Unifying recommendation and active learning for human-algorithm interactions

Scott Cheng-Hsin Yang[1], Jake Alden Whritner[1],
Olfa Nasraoui[2] & Patrick Shafto[1]

1 Department of Mathematics & Computer Science, Rutgers University–Newark
2 Department of Computer Engineering and Computer Science, University of Louisville

CogSci 2017

21st century online shopping

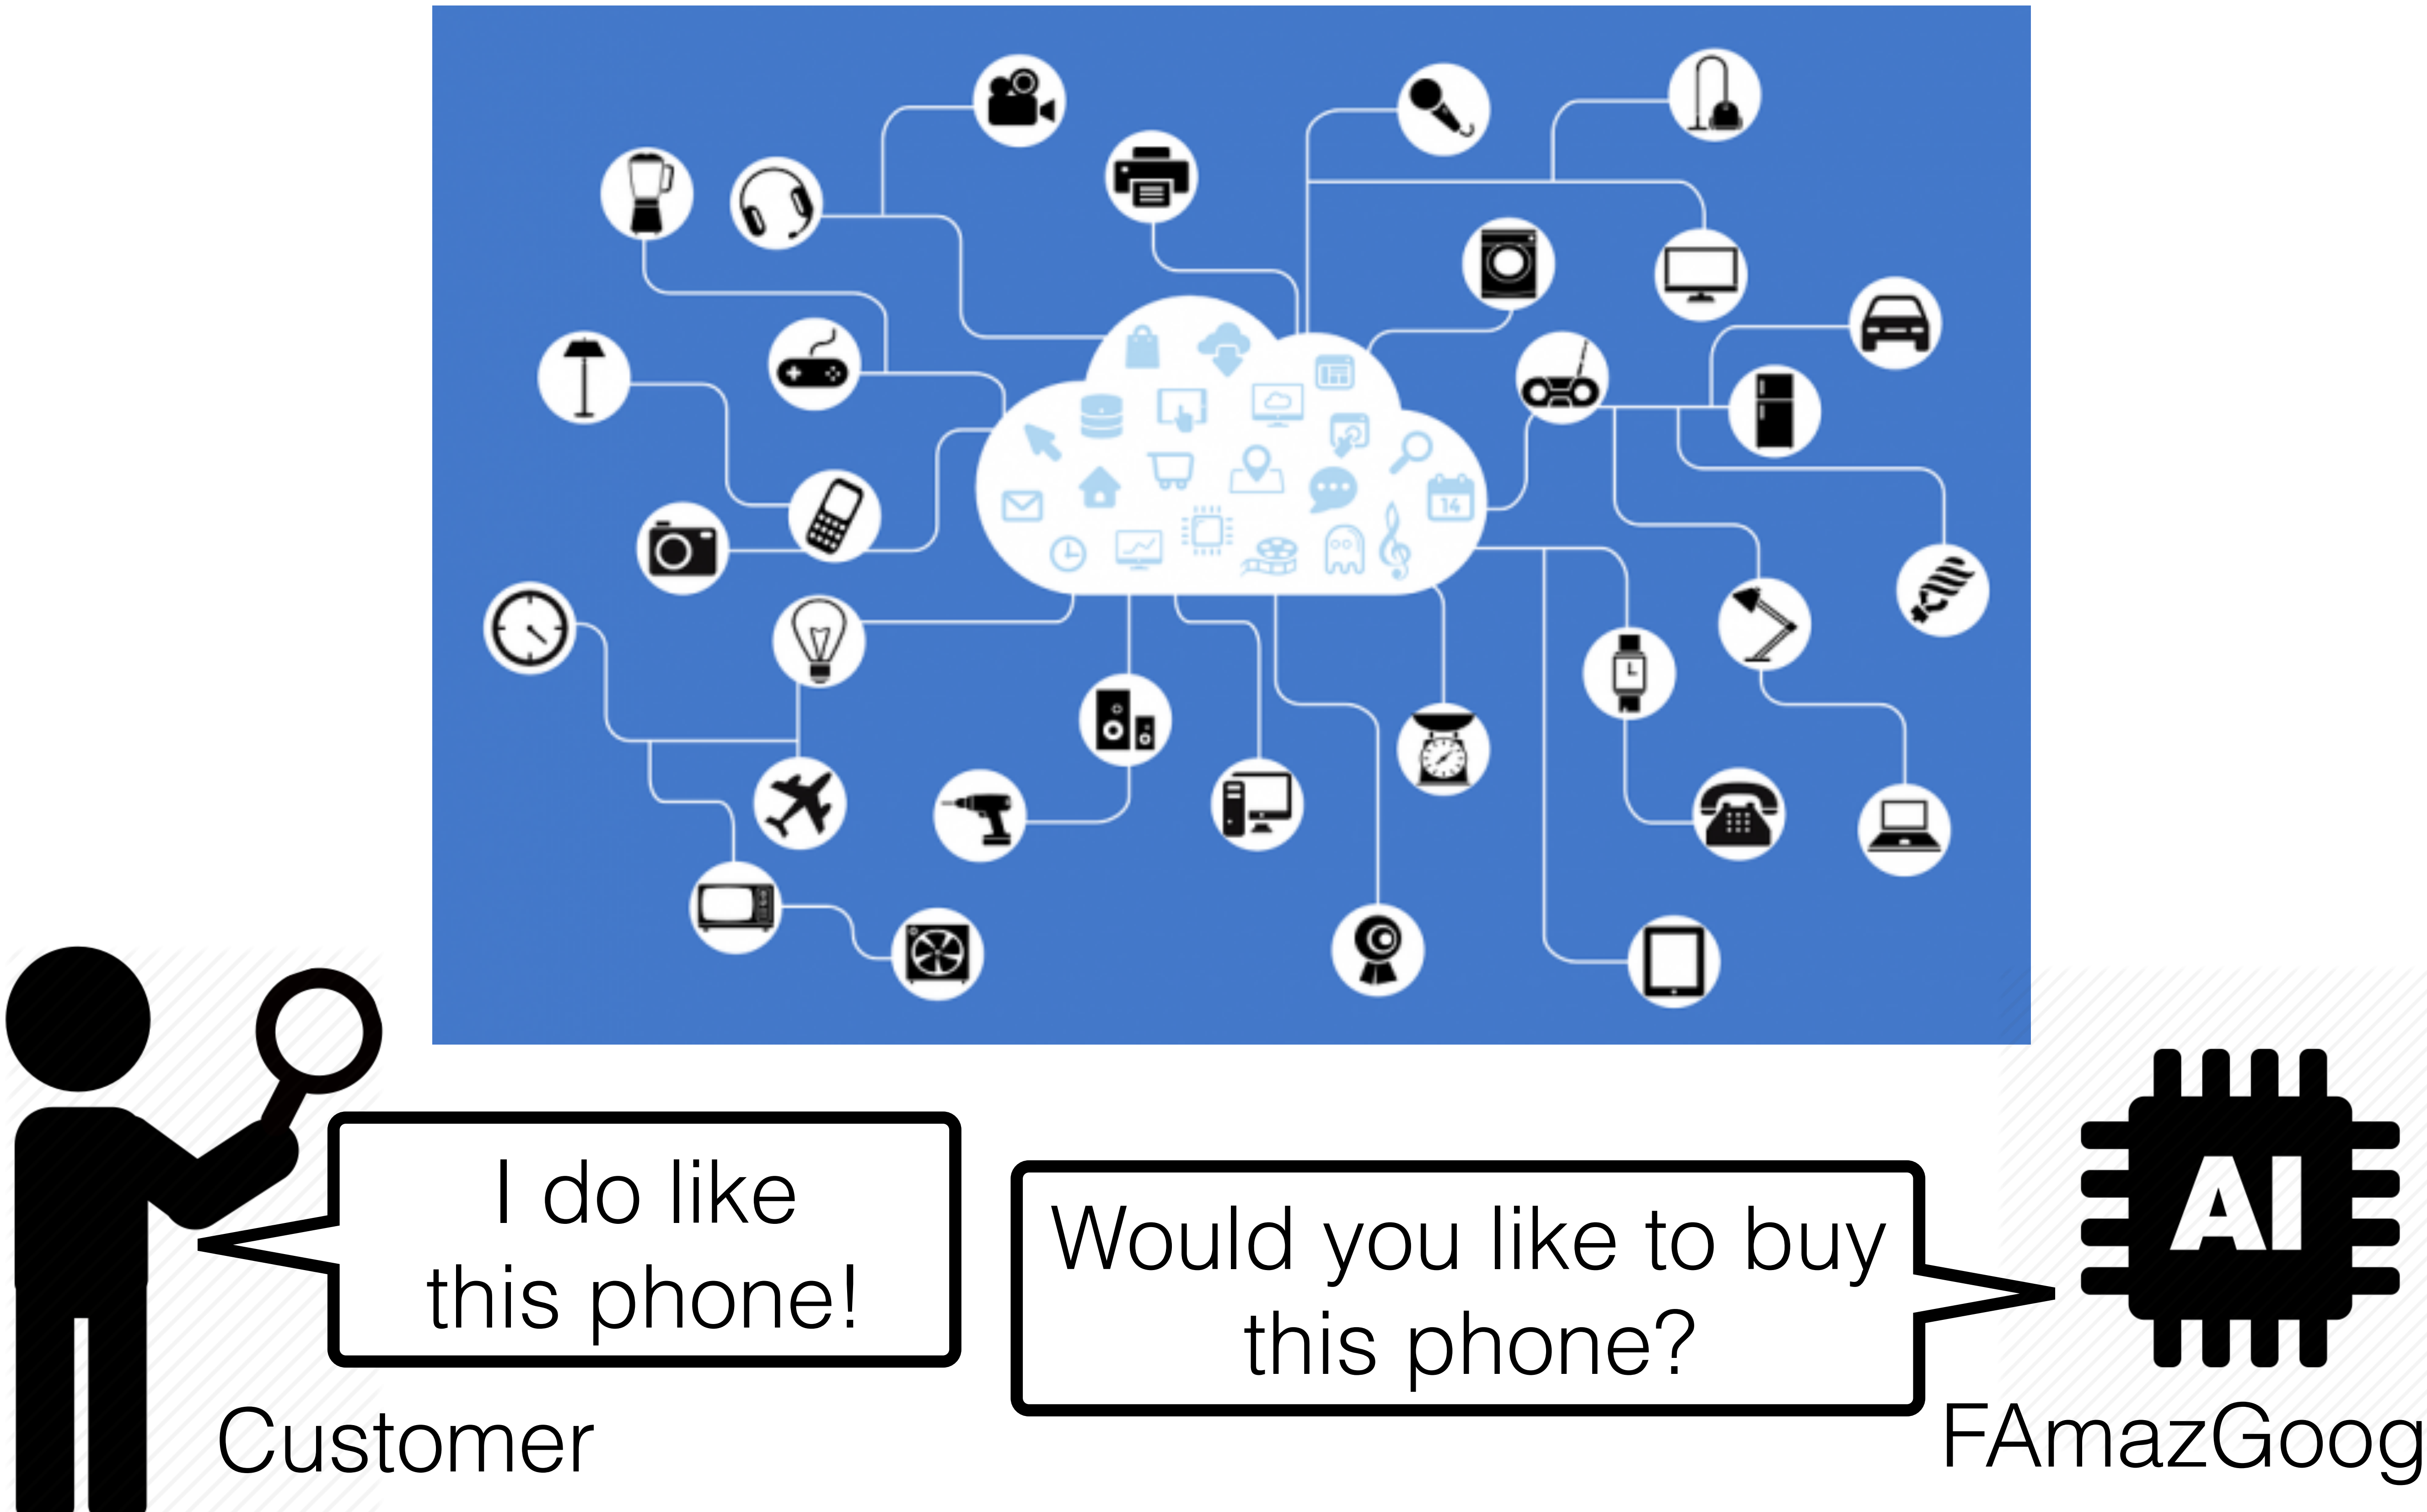

Problem

Active learning:

- **Goal**: figure out customers' preferences
- **Way**: test user's preference on items that the algorithm is uncertain how the user will like
- **Problem**: may show too many disliked items and hence drive customers away.

Recommender system:

- **Goal**: recommend items that customers will buy
- **Way**: recommend items similar to those that are known to be liked
- **Problem**: create "filter bubbles" that limit the customers to see only a restricted set of items.

Figuring out preferences vs. Recommending likable items

Exploration-exploitation tradeoff

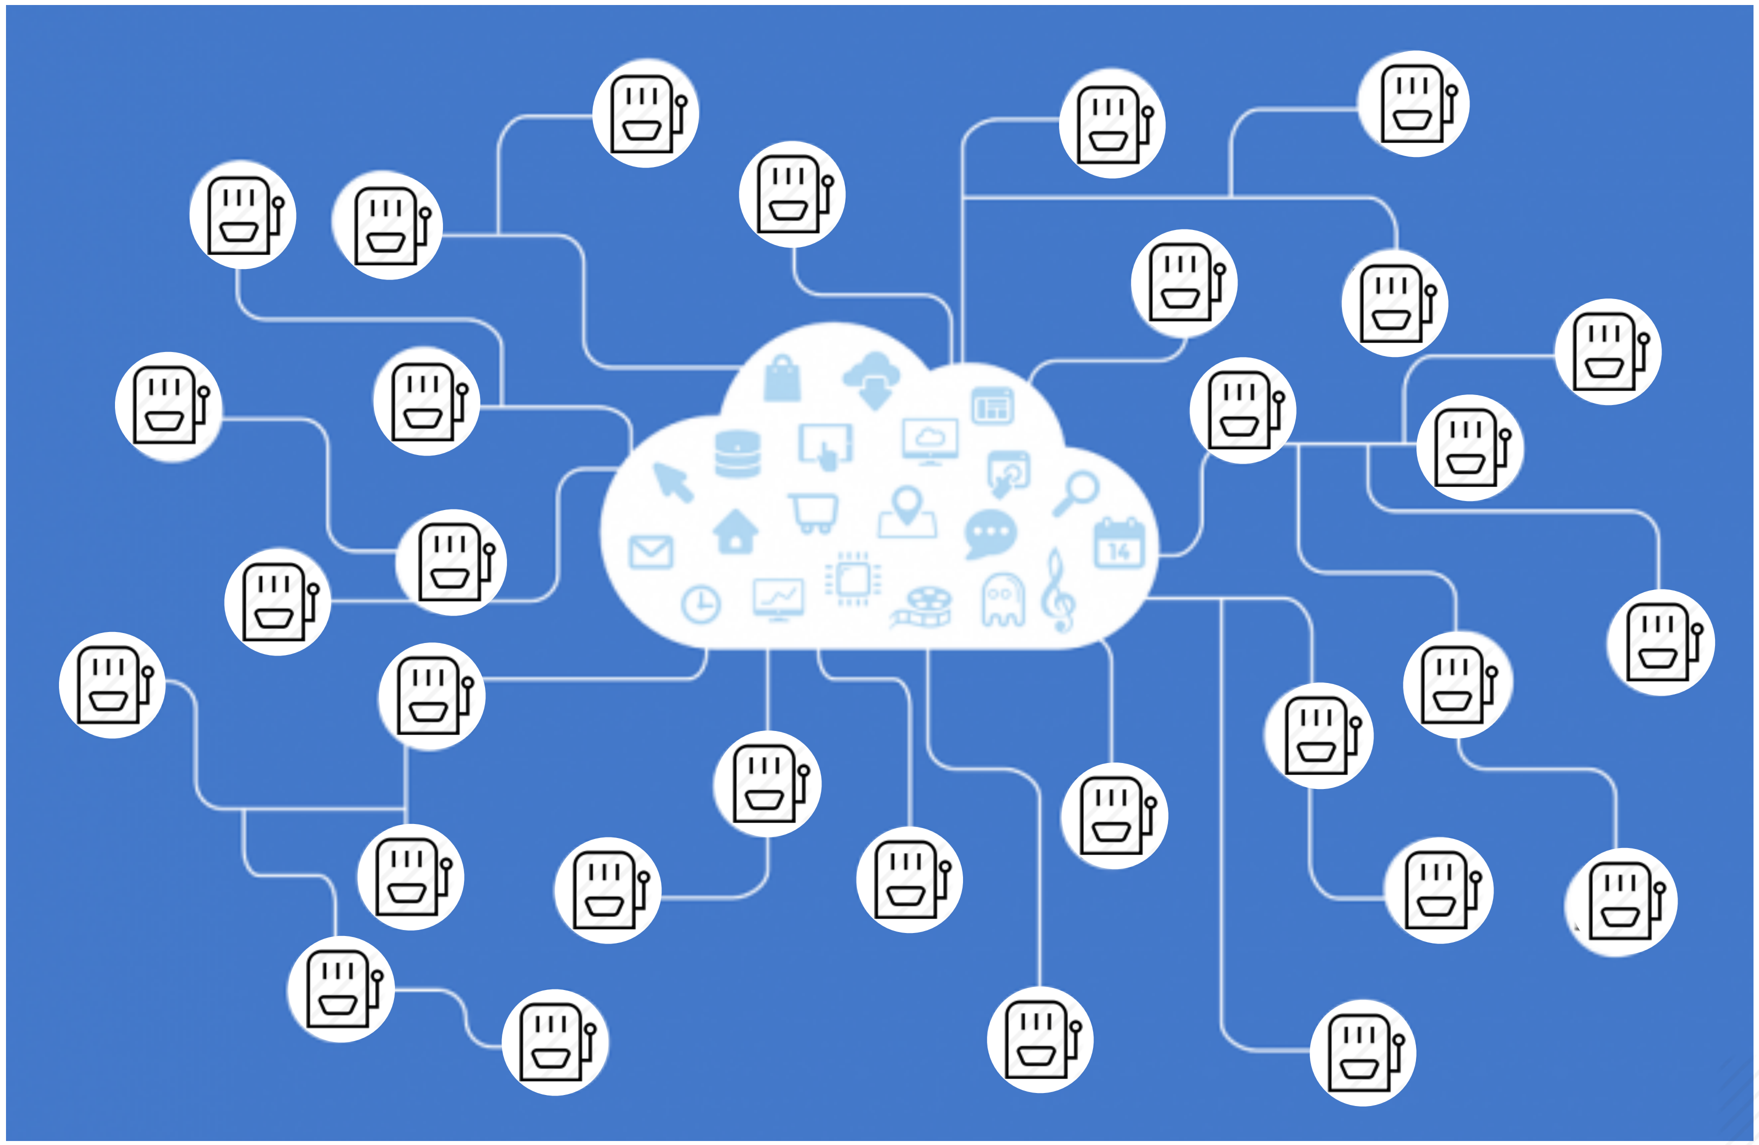

Customer

Should I stick to what I know to be OK, or should I risk trying something new to see if it is better?

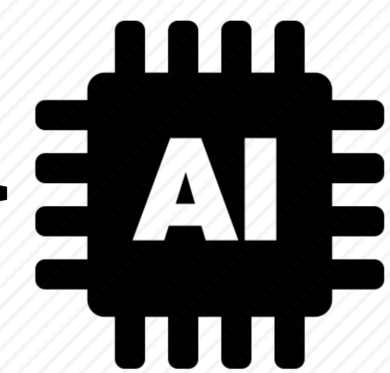

FAmazGoog

Cognitive science + Human-algorithm interaction

Specific Q: is there a way to overcome the trade-off?
General Q: given an algorithm, can we predict what the interaction will be like?

Human-algorithm interaction research (e.g., Pariser 2011, Baeza-Yates 2016):

- big data approach (e.g., collaborative filtering)
- uncontrolled decision factors

CogSci research (e.g., Bruner et al 1956, Shepard et al 1961):

- controlled decision factors
- traditionally no interaction with algorithms

CogSci + Human-algorithm interaction:

- human-algorithm interaction with controlled decision factors
- compare idealized responses with actual human responses

The framework

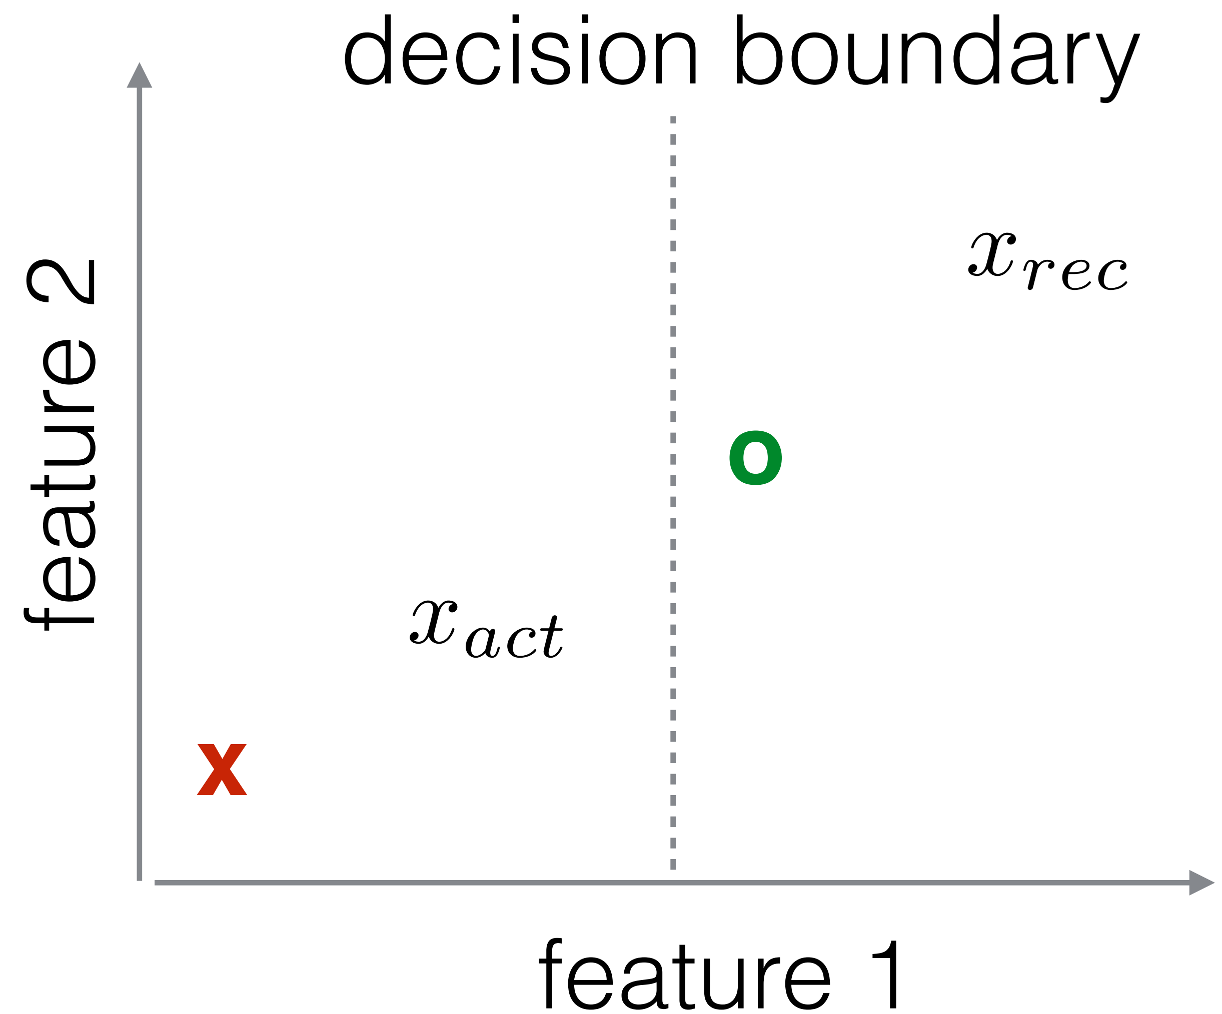

$$x_{rec} = \arg\max_{x^*} P(y = 1 | x^*, D)$$

$$x_{act} = \arg\min_{x^*} |0.5 - P(y = 1 | x^*, D)|$$

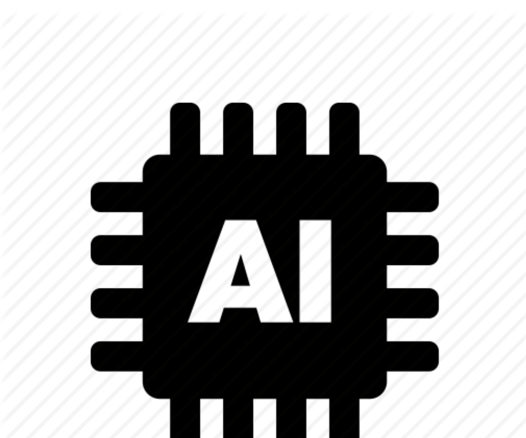

Active recommendation

$$x_{act} = \arg\min_{x^*} |0.5 - P(y=1|x^*, D)|$$

$$x_{rec} = \arg\max_{x^*} P(y=1|x^*, D) = \arg\min_{x^*} |1 - P(y=1|x^*, D)|$$

$$x_\alpha = \arg\min_{x^*} |\alpha - P(y=1|x^*, D)| \qquad \alpha \in [0.5, 1]$$

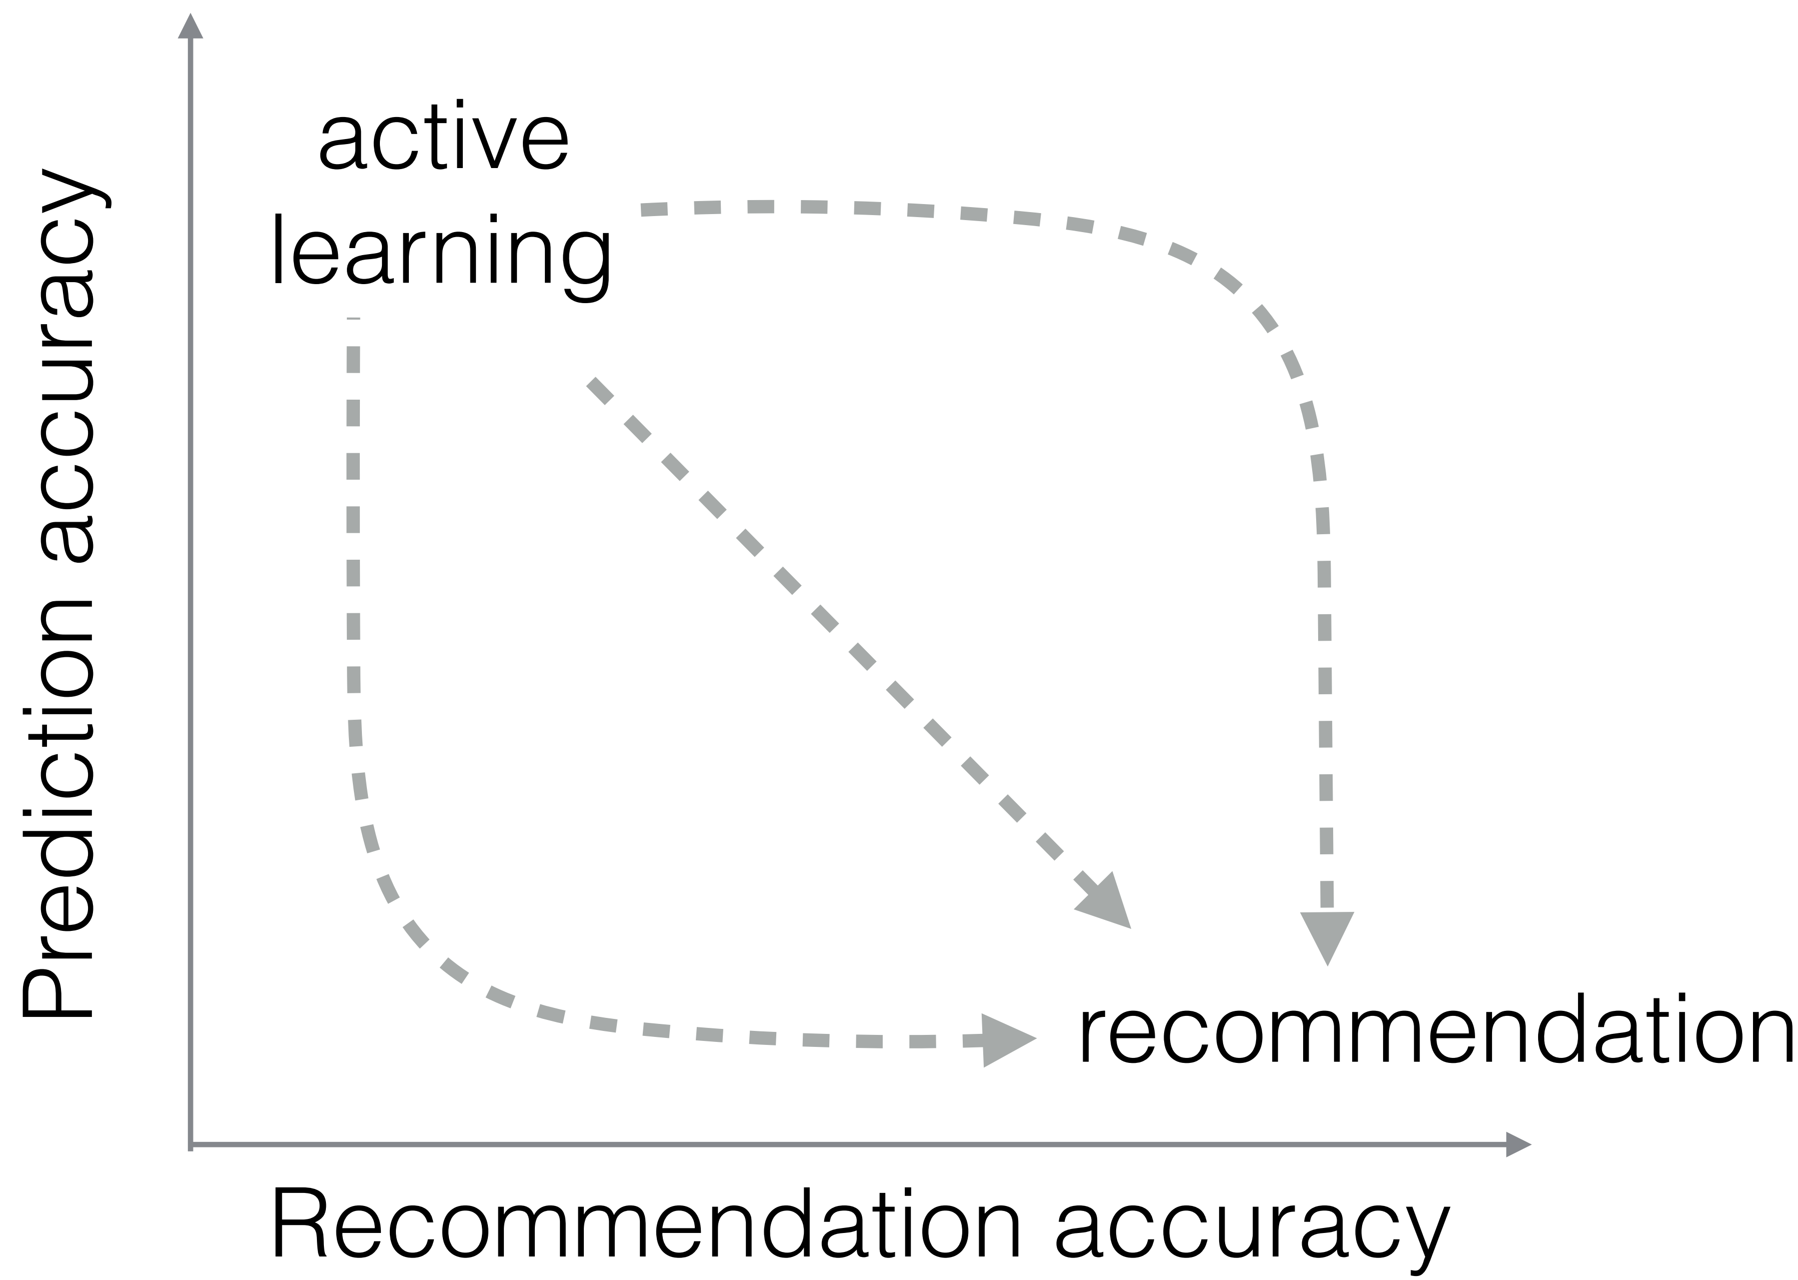

Experiment

Stimuli

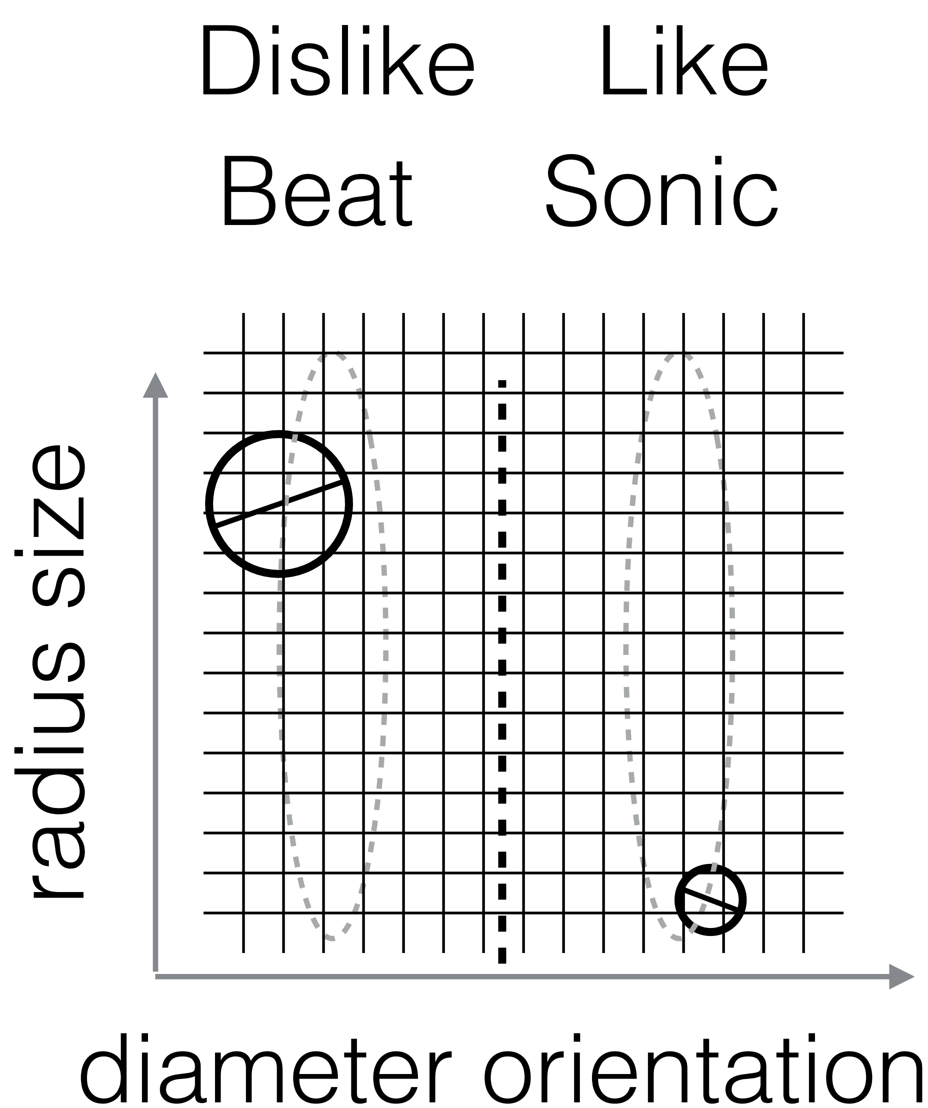

1. Training phase:
   - train subject to associate labels (Beat or Sonic) with stimuli
   - phase done when gets 19 out of the last 20 trials correct
2. Interaction phase:
   - instruct subject the preferred stimuli
   - naive algorithm chooses stimuli;
   - subject labels like/dislike;
   - algorithm updates setting
   - 20 trials
3. Check phase:
   - subject labels 20 stimuli sampled from a grid

Markant & Gureckis 2014

Conditions & subjects

- 6 interaction conditions:
  - random, α=0.5 (active), α=1 (recommend)
  - α=0.55, α=0.75, α=0.95 (active recommend)
- 30 subjects per condition
- Omit subject if check score < 18/20
  - ~ 4 subjects omitted per condition
- Consistency score: the fraction of the subject's responses in the interaction phase that matched the expected responses from the predefined boundary
  - Flip subjects like/dislike response if consistency score < 50%
  - ~ 3 subjects' responses flipped per condition

Results

Recommendation accuracy = the fraction of likes in the interaction phase.
Prediction accuracy = the fraction of correct model predictions, w.r.t. the true boundary, on 100 stimuli sampled from a grid in the feature space.

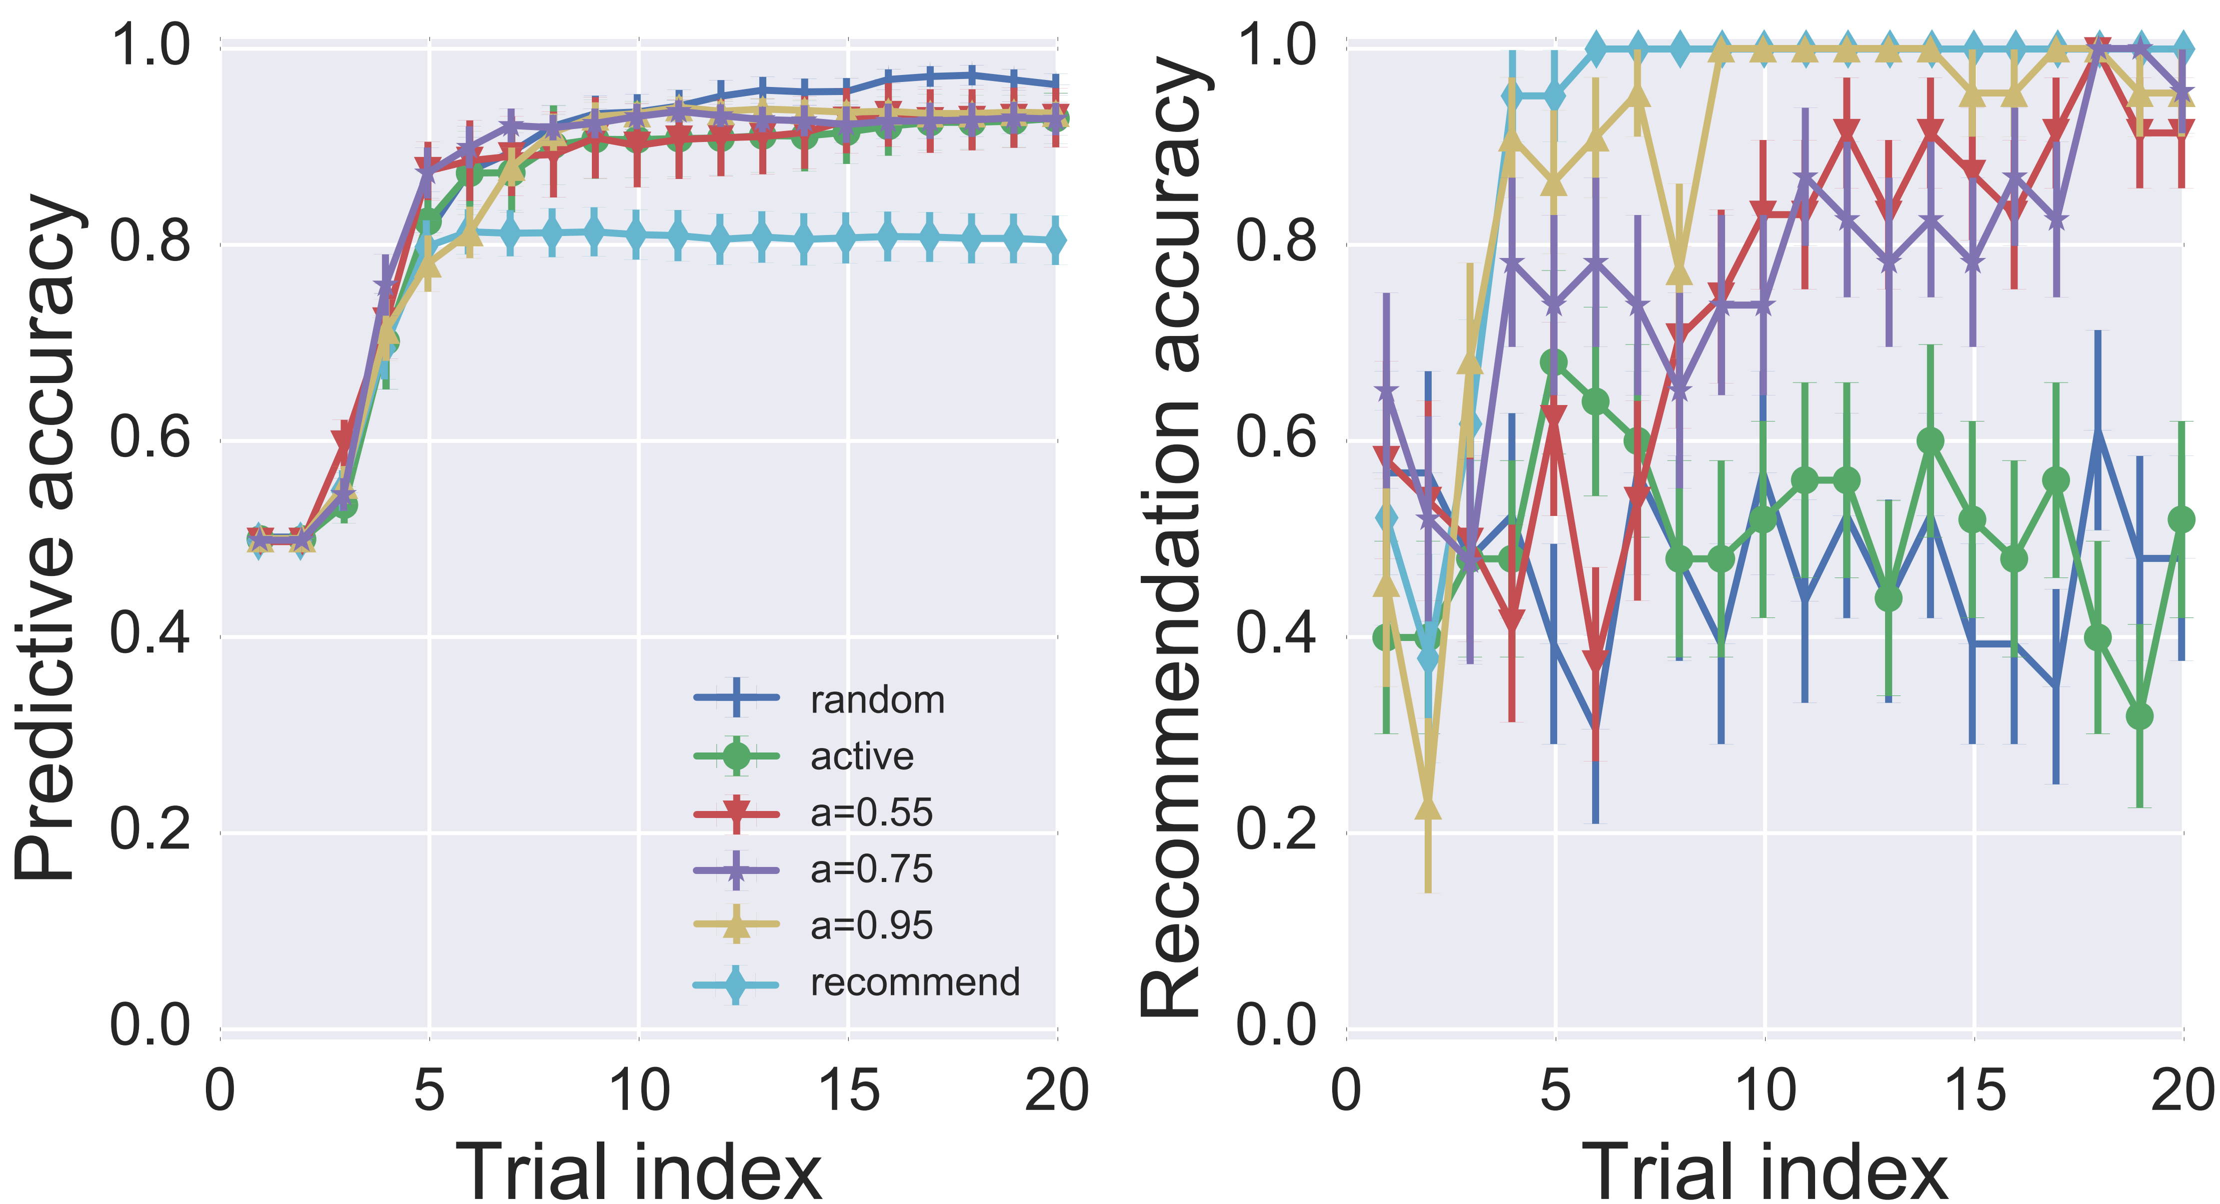

Active recommendation overcomes the tradeoff!

The distribution of interaction examples

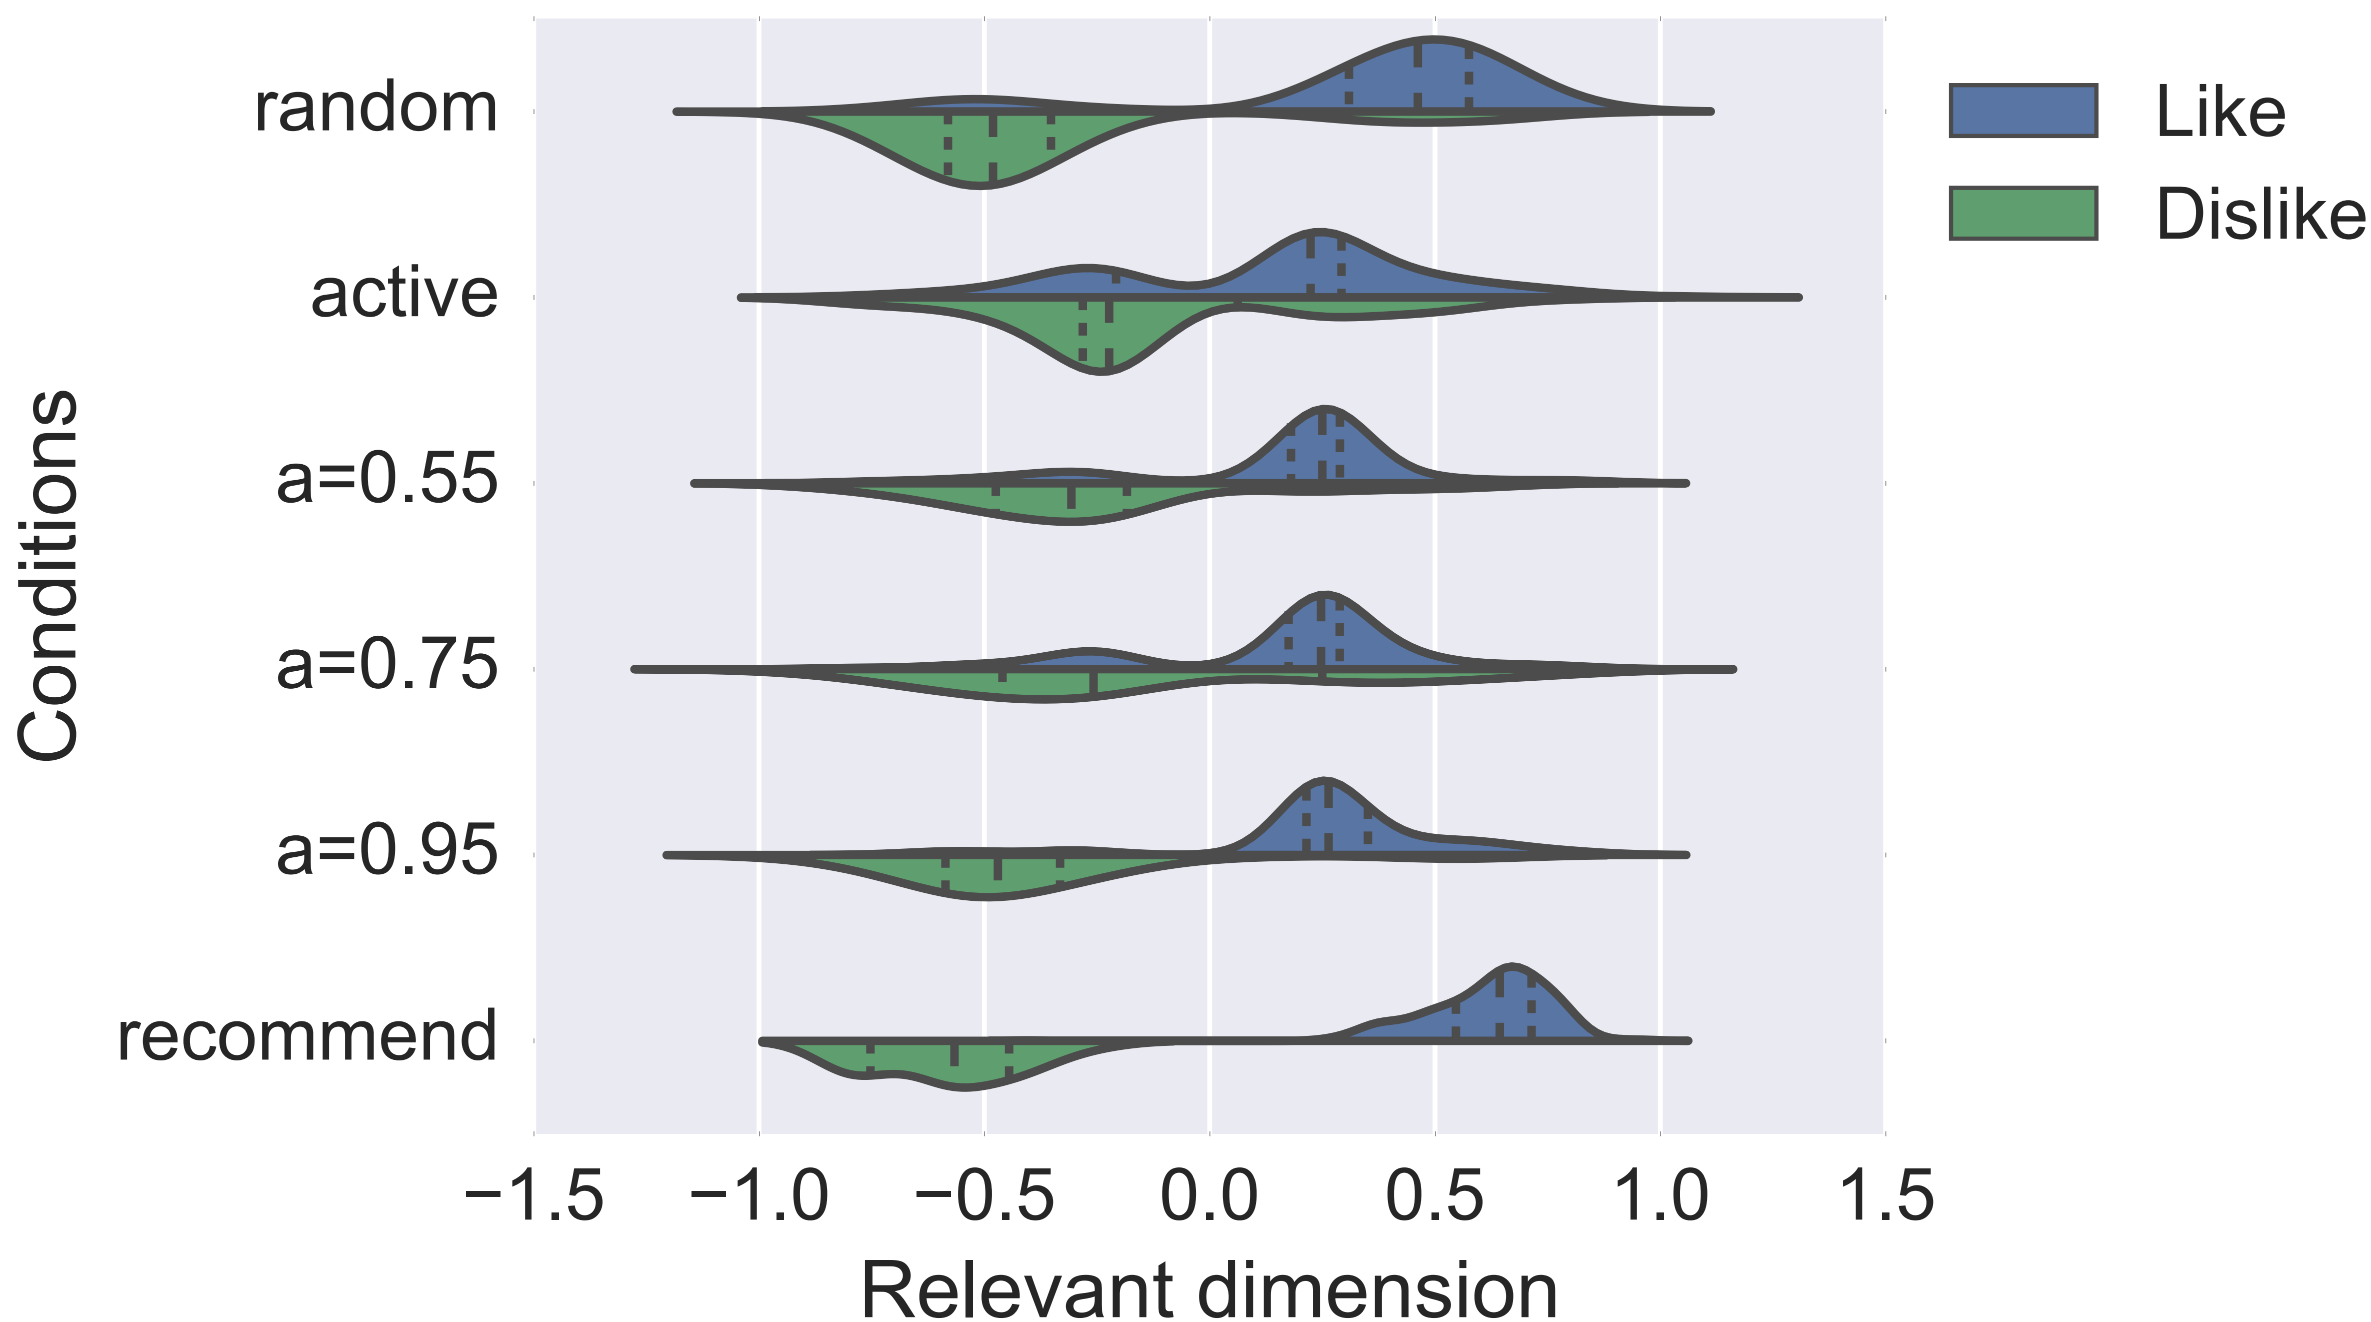

Active recommendation selects uncertain example ***within*** the relevant category.

Active recommendation

$x_{act} = \arg\min_{x^*} |0.5 - P(y=1|x^*, D)|$

$$x_{rec} = \arg\max_{x^*} P(y=1|x^*, D) = \arg\min_{x^*} |1 - P(y=1|x^*, D)|$$

$$x_\alpha = \arg\min_{x^*} |\alpha - P(y=1|x^*, D)| \qquad \alpha \in [0.5, 1]$$

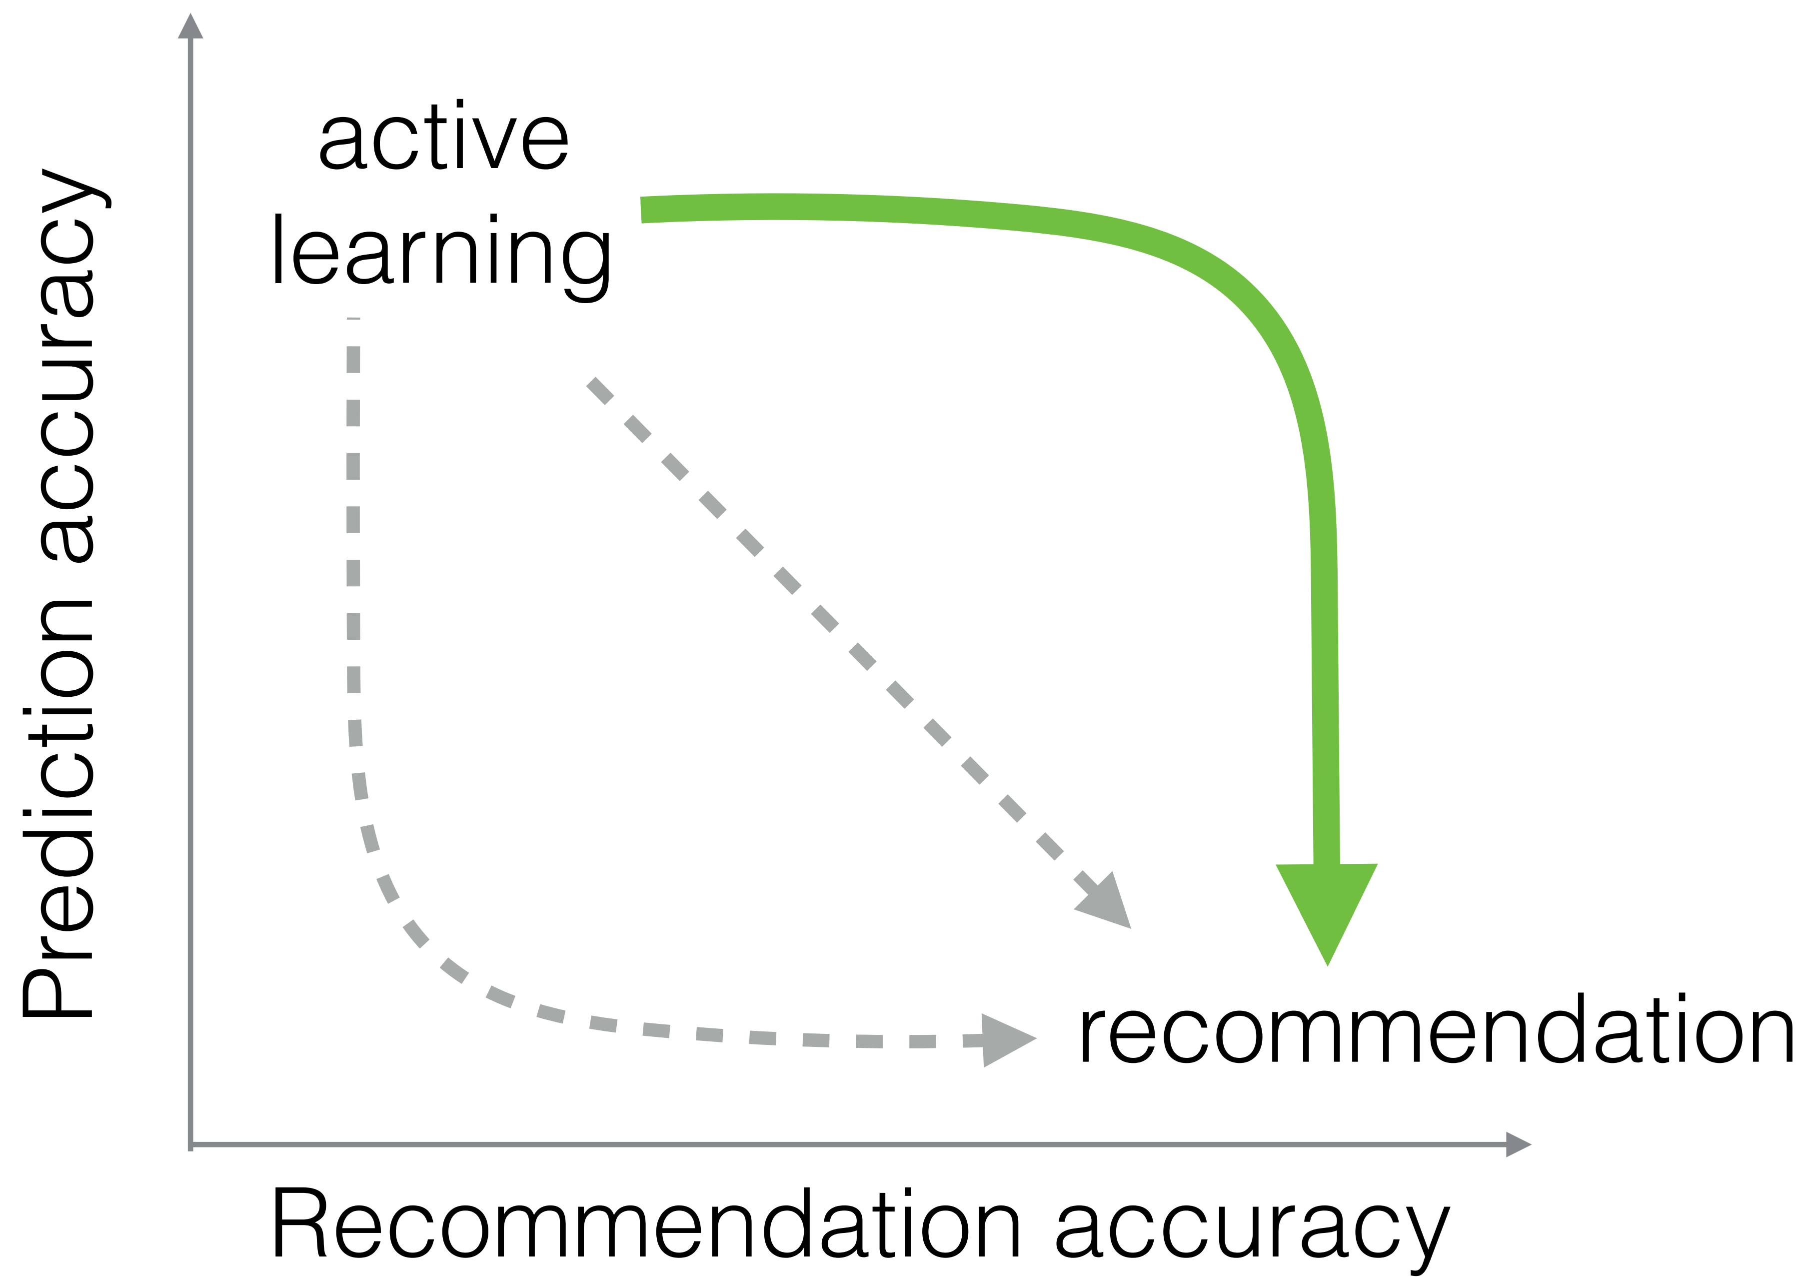

The effect of human variability

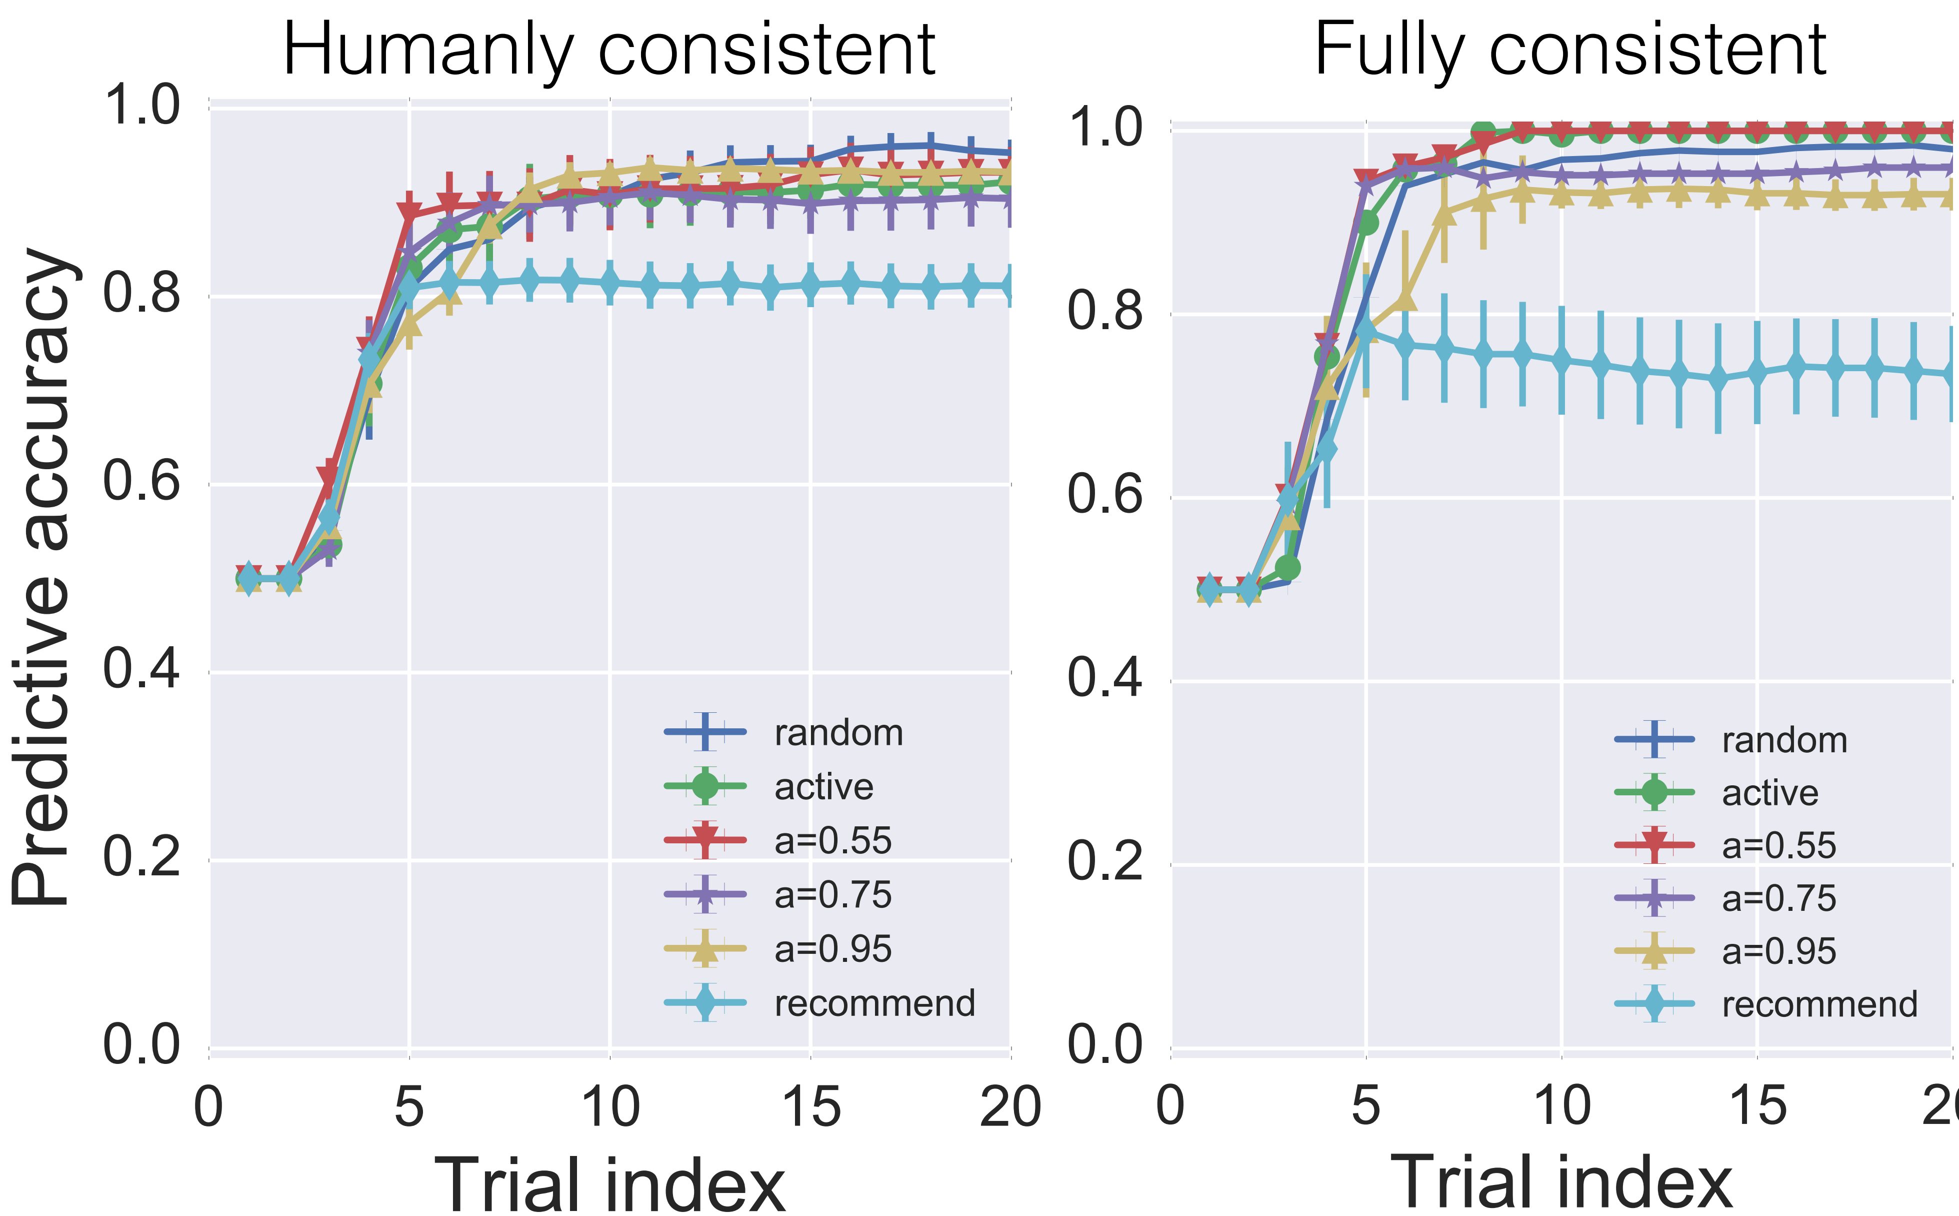

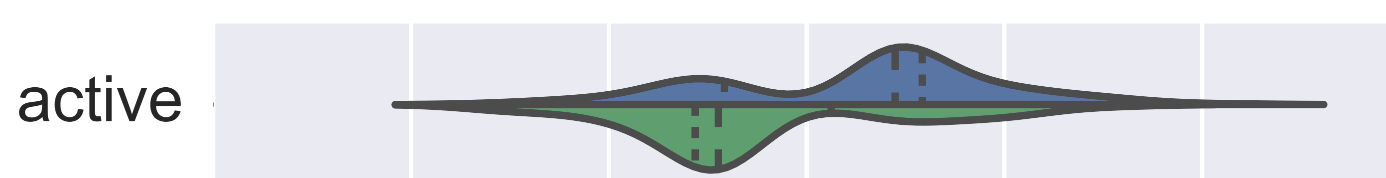

If look at only fully consistent subjects —> see strict ordering.
Noisy response close to the boundary —> imperfect prediction accuracy.

Active recommendation

$$x_{act} = \arg\min_{x^*} |0.5 - P(y=1|x^*, D)|$$

$$x_{rec} = \arg\max_{x^*} P(y=1|x^*, D) = \arg\min_{x^*} |1 - P(y=1|x^*, D)|$$

$$x_\alpha = \arg\min_{x^*} |\alpha - P(y=1|x^*, D)| \qquad \alpha \in [0.5, 1]$$

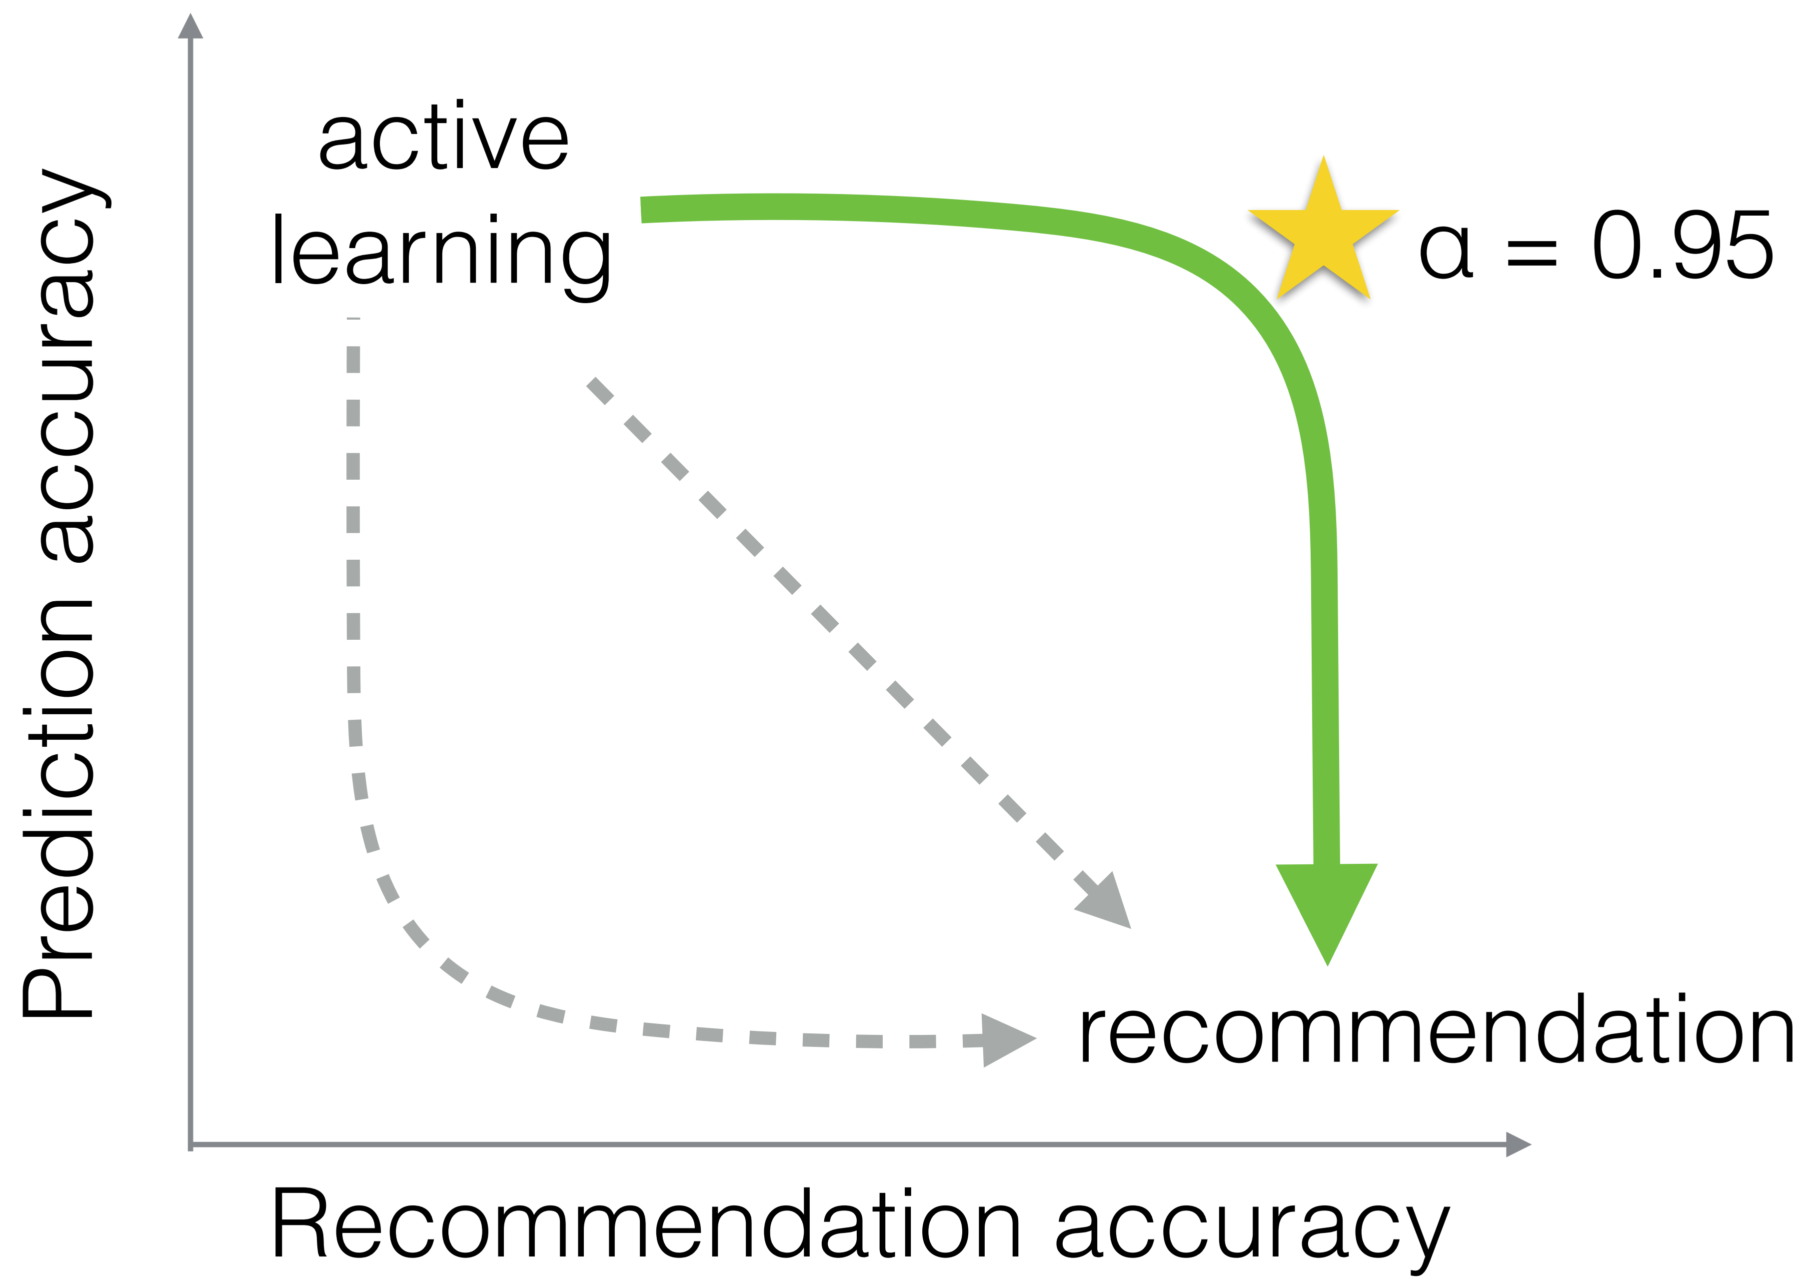

Conclusions

- Studied human-algorithm interaction as a cognitive concept learning experiment.
- Formalized a unification for recommendation and active learning.
- Challenge the explore-or-exploit dichotomy.
- Showed a case when the tradeoff doesn’t really exist.
- Active recommendation can overcome the tradeoff by selecting uncertain example within the relevant category.

Acknowledgments

Jake Whritner

Pat Shafto

Olfa Nasraoui

RUTGERS NEWARK

UNIVERSITY of LOUISVILLE

NSF

The core idea

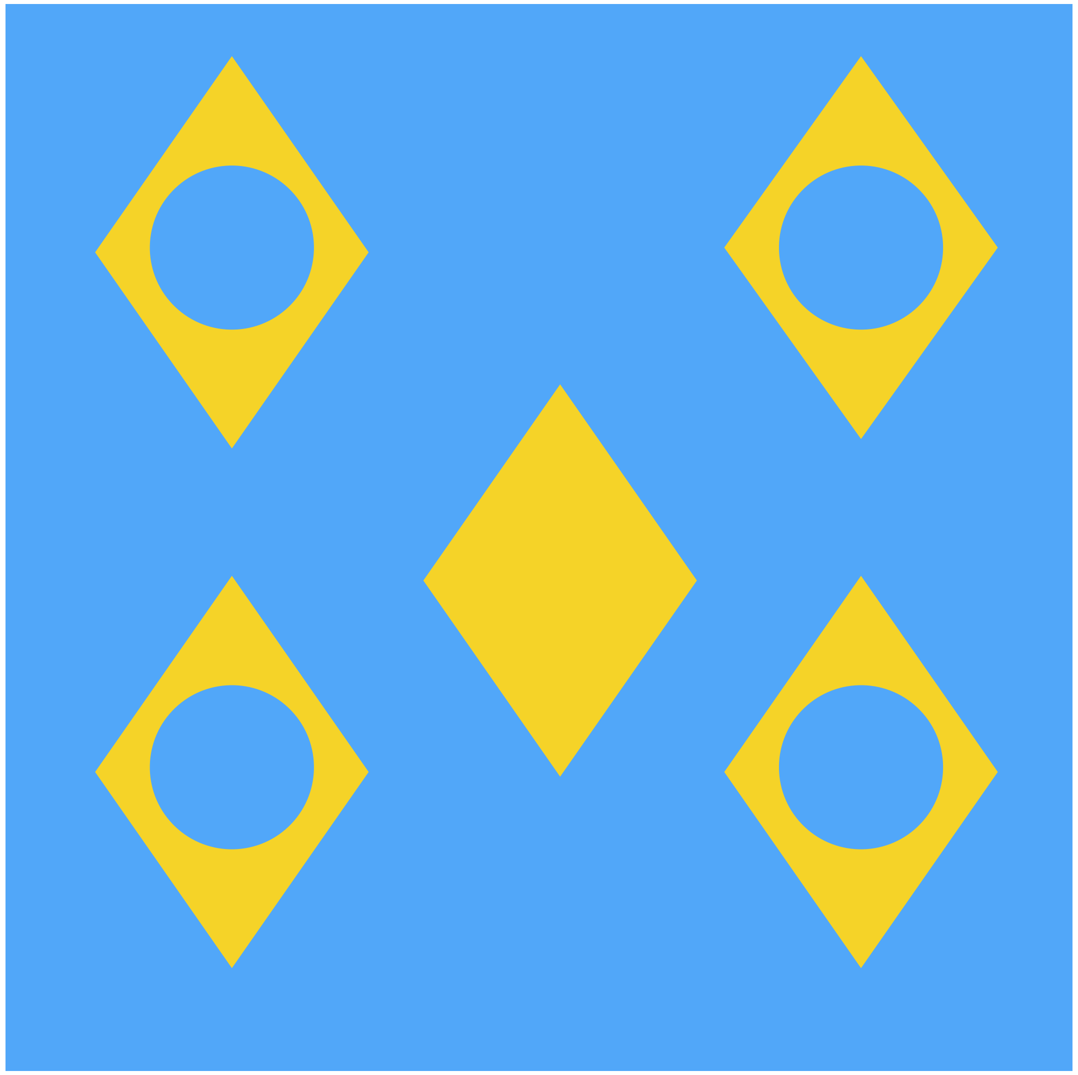

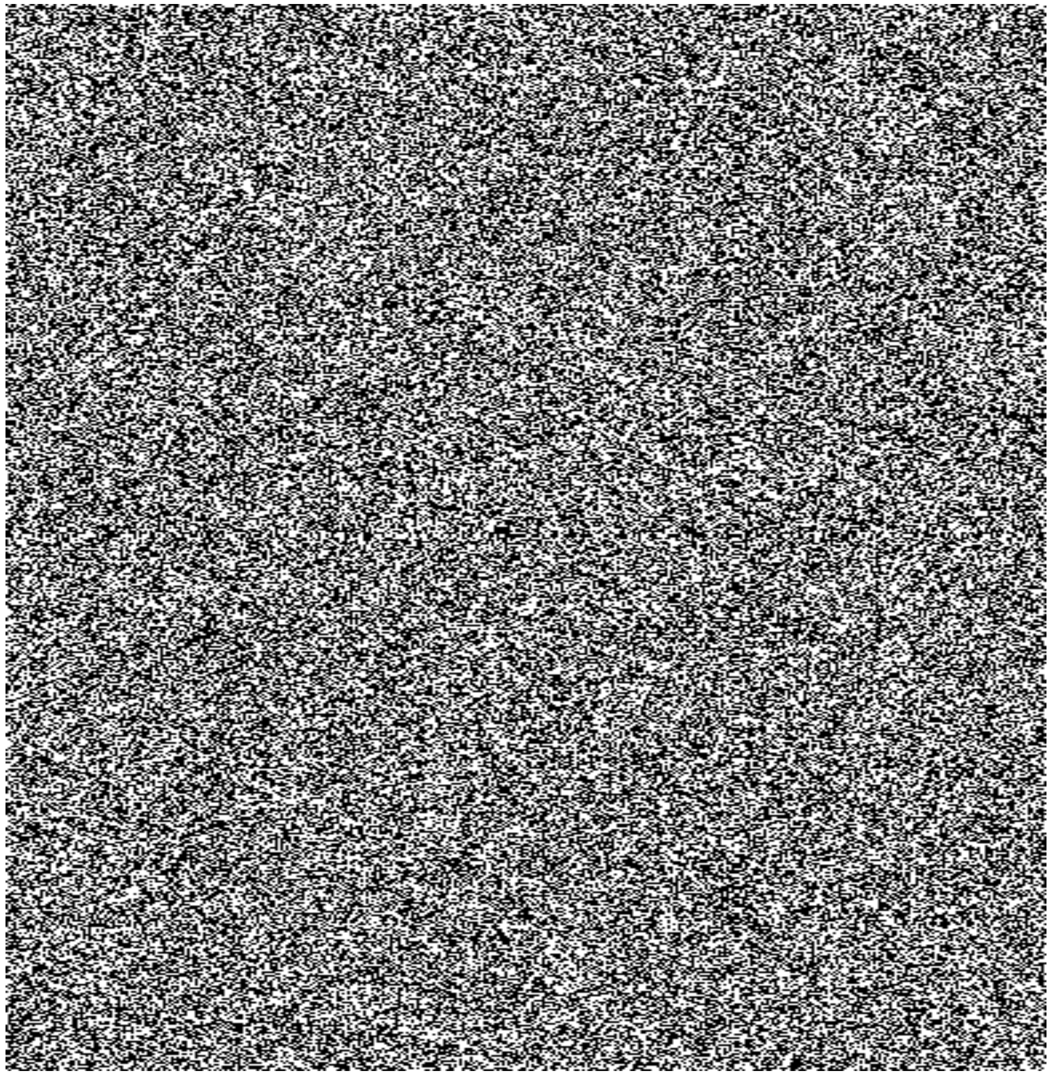

Active recommendation bypasses the tradeoff if the model captures the global and local structure.