AI+HCI

MICHAEL BERNSTEIN
STANFORD UNIVERSITY

Announcements

- Project fair on Thursday
- Shorter lecture and discussion to earn us back enough time for the project fair

- Jeff Hancock joining next Tuesday while Michael is at CSCW

Course Overview

INTRO	week 1	Intro to Interaction; Intro to Social Computing
	week 2	Intro to Design; Interaction
DEPTH	week 3	Interaction; Social Computing
	week 4	Social Computing
	week 5	Design
BREADTH	week 6	AI+HCI; Media
	week 7	Foundations
	week 8	Access; Programming
	week 9	Collaboration; Visualization
	week 10	Education; Critiques of HCI

People: where the AI hits the road

MIT Personal Robotics Group

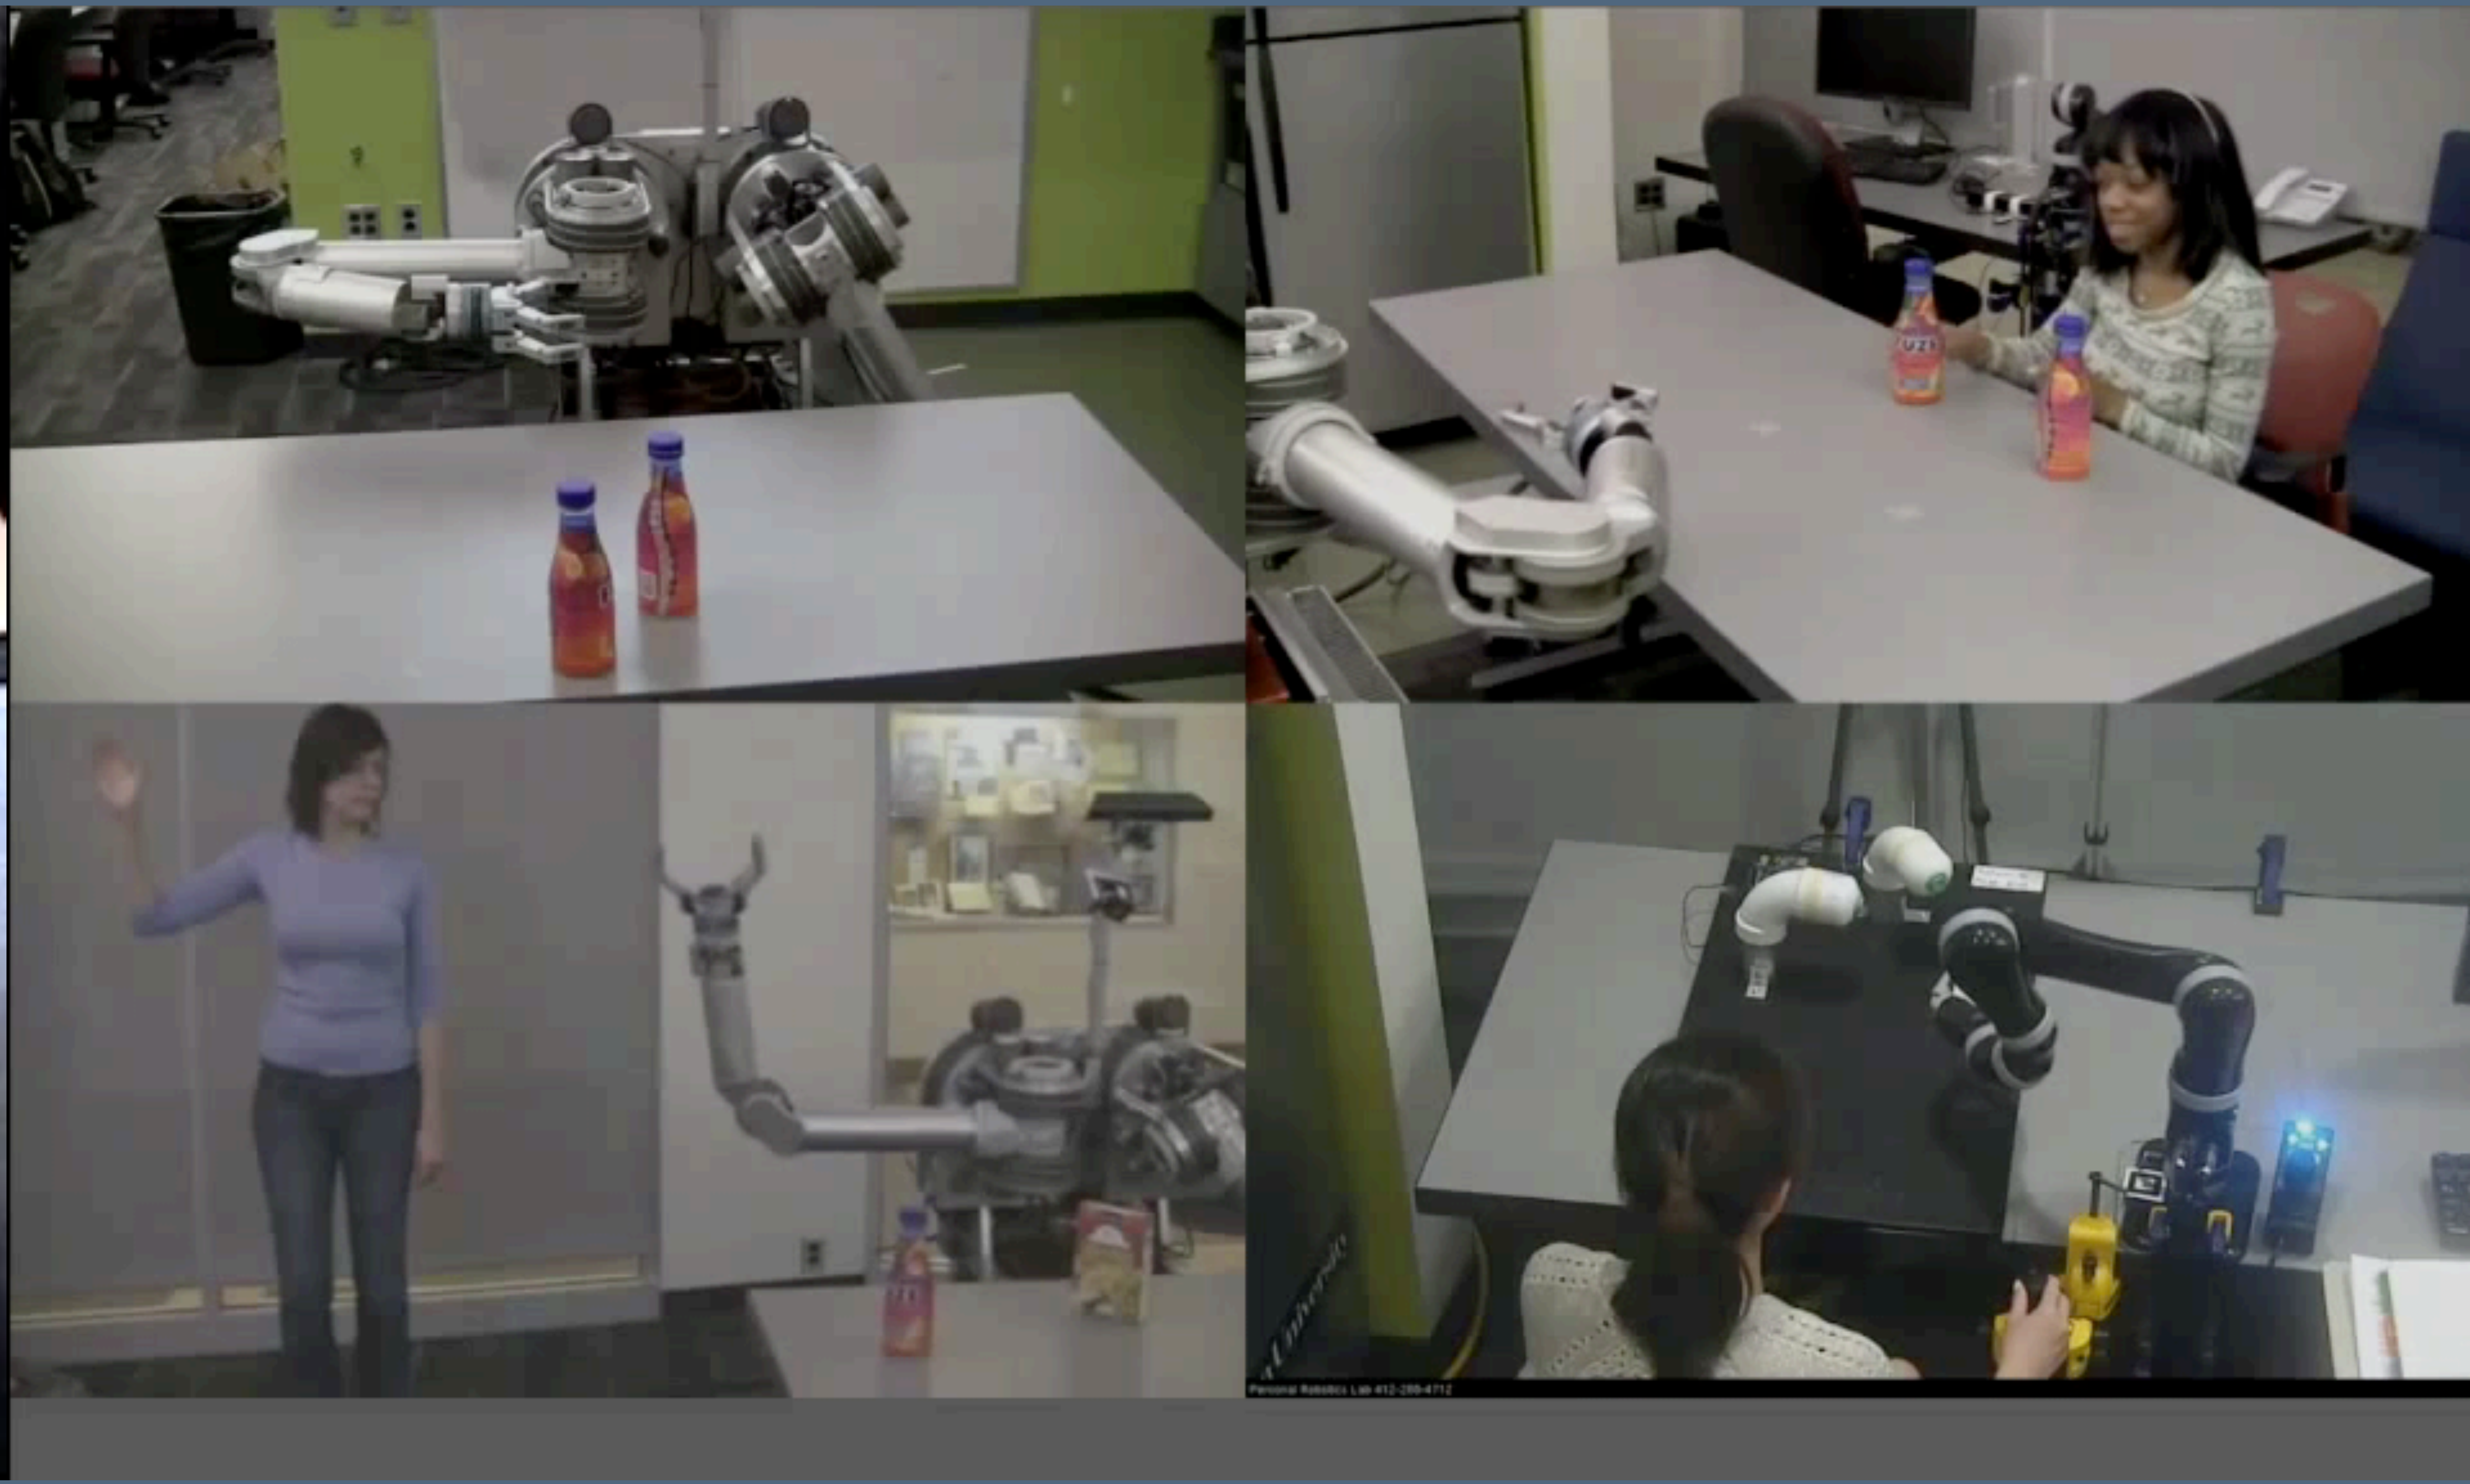

UC Berkeley InterACT laboratory

Interactive AI on live TV

Interactive AI in everyday use

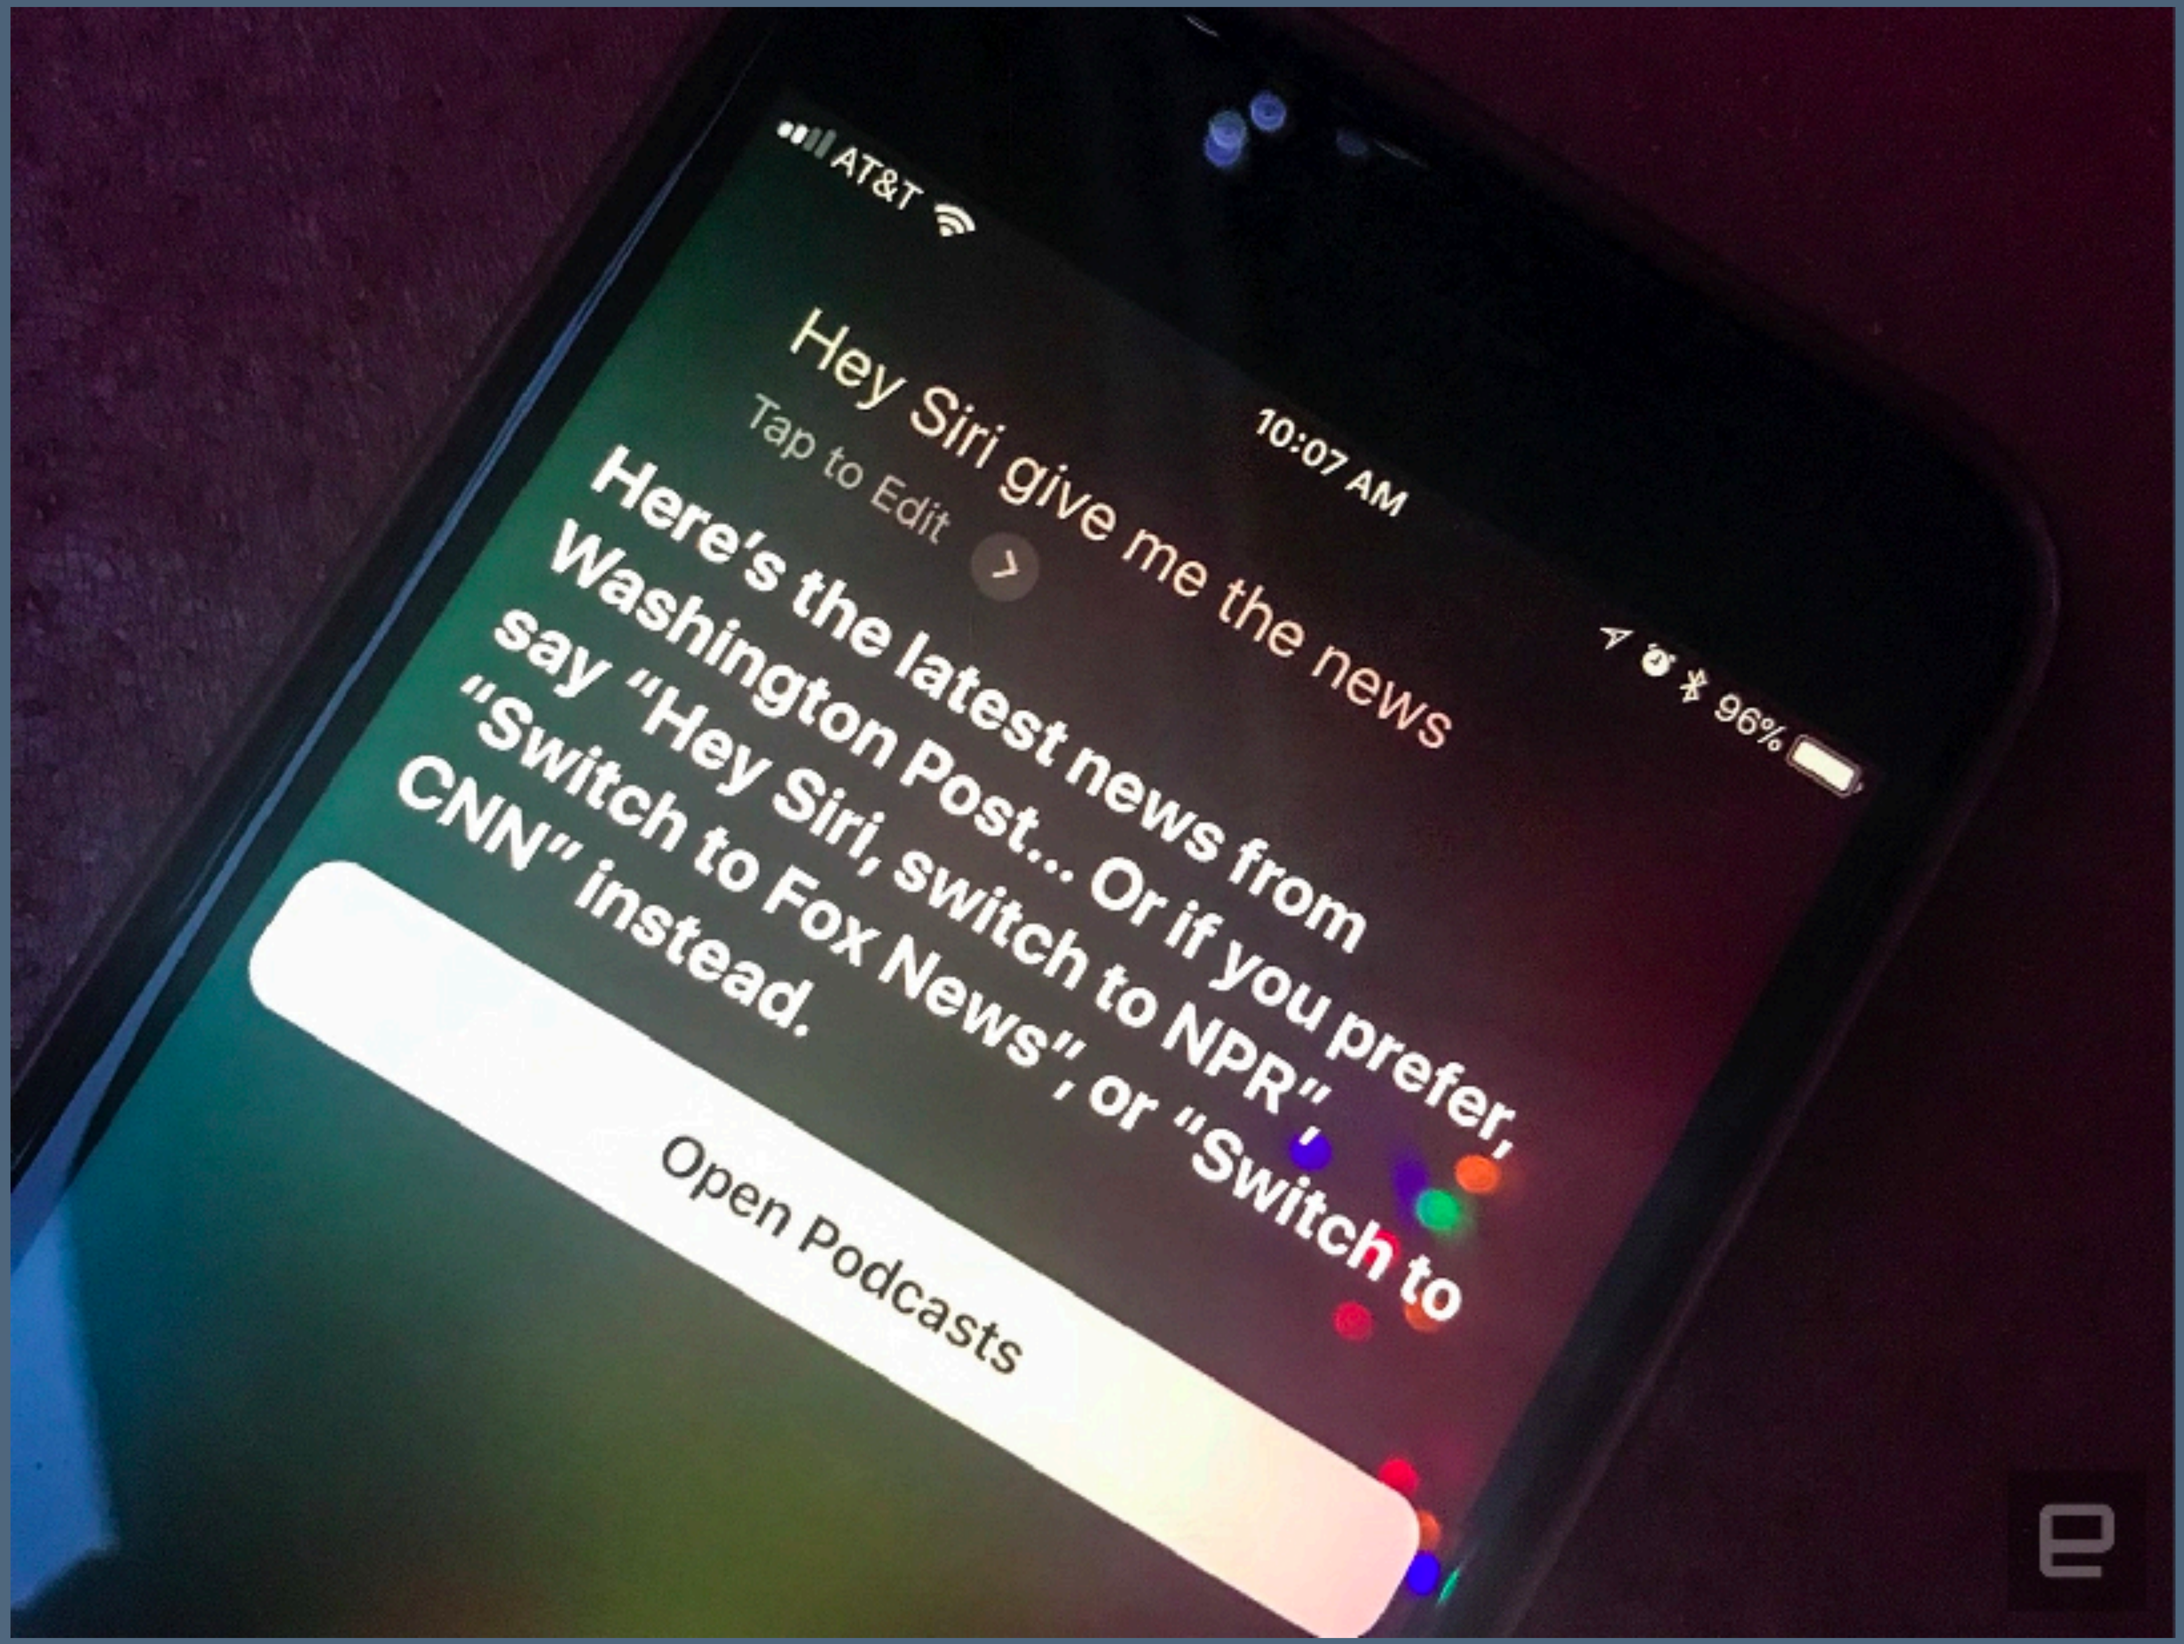

Siri, image from Engadget

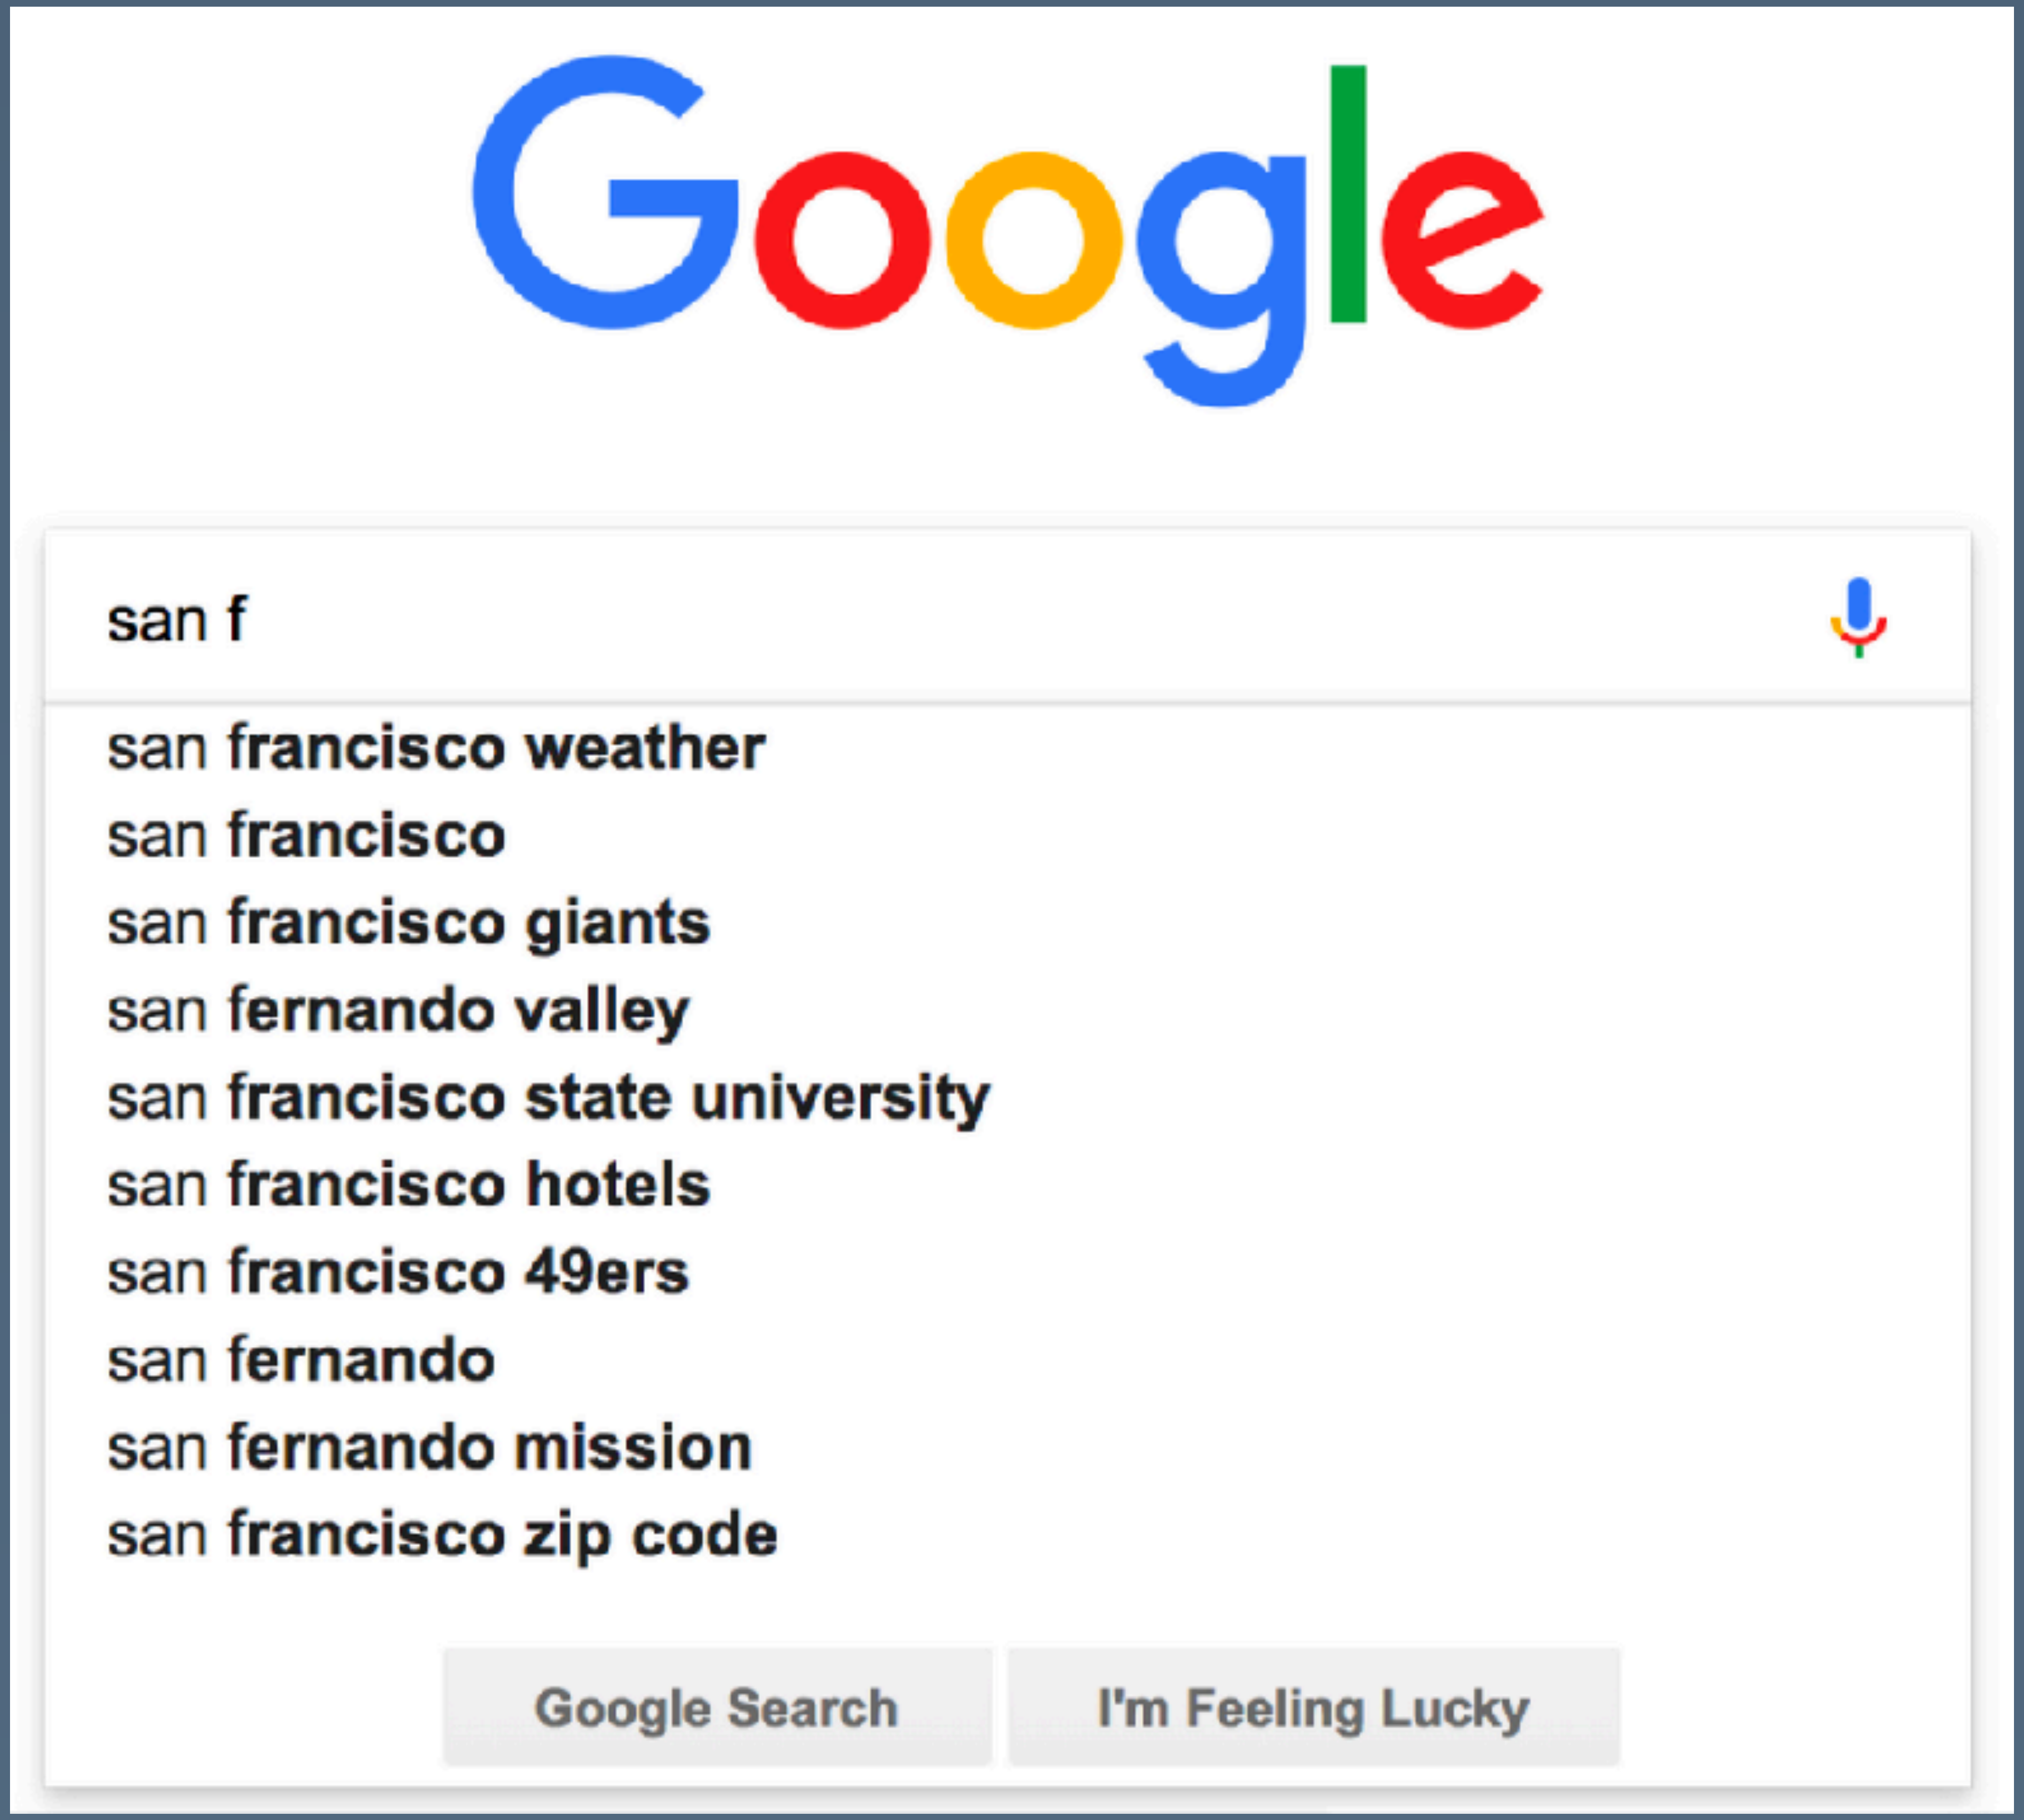

Google Autocomplete

Automation is not the answer

Intelligence Augmentation

A reaction to:

“AI will take over human jobs”

Engelbart Institute

AUGMENTING HUMAN INTELLECT: A CONCEPTUAL FRAMEWORK

Prepared for:

DIRECTOR OF INFORMATION SCIENCES
AIR FORCE OFFICE OF SCIENTIFIC RESEARCH
WASHINGTON 25, D.C.

CONTRACT AF 49(638)-1024

By: *D. C. Engelbart*

STANFORD RESEARCH INSTITUTE

MENLO PARK, CALIFORNIA

*SRI

SRI

Artificial Intelligence

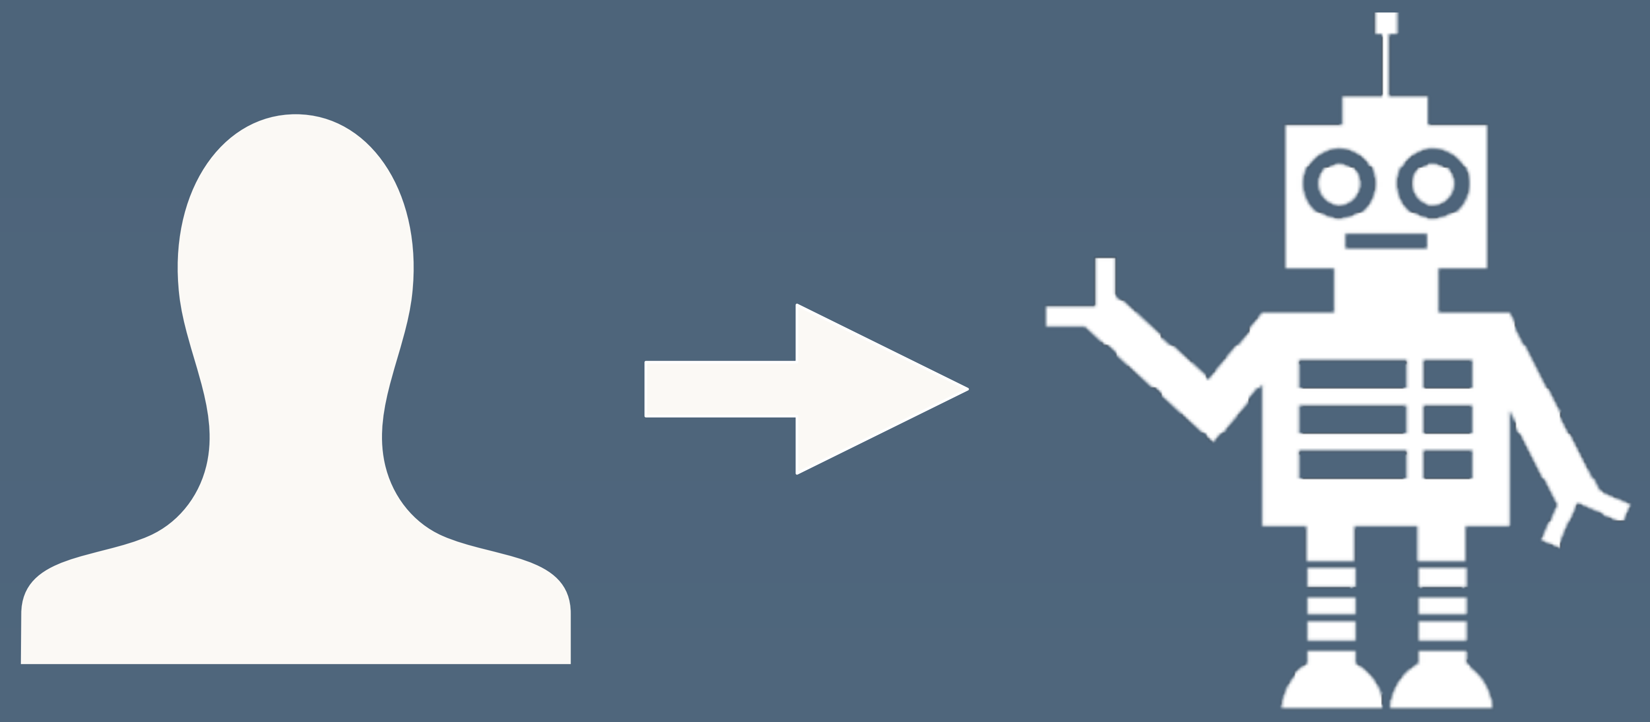

Replace human intelligence with artificial intelligence

Intelligence Augmentation

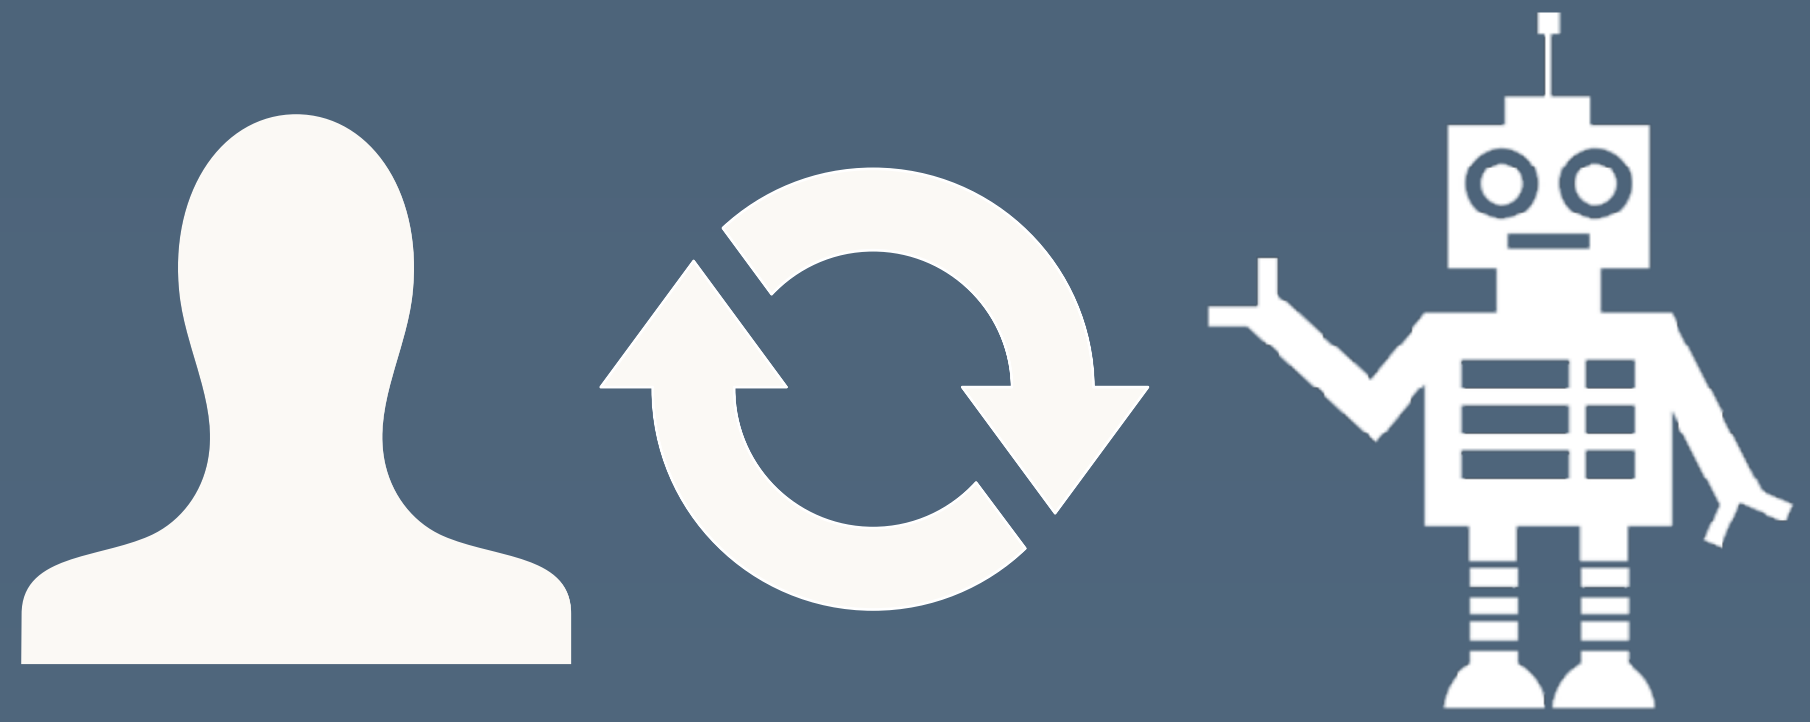

Augment human intelligence with artificial intelligence

Modeling Uncertainty and Error

Software agents

- Delegate to proactive software and artificial intelligence

Pattie Maes, MIT Media Lab

Direct manipulation

- Users should always have full control

Ben Shneiderman, U. Maryland

A problem has been detected and Windows has been shut down to prevent damage to your computer.

The problem seems to be caused by the following file: kbdhid.sys

MANUALLY_INITIATED_CRASH

If this is the first time you've seen this stop error screen, restart your computer. If this screen appears again, follow these steps:

Check to make sure any new hardware or software is properly installed. If this is a new installation, ask your hardware or software manufacturer for any Windows updates you might need.

If problems continue, disable or remove any newly installed hardware or software. Disable BIOS memory options such as caching or shadowing. If you need to use safe mode to remove or disable components, restart your computer, press F8 to select Advanced Startup Options, and then select Safe Mode.

Technical Information:

*** STOP: 0x000000e2 (0x00000000, 0x00000000, 0x00000000, 0x00000000)

The problem

- Unlike traditional interfaces, introducing an AI into a system creates an element of uncertainty
- Will it understand you correctly? Will it make the correct inferences?
- How do you design a system that can be robust to these kinds of errors?

Mixed-initiative interaction

Mixed-initiative: combine the best of both worlds

YOU READ THIS

[Horvitz CHI '99]

- Utility-based calculation:
  - u(A,G) = (positive) utility of taking an automated action when the goal is correctly guessed
  - u(A,¬G) = (negative) utility of taking the same action when the goal is incorrectly guessed
  - u(¬A,G) and u(¬A,¬G) similarly

	Desired goal	Not desired goal
Take action	u(A,G)	u(A,¬G)
No action	u(¬A,G)	u(¬A,¬G)

Mixed-initiative: utility calculation

[Horvitz CHI '99]

- What's the expected value of taking action?
  - $\left(P(G) \cdot u(A, G)\right) + \left(P(\neg G) \cdot u(A, \neg G)\right)$
- What's the expected value of taking no action?
  - $\left(P(G) \cdot u(\neg A, G)\right) + \left(P(\neg G) \cdot u(\neg A, \neg G)\right)$

	Desired goal	Not desired goal
Take action	u(A,G)	u(A,¬G)
No action	u(¬A,G)	u(¬A,¬G)

Mixed initiative: visually

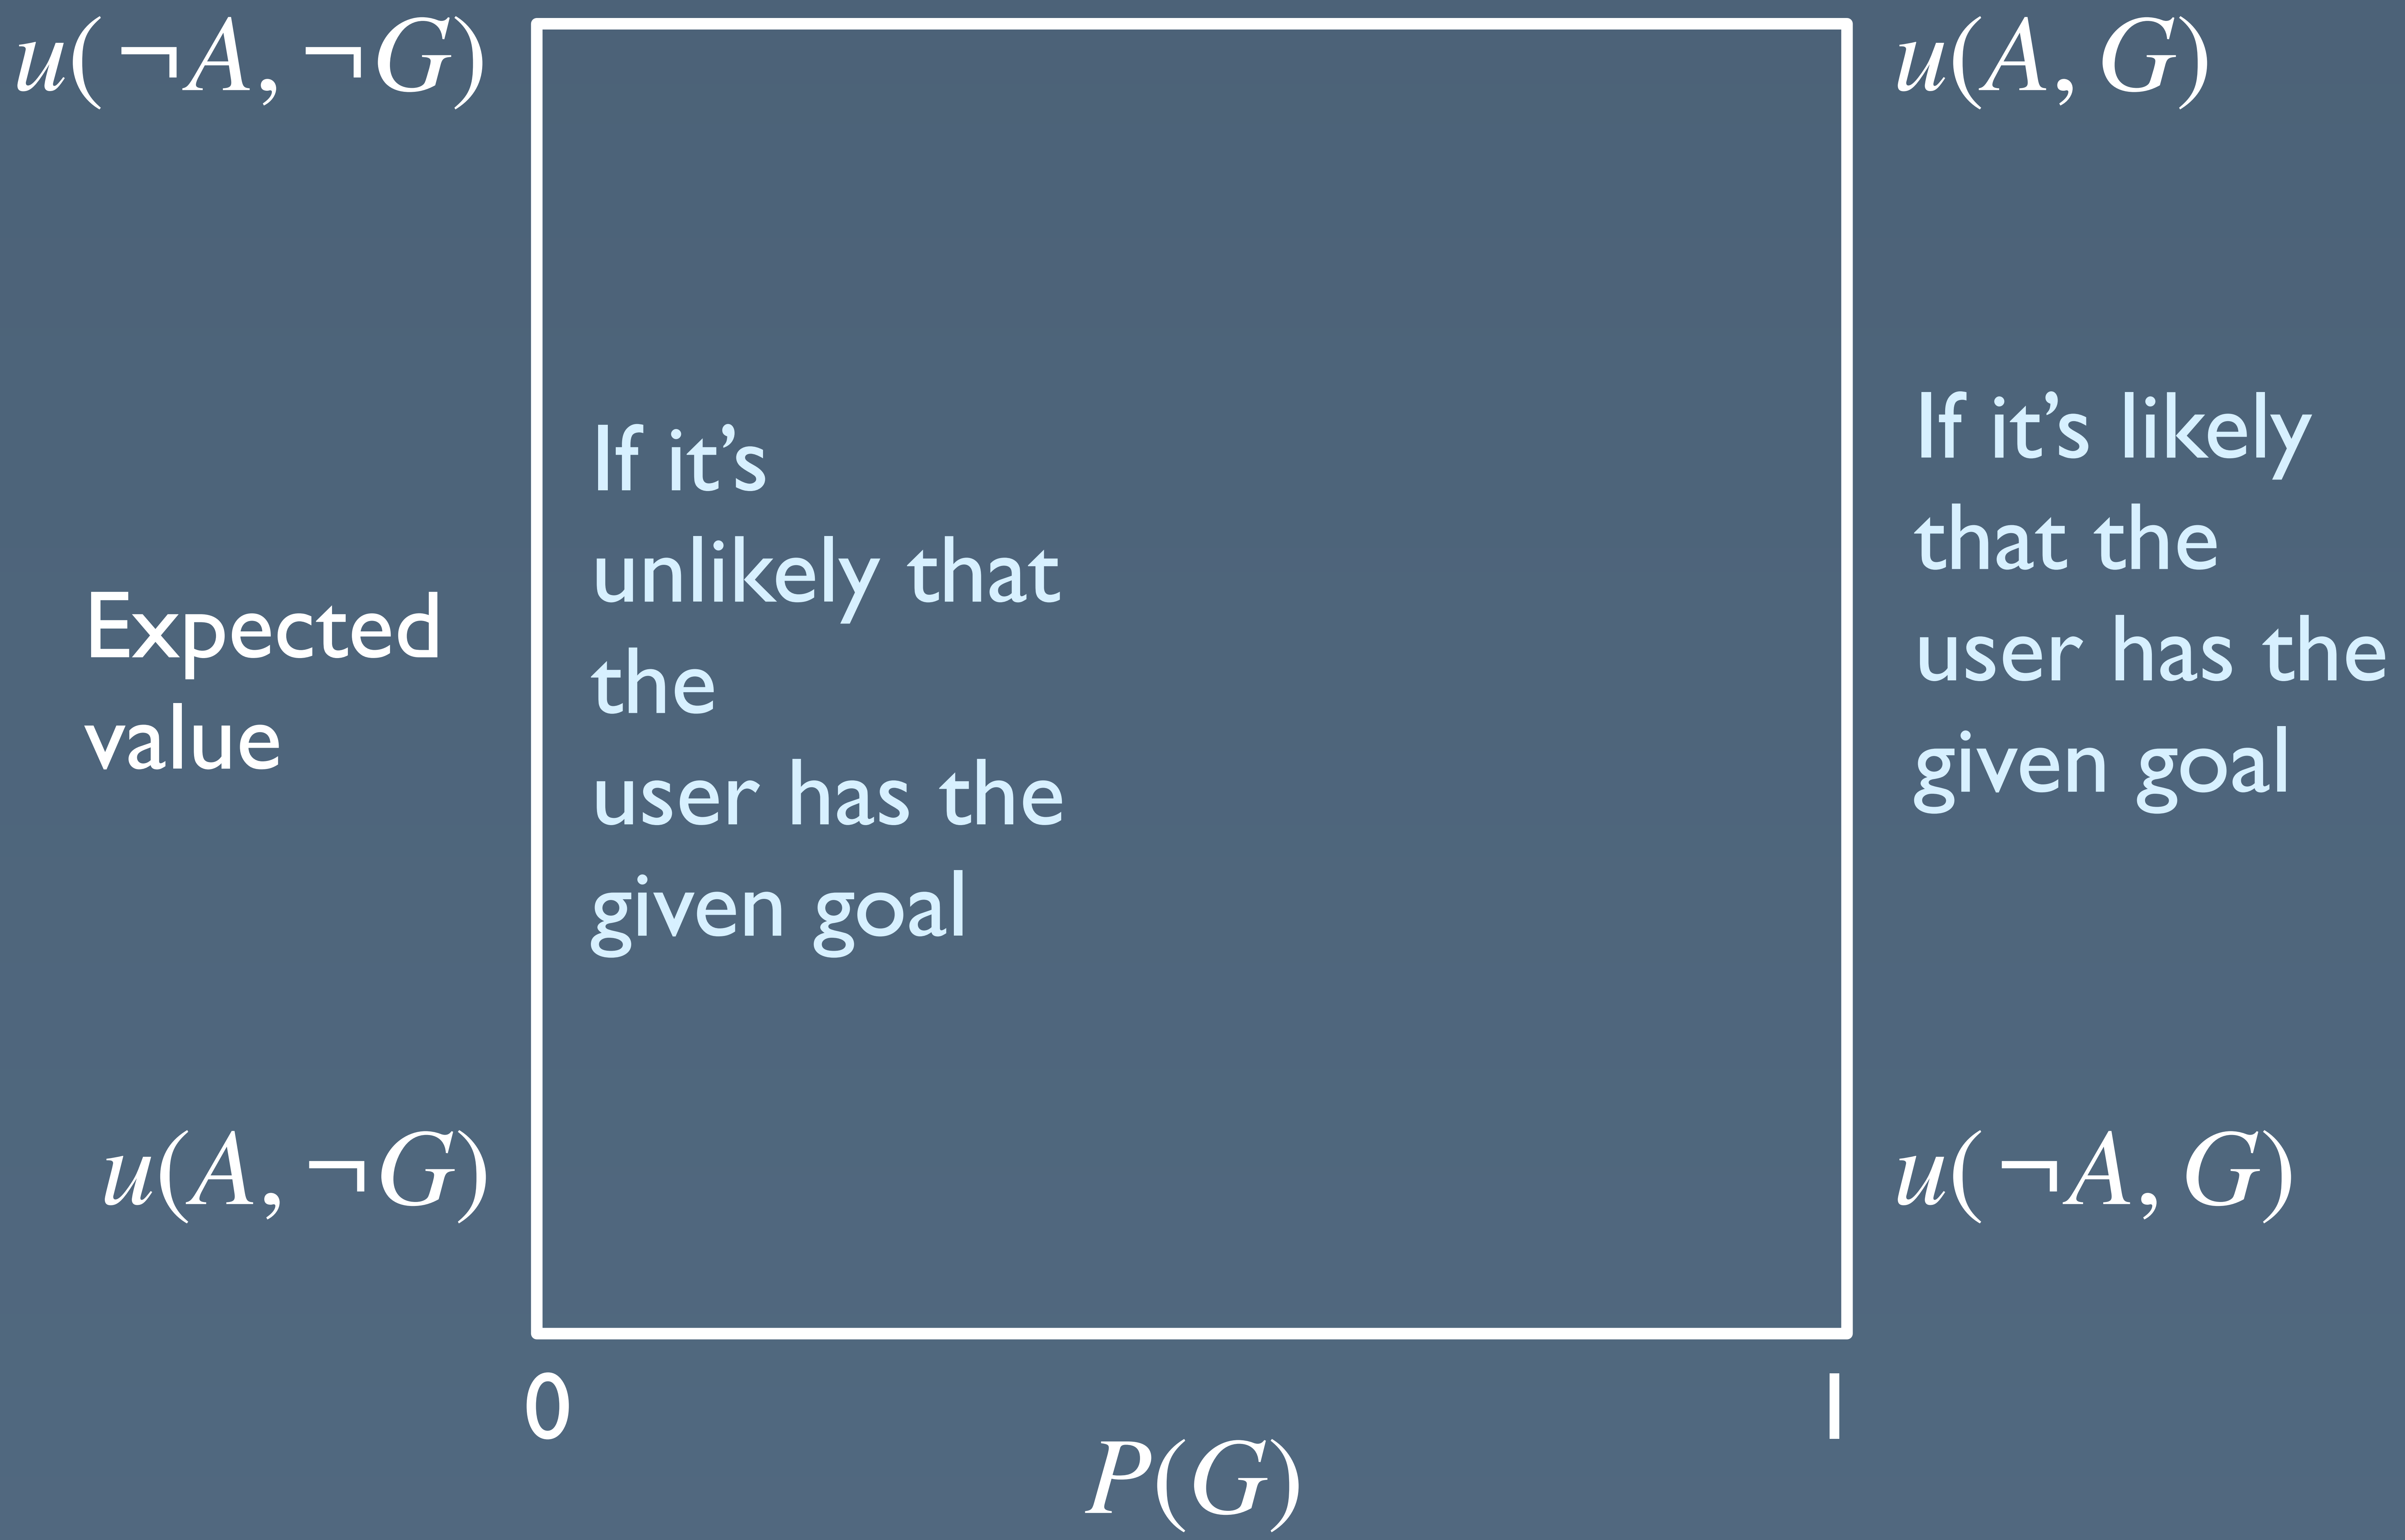

Mixed initiative: visually

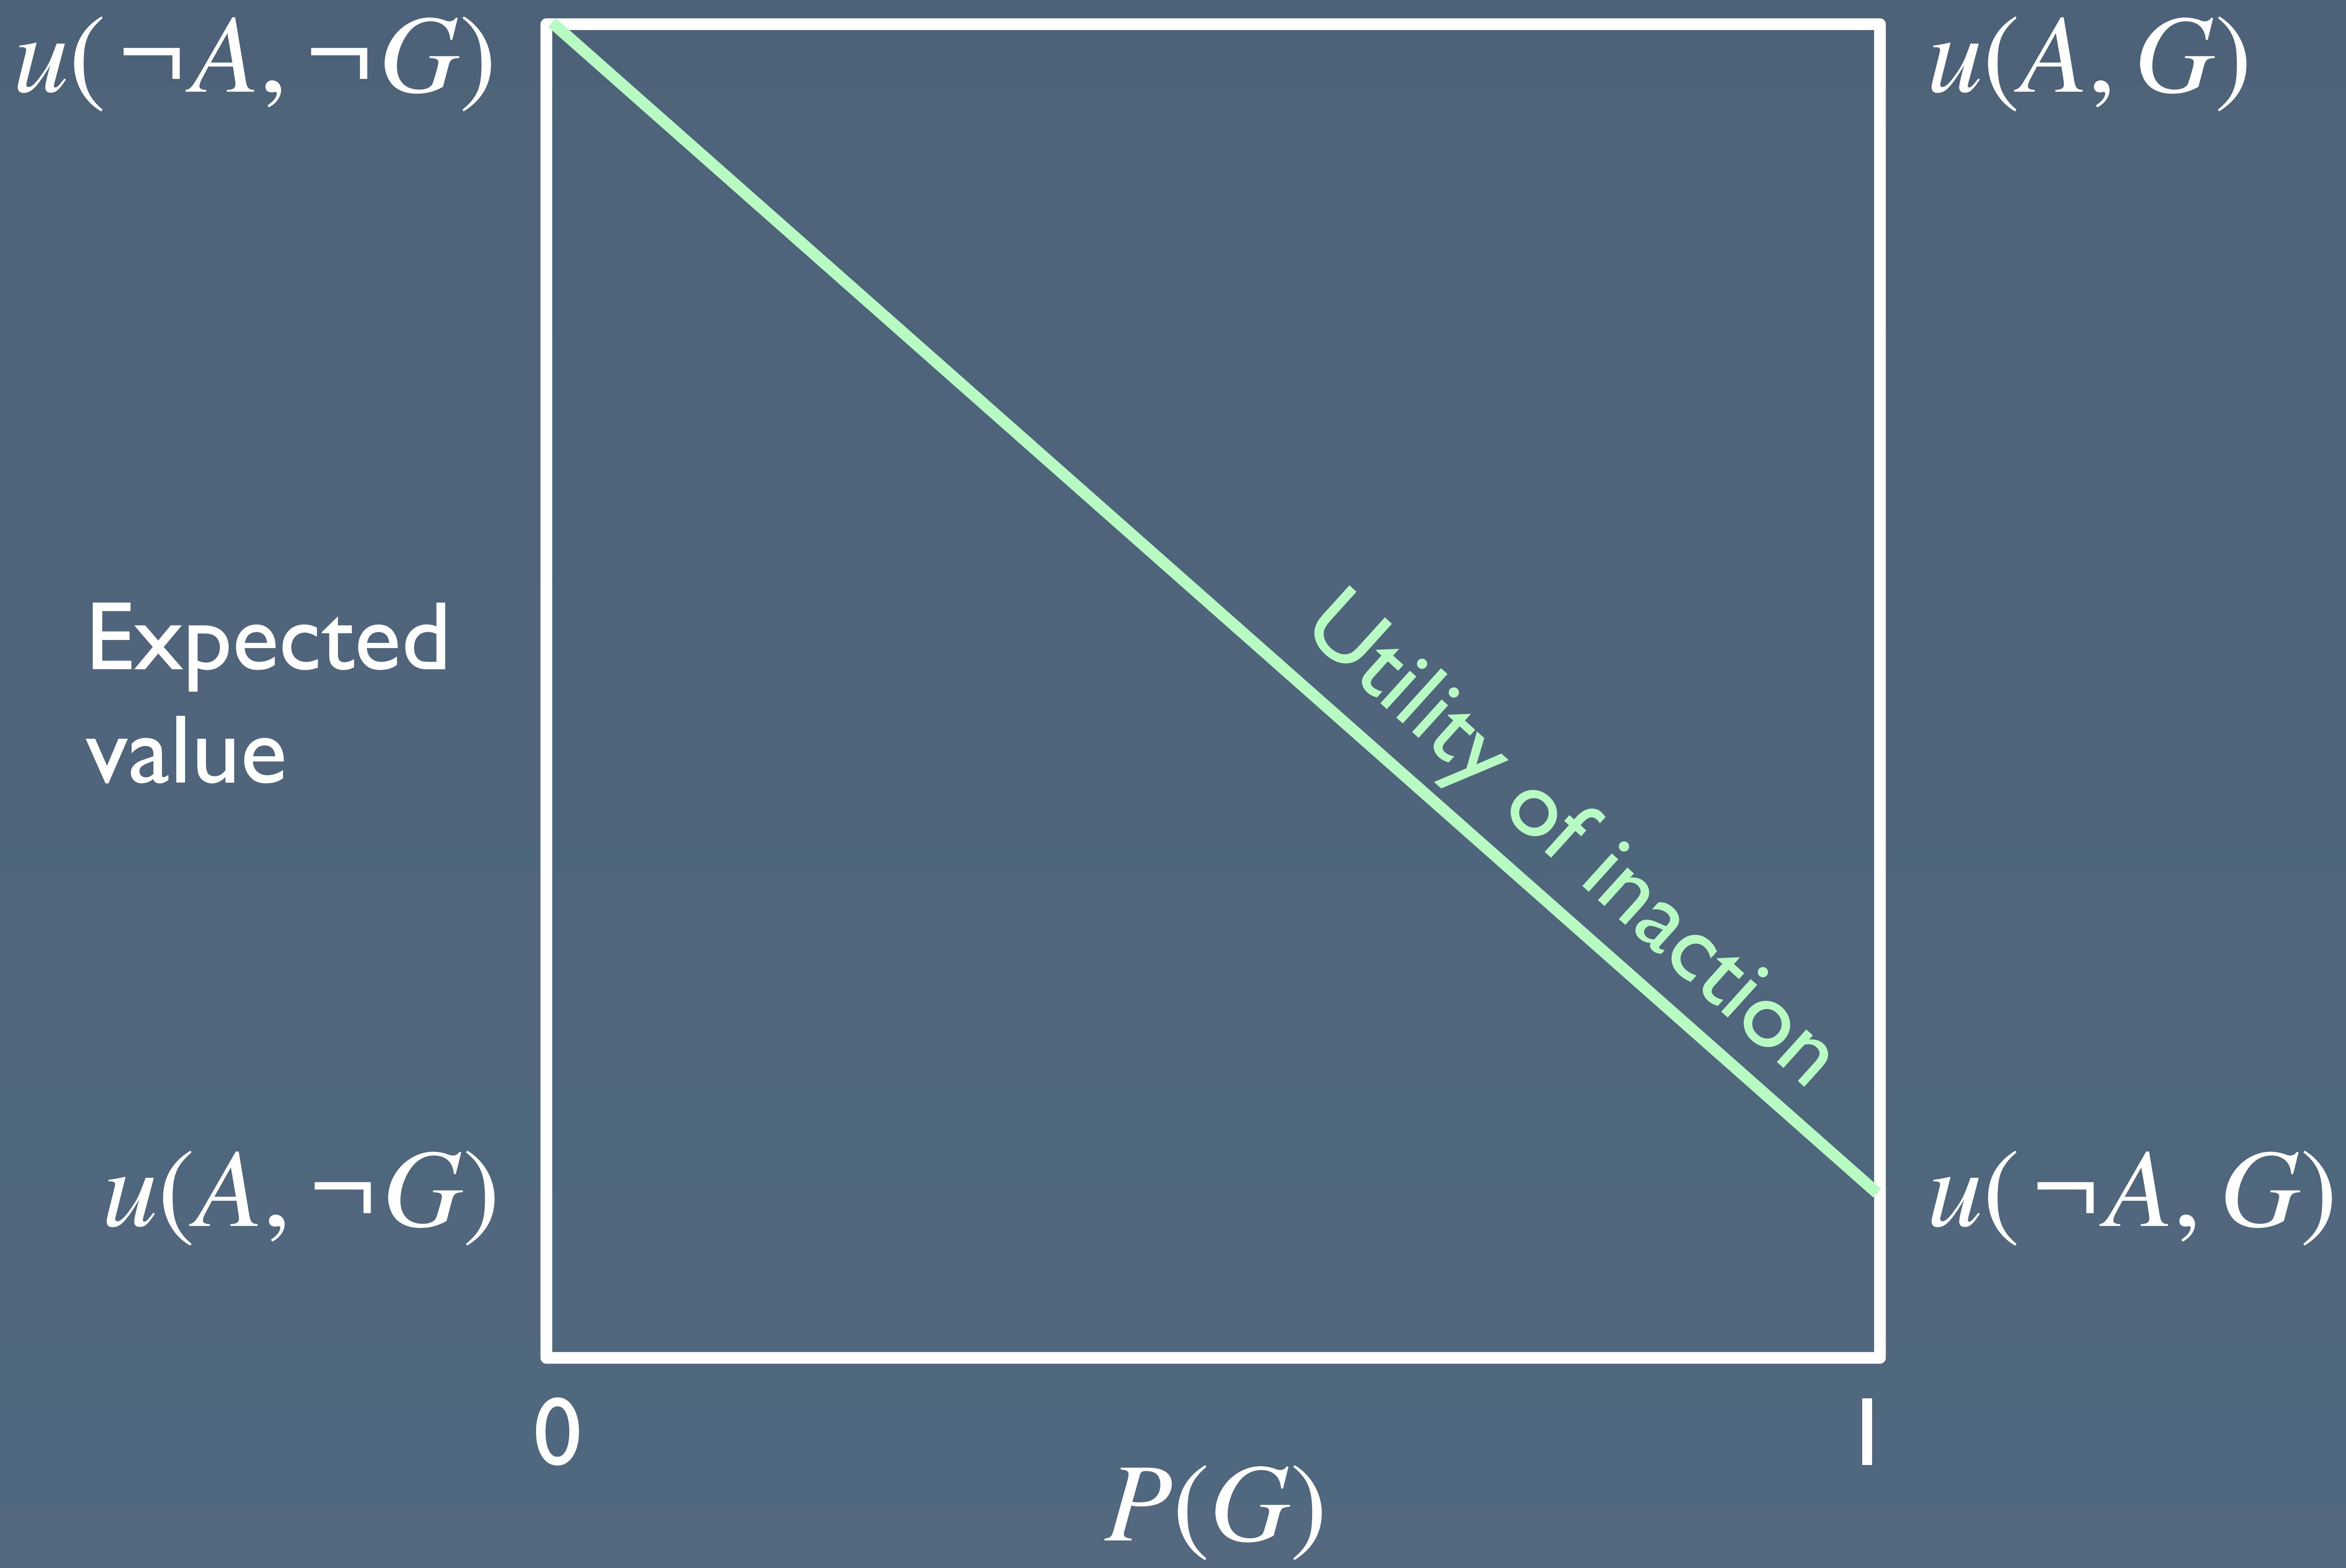

Mixed initiative: visually

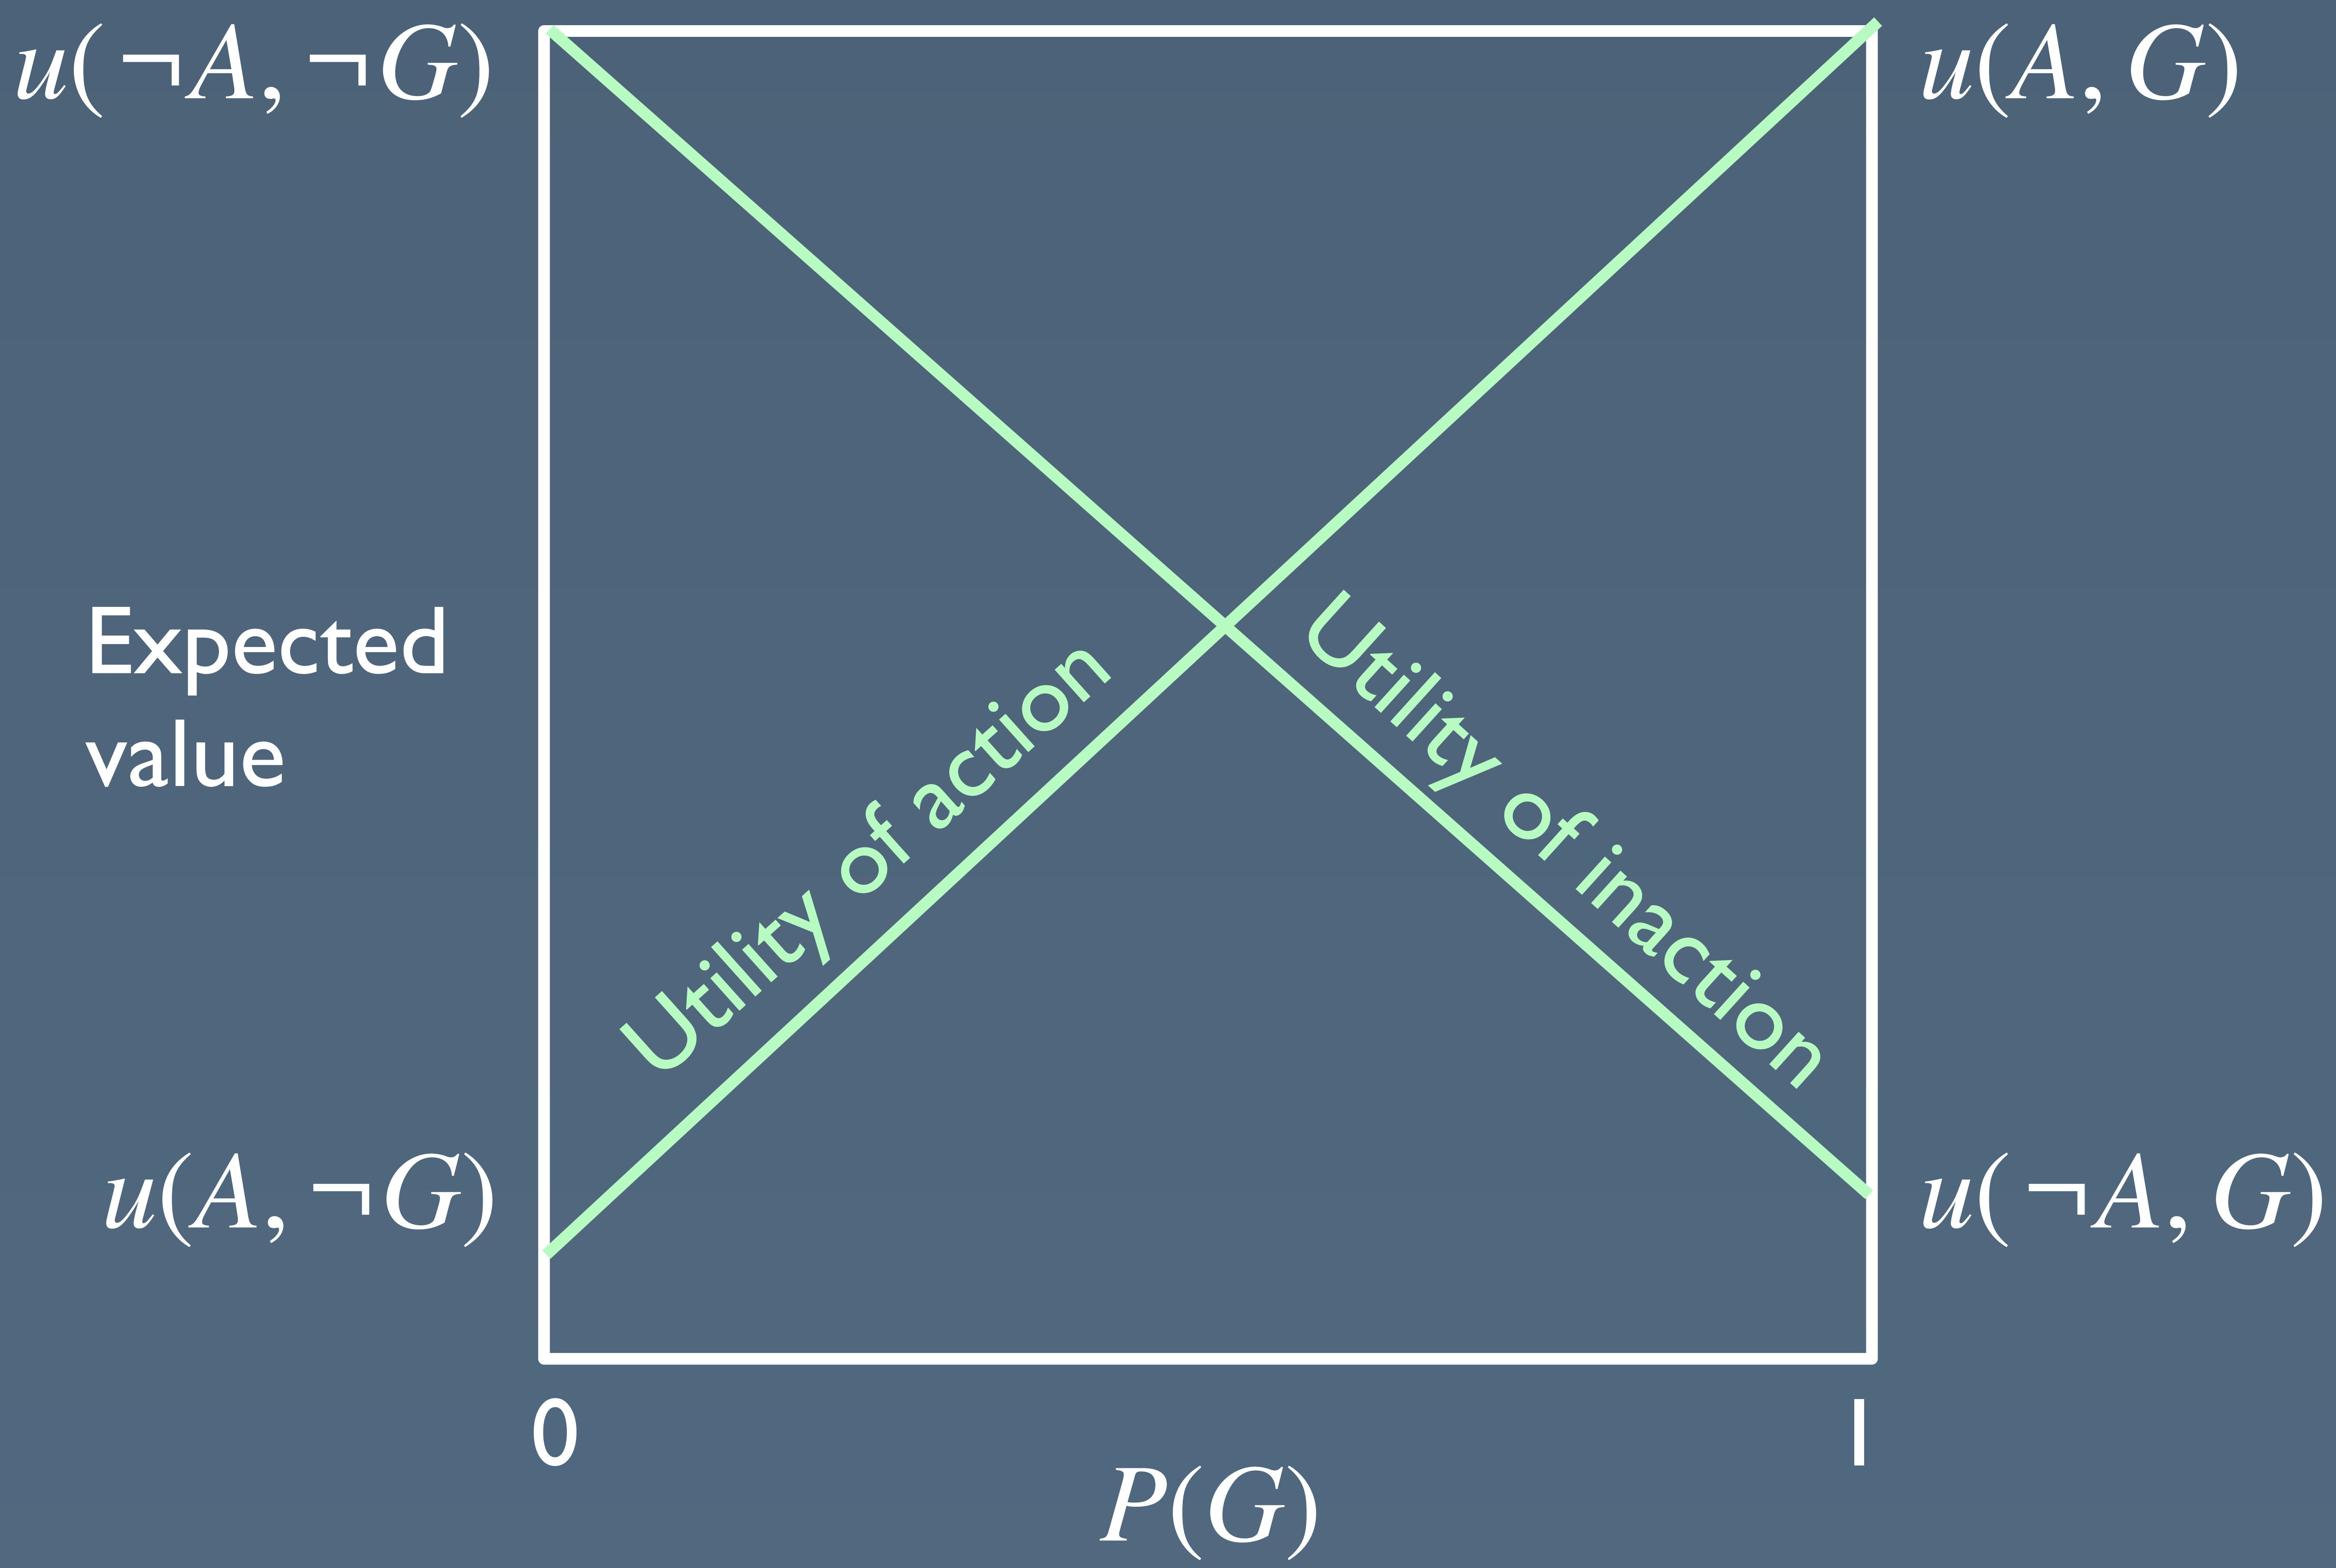

Mixed initiative: visually

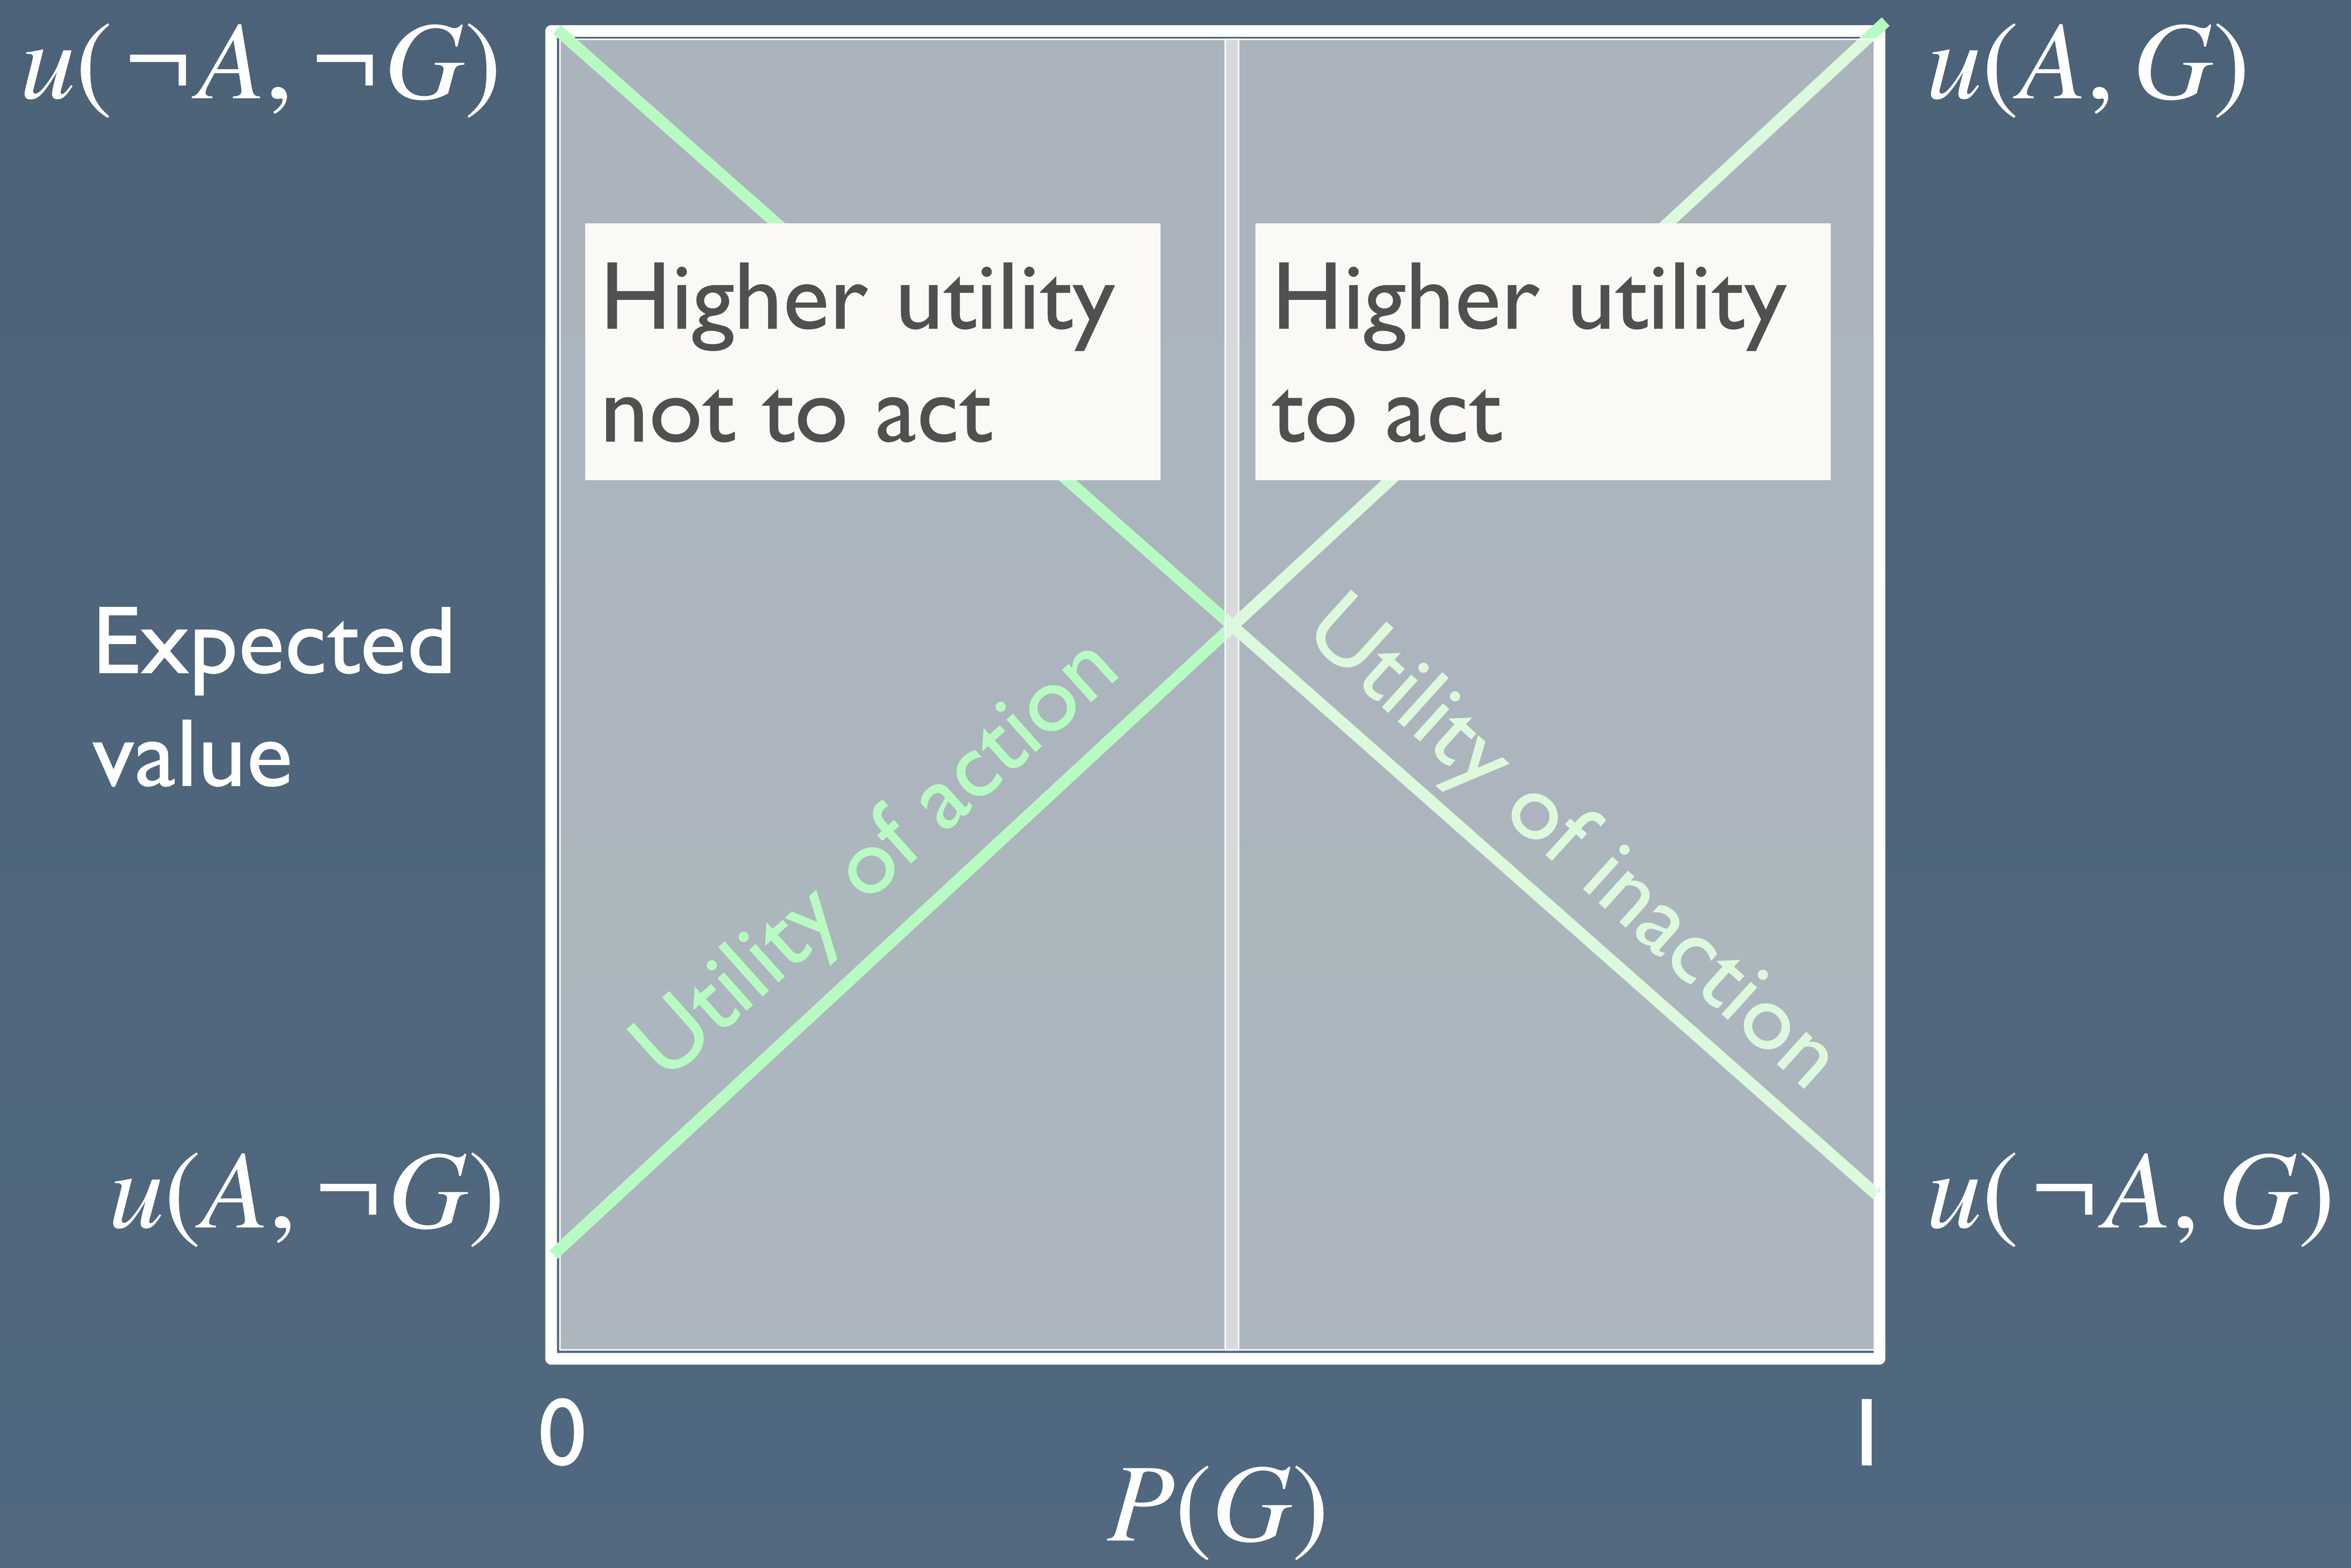

What if making an error is costly?

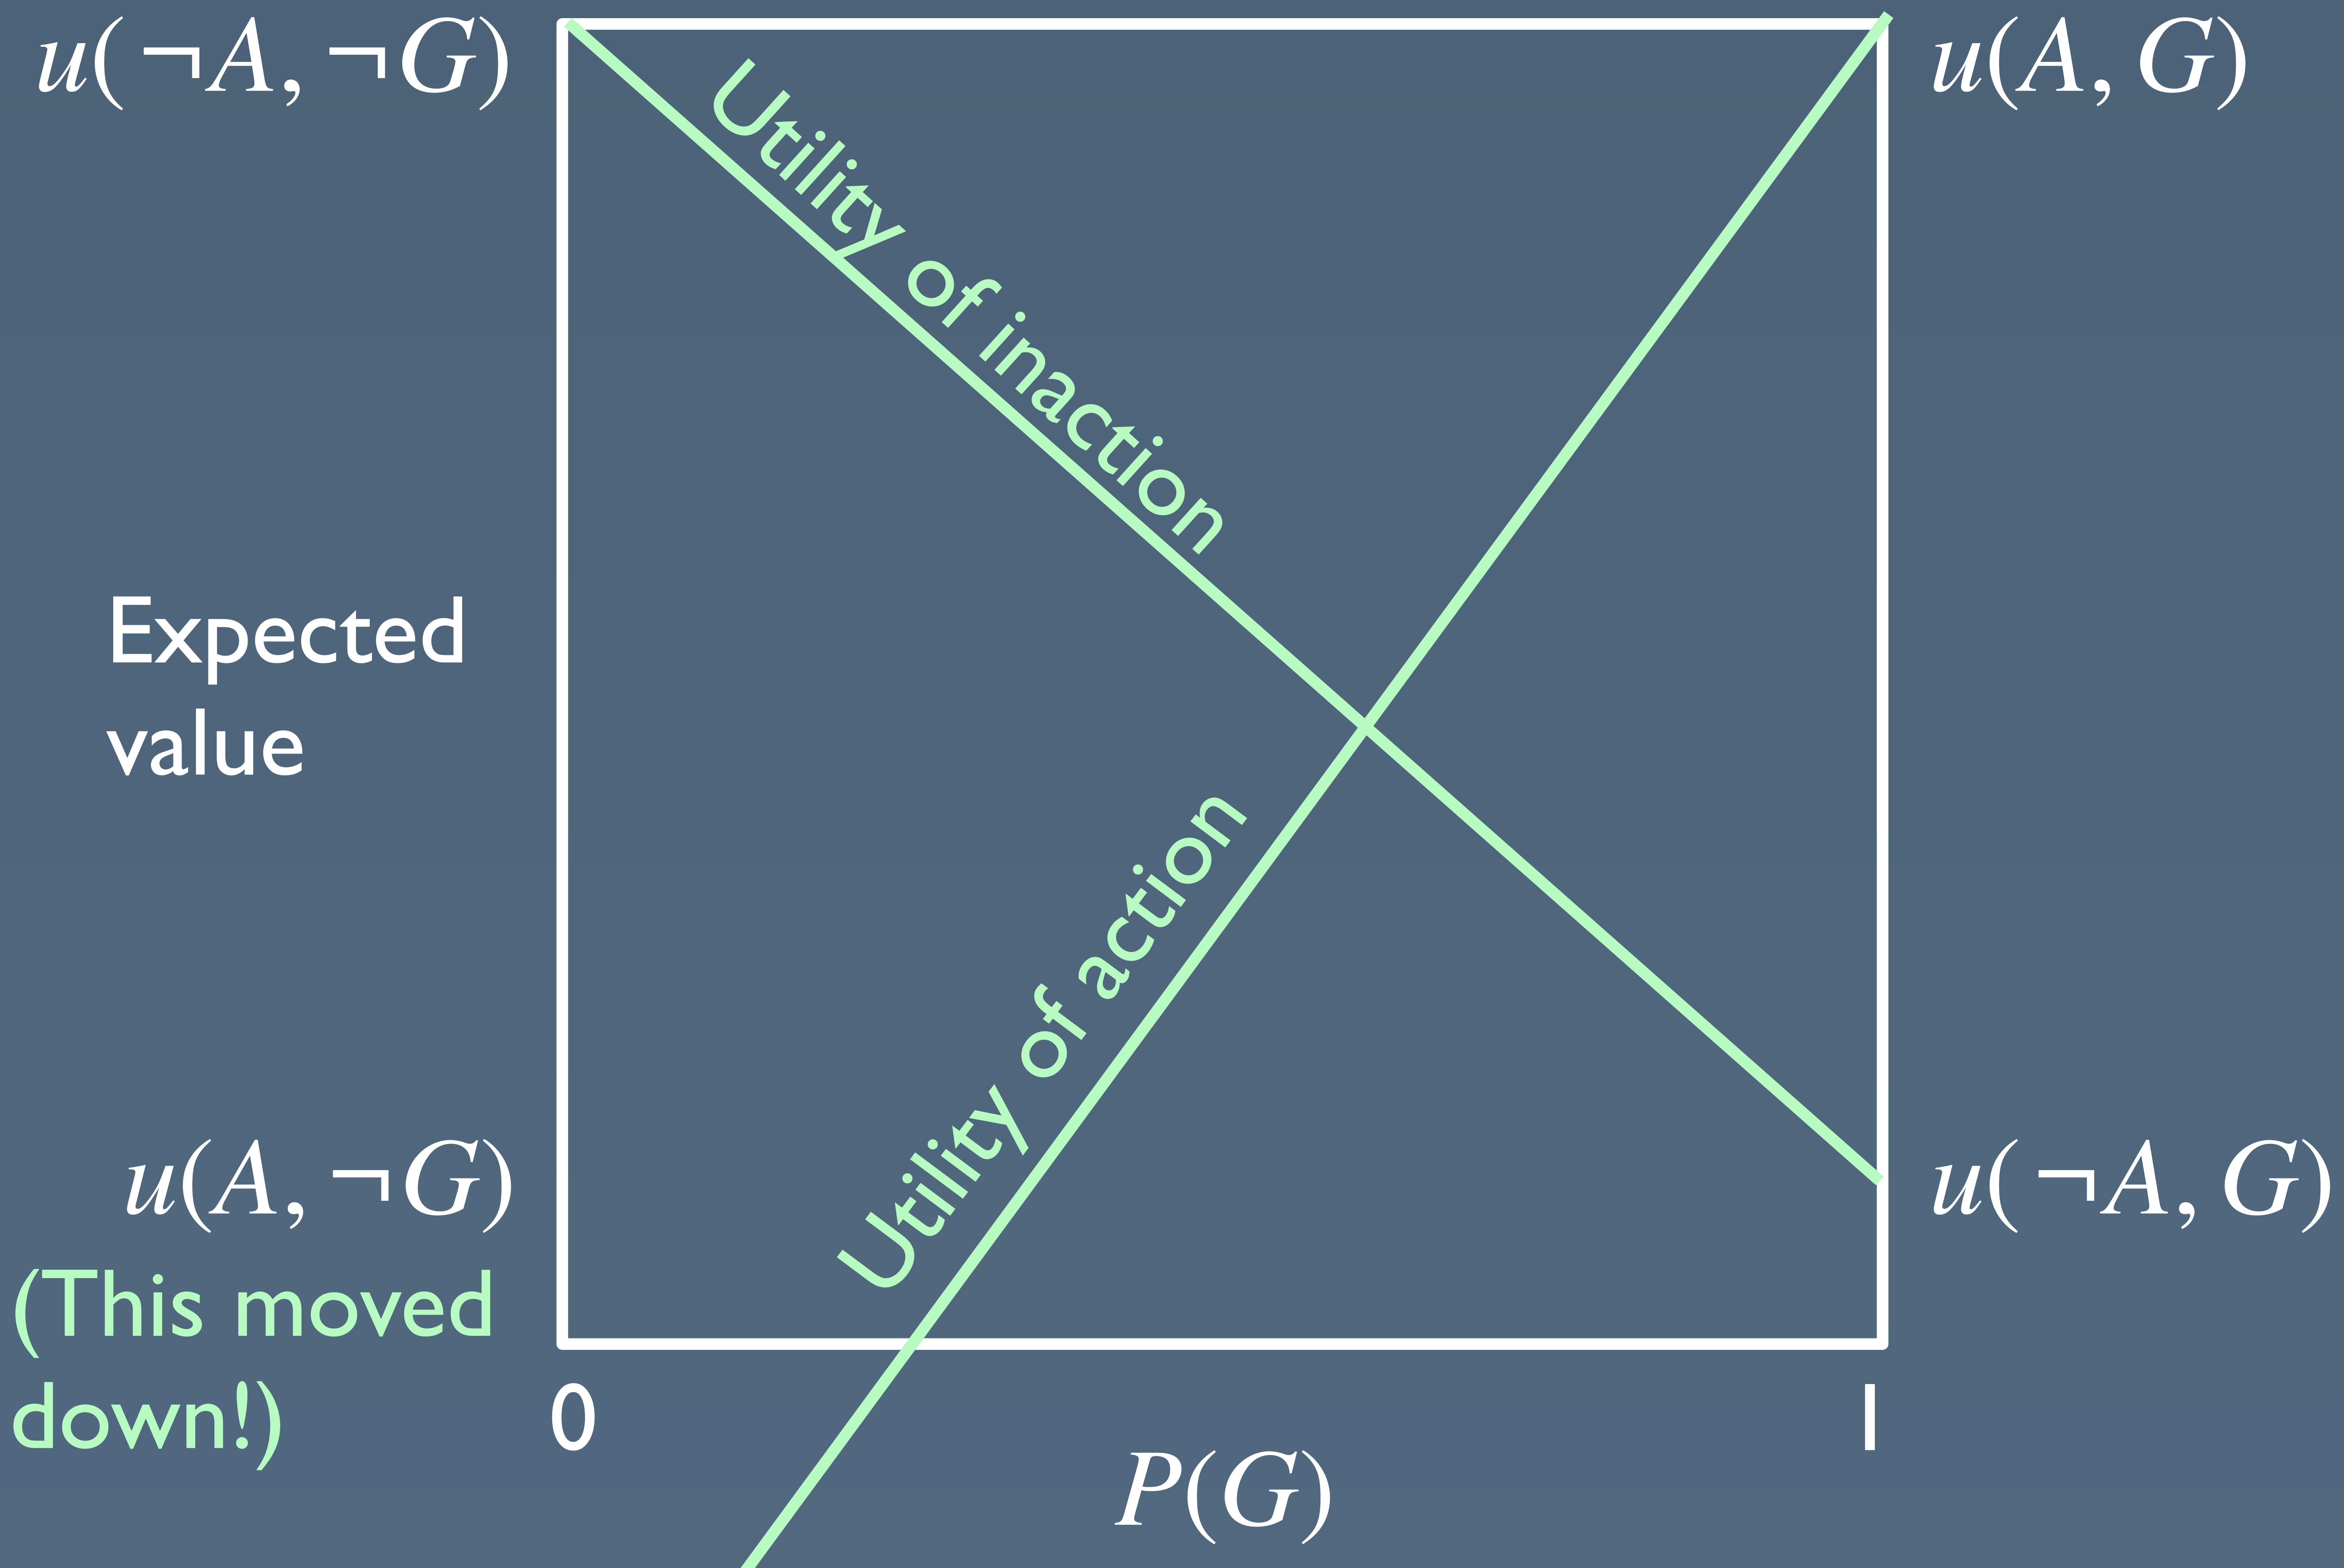

What if making an error is costly?

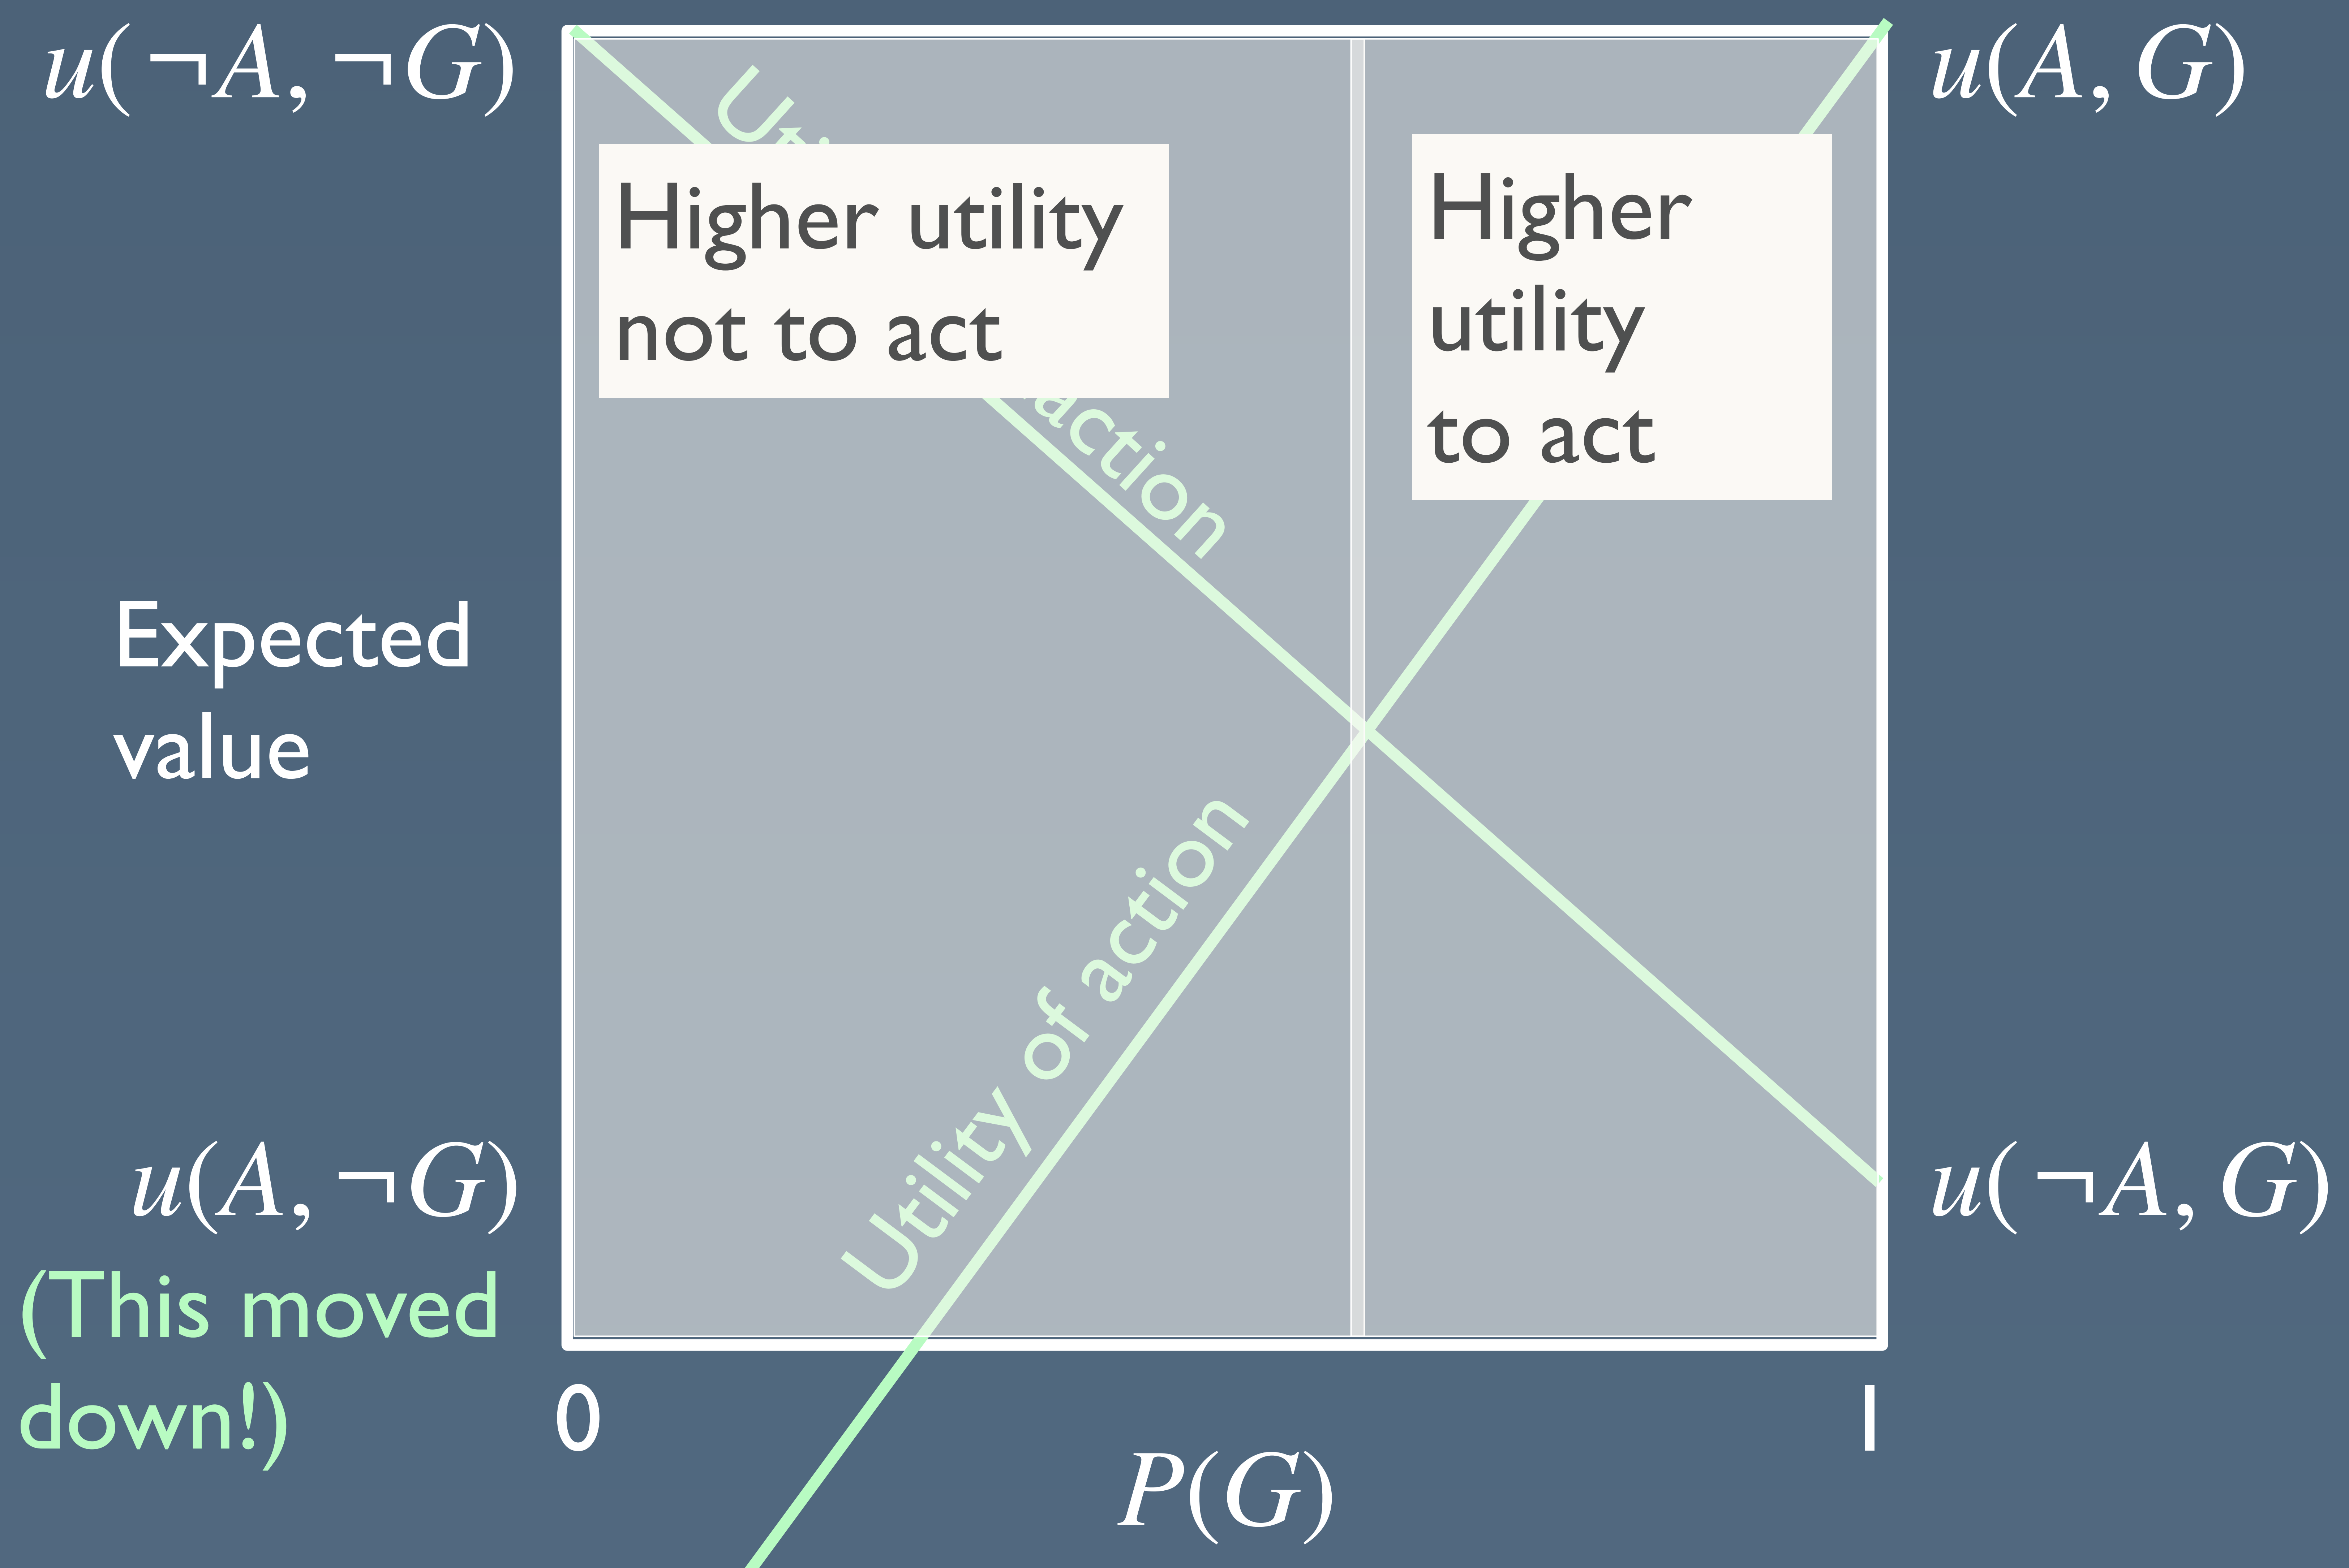

Now we only take action if we are even more certain that we correctly estimated the user's goal

What if we ask the user?

Asking often carries lower risk, but also lower utility

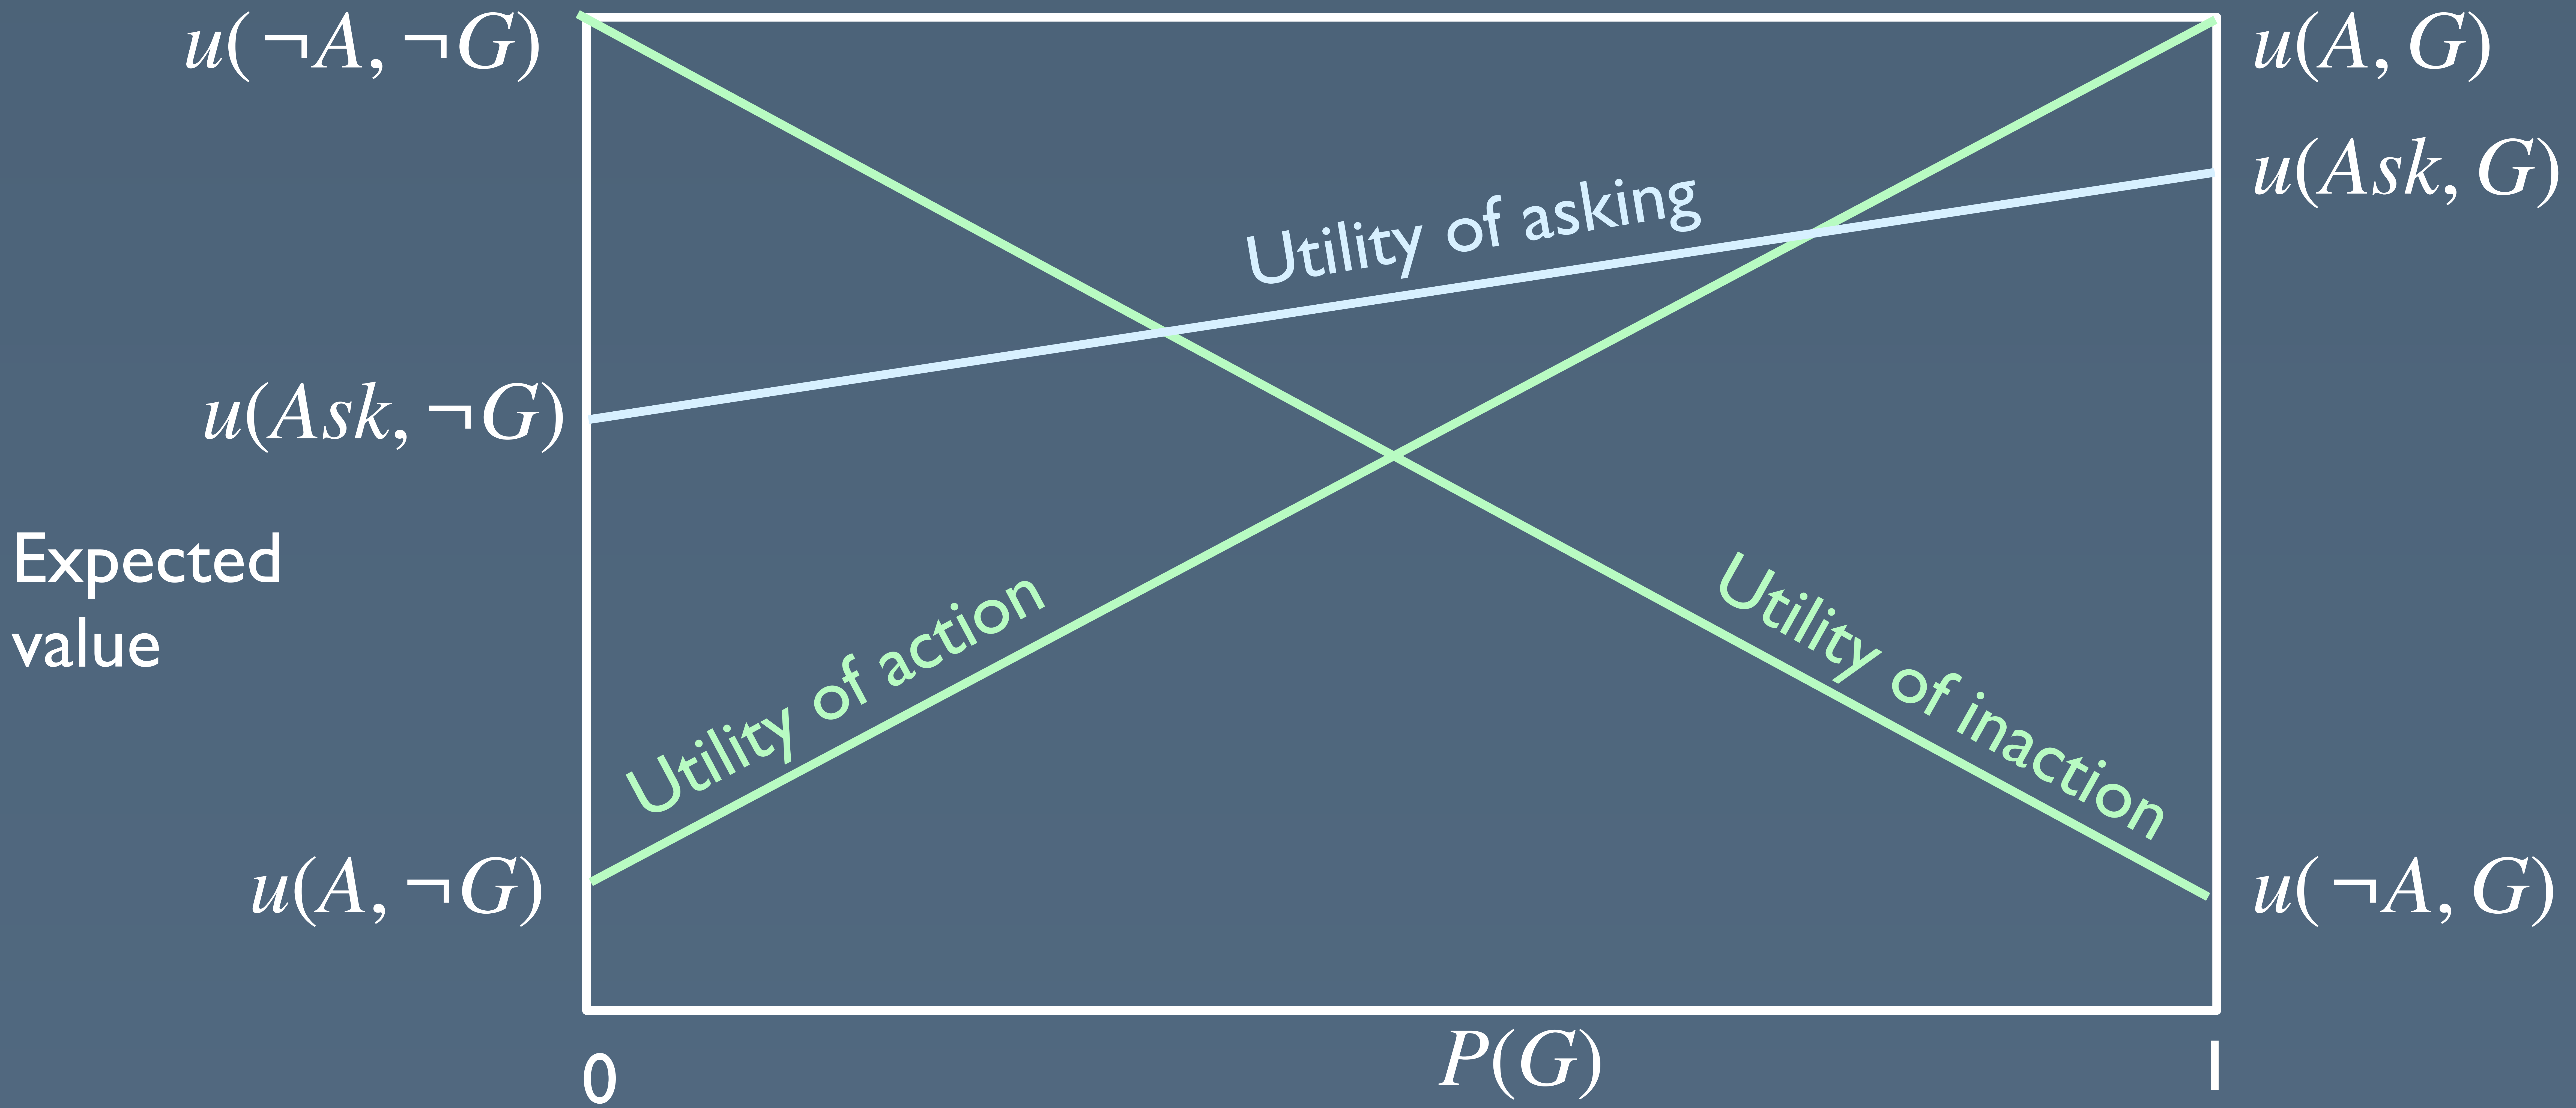

What if we ask the user?

Asking often carries lower risk, but also lower utility

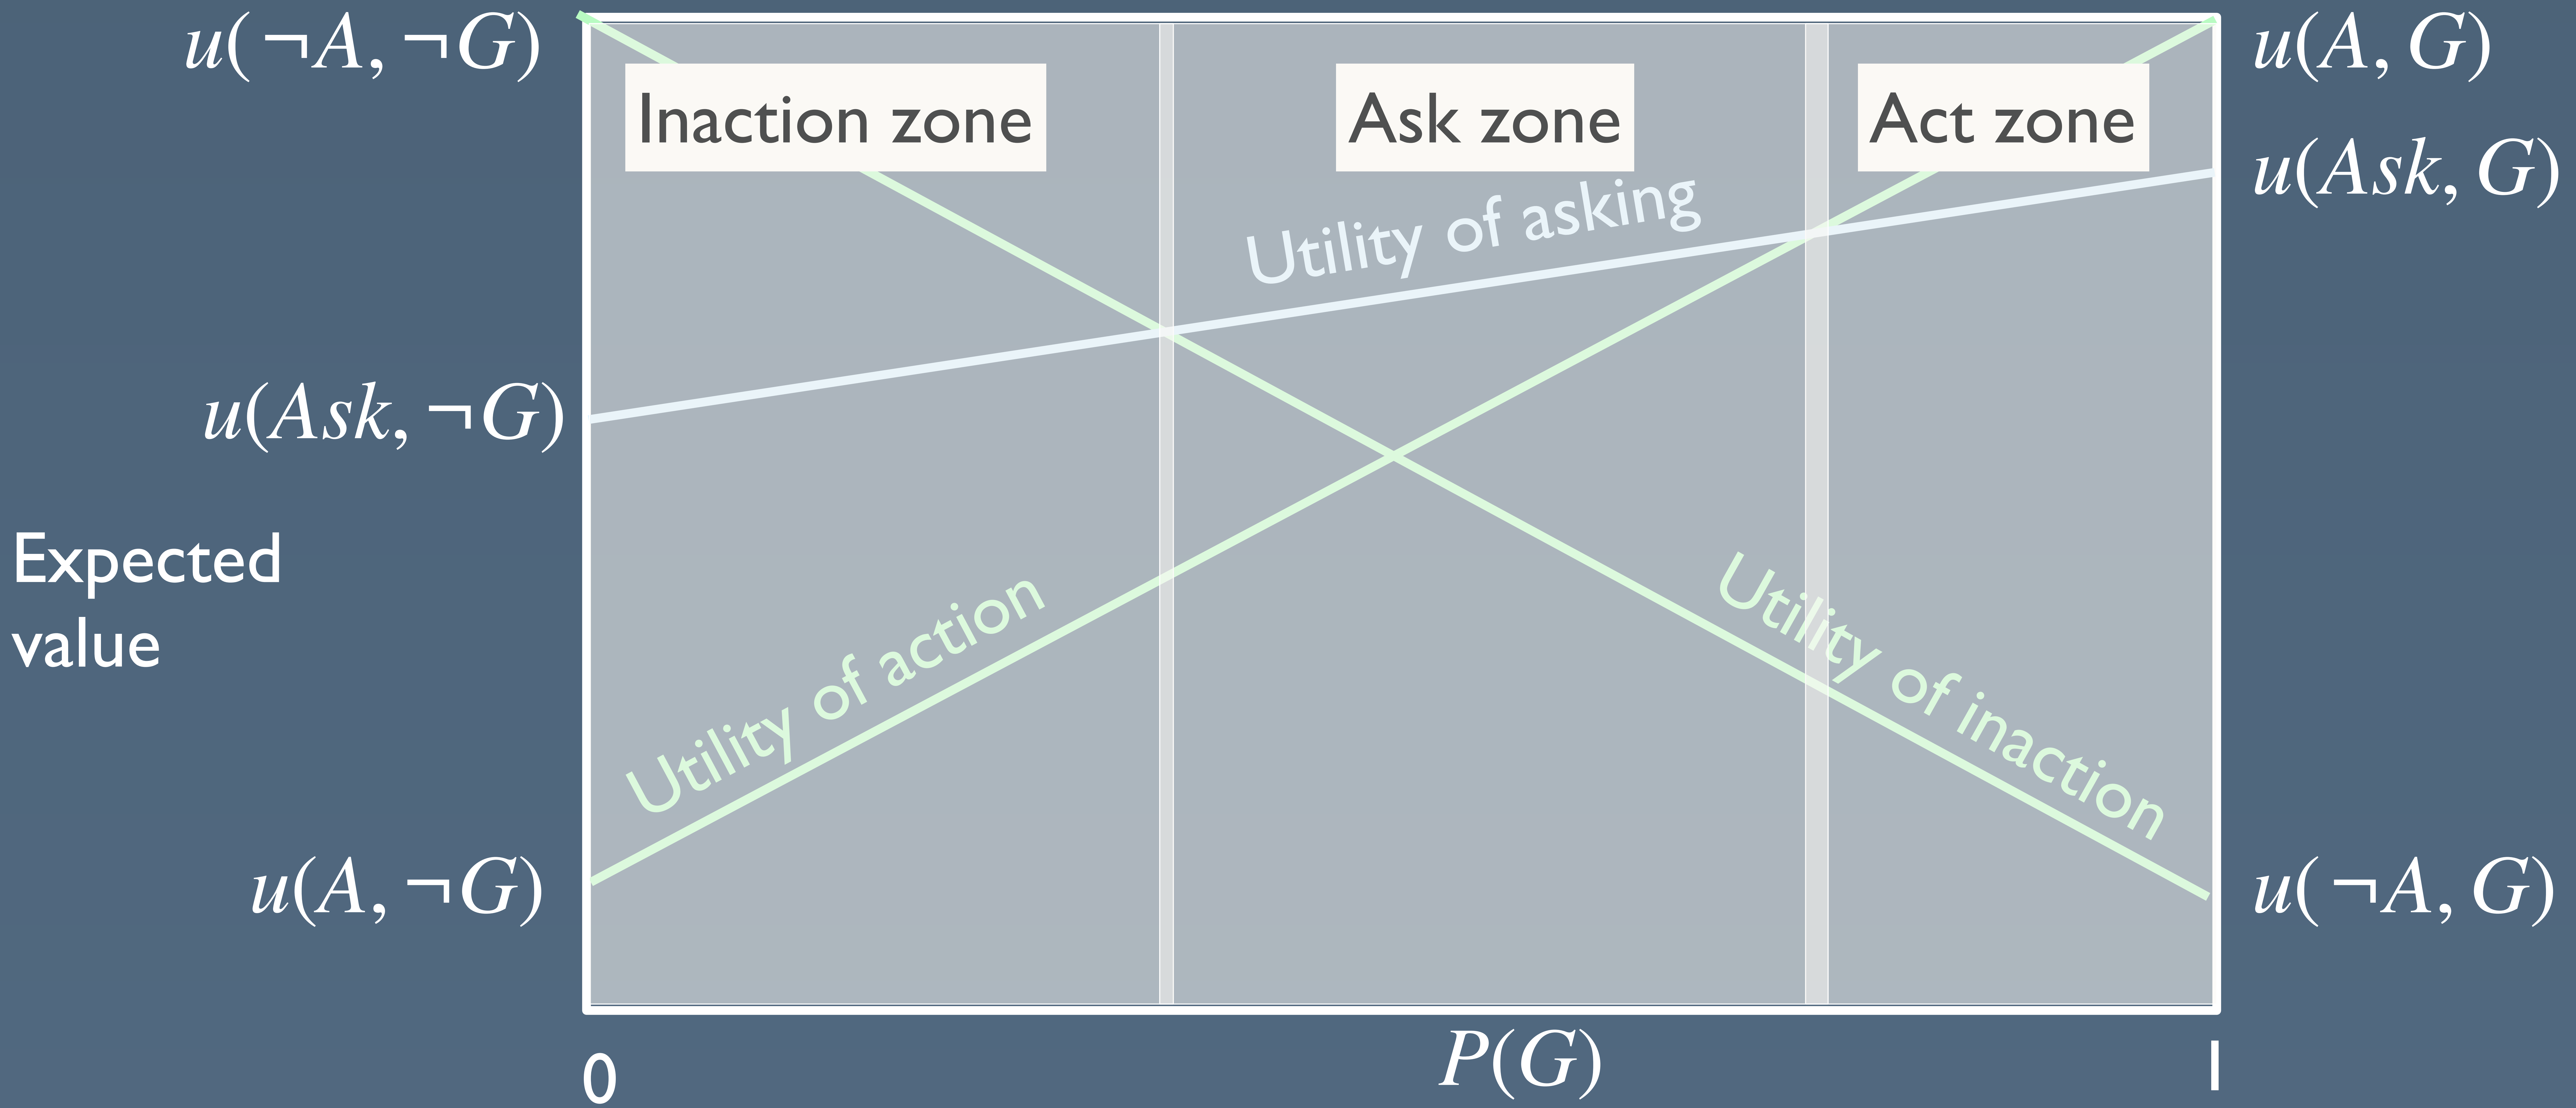

So, when does this screw up?

- When the system cannot accurately assess the probability of the user having the goal P(G)

  or

- When the utilities are not correctly estimated
  e.g., too high a utility for asking if the user doesn't have the goal G. "Are you writing a letter right now?"

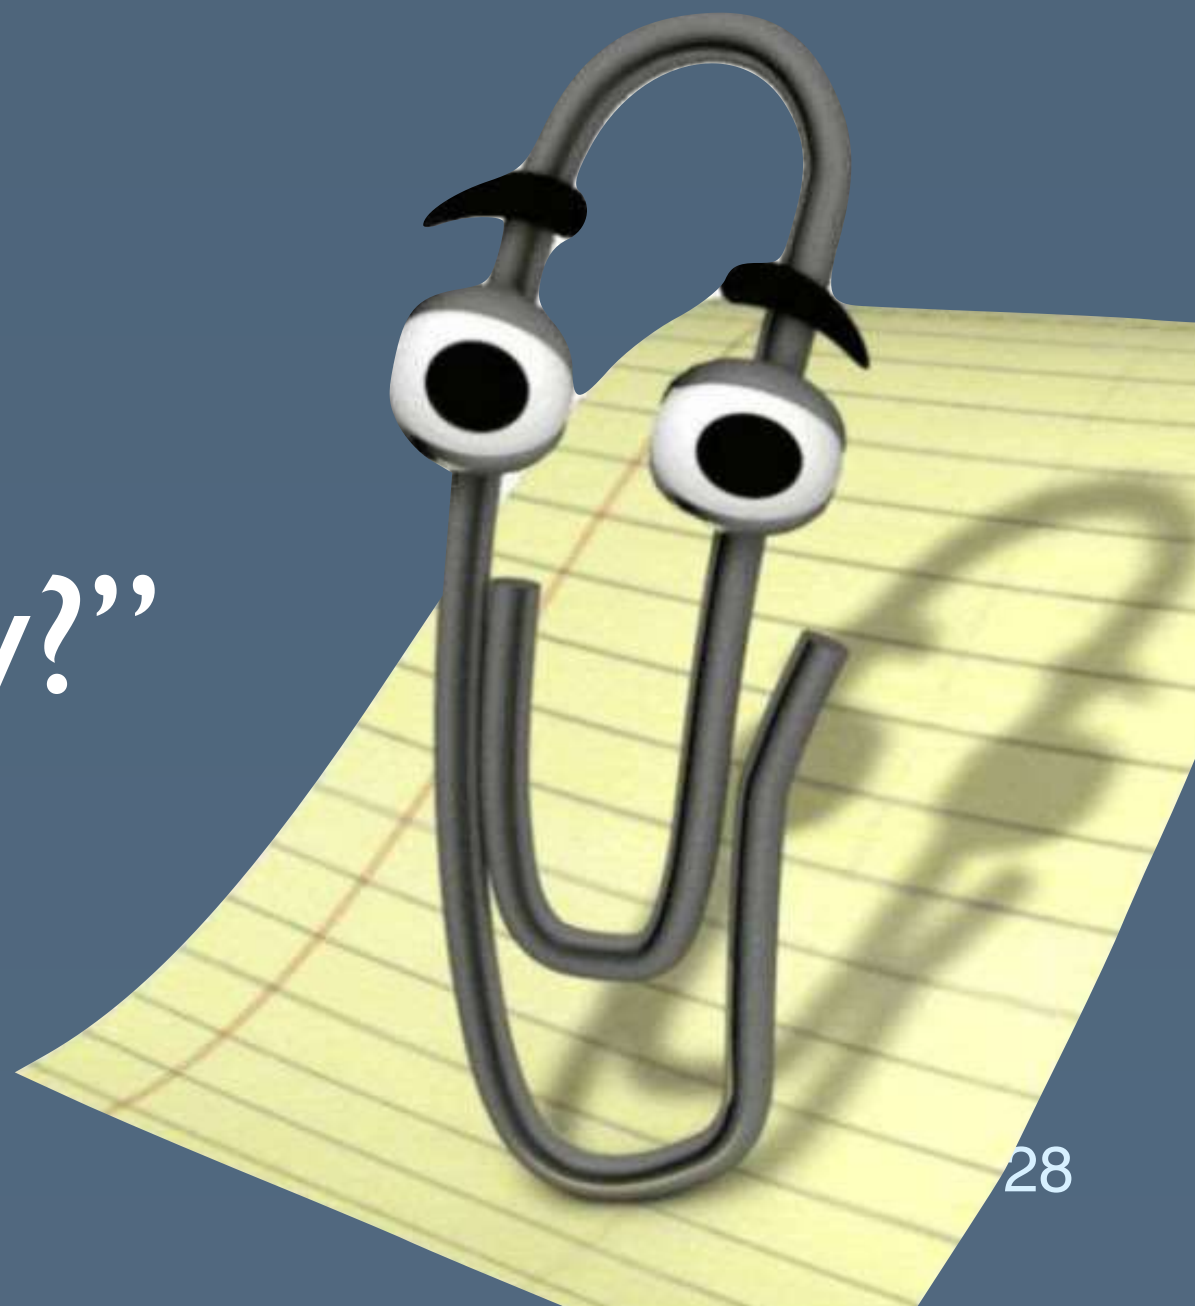

Interactive Machine Learning

Now: diving into the ML models

- AI systems are built on the back of machine learning models
- What lessons can we apply to make these models more powerful and effective?

- Assumed here: a basic knowledge of machine learning

What is your black box learning?

- Machine learning models are powerful, but opaque and unintelligible
- Difficult to predict, design, and debug

- This produces nonintuitive behavior:

[Goodfellow, Shlens, Szegedy 2014]

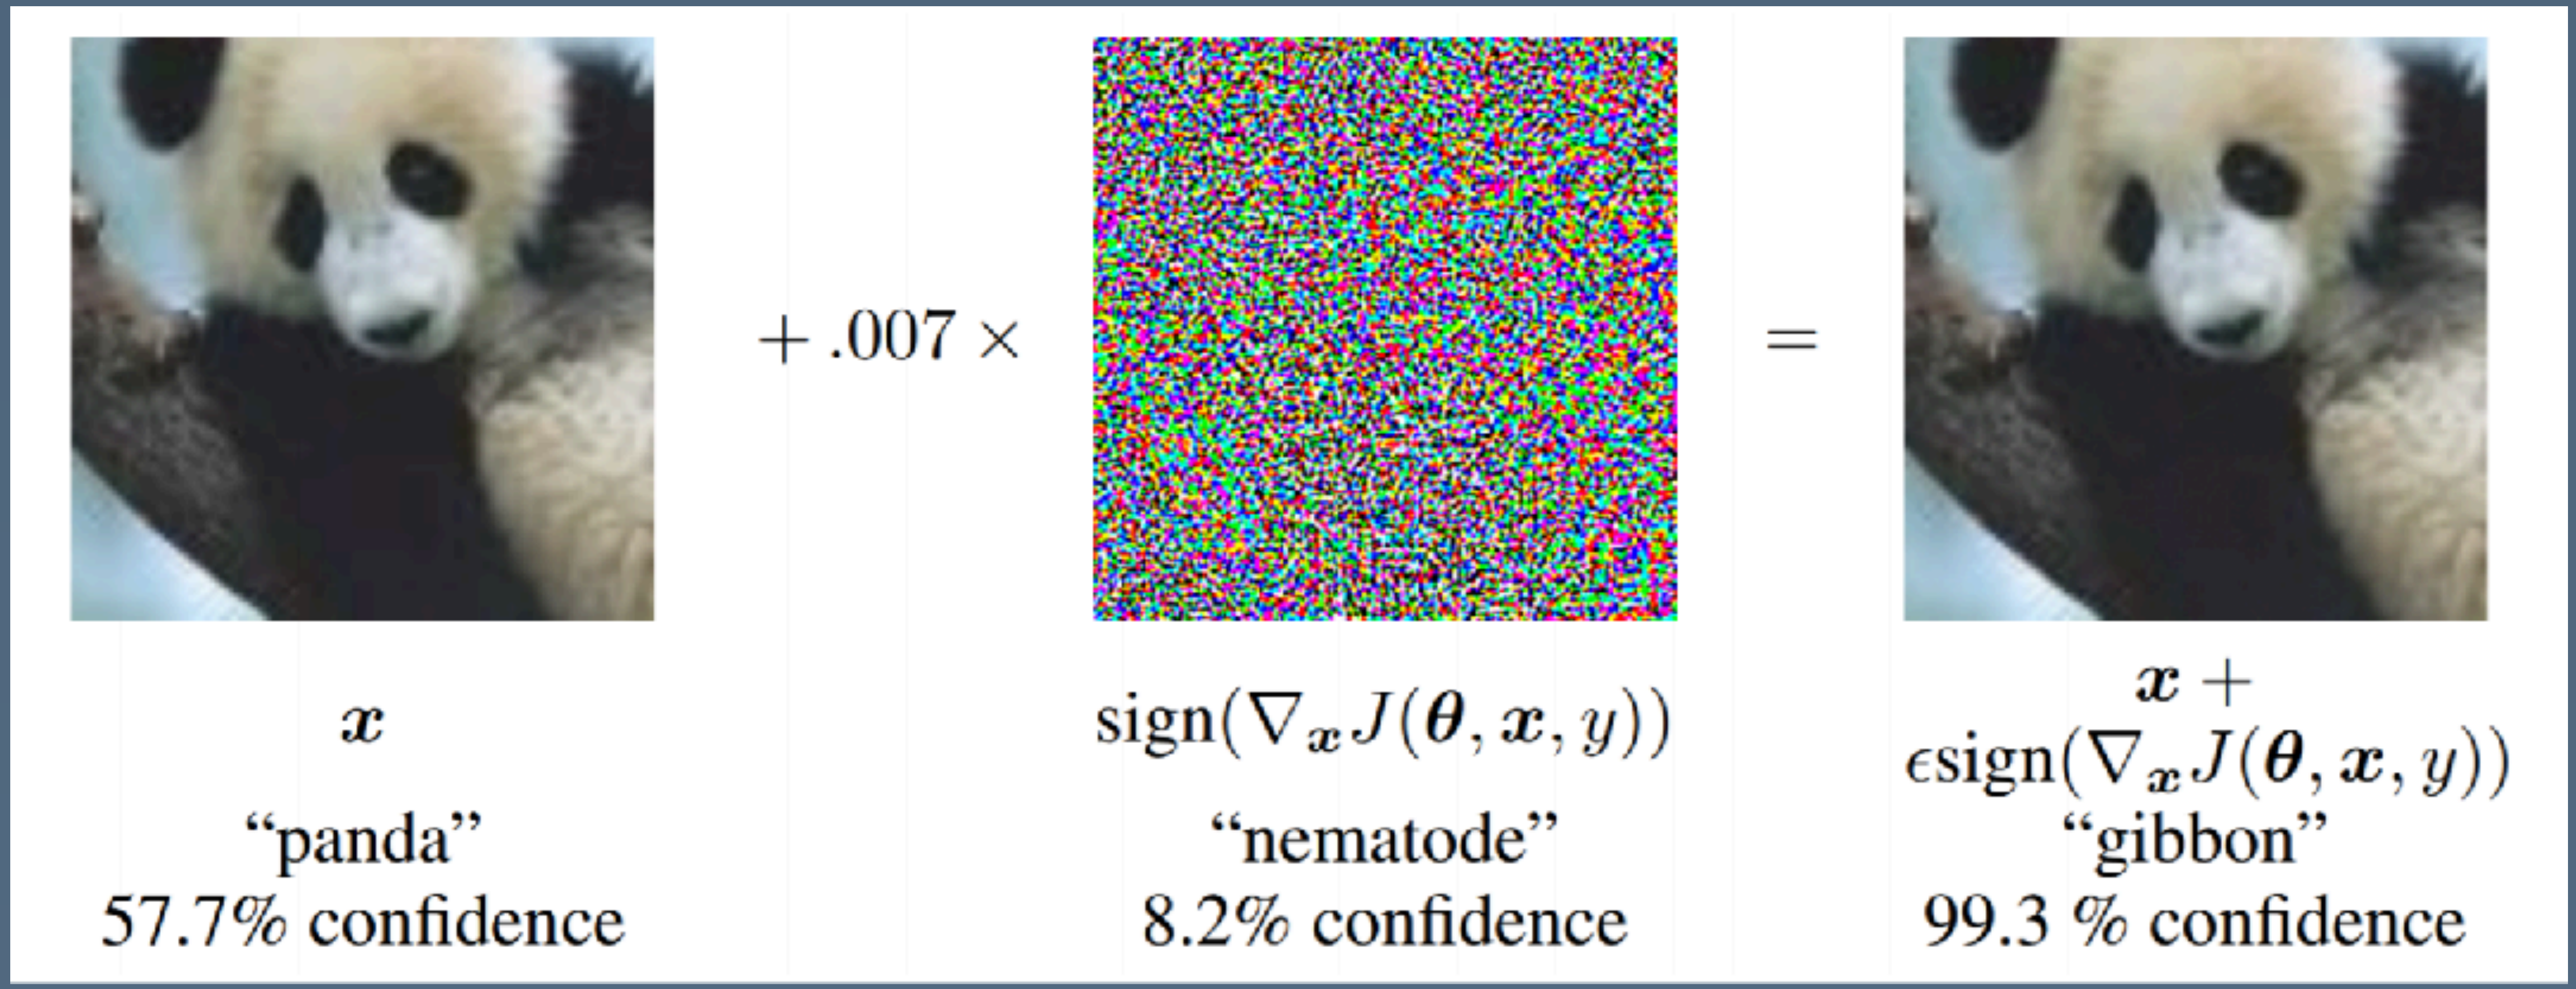

Goal: intelligibility in ML models

- A model is *intelligible* to the extent that a human can predict how a change to model's inputs will change its output. [Weld 2018]

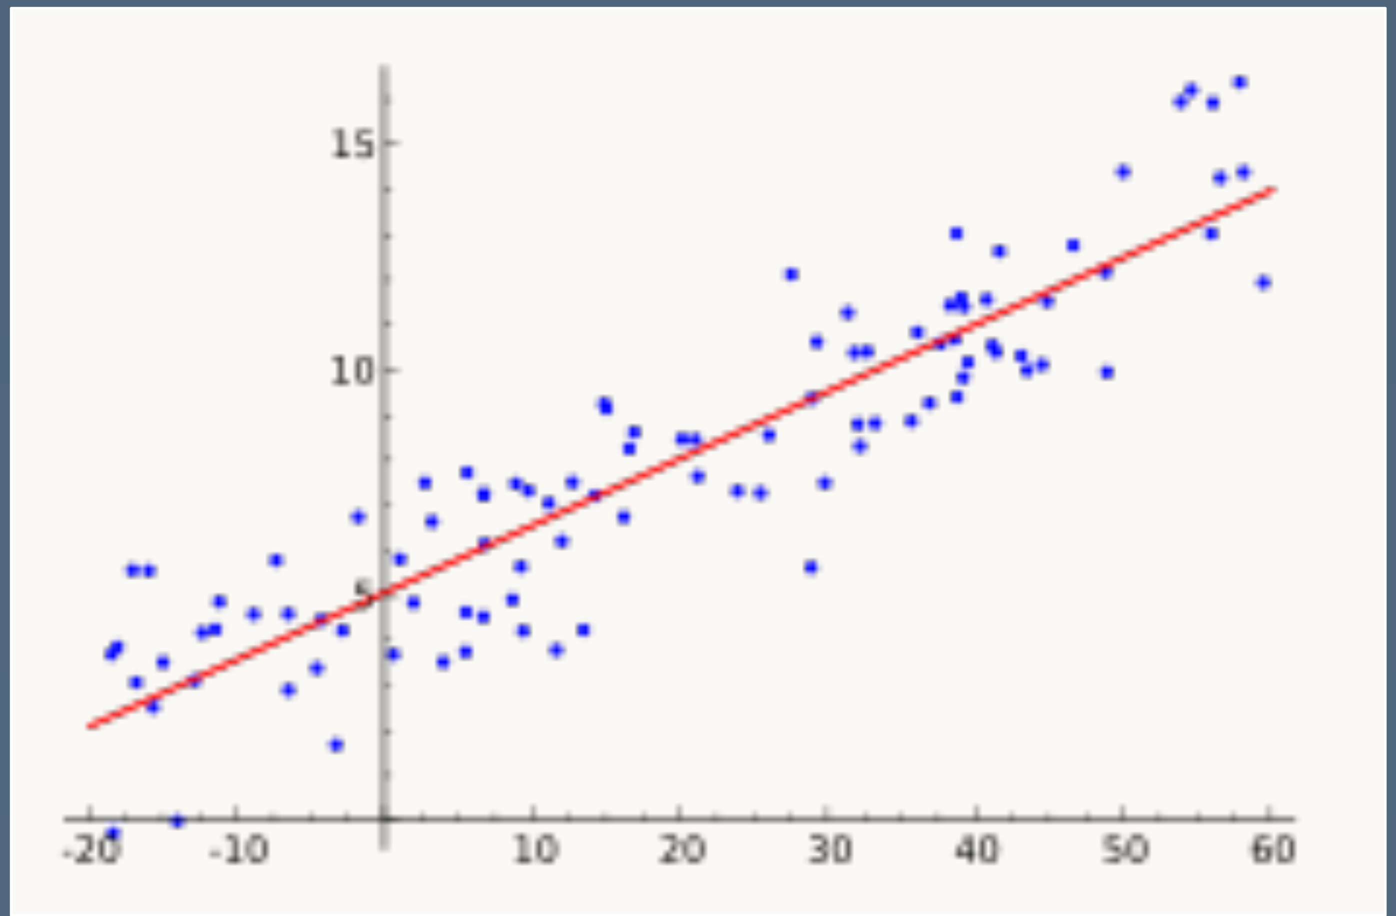

Linear relationship: intelligible
[image from Wikipedia]

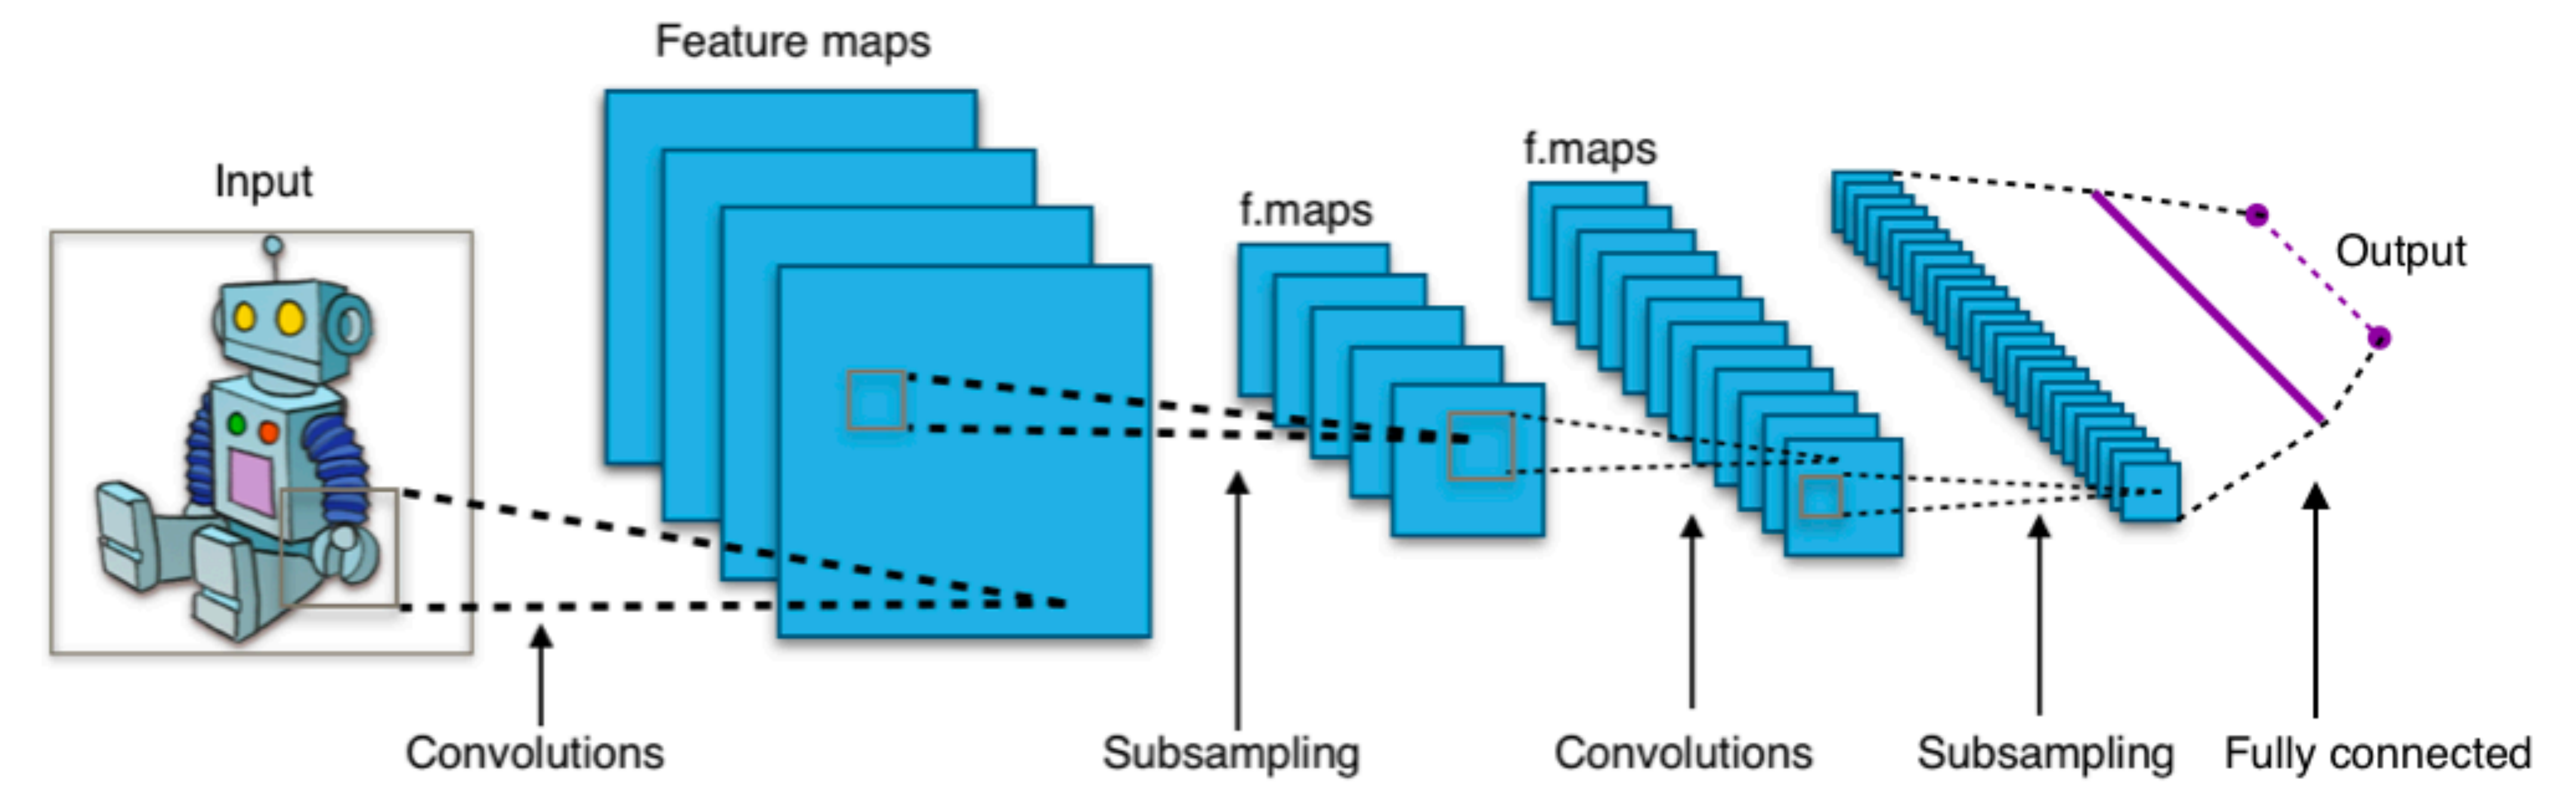

Convolutional neural networks: far less intelligible
[image from Wikipedia]

What does the model “see”?

[Olah et al. 2018]

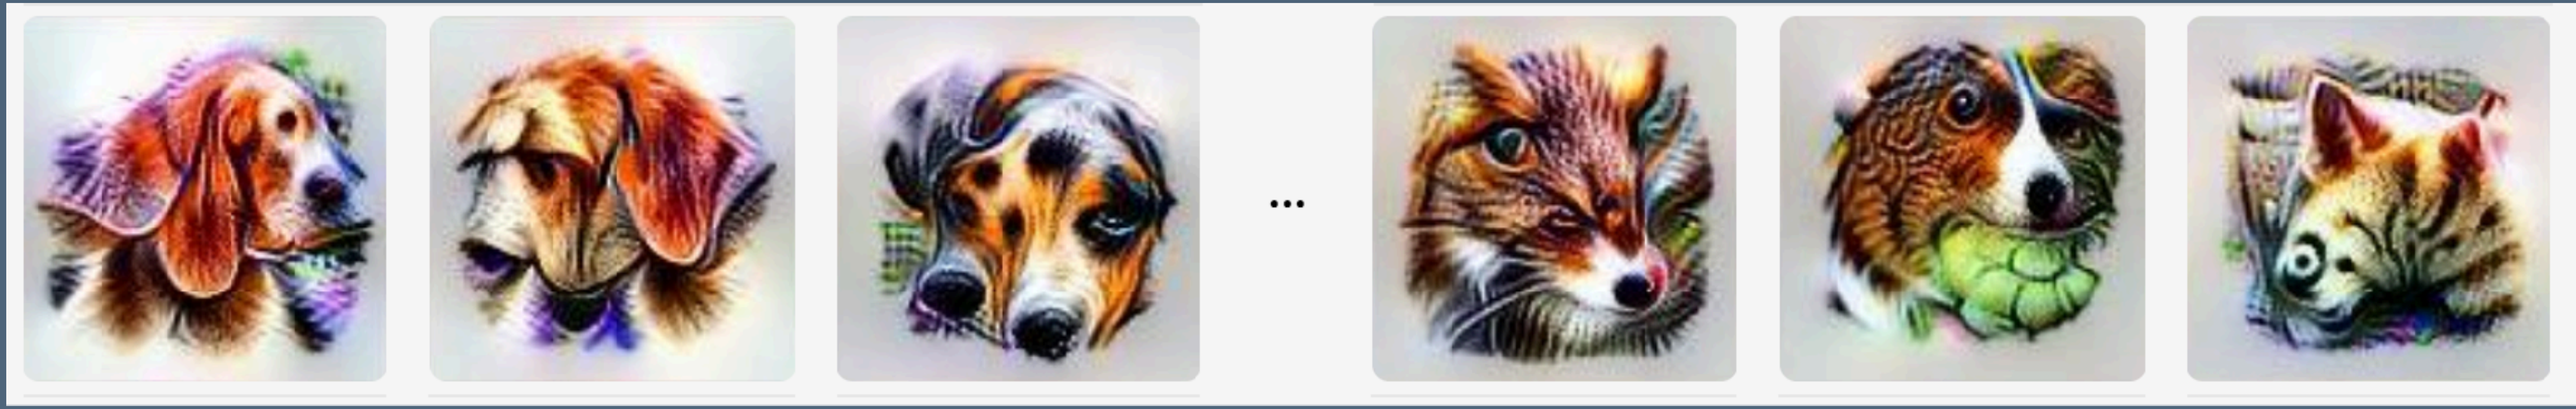

- Above: labrador retriever (left) vs. tiger cat (right)
- If it requires a person to predict its behavior, ML systems that are intelligibility require that people can “see what they see” to a certain extent

One approach: simplify the model

- Sometimes you can get most of the performance with far higher intelligibility

GA²M

[Lou et al. 2013, Caruana et al. 2015]

- A model that learns all features of the form: $\sum f_i(x_i) + \sum f_{ij}(x_i, x_j)$
- In other words: the system learns univariate and bivariate relationships between the input features and the outcome

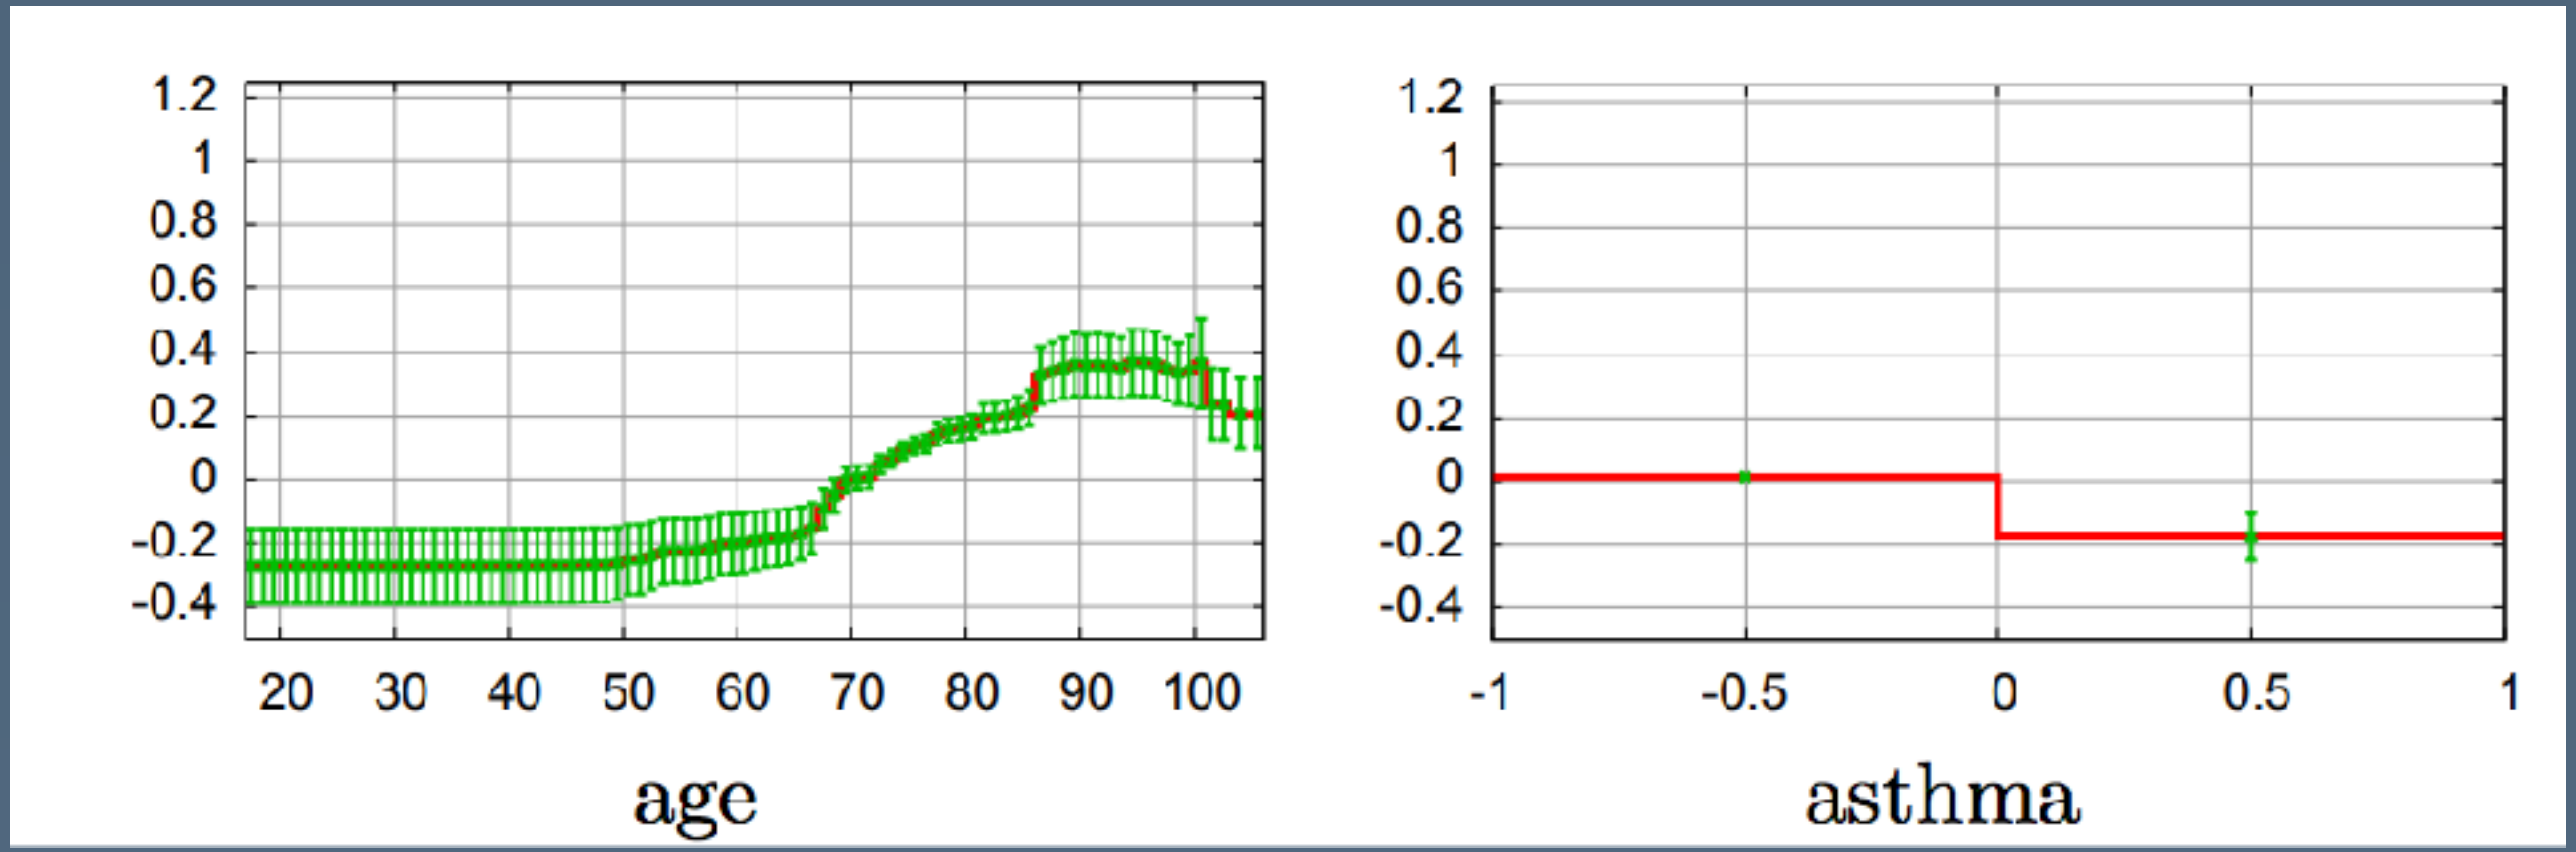

Univariate features predicting pneumonia risk

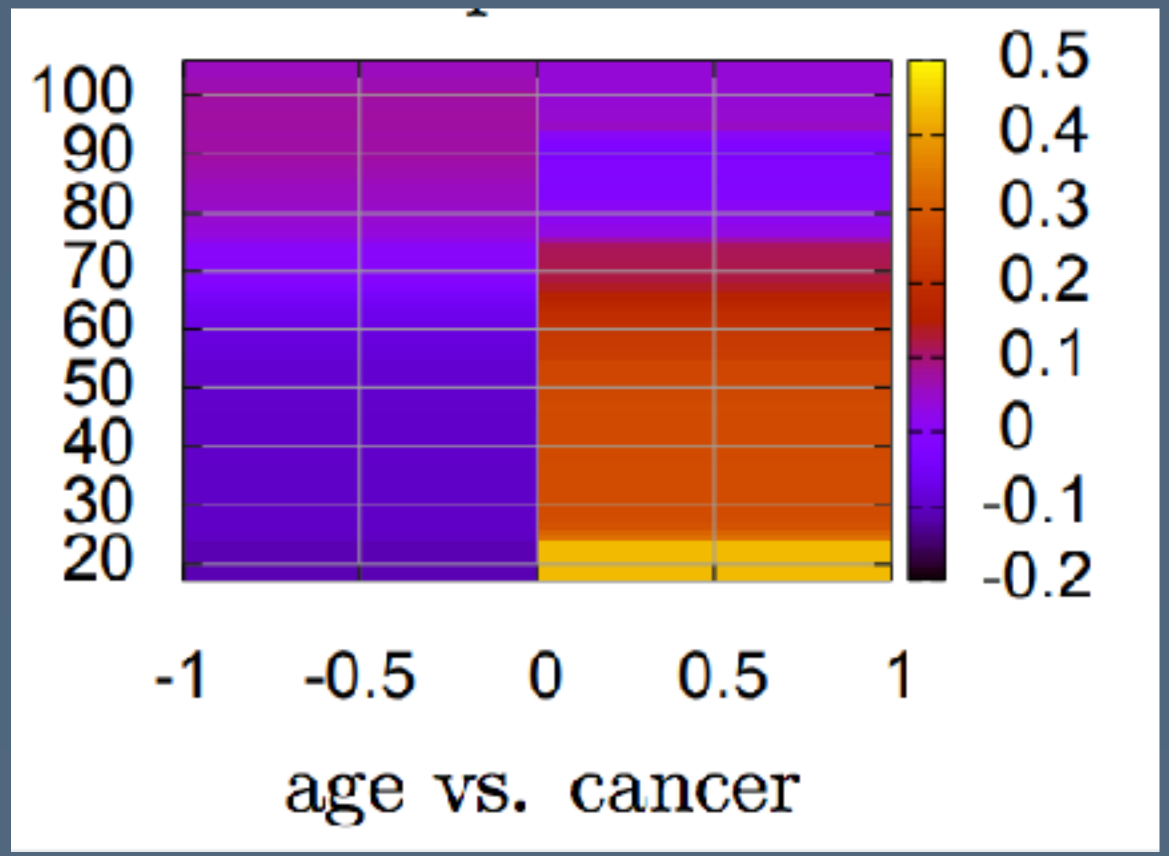

Bivariate interaction predicting pneumonia risk (note particular risk of young people with cancer)

LIME local explainers

[Ribeiro et al. 2016]

- Often the learned decision boundary (red vs. blue) is nonlinear →
- Suppose you were trying to explain the bright red cross example. What would you do?

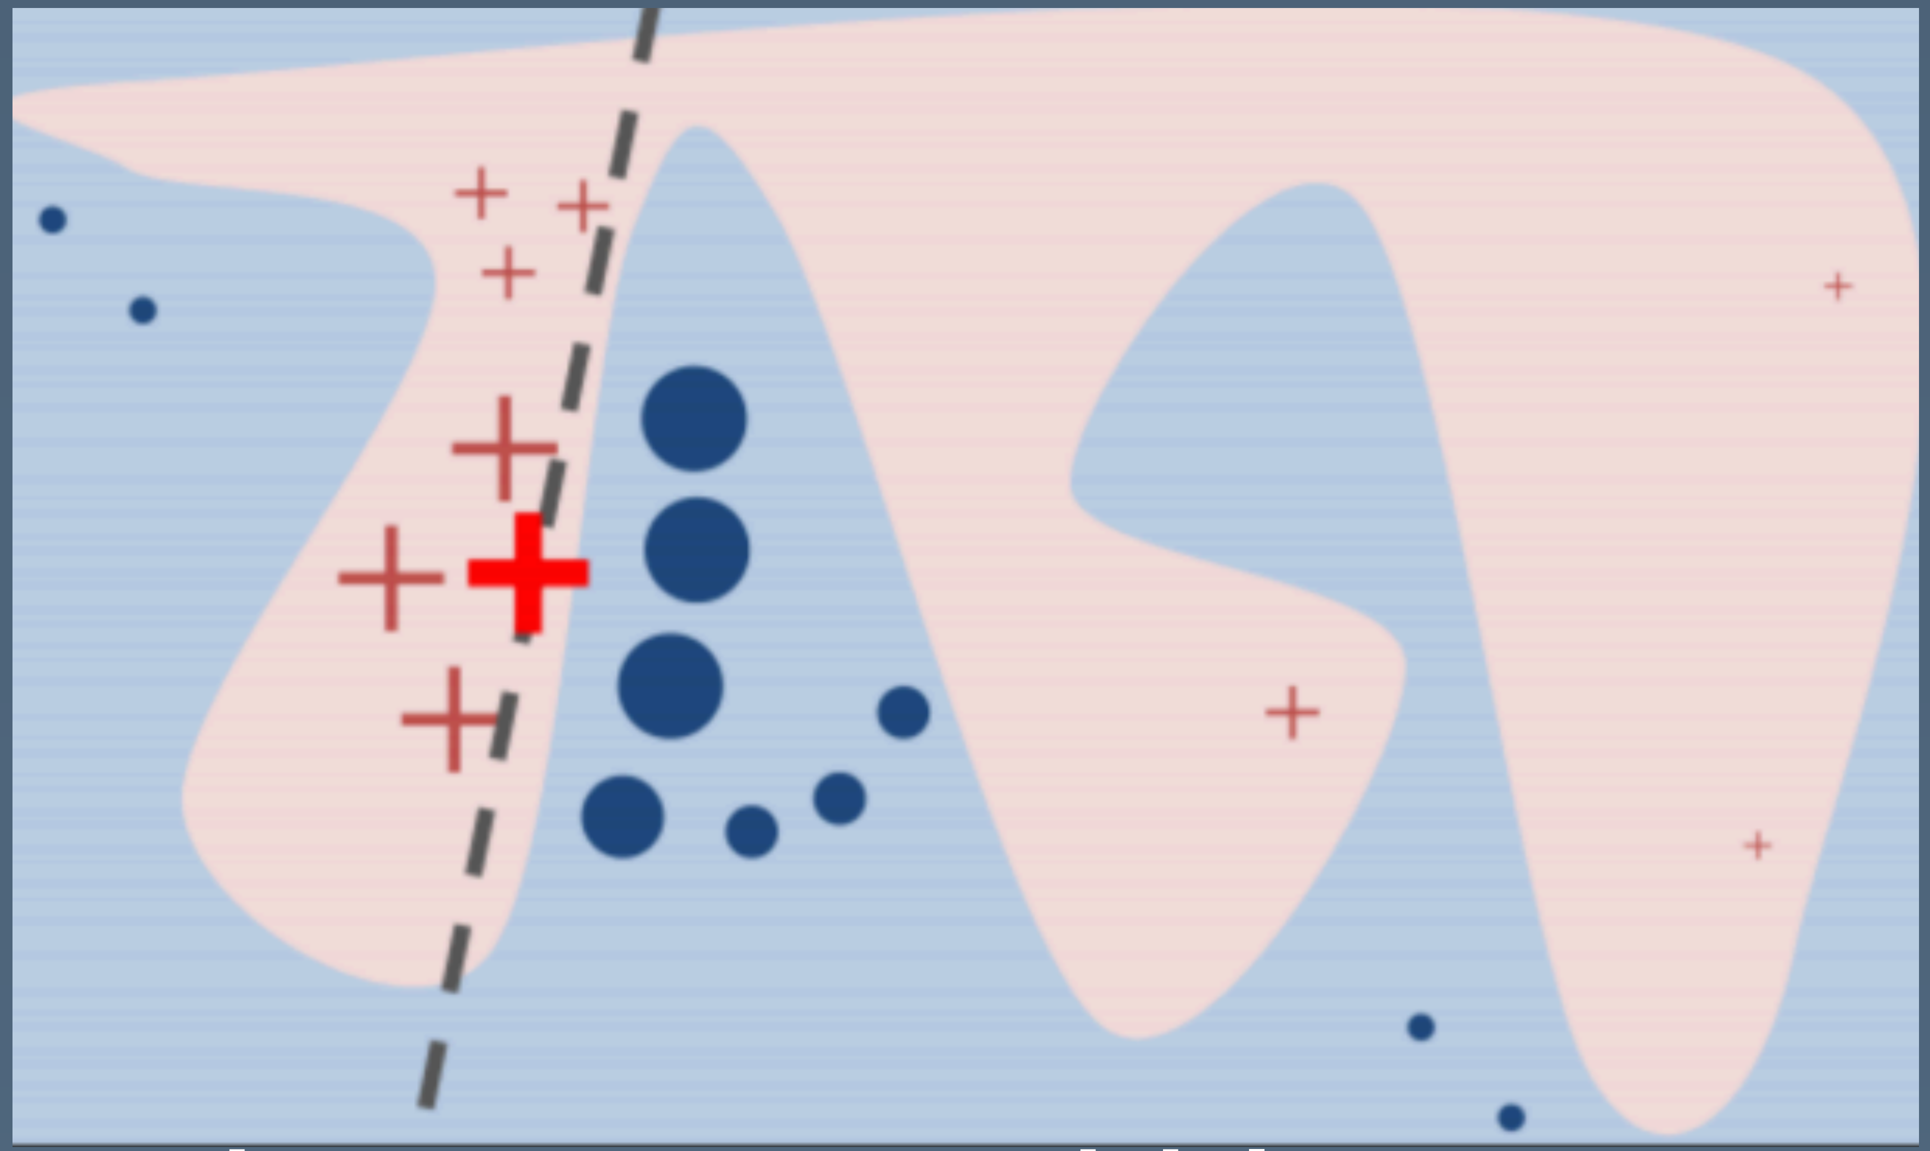

- LIME: sample other points nearby (large red crosses and blue circles), weigh them in proportion to their proximity to the red cross, and learn a linear separator for them.
  - This is *not* an accurate representation of the whole model! But still useful in local explanation.

Influence functions

[Koh and Liang et al. 2017]

- Mathematical approach that traces a prediction back to the most influential training points that produced the prediction

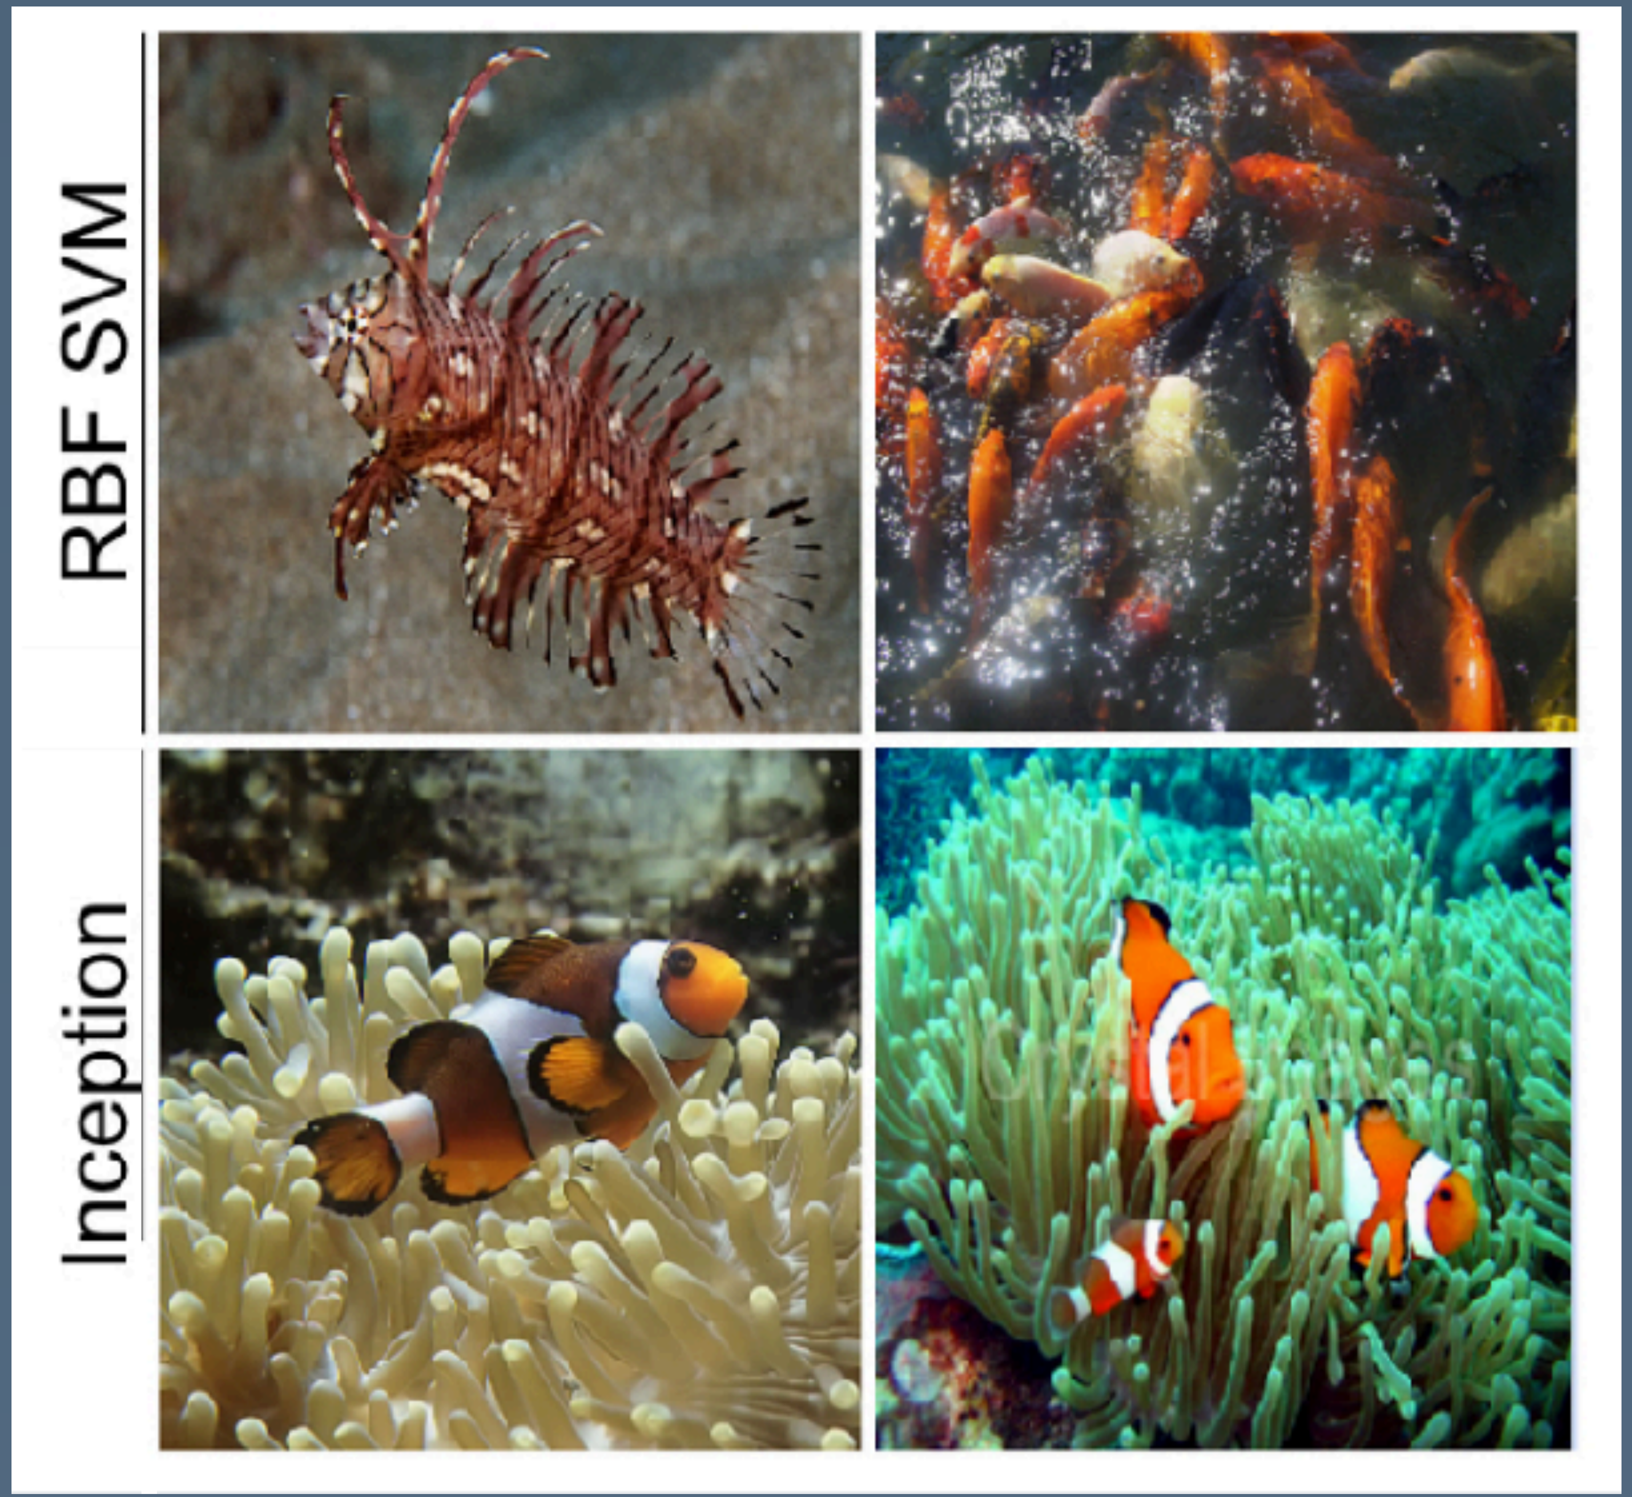

The dilemma

[Weld and Bansal 2018]

- Any model simplification is a lie!
- But any non-simplification is unintelligible.

- Recommendation: draw on psych research to guide explanation
  - Make explanations contrastive: "Why did you recommend Movie X?" (implicitly: "Why didn't you recommend Movie Y?")
  - Necessary causes are better than sufficient ones
  - Use few conjuncts

Guiding users to train effective classifiers

Interactive training

[Fogarty et al. 2008]

- Allow users to keep training and re-training by drag-dropping instances into positive and negative classes as they go

Revising your training as you go

[Chang, Amershi and Kamar 2017]

- Facilitate concept evolution through a “could be” category that allows clustering into subcategories you can change labels for

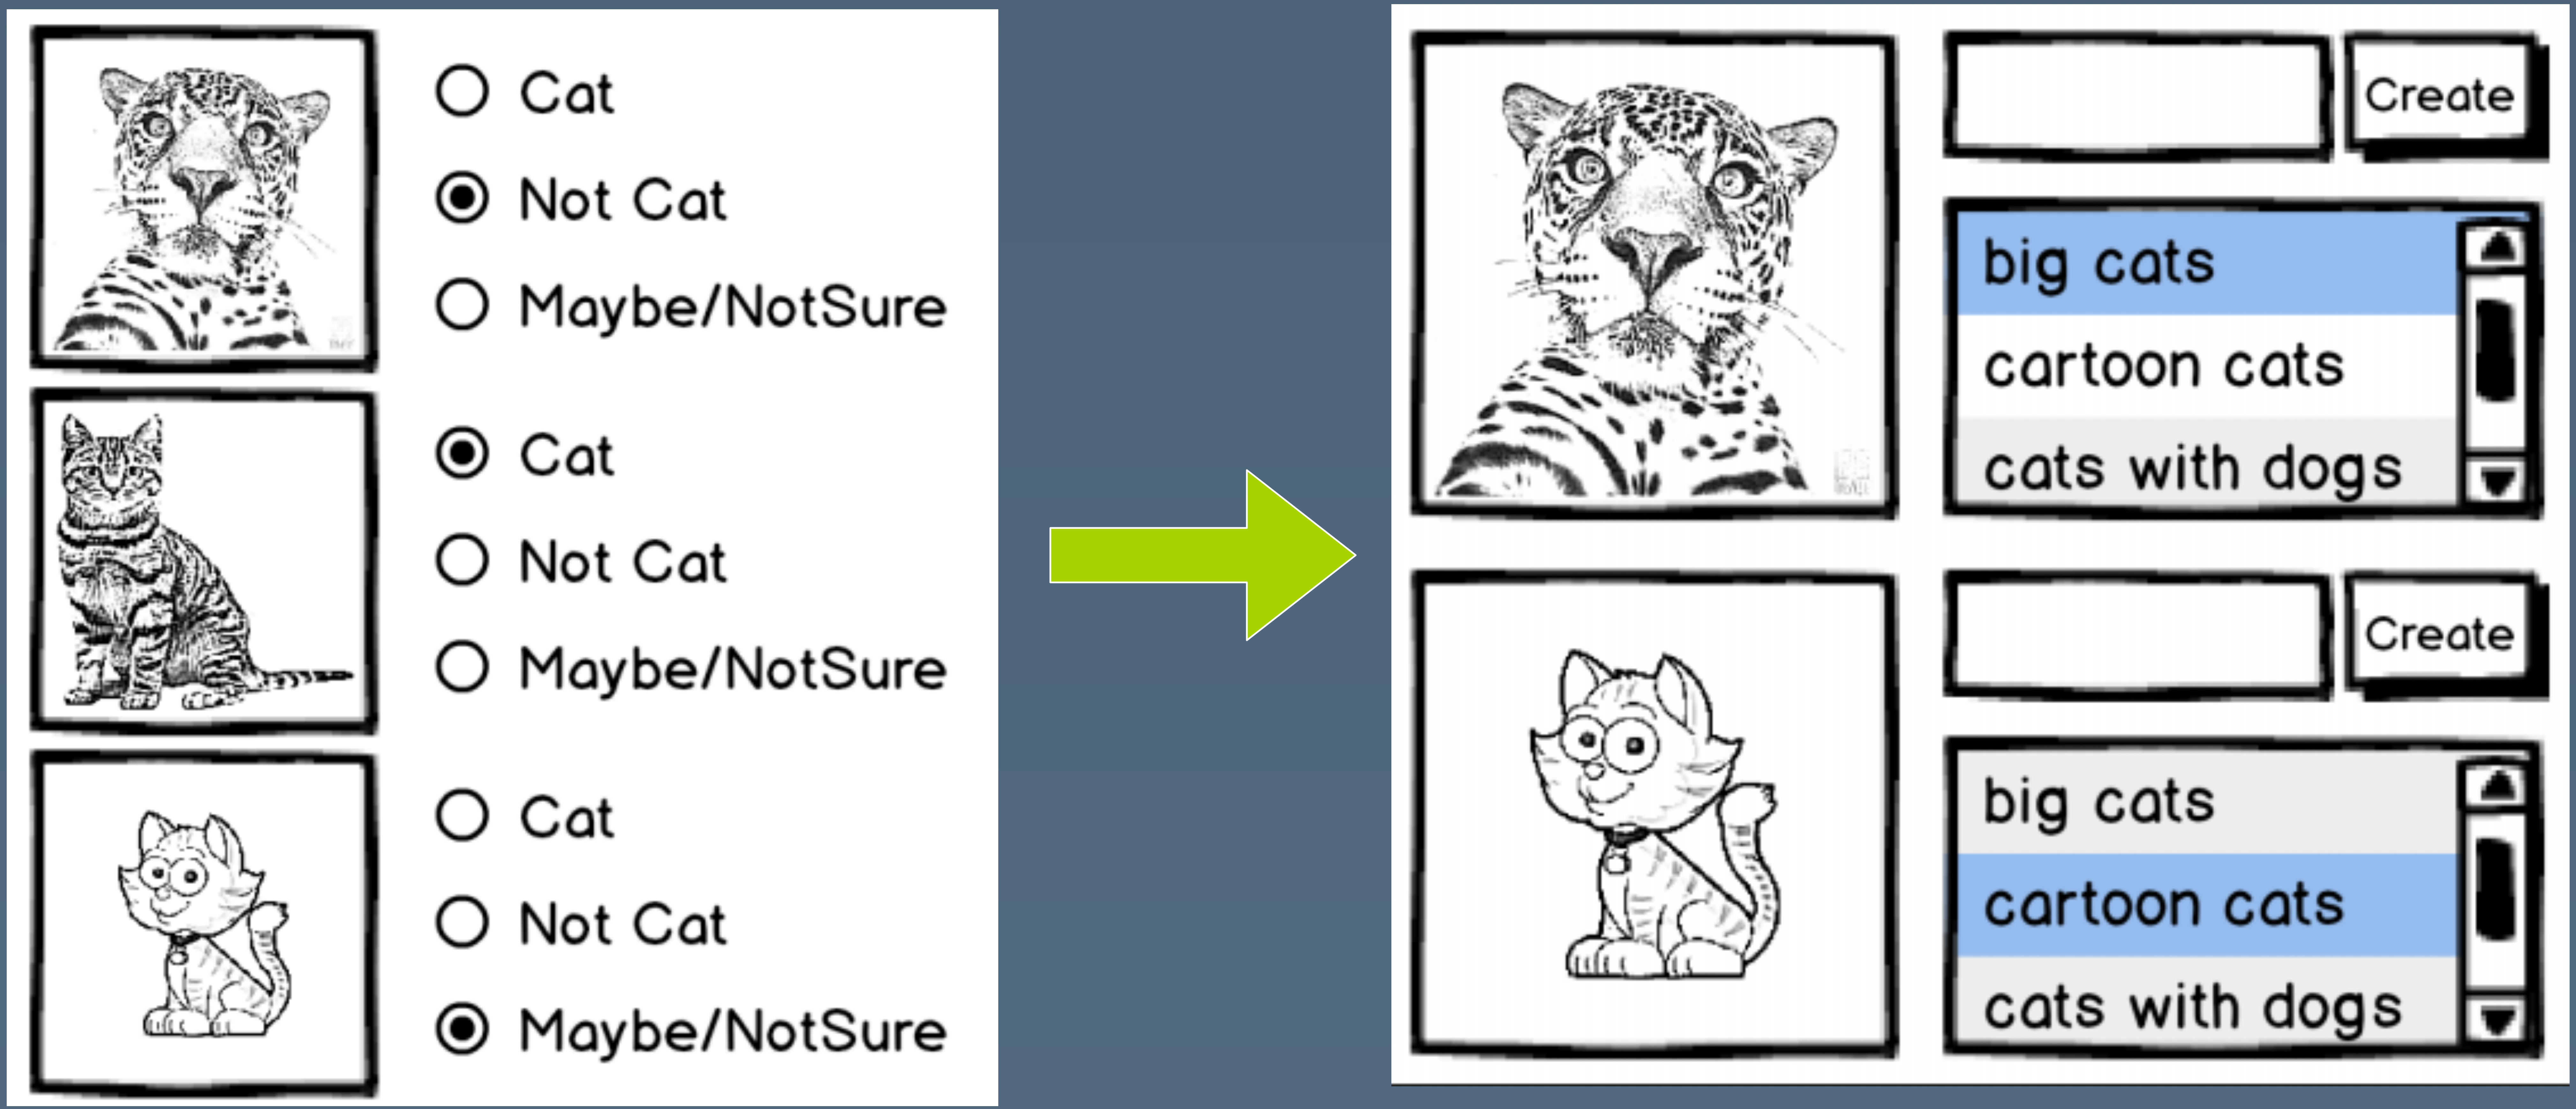

Play-along learning

[Fiebrink, Cook, and Trueman, ICMC '09]

- Create the output (sounds) you desire
- "Play along" and demonstrate the input that should generate that output

Developing intelligent software

If you wanted a smart doorbell...

- To automatically control entrance to your room
- To let in possible donors for your Stanford education

How would you train the system quickly?

Crayons: camera-based interaction

[Fails and Olsen, CHI '03]

Direct-manipulation training

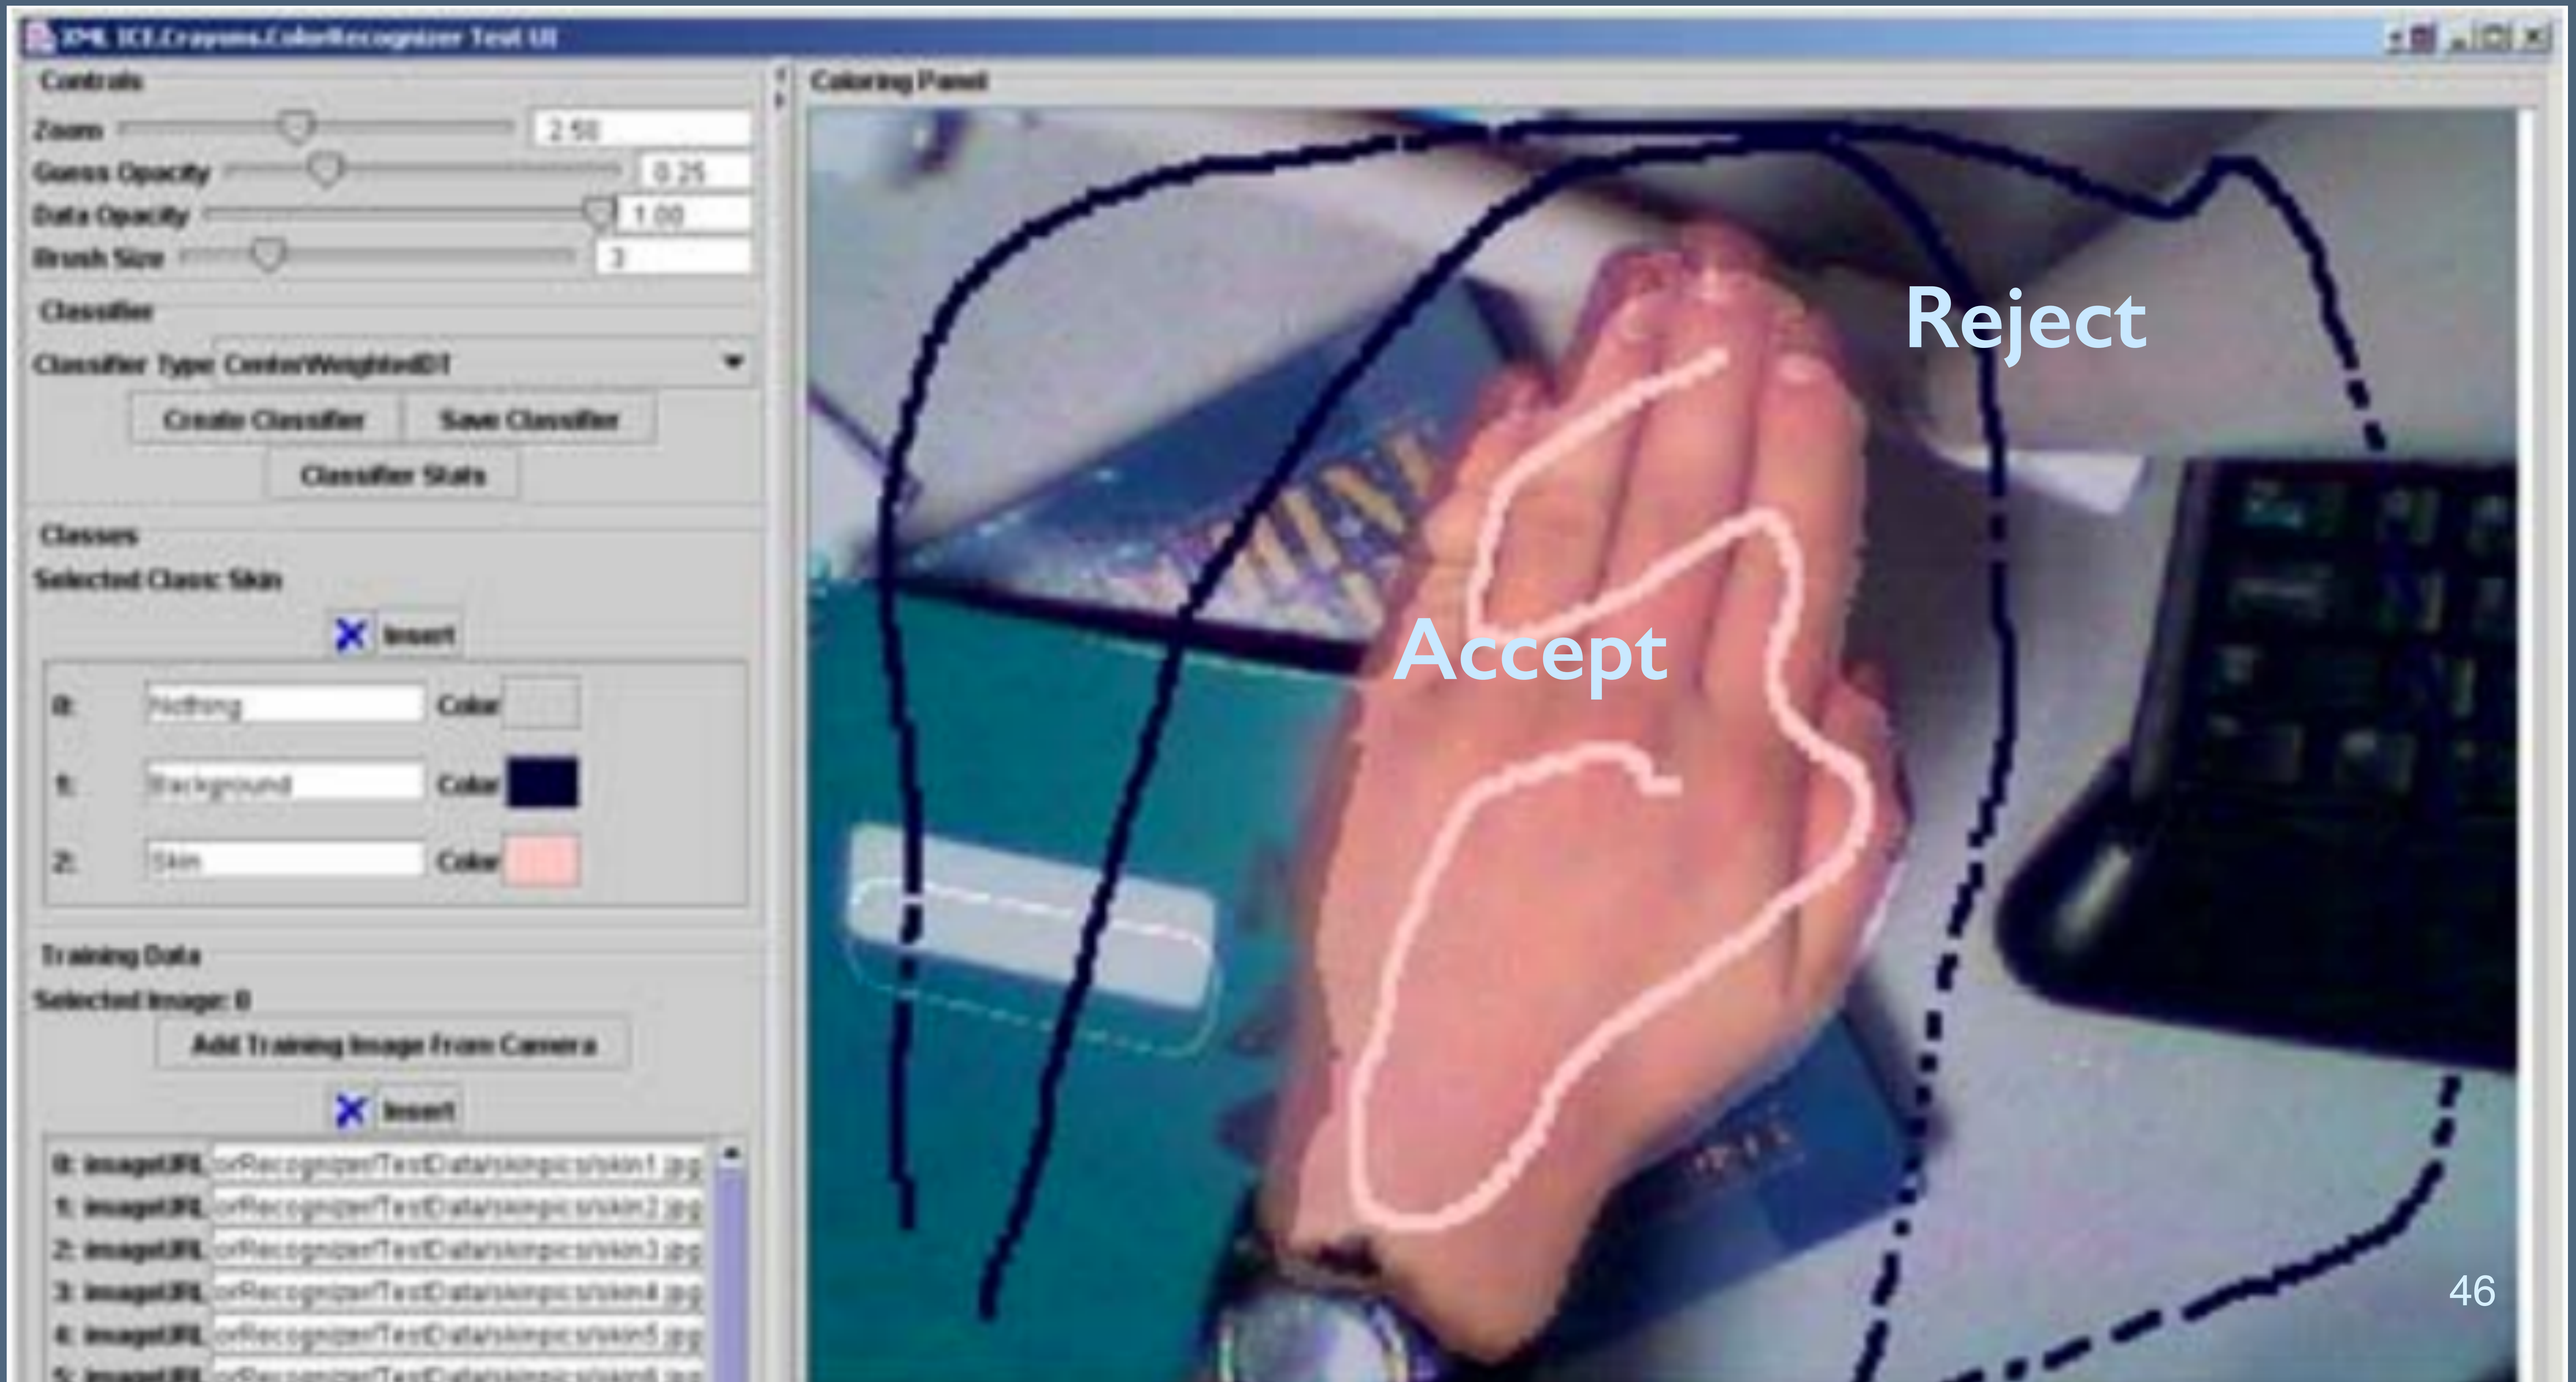

Development challenges with ML

[Patel et al., CHI '08]

- Software development benefits from modularity, but machine learning is iterative and nonlinear
- Difficulty understanding the statistical process underlying machine learning algorithms
- Evaluation of progress is difficult

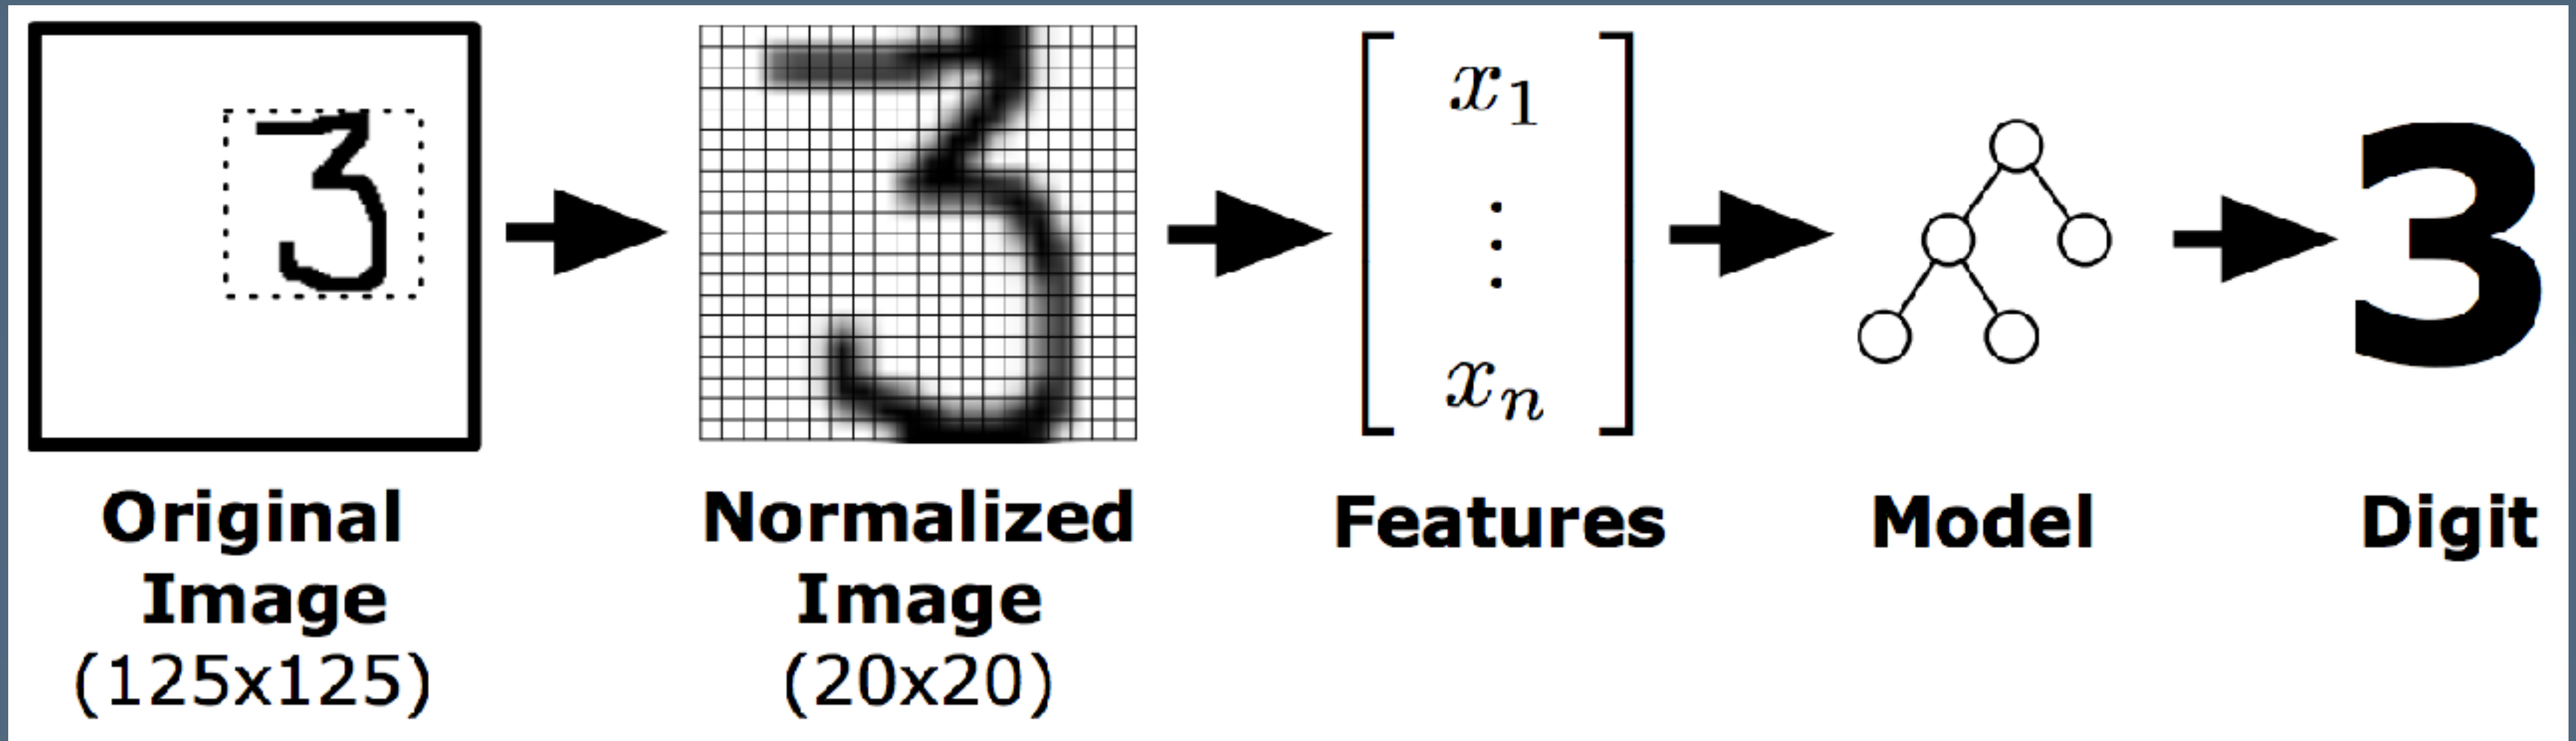

Papier-Mâché: toolkit support for tangible input

[Klemmer et al., CHI '04]

- Monitoring window, wizard-of-oz input, listeners, designed and evaluated as a user interface

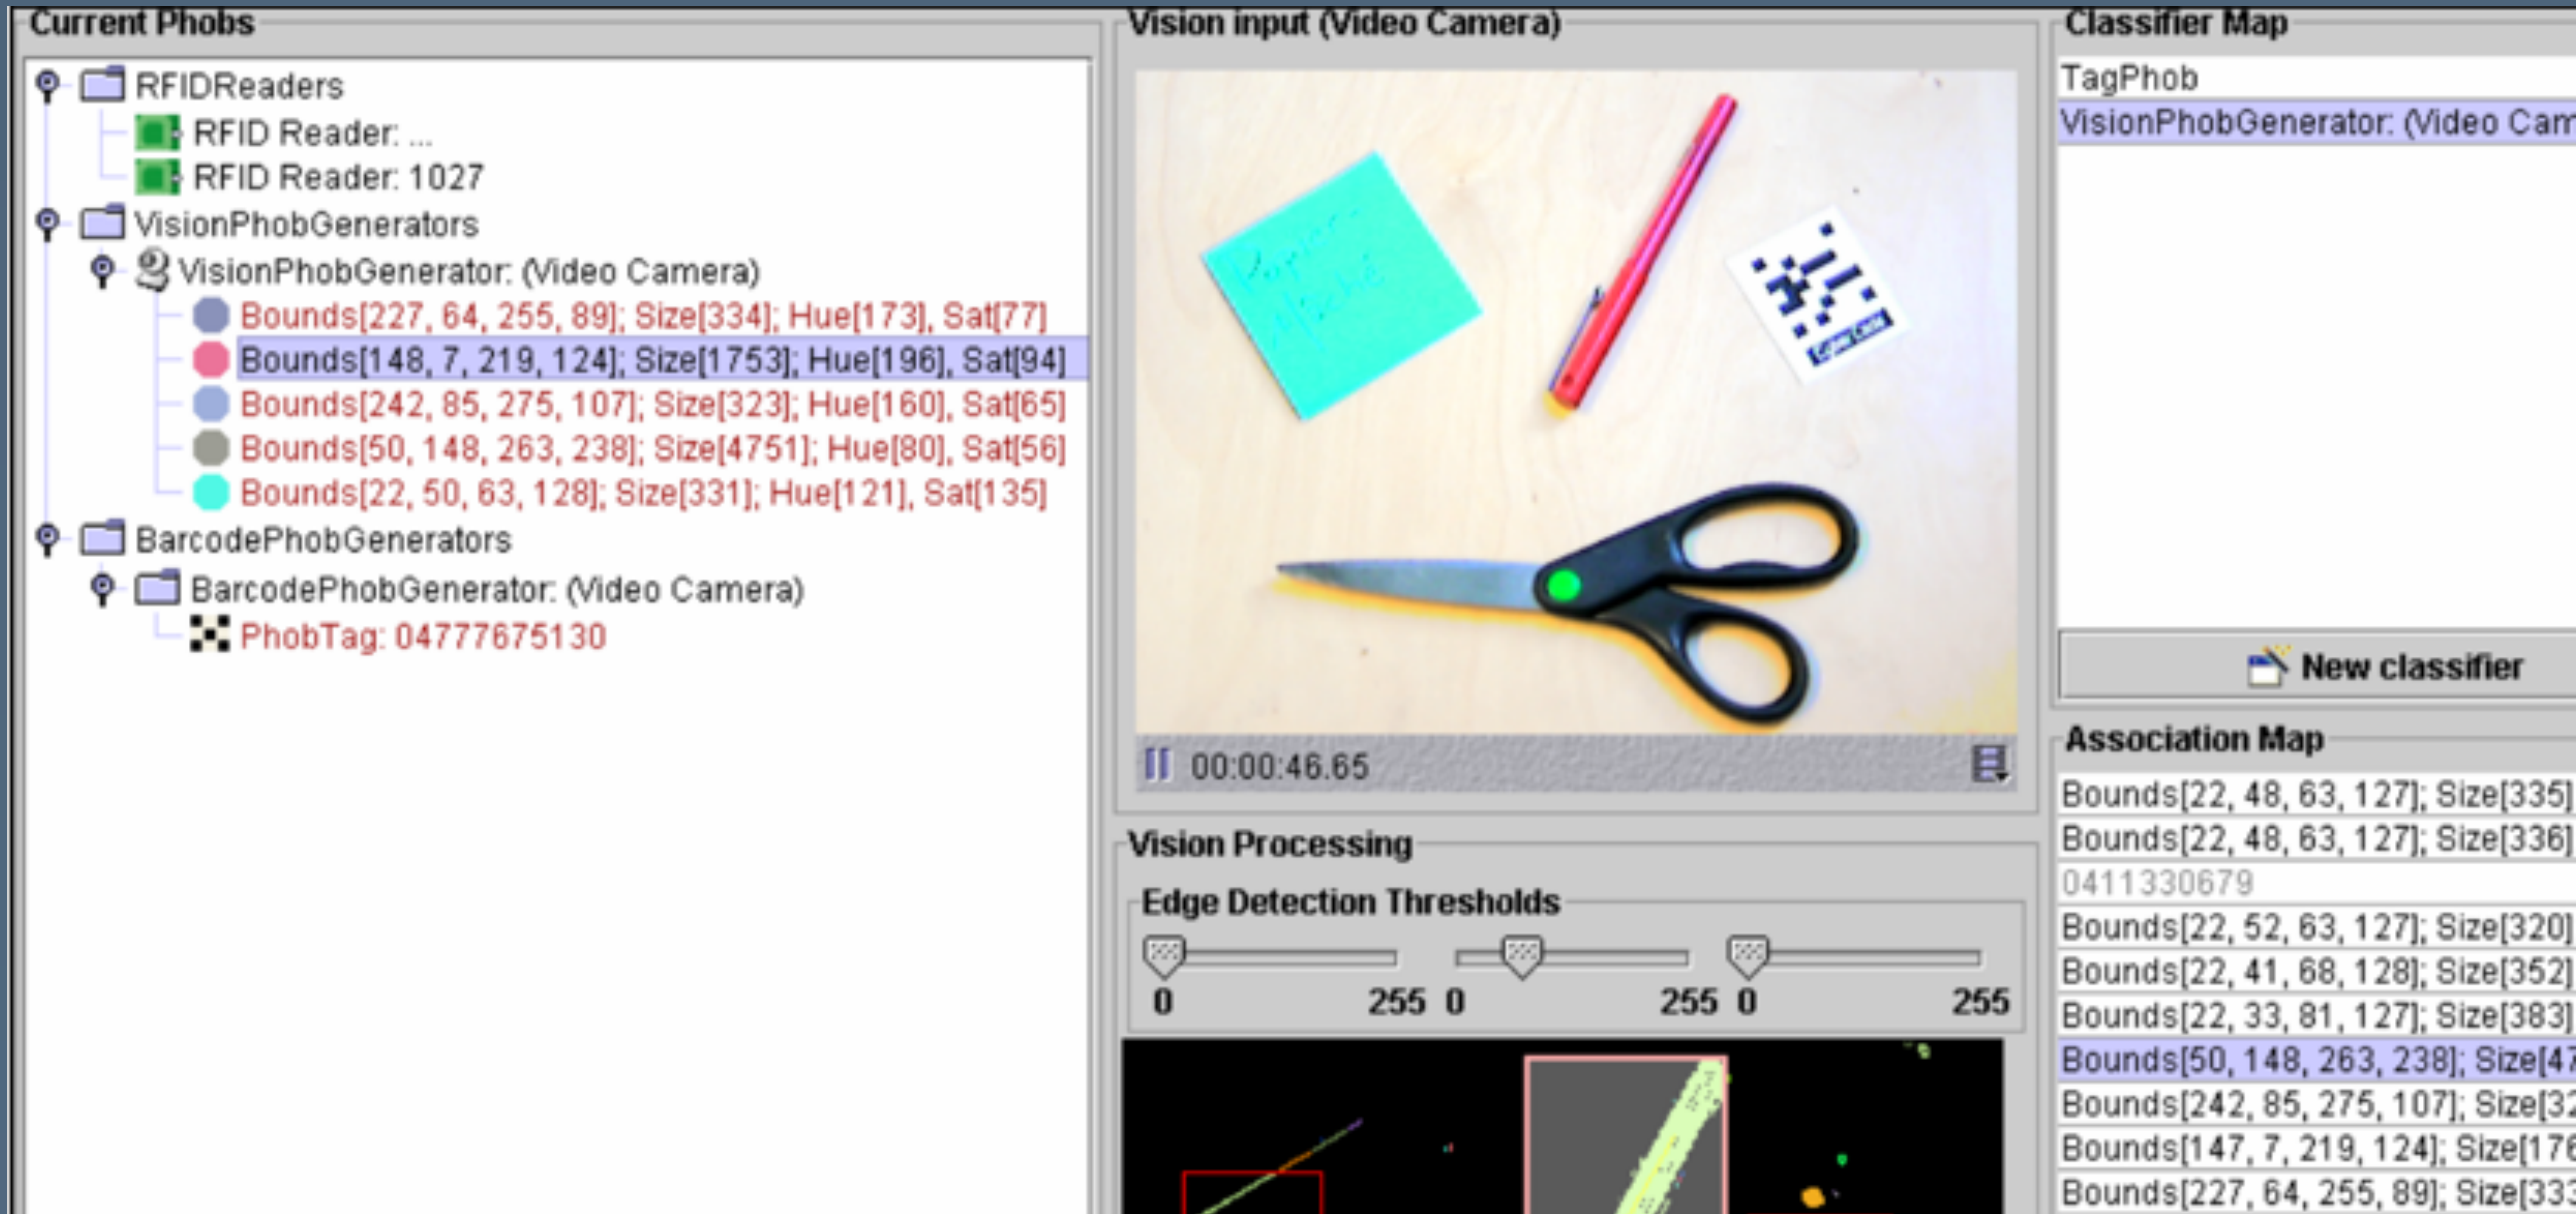

IDE support for ML development

[Patel et al., UIST '10]

- Explicit support for each step: feature extraction, model generation, training and testing

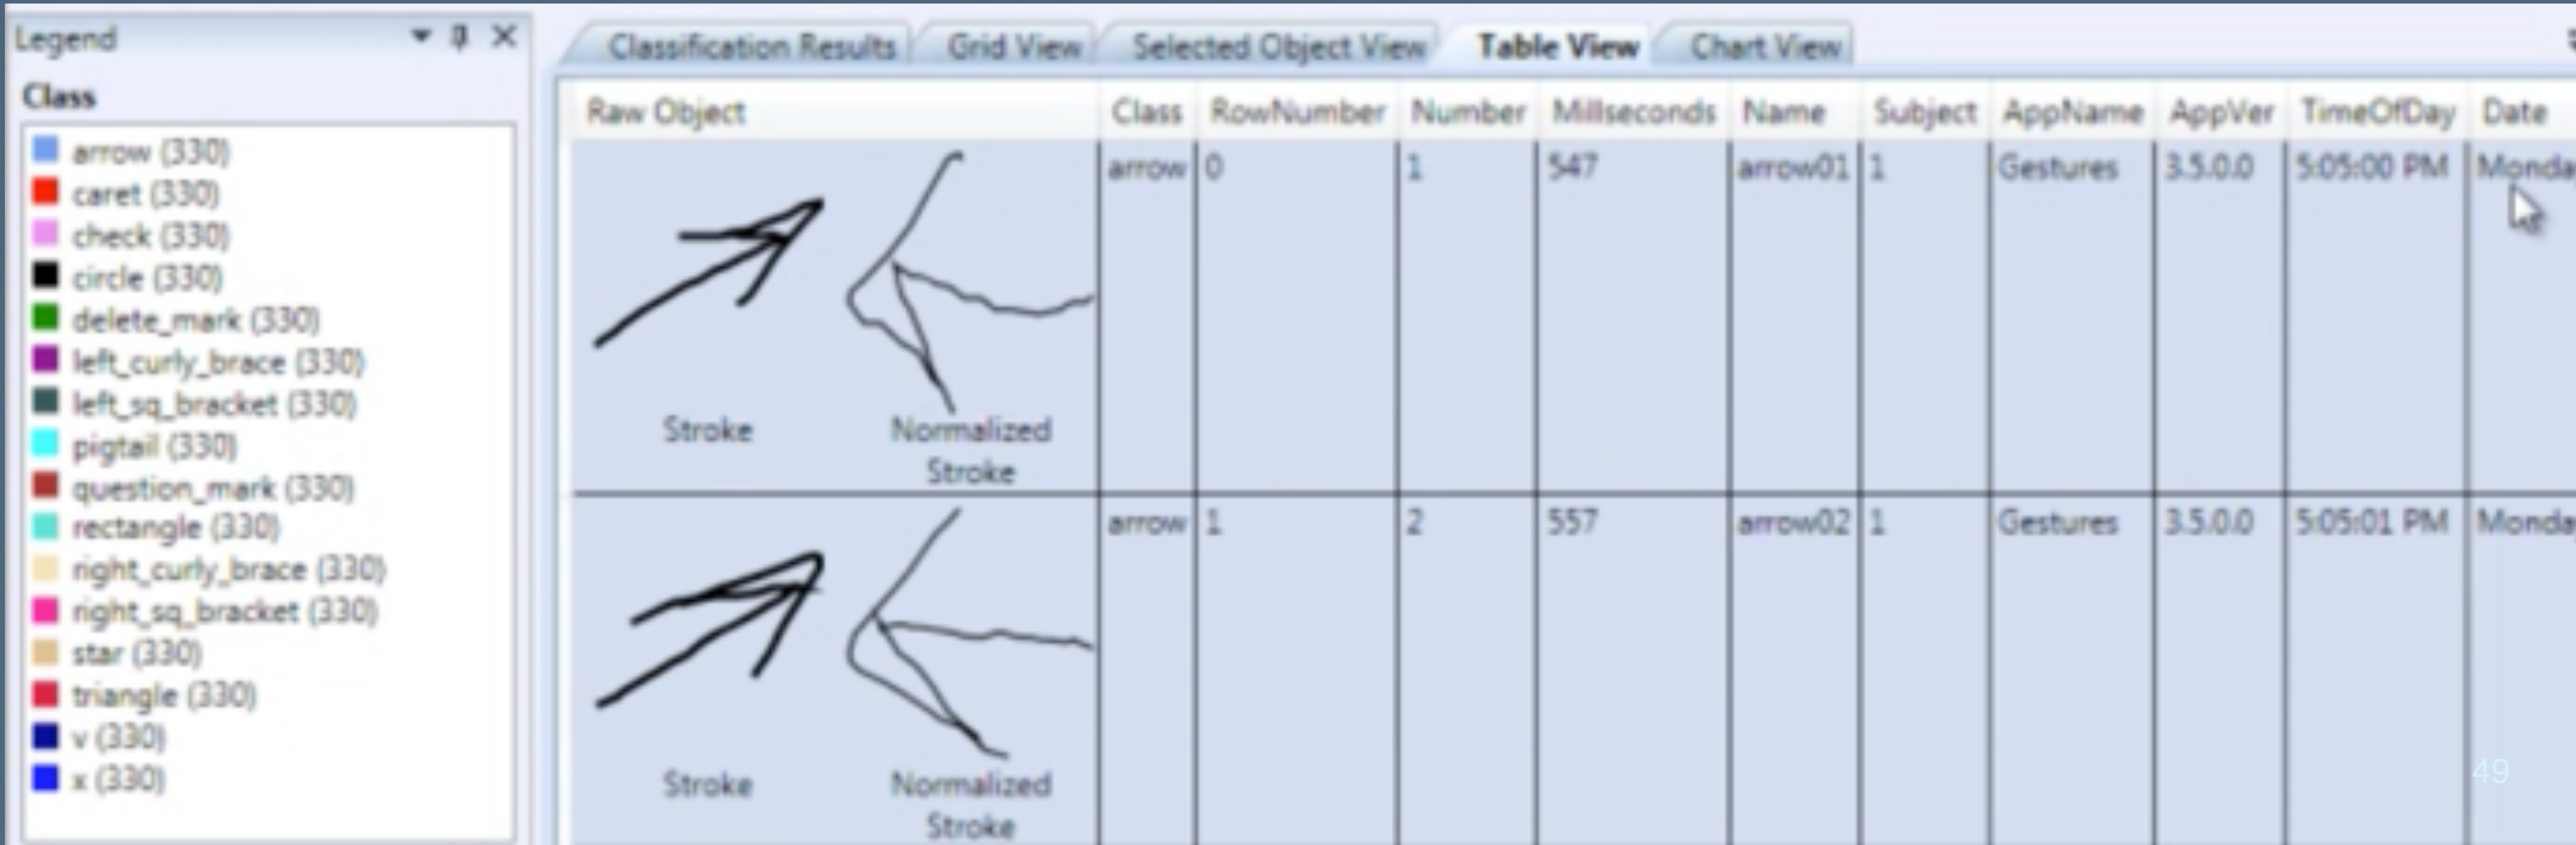

AI-driven design

Does your design look a bit like this?

AI-driven design

- Learn design patterns from high-quality designs, and make it straightforward to apply those design patterns to your own designs

Retargeting designs

- "Can I borrow your design?" for the web
- Structured tree mapping algorithm
  - Roughly: costs associated with violating ancestry and sibling relationships in creating a mapping

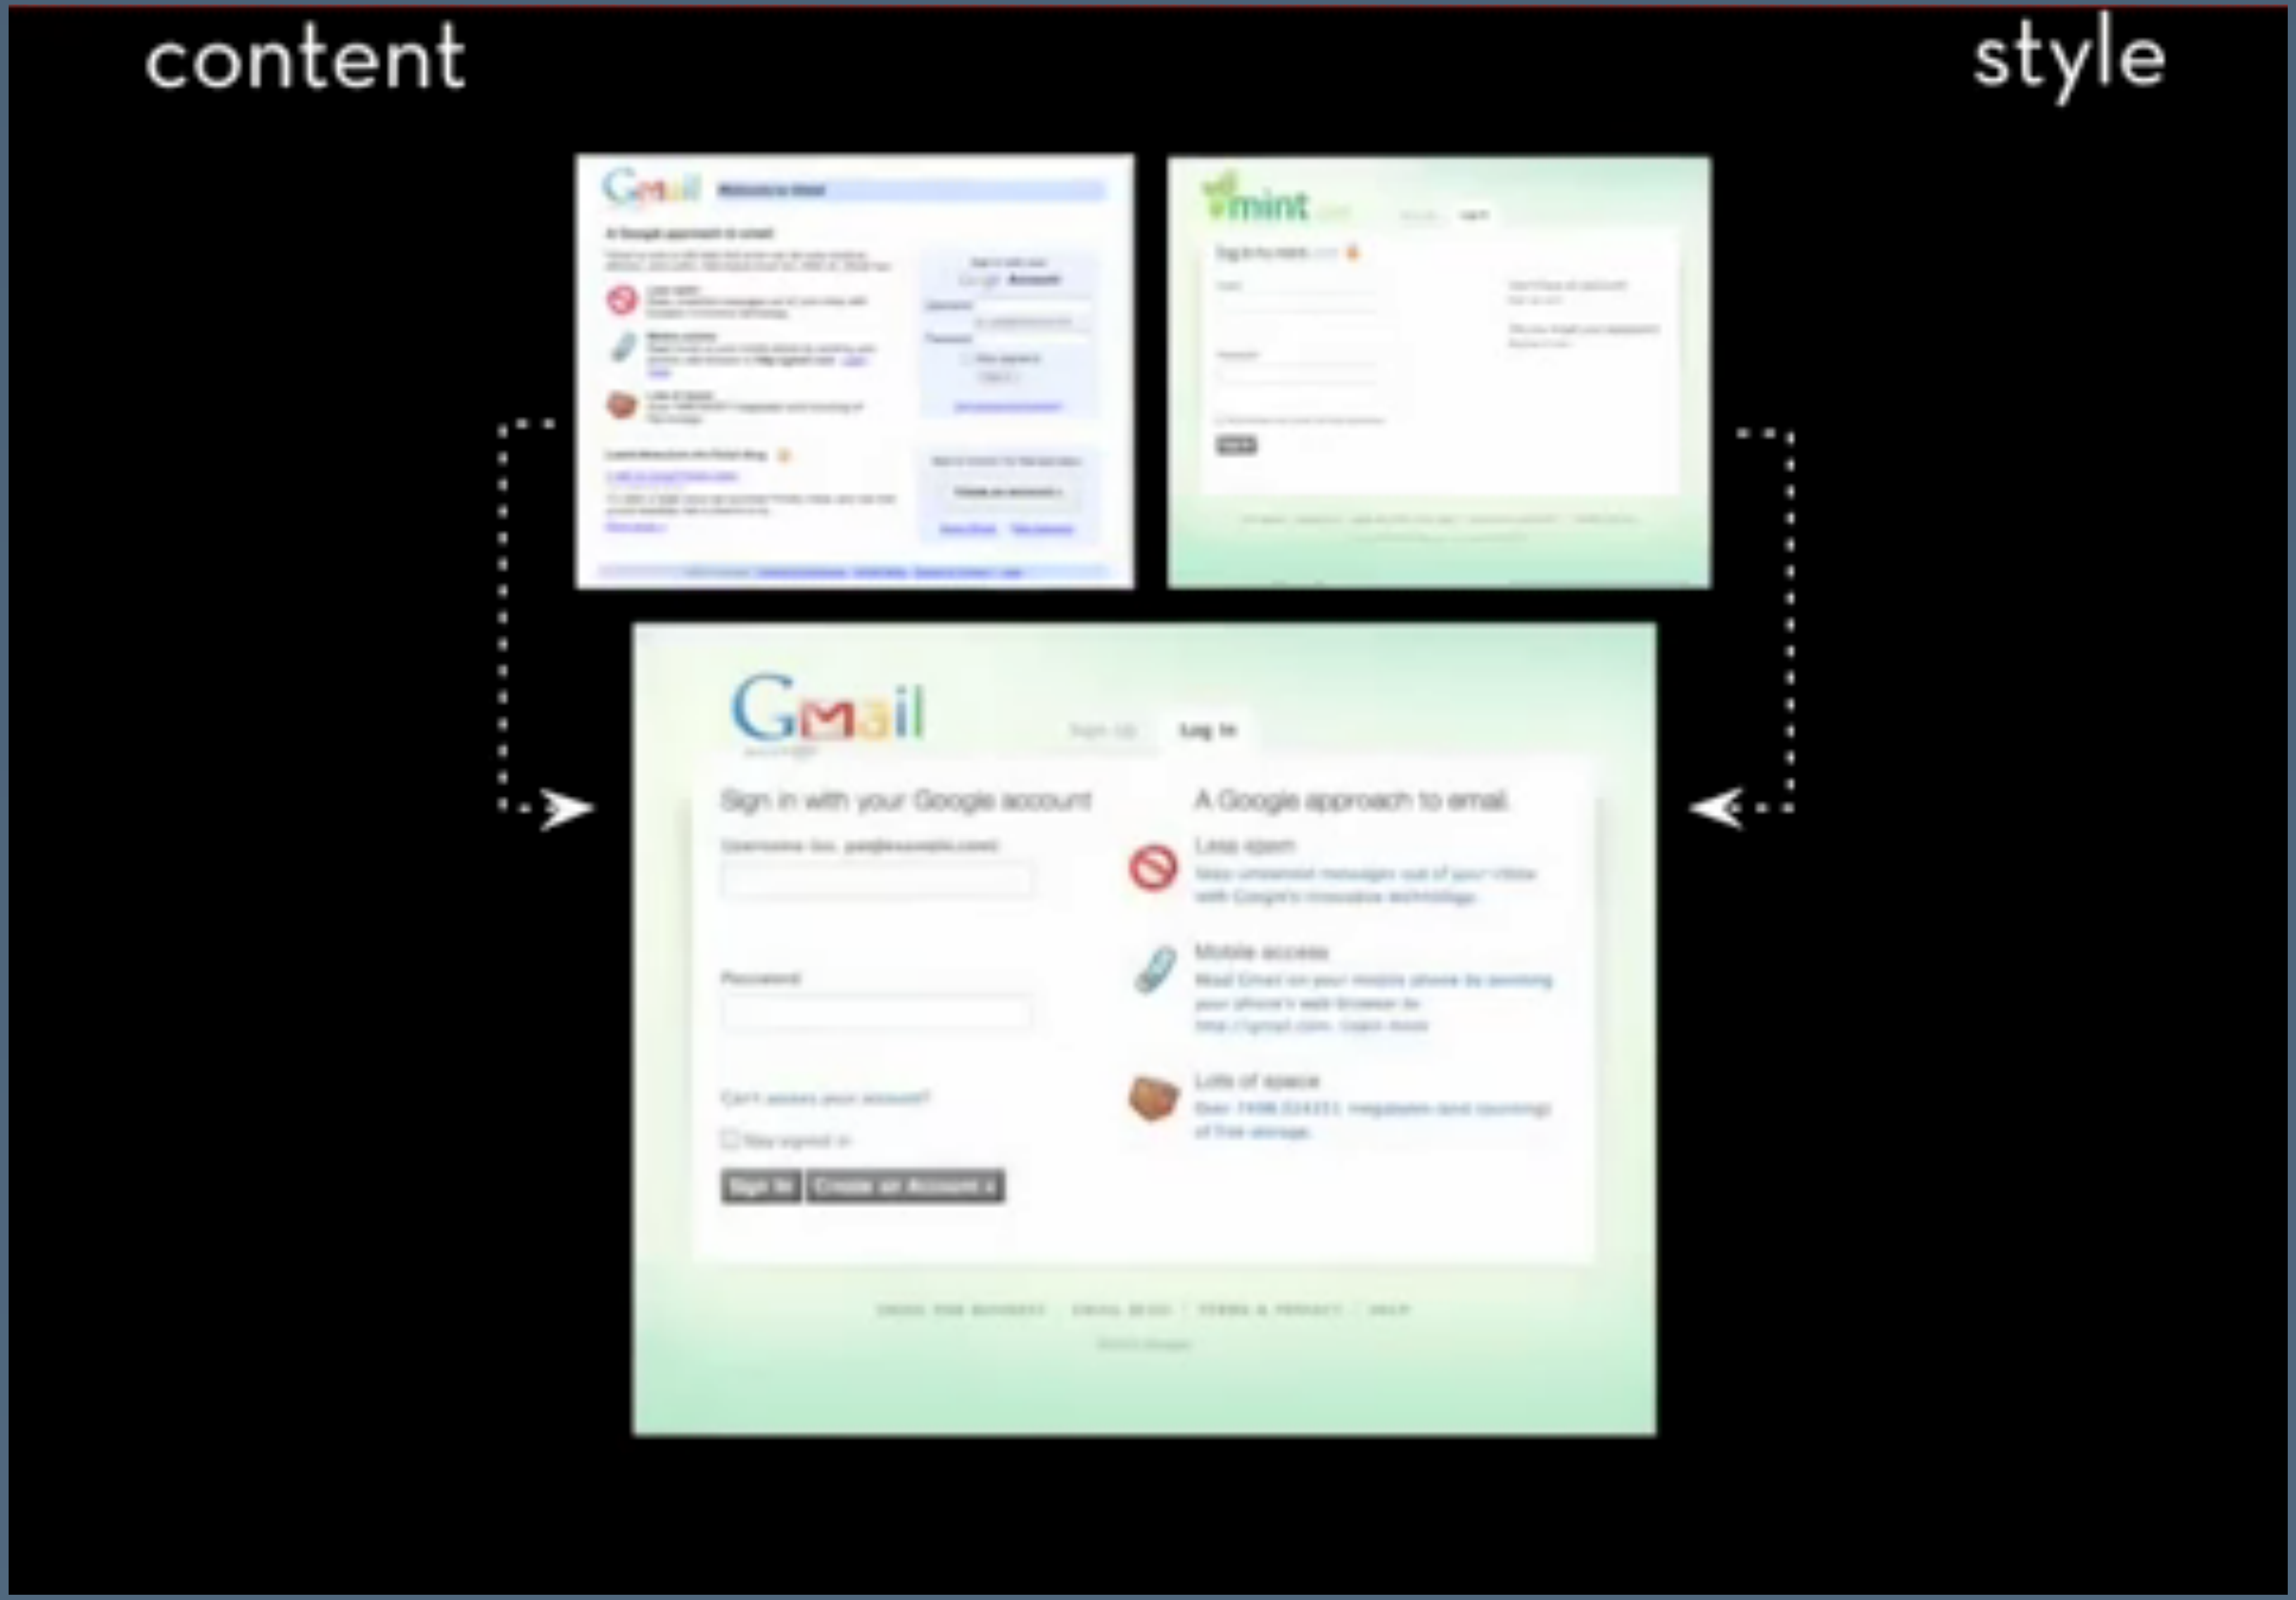

Webzeitgeist

[Kumar et al., CHI '13]

- Crawl the web and index large-scale design elements
- Main idea: what happens if we start data mining designs, rather than user behavior?

Human-robot interaction: enforcing social norms

[Porfirio et al., UIST '17]

Adaptive interfaces

[Gajos and Weld, IUI ‘04]

- Reactive design: remaps to output affordances
- Minimize a cost function derived from navigating between widgets in user traces

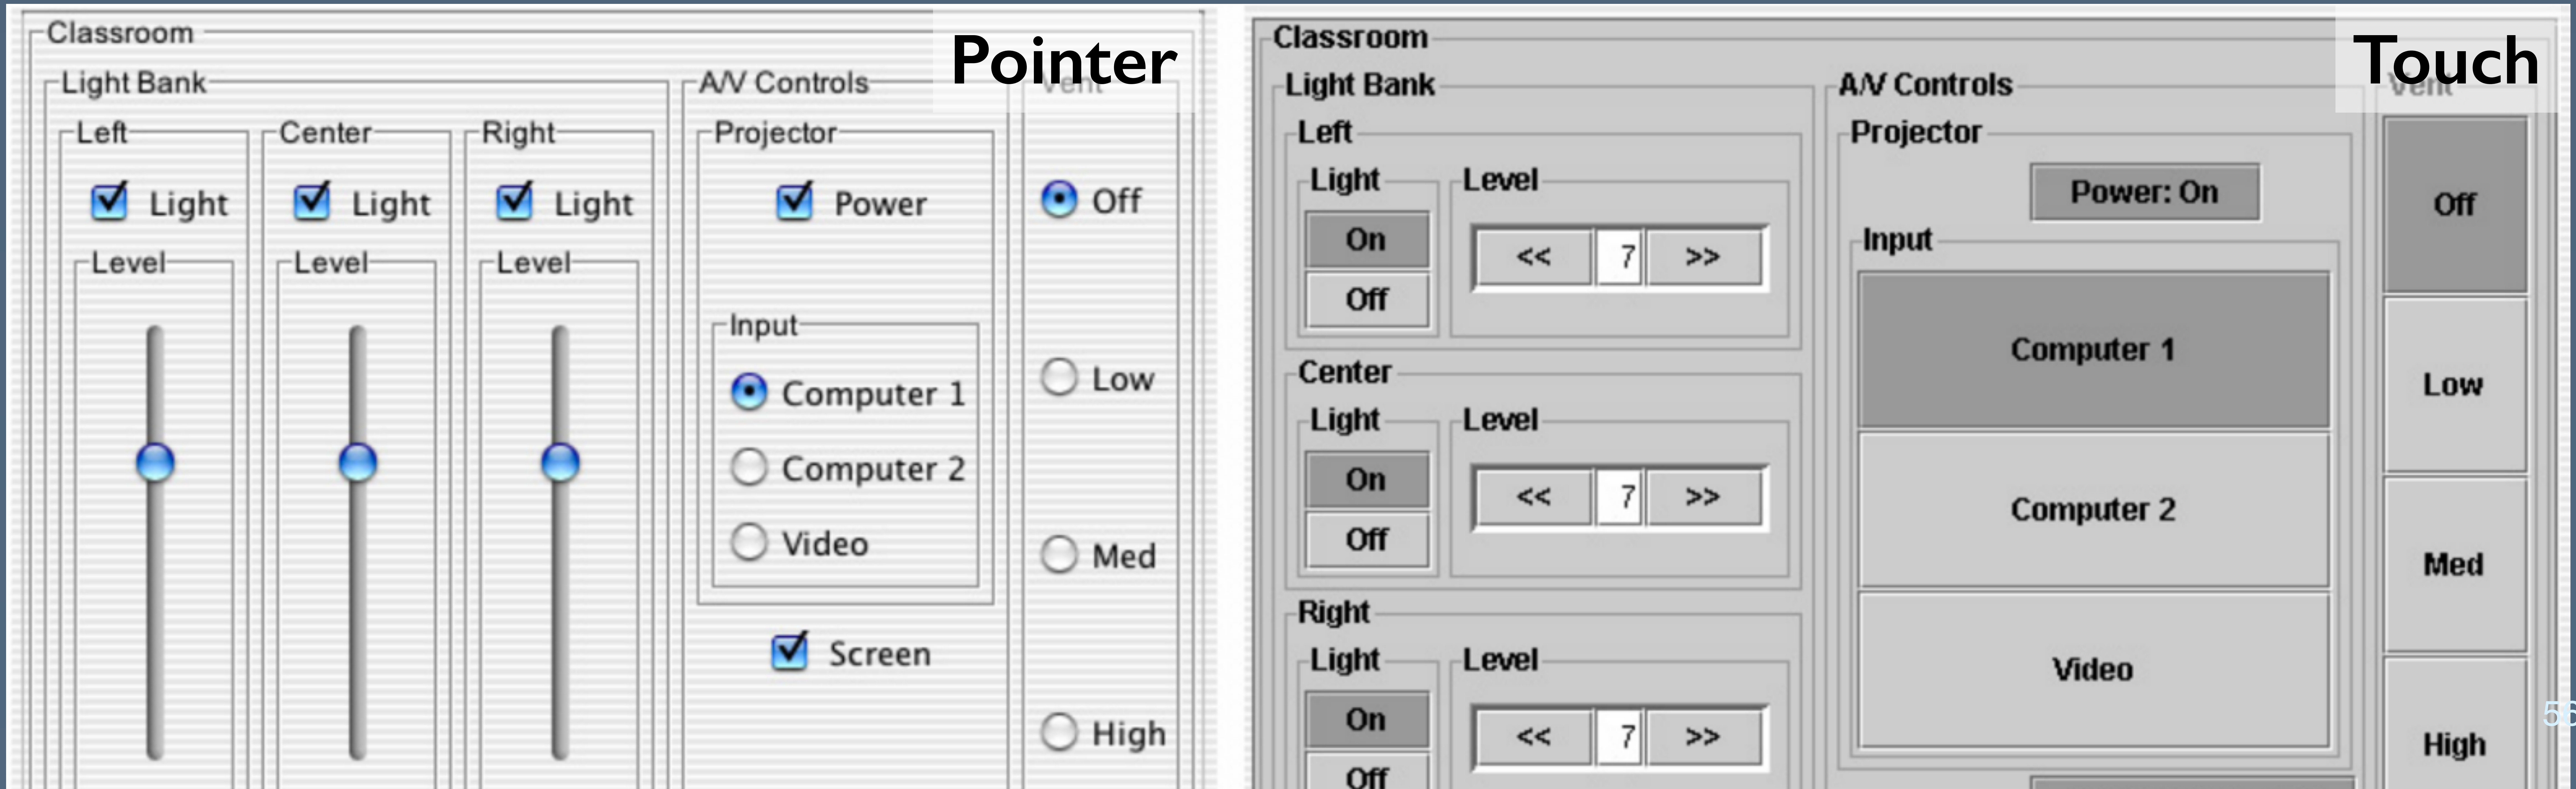

We see this all over…

- …as in papers we already covered!

Sketch the interaction to produce working systems

- SILK [Landay, CHI '96]

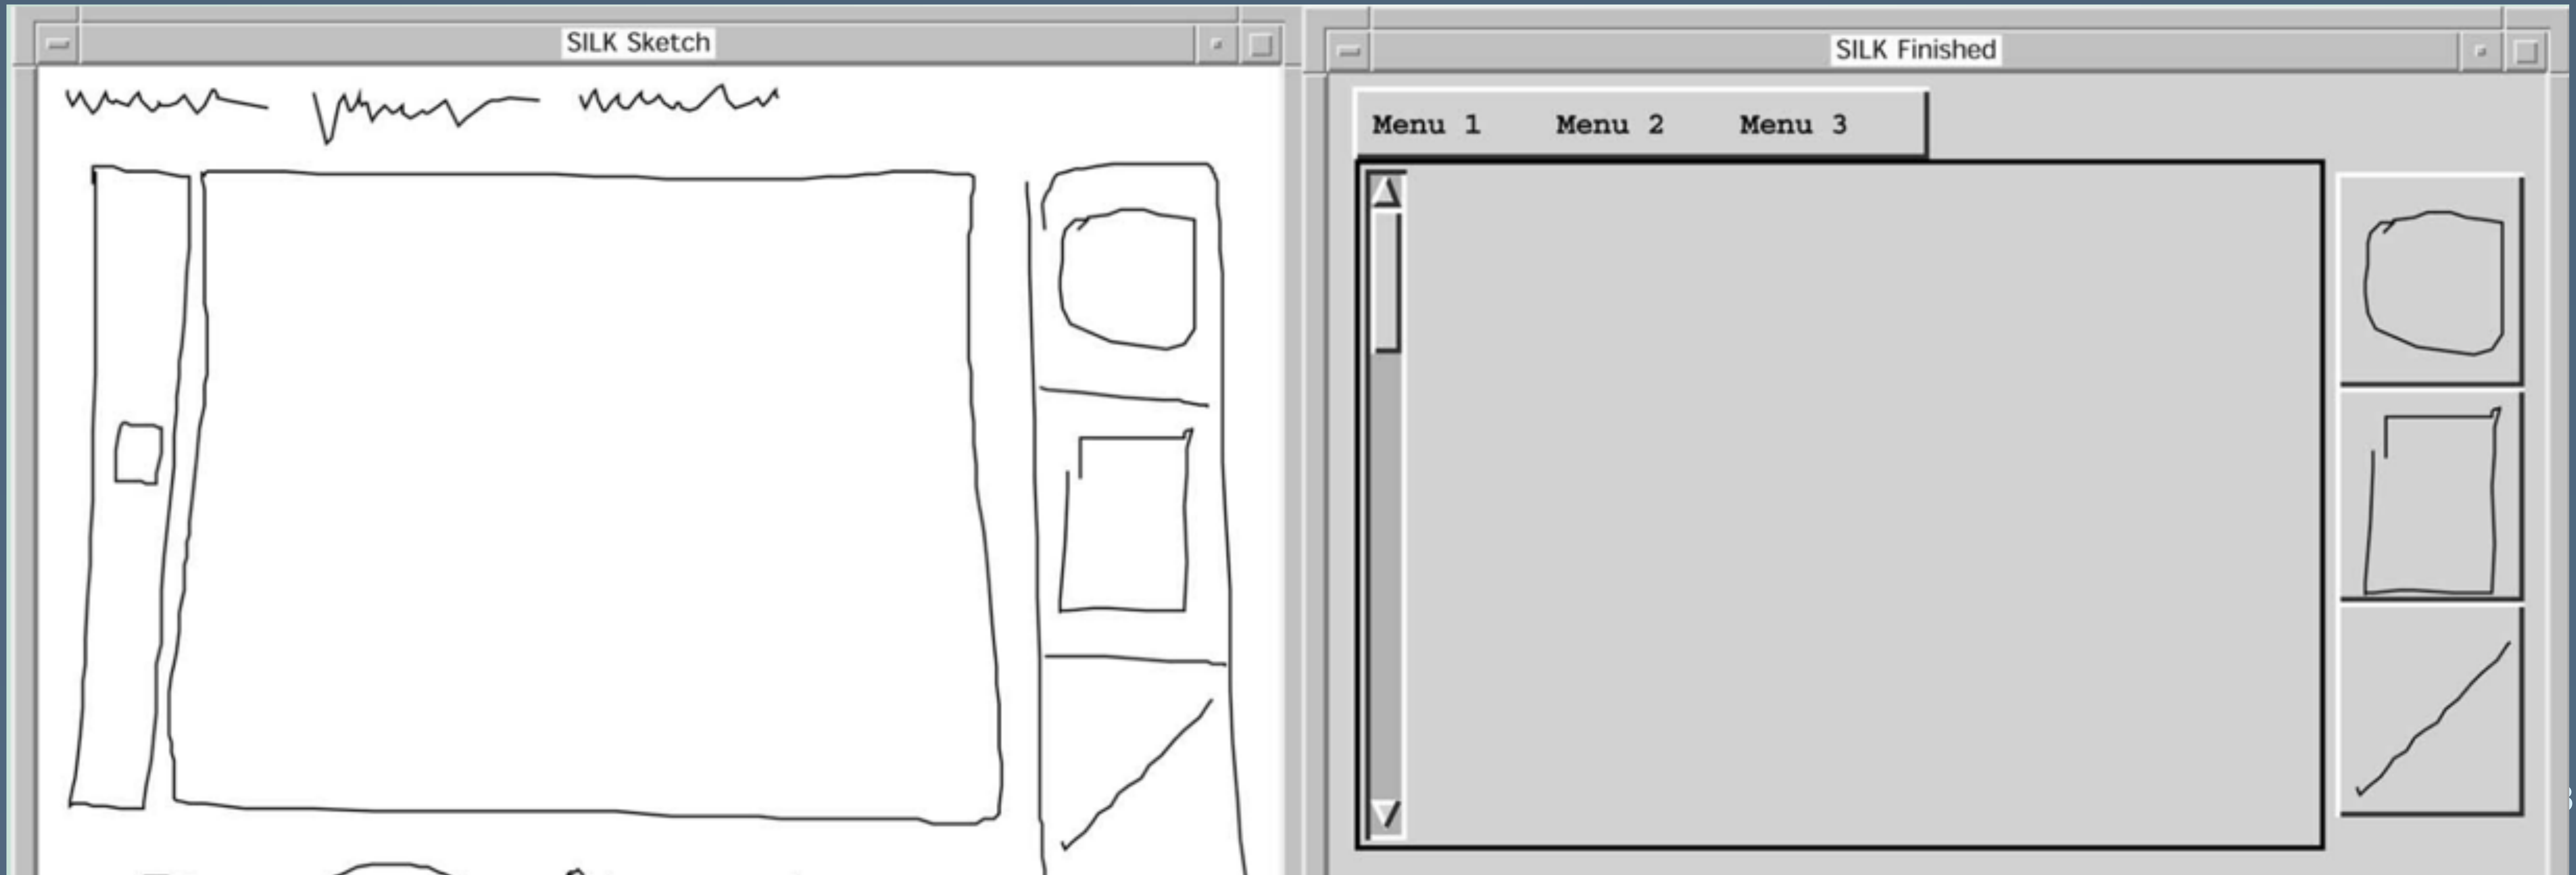

Quantifying Visual Preferences

[Reinecke and Gajos CHI 2014]

- LabInTheWild data via a quiz about which web sites you like

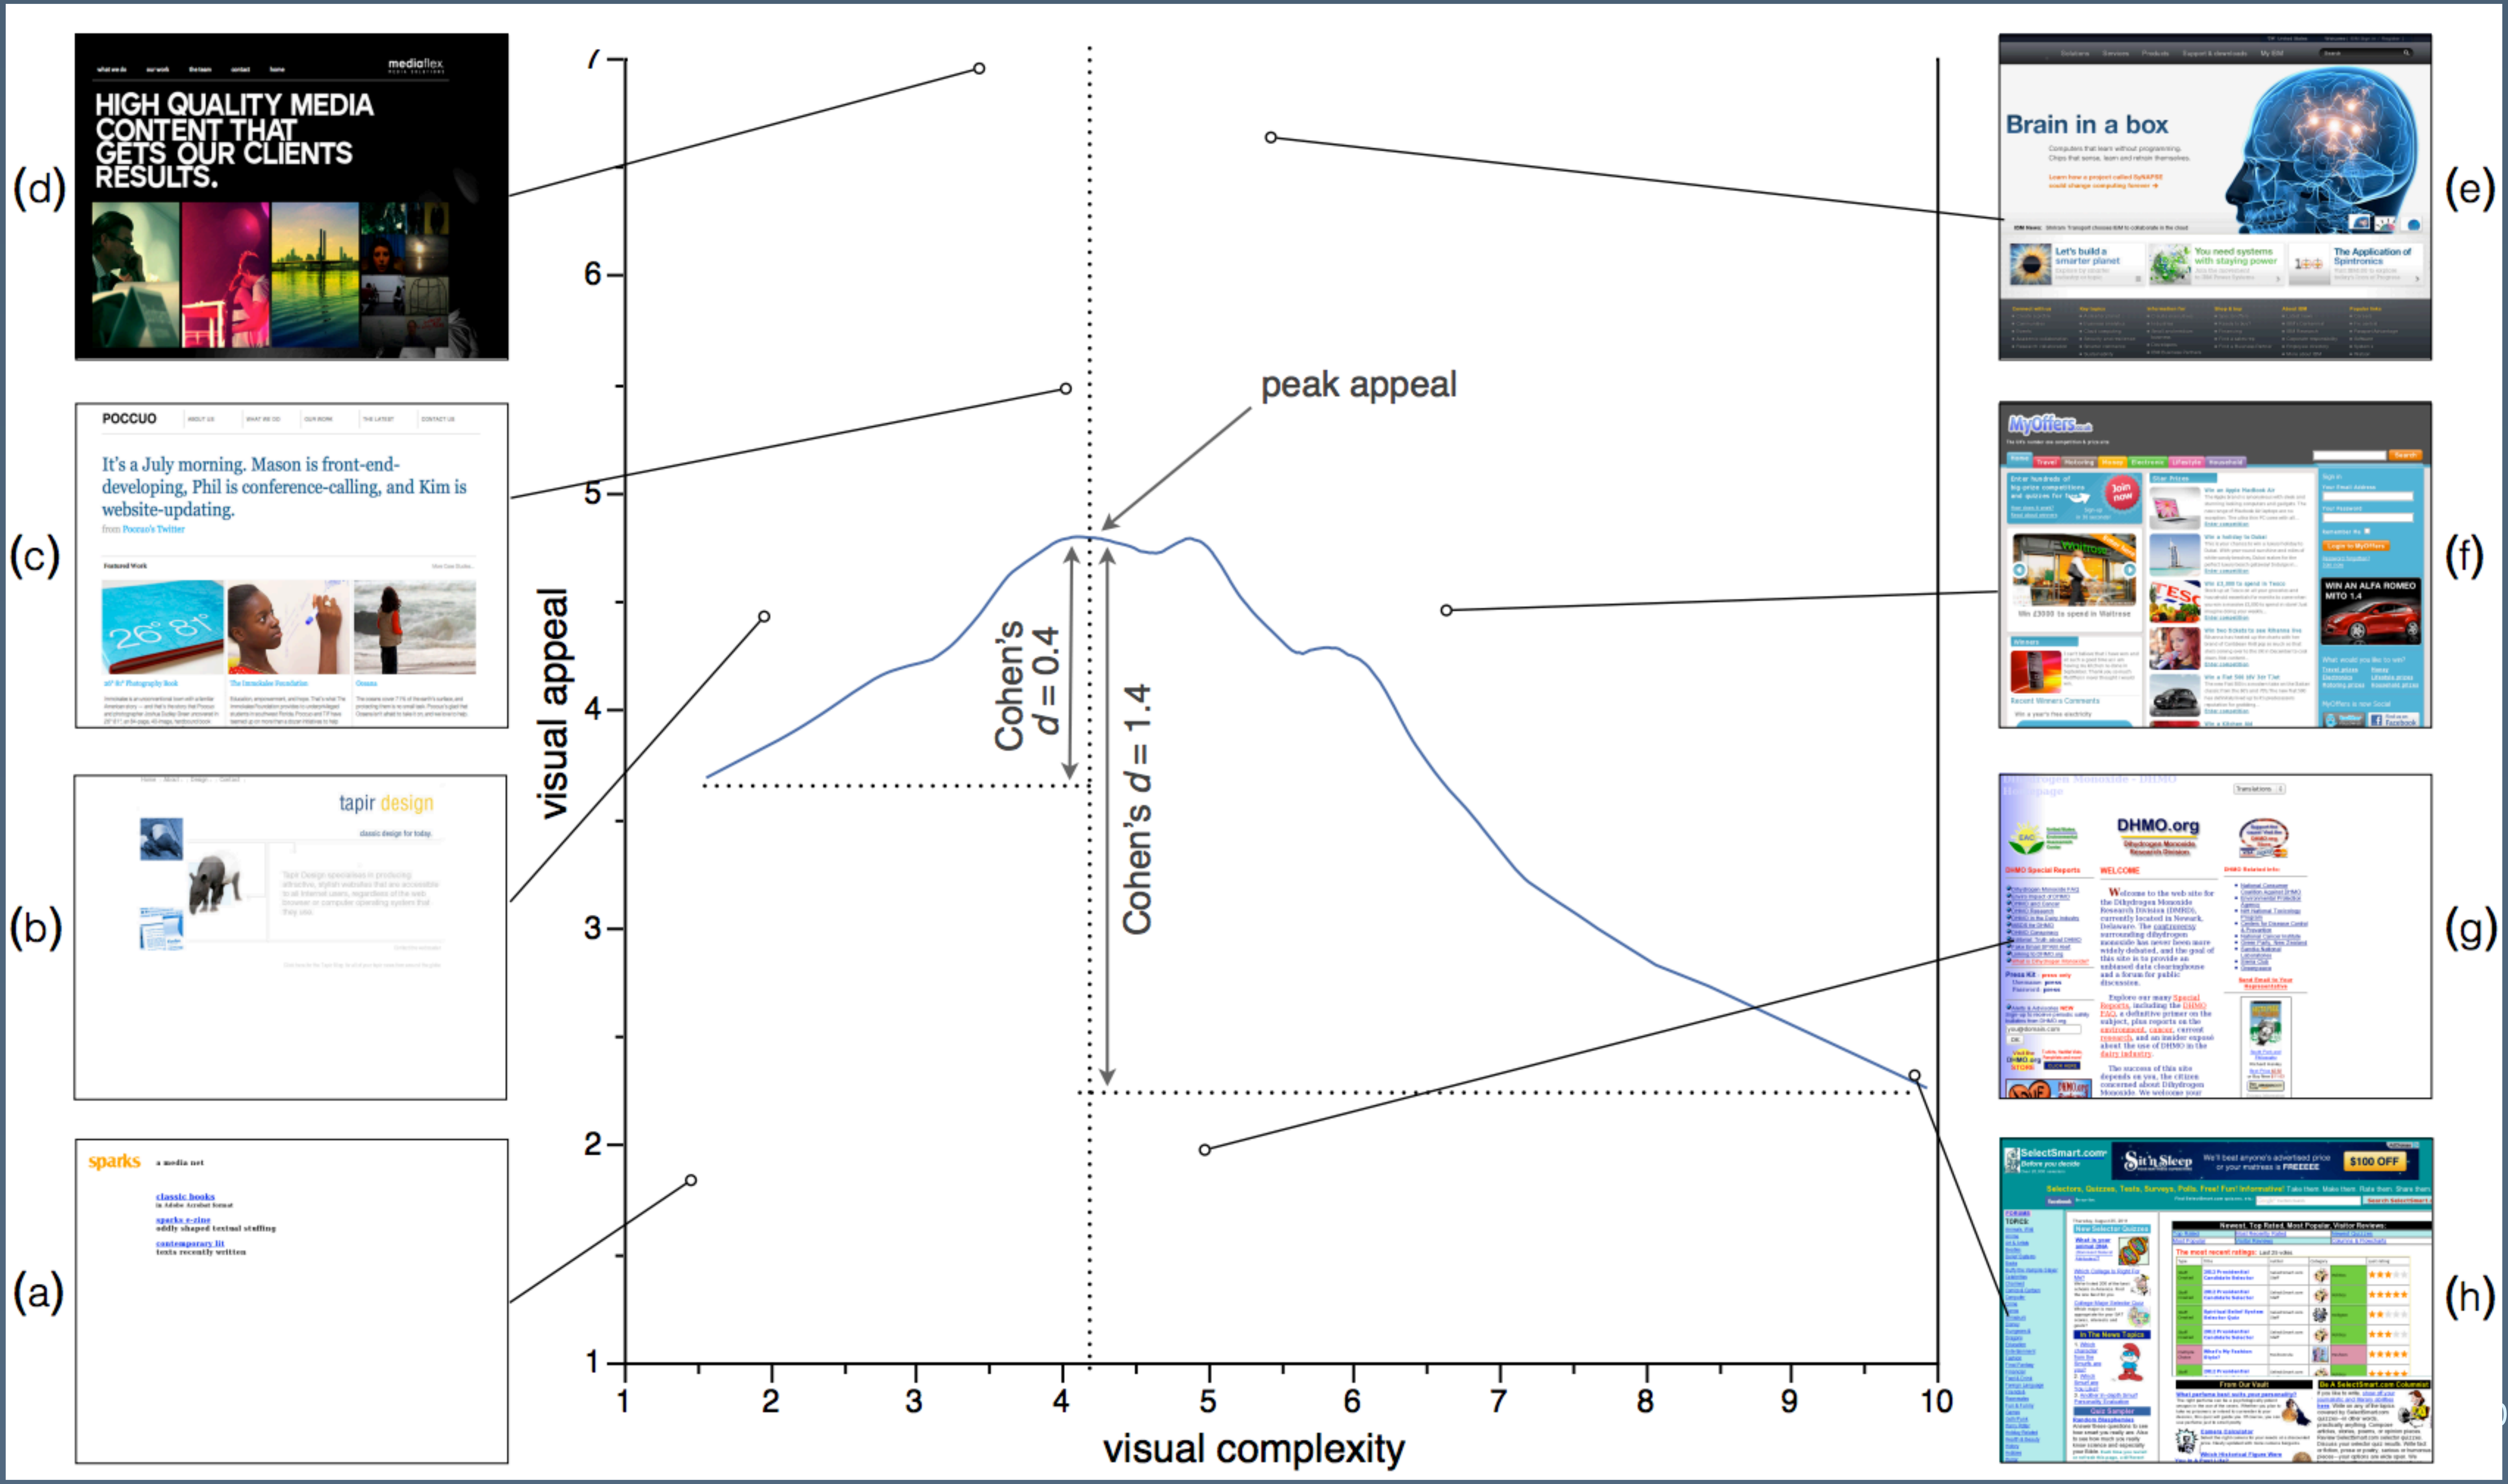
(a)
(b)
(c)
(d)
(e)
(f)
(g)
(h)
visual appeal
visual complexity
peak appeal
Cohen's d = 0.4
Cohen's d = 1.4
1
2
3
4
5
6
7
8
9
10
HIGH QUALITY MEDIA CONTENT THAT GETS OUR CLIENTS RESULTS.
POCCUO
It's a July morning. Mason is front-end-developing, Phil is conference-calling, and Kim is website-updating.
tapir design
sparks
Brain in a box
MyOffers
DHMO.org
SelectSmart.com

DesignScape: interactive layout

[O'Donovan, Agarwala, and Hertzmann CHI '15]

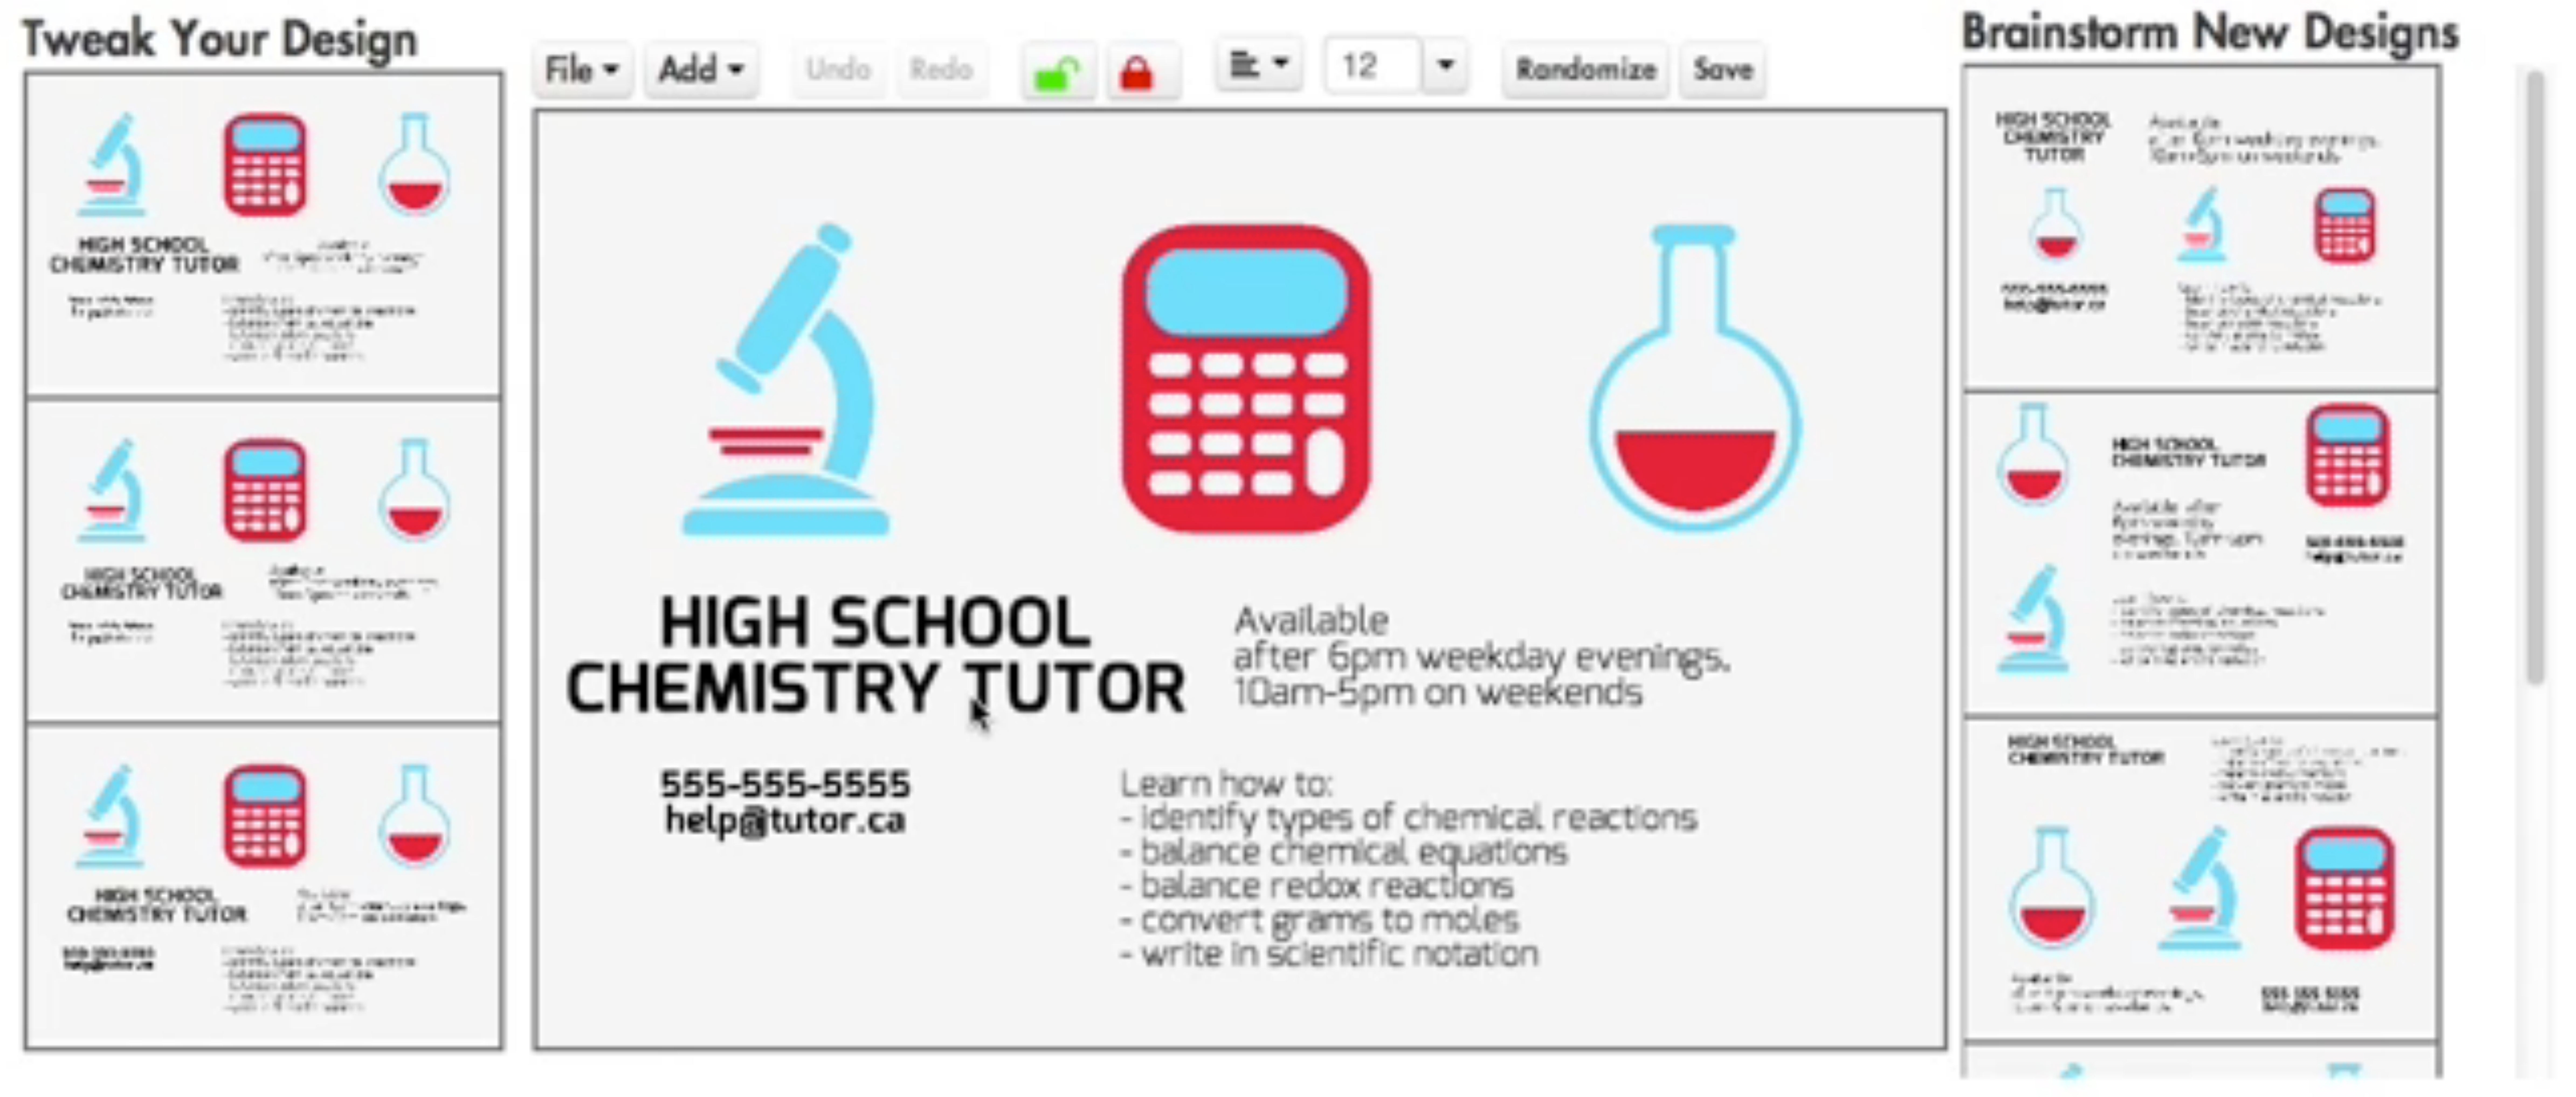

Learning Visual Importance

[Bylinskii et al., UIST '17]

What to take away?

- AI can…
  - Help identify effective designs
  - Help map your design onto a target design
  - Predict how people will react to your design

Voice, gesture, and vision-based interaction

Interaction off the desktop requires intelligence

Amazon Echo
+ Siri

FitBit + Apple Watch

Nest thermostat

AI is now a component of many sensing pipelines

- …as in papers we already covered!

Activity recognition

- Sense the user's physical state by using minimally invasive sensors
- For example, wearing five 2d accelerometers and predicting tasks like walking, watching TV, reading, eating...

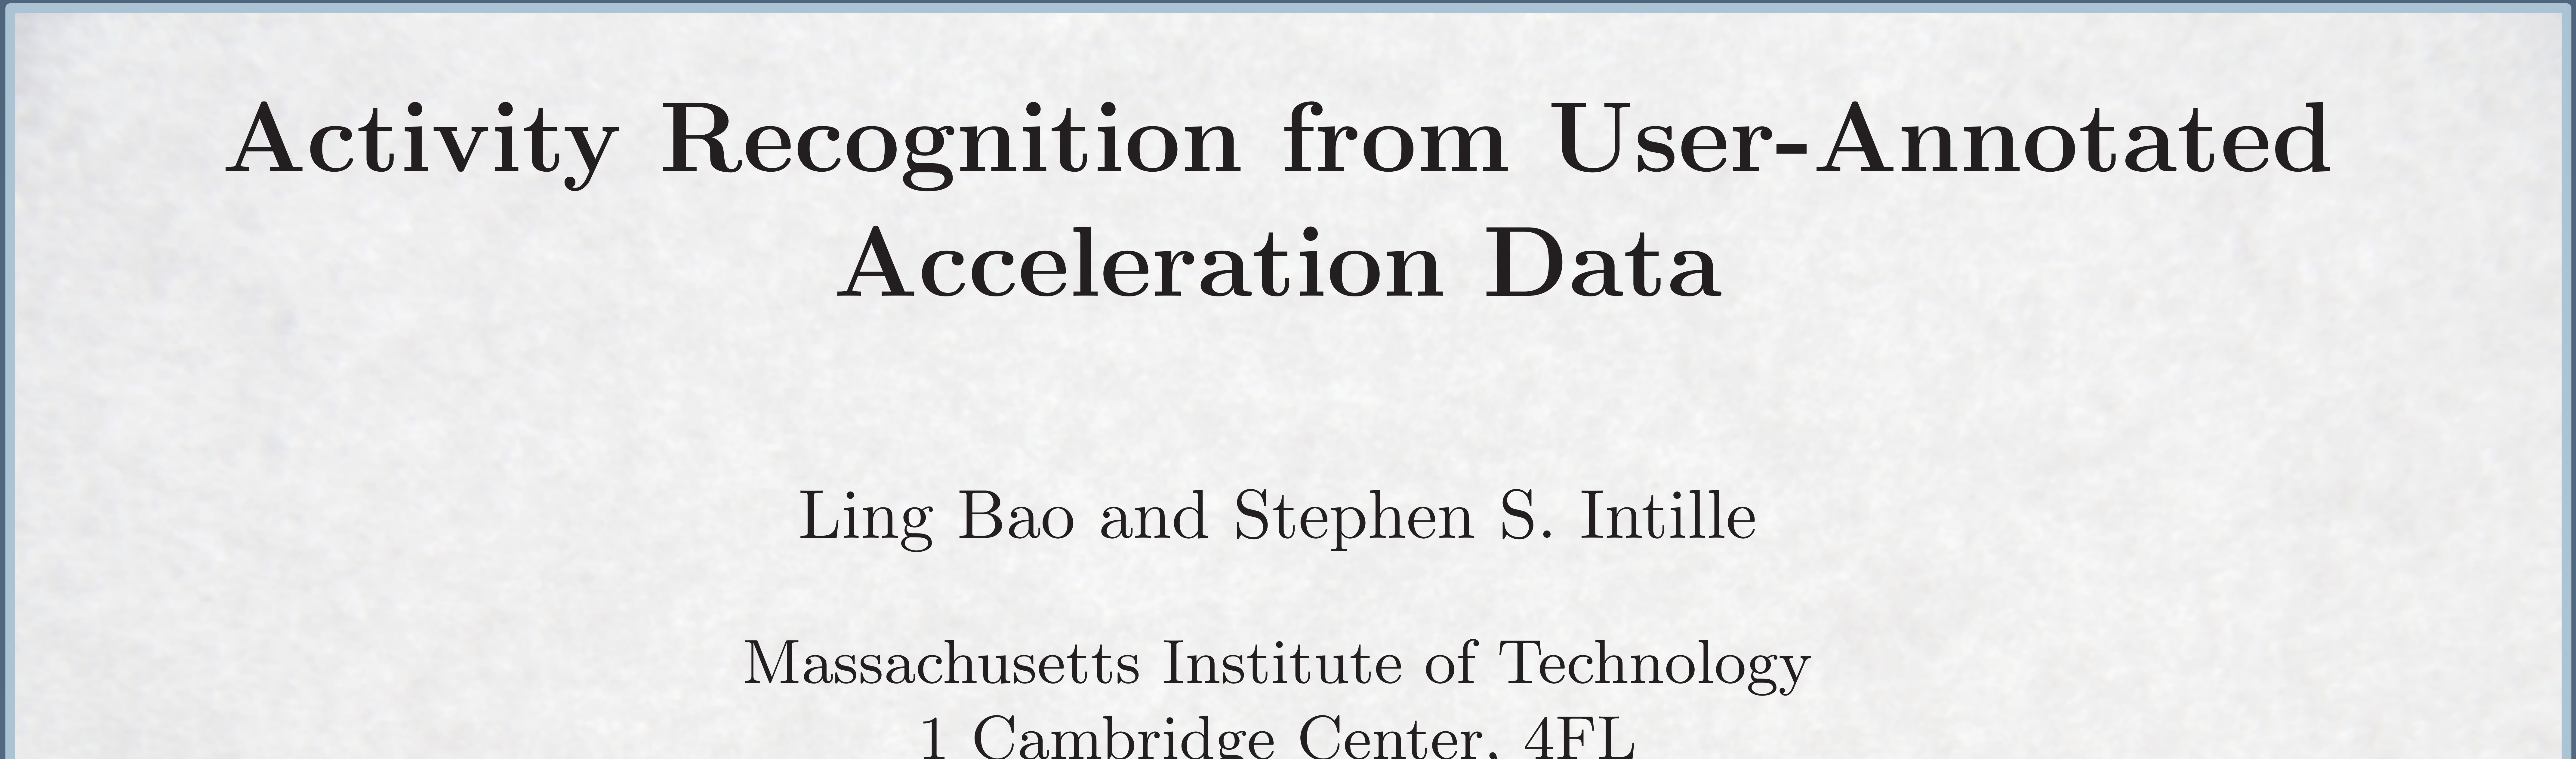

Activity Recognition from User-Annotated Acceleration Data

Ling Bao and Stephen S. Intille

Massachusetts Institute of Technology
1 Cambridge Center, 4FL

Patel et al. At the Flick of a Switch: Detecting and Classifying Unique Electrical Events on the Residential Power Line. Ubicomp '07.

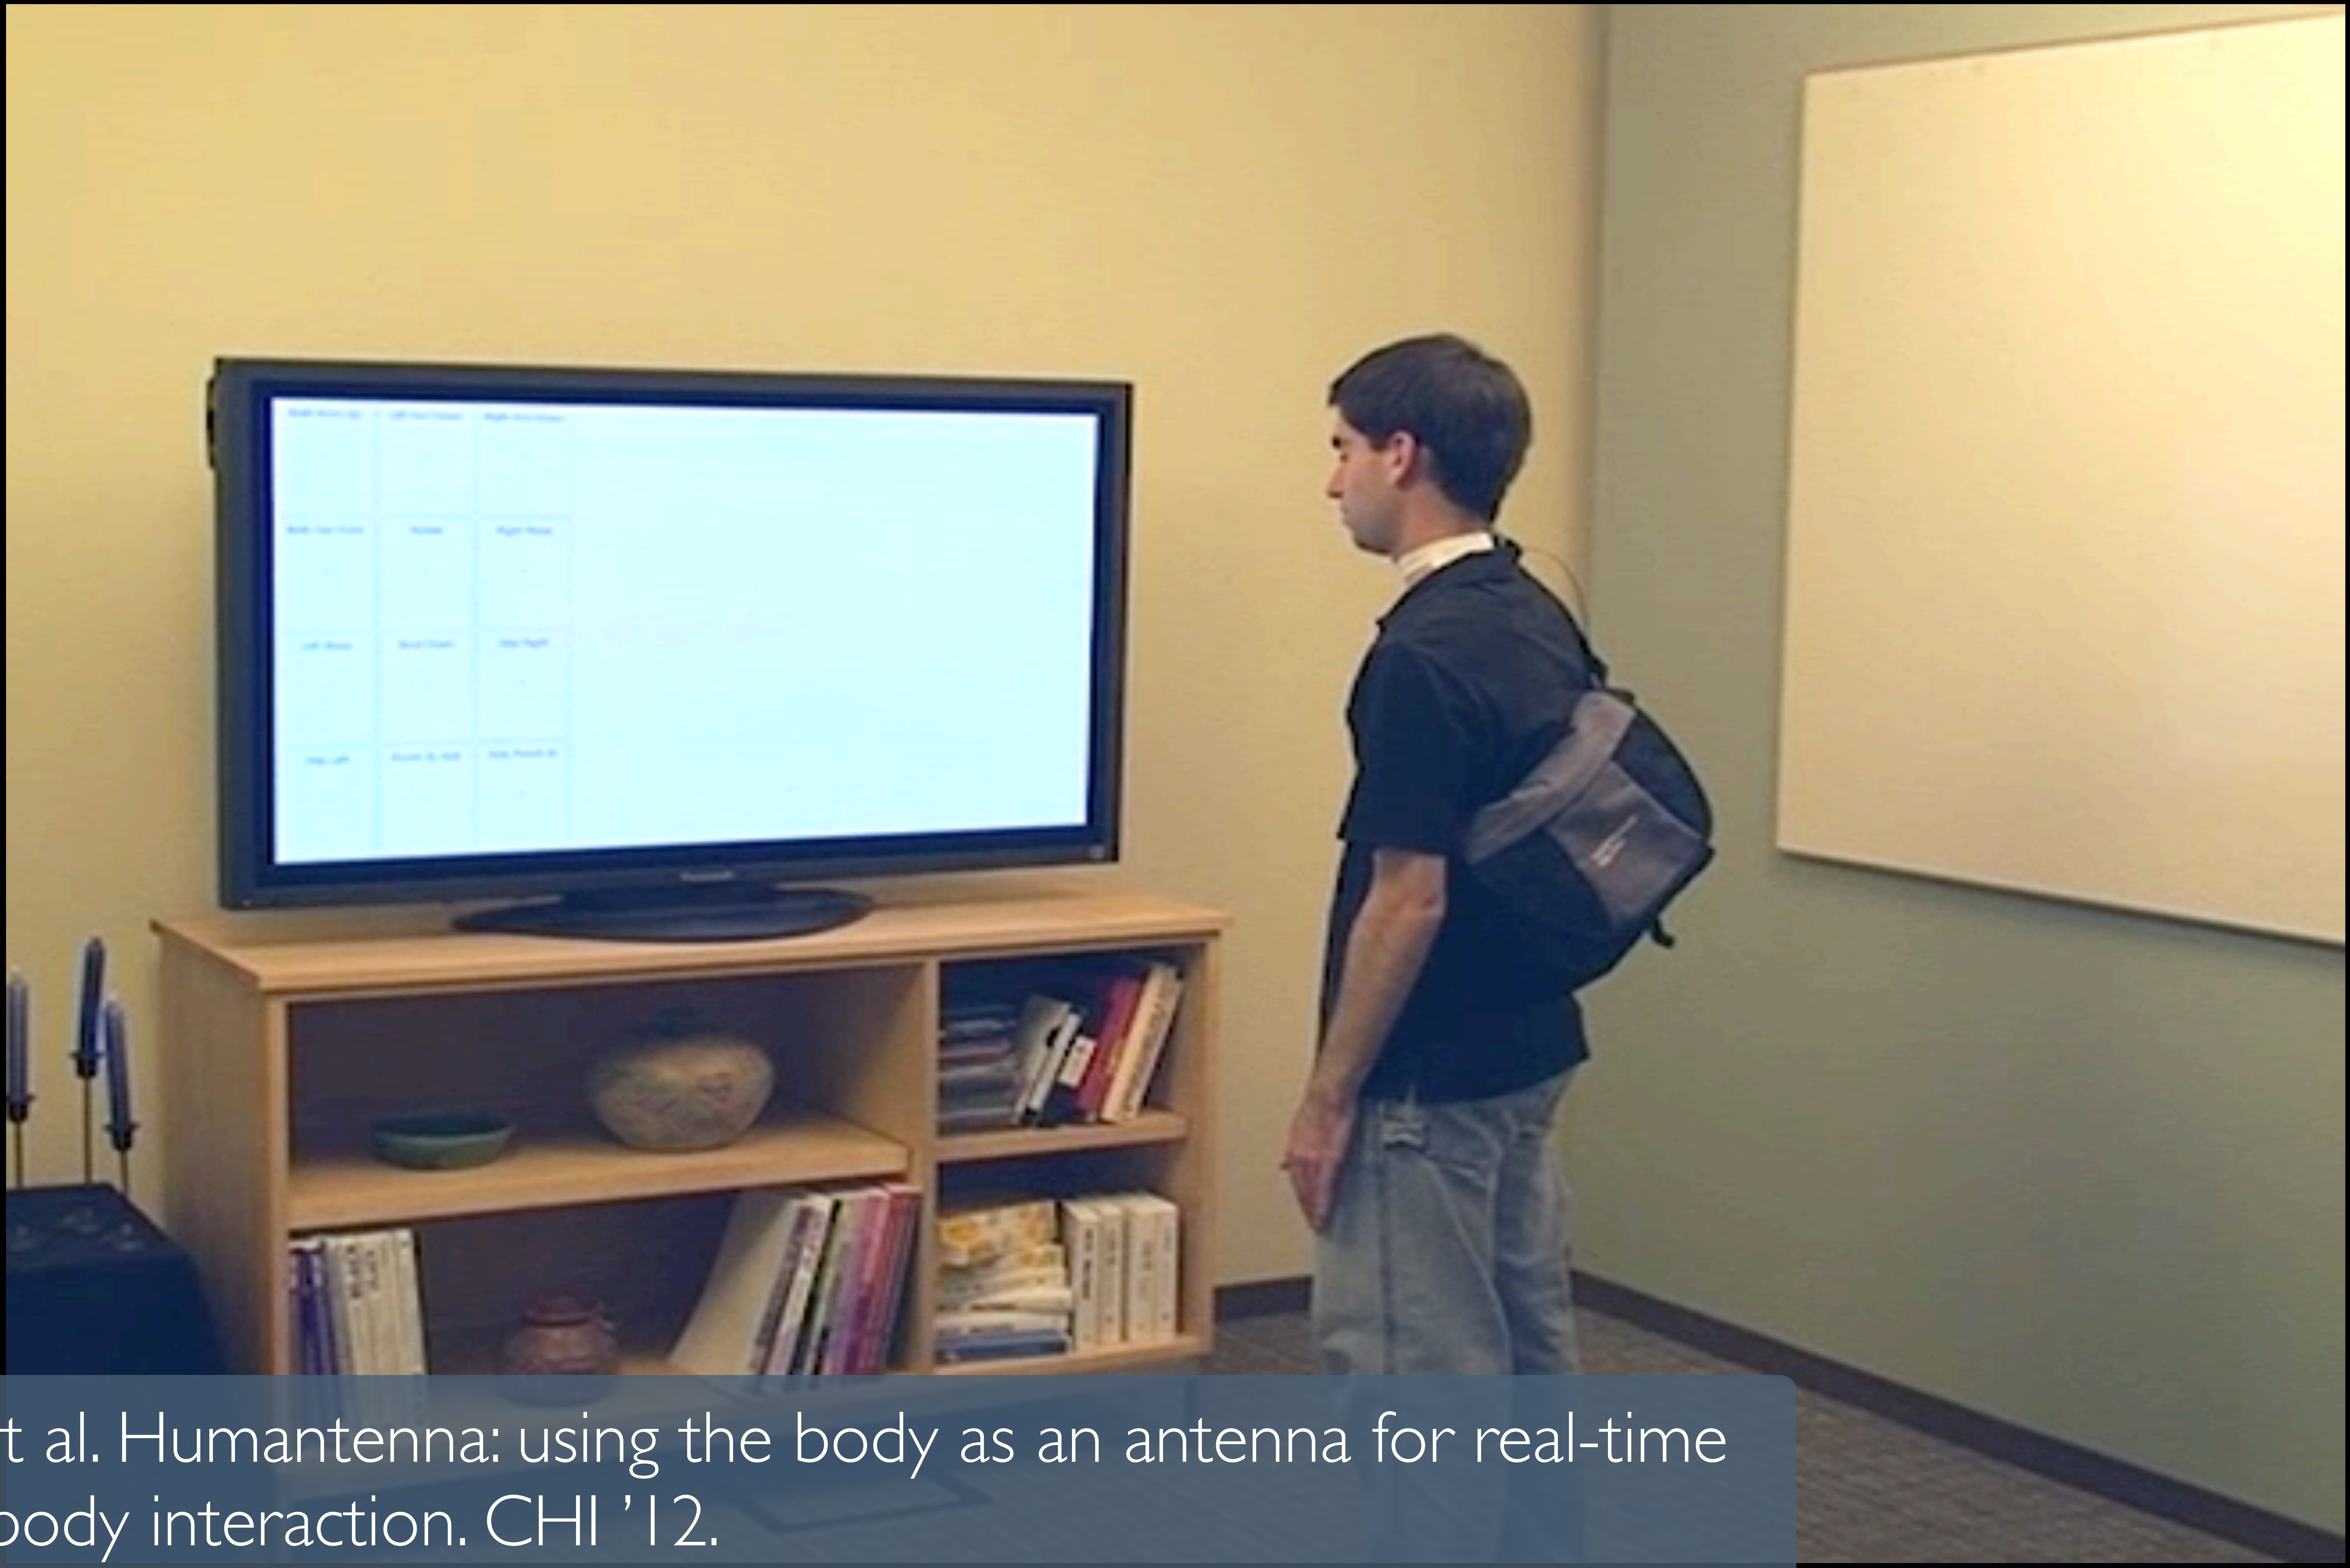

Cohn et al. Humantenna: using the body as an antenna for real-time whole-body interaction. CHI '12.

Harrison, Morris, Tan. Skinput: Appropriating the Body as an Input Surface. CHI '10.

Harrison, Benko, Wilson. Omnitouch: Wearable Multitouch Interaction Everywhere. UIST '11.

EM-Sense

Touch Recognition of Uninstrumented, Electrical and Electromechanical Objects

Gierad Laput Chouchang Yang Robert Xiao
Alanson Sample Chris Harrison

Carnegie Mellon University
Disney Research

Laput, G. et al. 2015. EM-Sense: Touch Recognition of Uninstrumented, Electrical and Electromechanical Objects. UIST '15.

Acoustics

Laput et al. Acoustruments: Passive, Acoustically-Driven Interactive Controls for Hand Held Devices. UIST '15.

Multimodal interaction

Using simultaneous inputs

- Sensor fusion can help disambiguate multiple noisy signals

Put That There

Speech N-best	Gesture N-best	Multimodal N-
Zoom in	Checkmark	Zoom out
Show info		
Show all		
Zoom out		

Quickset

[Oviatt, CHI '99]

Reflections

- AI is a powerful tool, but brings massive user interaction problems as a result of the uncertainty it introduces for the user
- Smart interaction design can hide or manage that uncertainty.
- "Don't let your AI write a check that your UI can't cash;
  Don't let your UI write a check that your AI can't cash."
  - Eytan Adar, University of Michigan