0

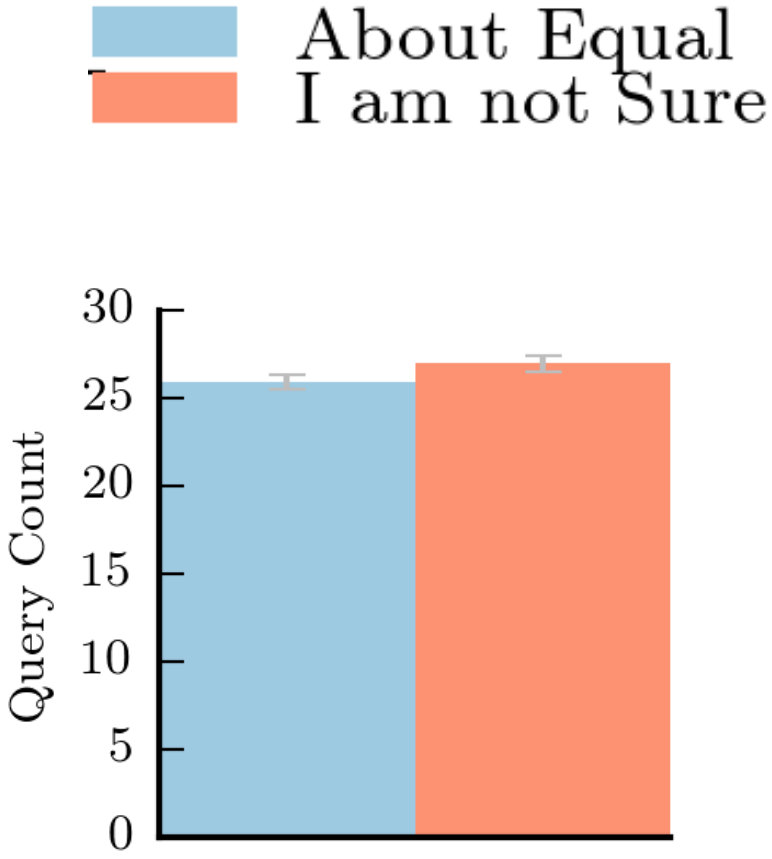
About Equal
I am not Sure
Query Count
30
25
20
15
10
5
0

I am Not Sure

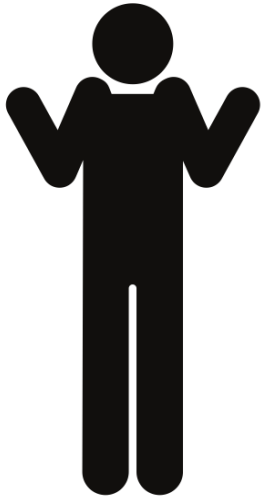

About Equal

Source: Noun Project

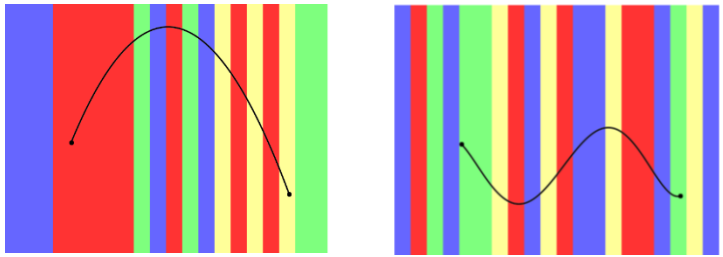

vs.

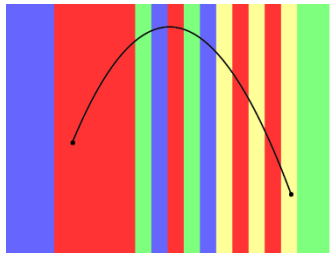

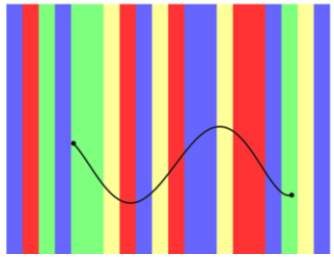

Left

Right

Left

About Equal

Right

Forced Preference

Uncertain Preference

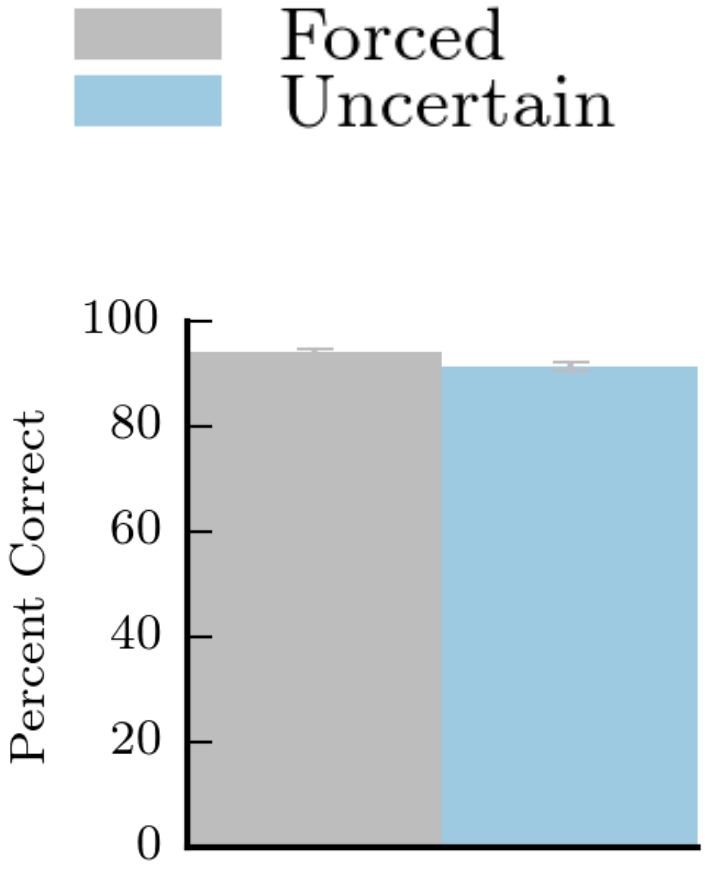
Forced
Uncertain
Percent Correct
100
80
60
40
20
0

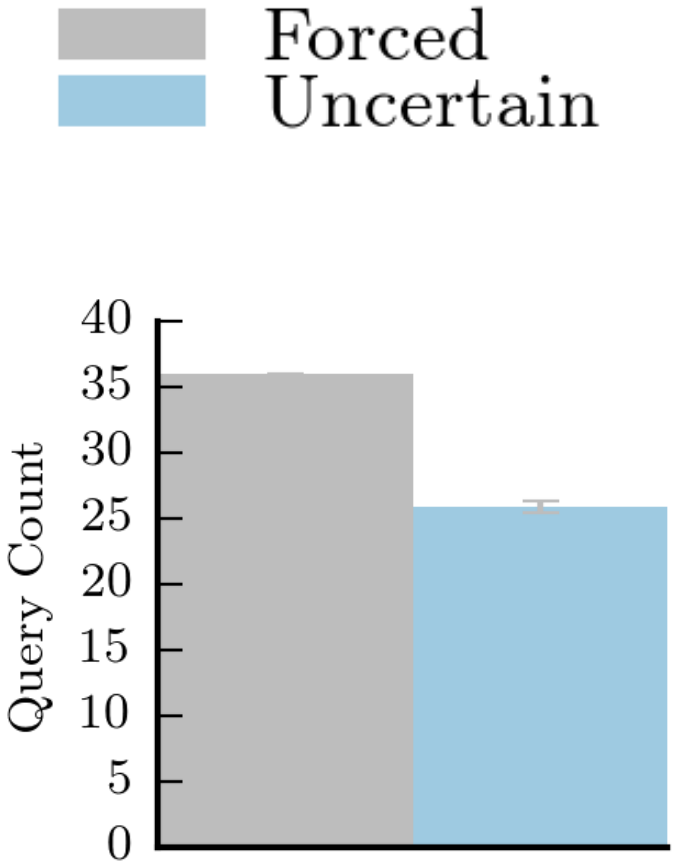
Forced
Uncertain
Query Count
40
35
30
25
20
15
10
5
0

Key Insight

The user's uncertainty is informative about their cost function

Active Comparison Based Learning Incorporating User Uncertainty and Noise

Rachel Holladay
Shervin Javdani
Anca Dragan
Siddhartha Srinivasa

www.personalrobotics.ri.cmu.edu