What is (and isn't) HCI research?

MICHAEL BERNSTEIN
CS 376

Announcements

- Project brainstorm 1 due at 11:59pm Friday
- Group mixer at 5:40pm Monday

Why are we here?

- This is a good class for you if: you are looking to get engaged in HCI research, or want to deepen your understanding of it

- HCI Research is a graduate-level research seminar course, not your typical HCI project course. It requires mastery of HCI concepts or concepts in adjacent fields.

What is HCI research?

Research vs. practice

- Research introduces a fundamental new idea into the world of human-computer interaction.

- An idea articulates a high-level approach to design, or a social scientific insight.
  - Examples from Monday: making bits "tangible", sensing exercise activity using accelerometers, embedding interfaces into clothing, projecting interfaces and using a depth sensor to detect interaction

Research vs. practice

- While they are situated in a particular context, ideas are generalizable and can be applied to new situations.

How do I know?

For design and engineering ideas

- Ask yourself: is it possible to solve this problem using a set of techniques that is already known?
  - If so, it is not research.
  - If not, it is more likely to be research.
- Ask yourself: has this technique been introduced in other HCI contexts?
  - If so, it is not research.
  - If not, it is more likely to be research.

How do I know?

For design and engineering ideas

- Ask yourself: is the problem one that is known to the HCI community?
  - If so, it is not research.
  - If not, it is more likely to be research.
- A good idea may be old news! (Ex: Apple Watch)

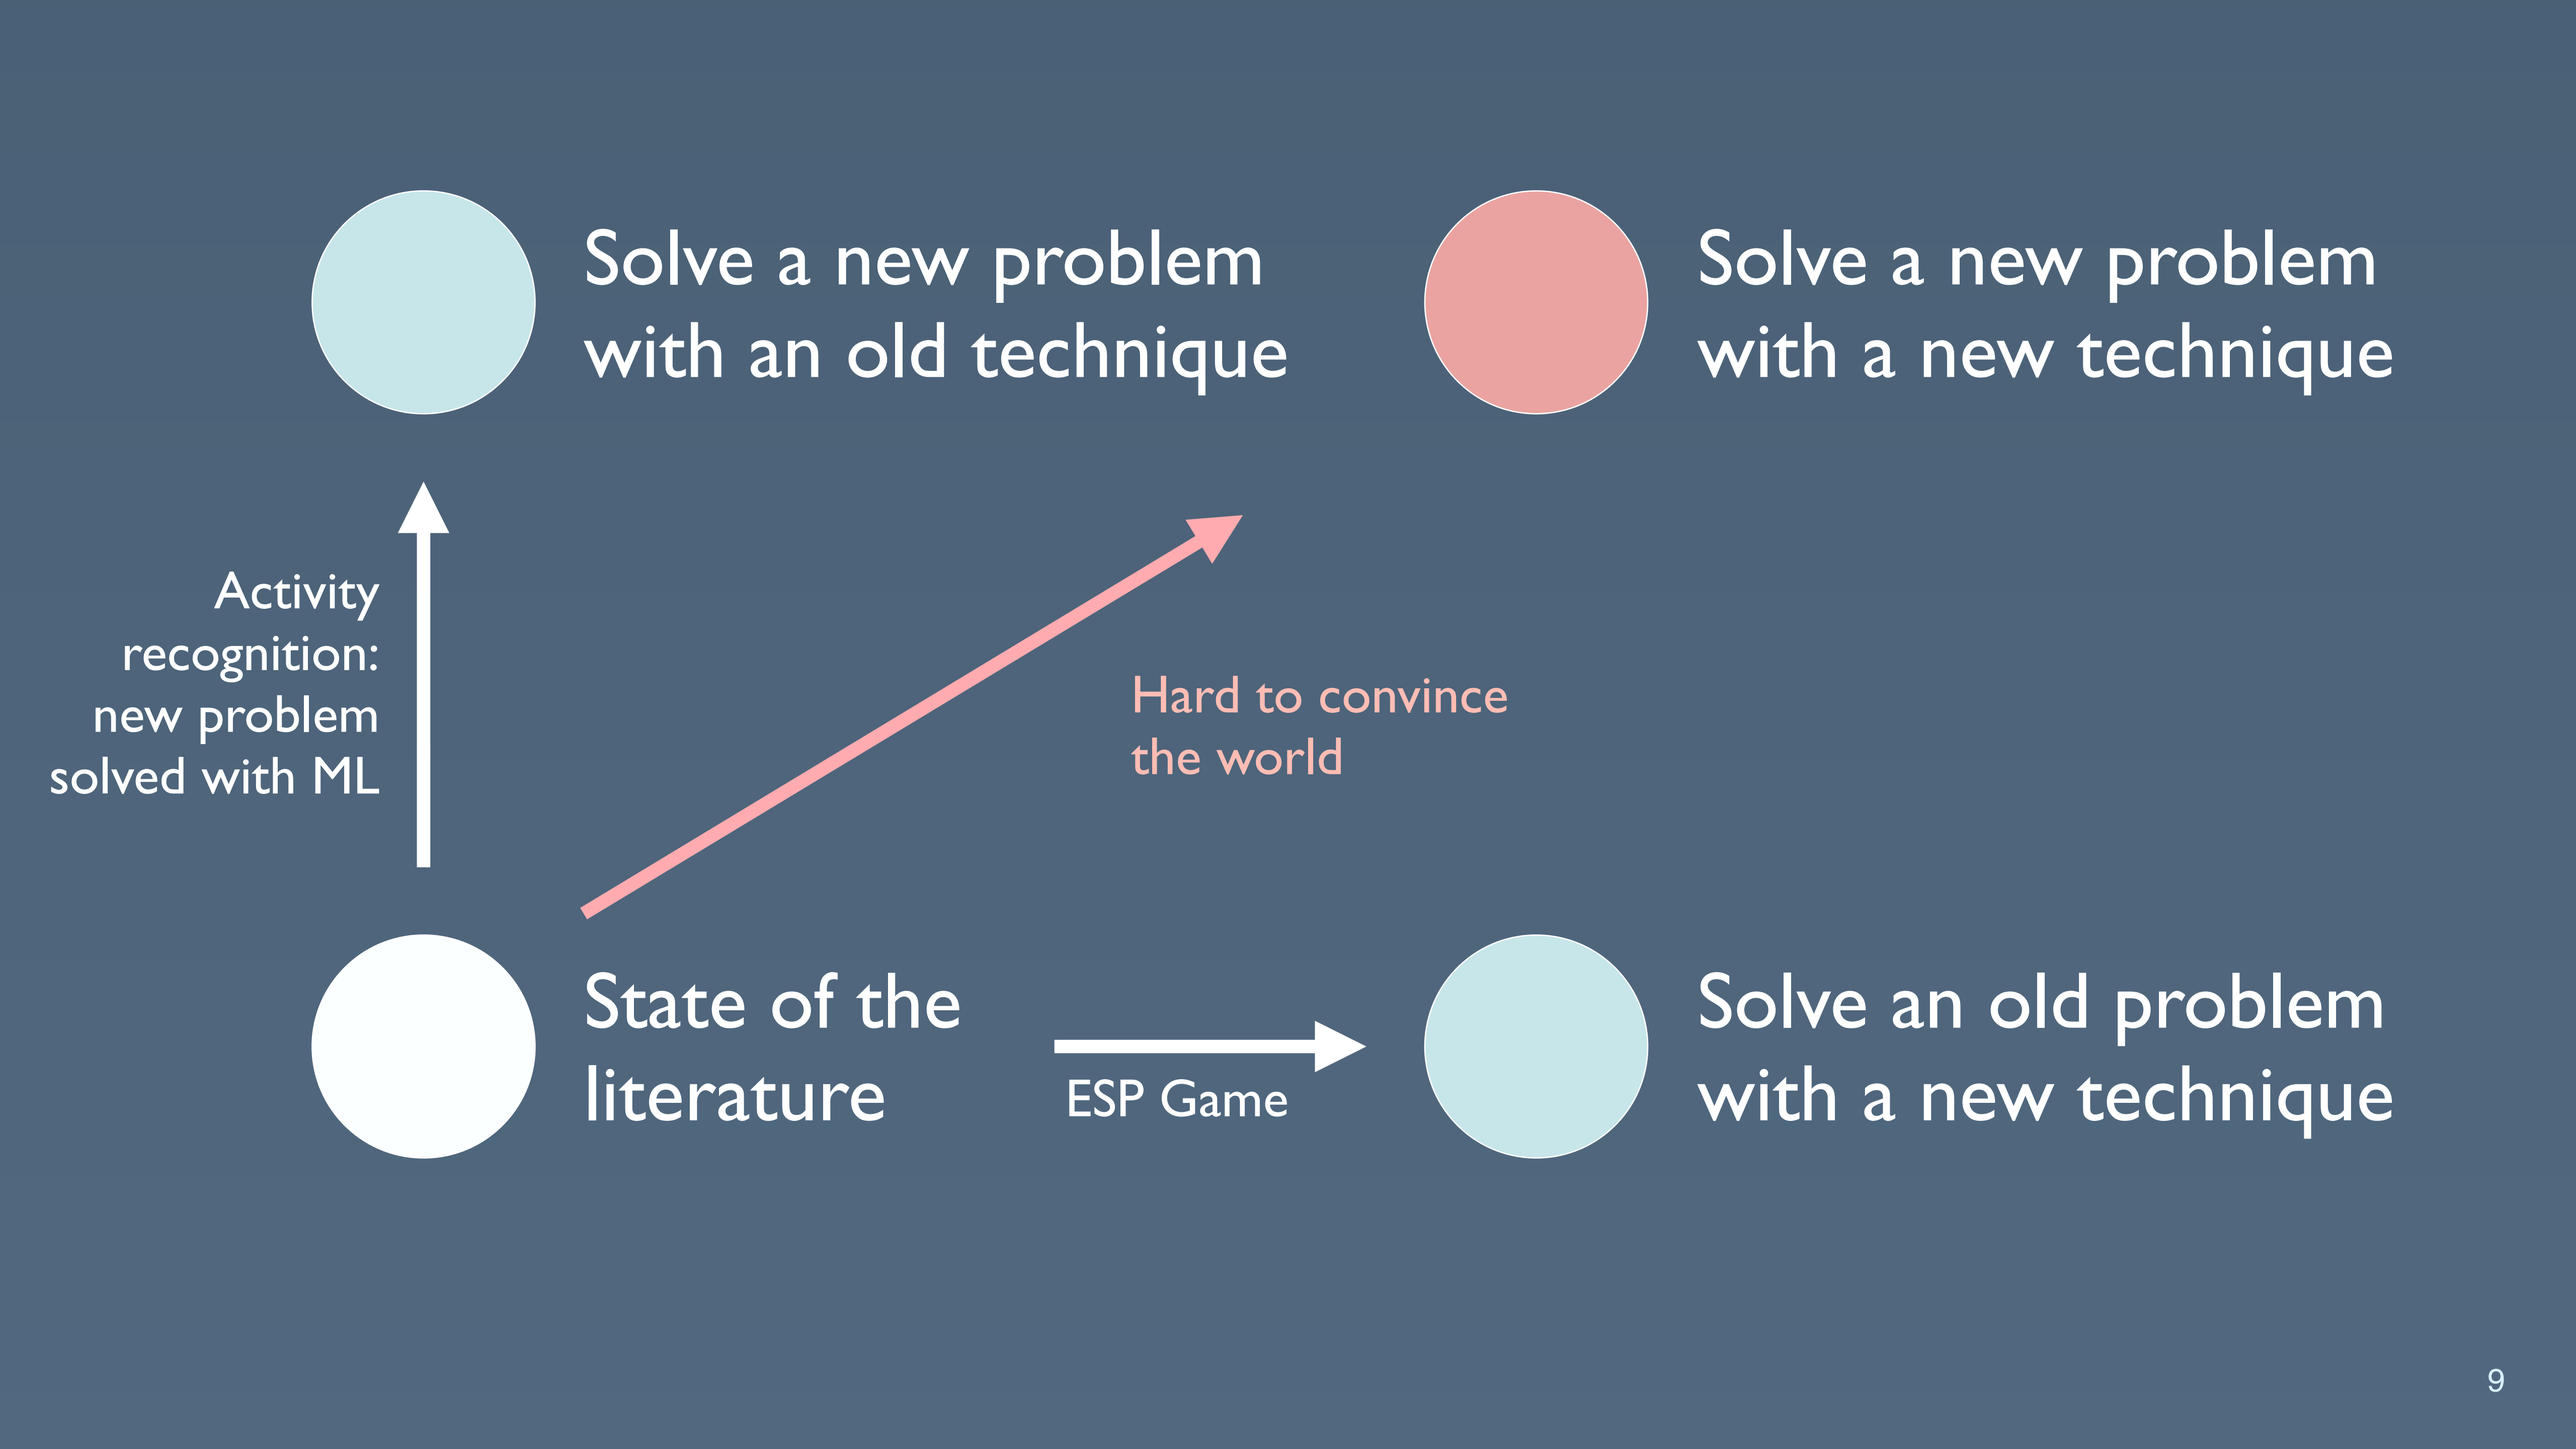
Solve a new problem with an old technique
Solve a new problem with a new technique
Activity recognition: new problem solved with ML
Hard to convince the world
State of the literature
ESP Game
Solve an old problem with a new technique

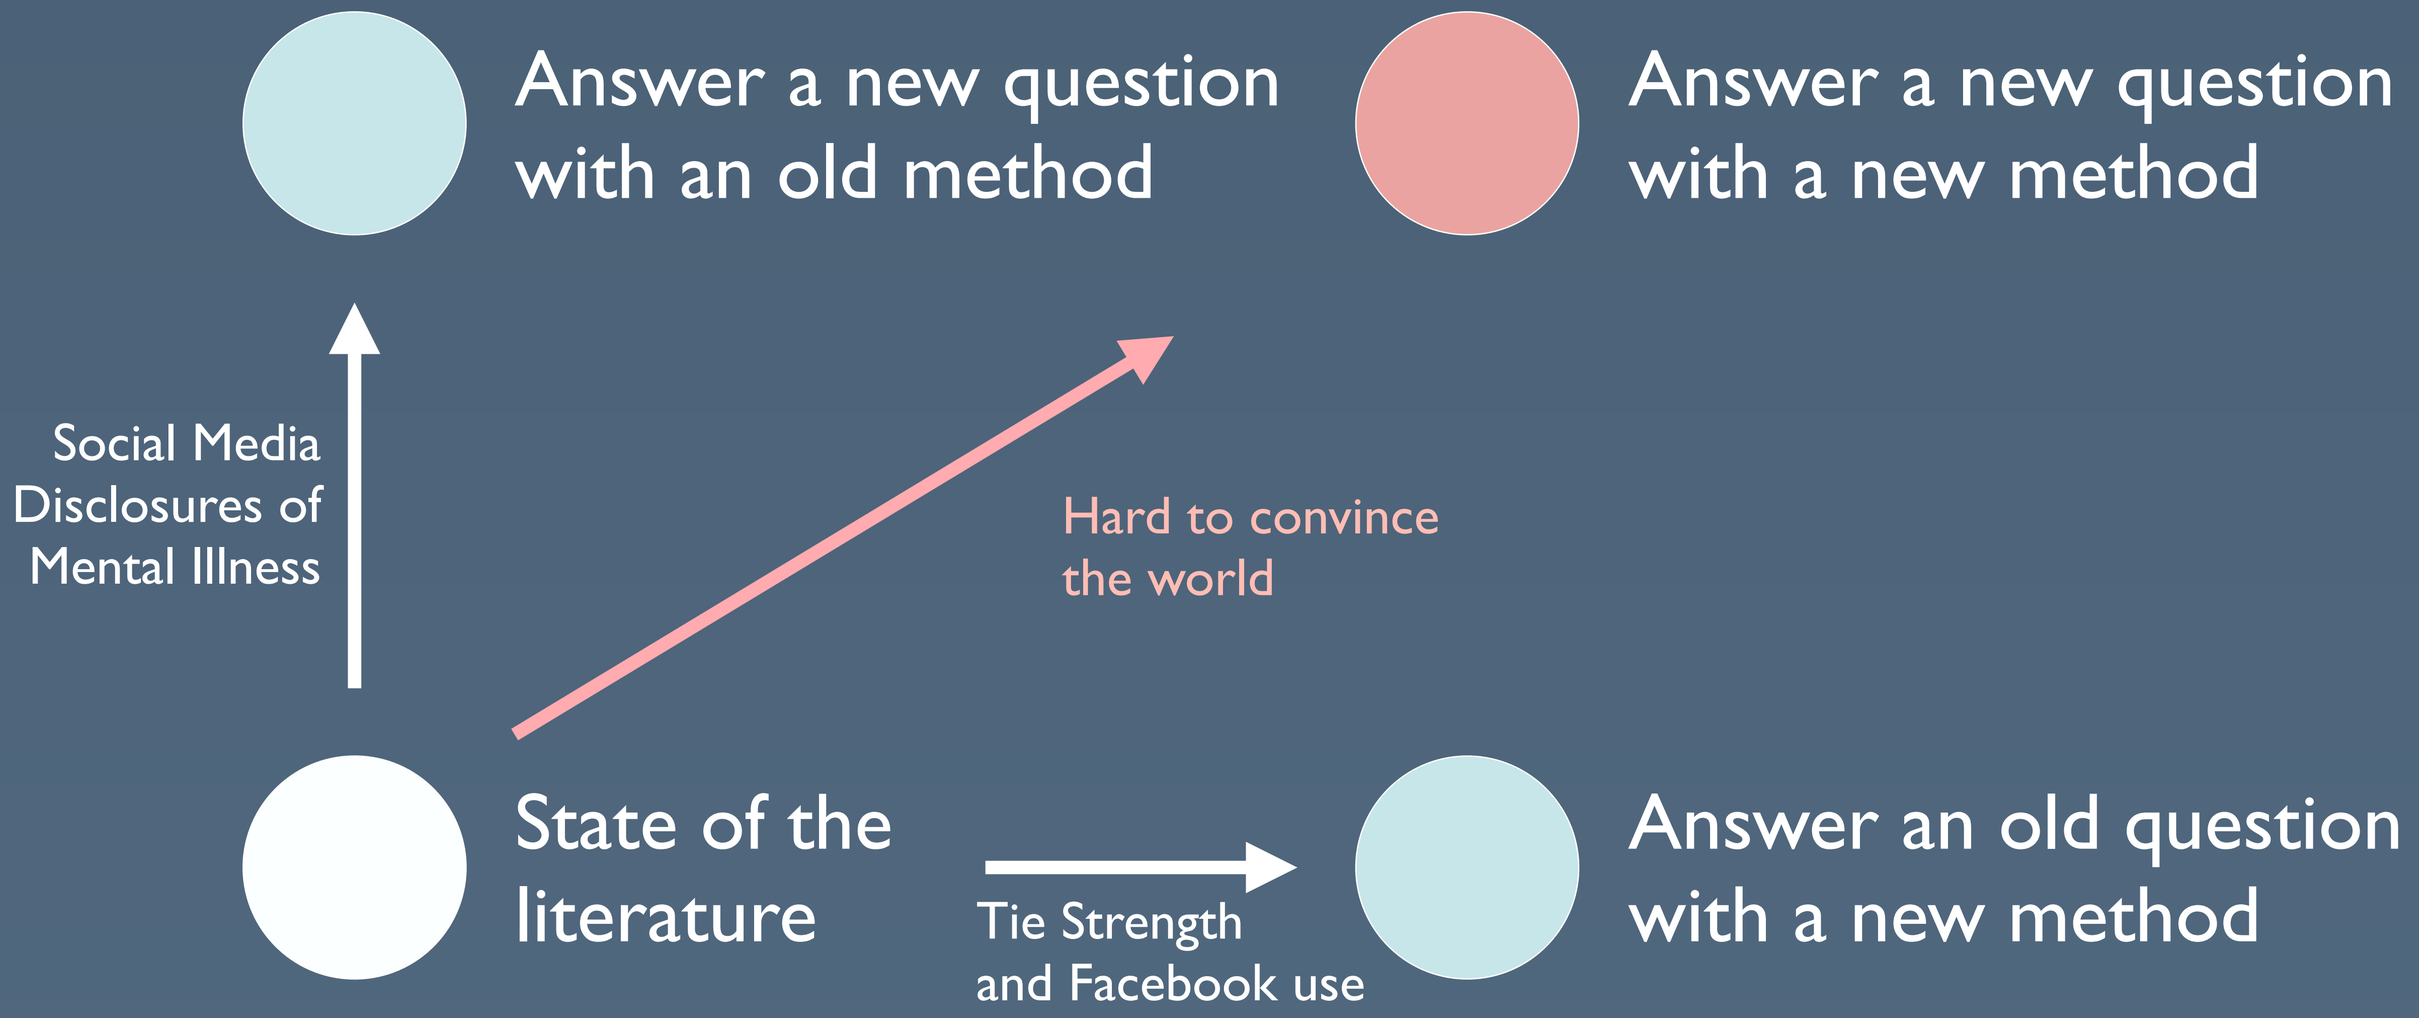
Answer a new question with an old method
Answer a new question with a new method
Social Media Disclosures of Mental Illness
Hard to convince the world
State of the literature
Tie Strength and Facebook use
Answer an old question with a new method

How do I know?

For social science ideas

- Ask yourself: is this phenomenon describable or is this question answerable using our existing social scientific knowledge?
  - If so, it is not research.
  - If not, it is more likely to be research.

- A good idea may be old news! (Ex: People using Wikipedia a lot but rarely contribute content — social loafing and diffusion of responsibility)

Examples

“Location sensing to autoshare shopping habits.”

Could be research if:

- nobody has ever proposed shopping as a problem
- your solution generalizes to other problems and has never been demonstrated
  - e.g., sensing location based on smell
  - e.g., public shaming to change behavior

Probably not research if:

- you are applying a solution that we know about already to a problem that we know about already

“A mirror to show me how I’d look if I lost weight”

Could be research if:

- nobody has ever studied how people use technology to envision health outcomes
- your solution generalizes to other problems and has never been demonstrated before (e.g., a model that generates realistic weight loss alterations)

Probably not research if:

- you are applying a solution we know about already to a problem that we know about already
  - e.g., this is solely a user-centered design project
  - e.g., you are not contributing a new technique or domain

“Researching the new hot app SnortChat.”

Could be research if:

- SnortChat exemplifies an interesting point in the design space, and we use it to understand that axis or design space
- Theories suggest that SnortChat should work one way or should not succeed, but it’s the opposite.

Probably not research if:

- You have trouble articulating what broader design choice SnortChat is an example of
- We have studied applications like SnortChat in the past, and SnortChat works the same way
- You have to put the word “researching” in the title

“I’m doing research already!”

- You have two options for your final project:
- The “Macro” option
  - Continue on your research path with the faculty member
  - Write up the overall project as your final project submission
- The “Micro” option
  - Carve out a sub-research problem of the larger project, or a riff on the project, and tackle it end-to-end within the scope of the class
- Either way, submit the idea brainstorms with your team. The point of the assignment is to train you to articulate research concepts.

Social Computing

MICHAEL BERNSTEIN
CS 376

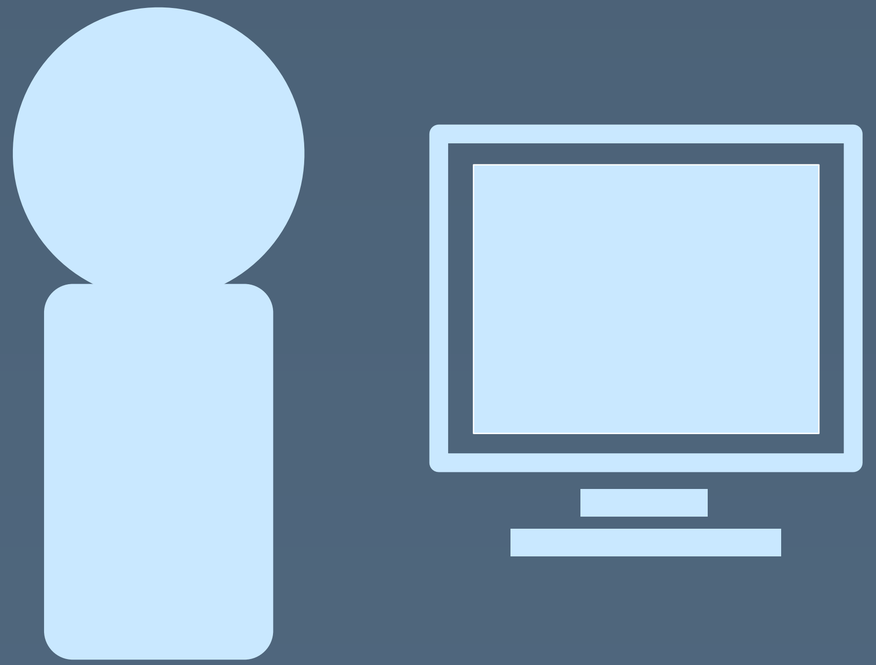

Human-computer interaction

Ubiquitous computing

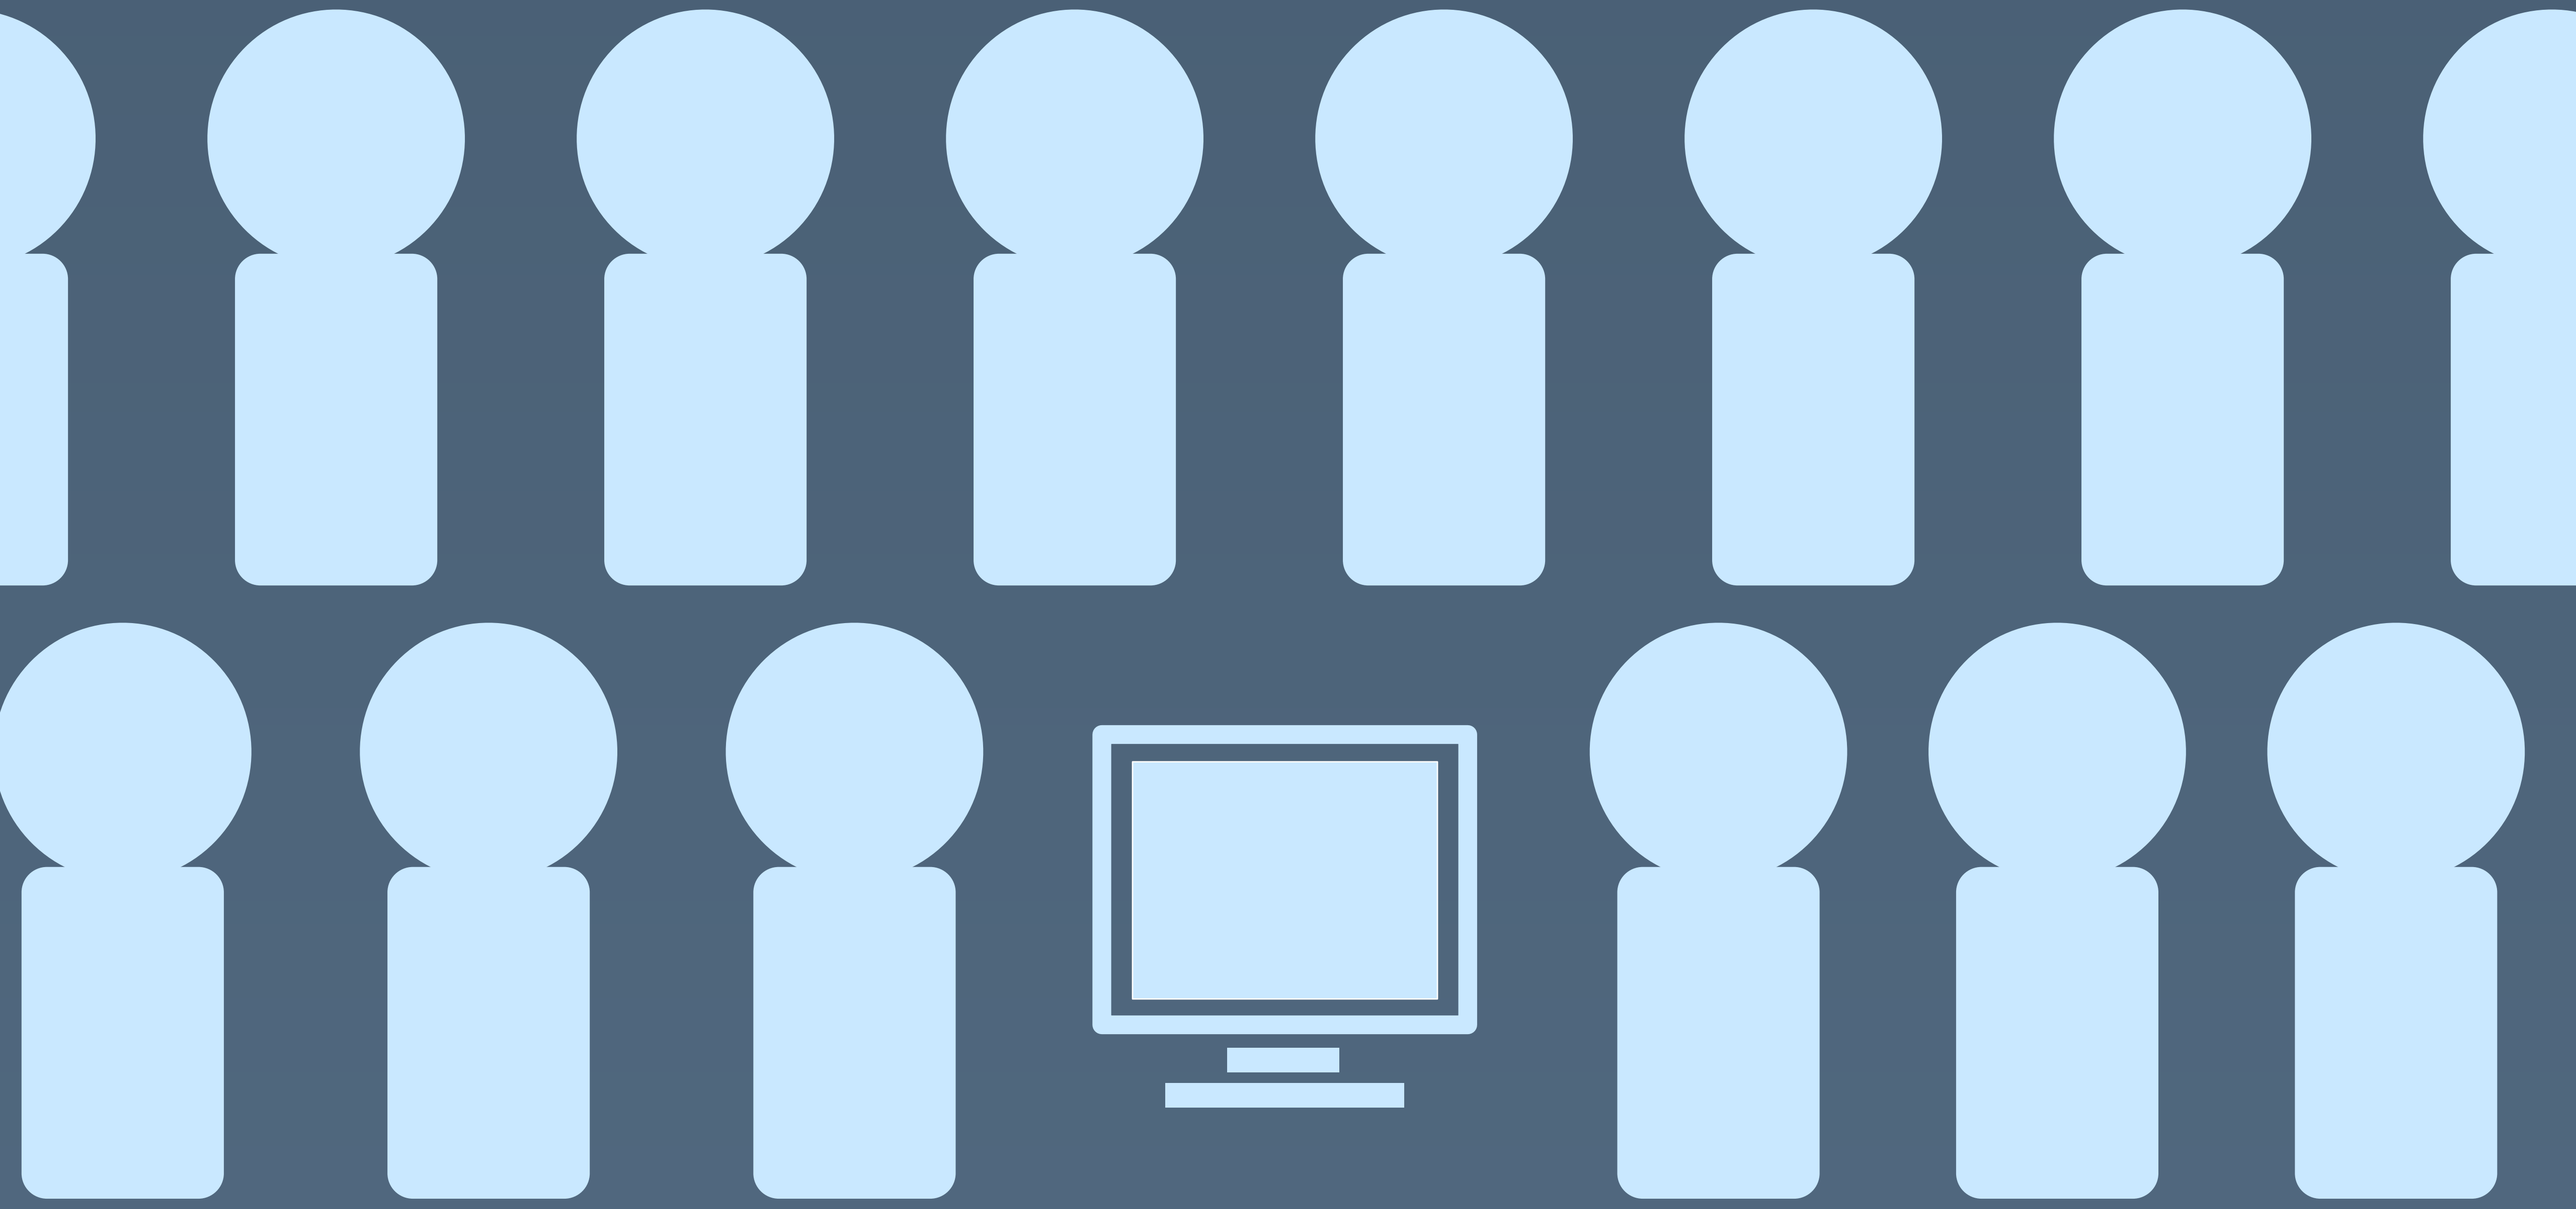

Social computing

Social computing goals

- Design new forms of large-scale human interaction
- Take advantage of the technology-mediated nature of the medium to understand human relationships
- Guide large groups of people to achieve complex, large-scale goals

The intellectual challenge of social computing design

- User-centered design perspective:
  - "The social-technical gap is the divide between what we know we must support socially and what we can support technically." [Ackerman 2000]
- Invention and design thinking perspective:
  - By lowering the transaction costs to connect with others, what kinds of unstated needs and new behaviors might the internet empower? [Shirky 2008]

The intellectual challenge of social computing science

- How has technology-mediated interaction changed our relationship with each other and with the world?
- By manipulating the technology platform, can we learn how people interact with each other?

Sociotechnical system

Emergent behaviors result from interactions
between social relationships and technological interventions.

Design

Online communities

- Reducing editor dropouts due to 'hazing' in Wikipedia [Halfaker, Geiger and Terveen, CHI '14]

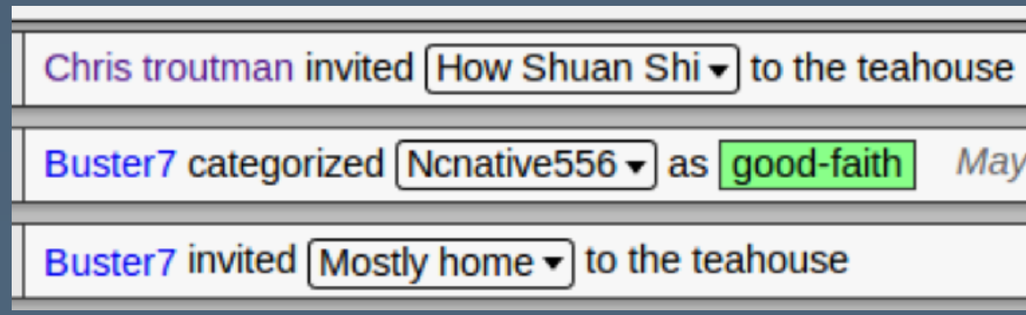

- Encouraging collective action online [Salehi et al. 2015]

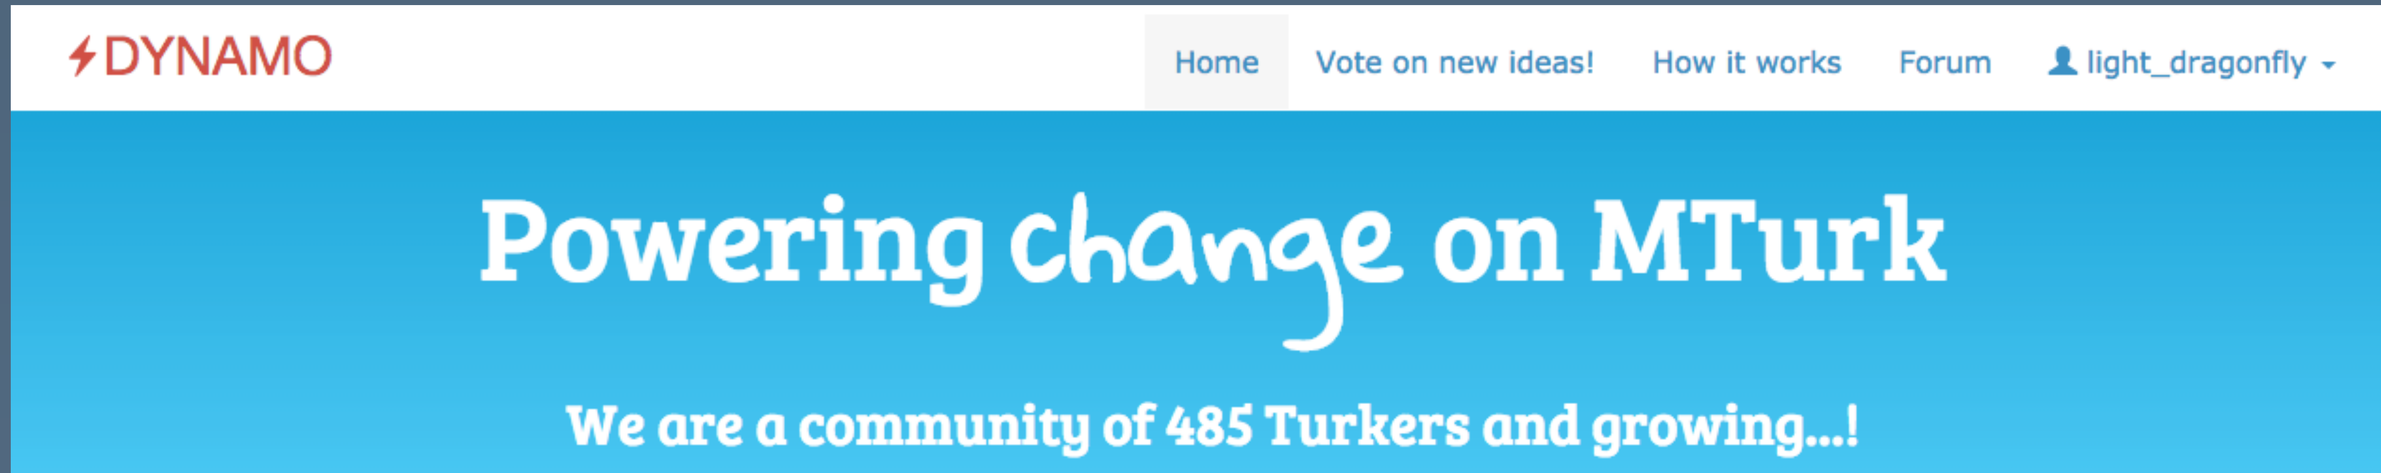

Social behavior as signal

- Learning from one user's behavior to predict another user's behavior
  - GroupLens, aimed at personalizing and filtering usenet
    [Resnick et al., CSCW '94]

- Designing peer assessment systems for MOOCs
  [Kulkarni et al., TOCHI '14]

Understanding

Social capital

- Collective benefits derived from involvement in social environments
- In other words: friends with benefits
- Bridging social capital
  - Social capital built up with a community or across groups (e.g., between any Stanford students)
- Bonding social capital
  - Social capital built up between close friends and family

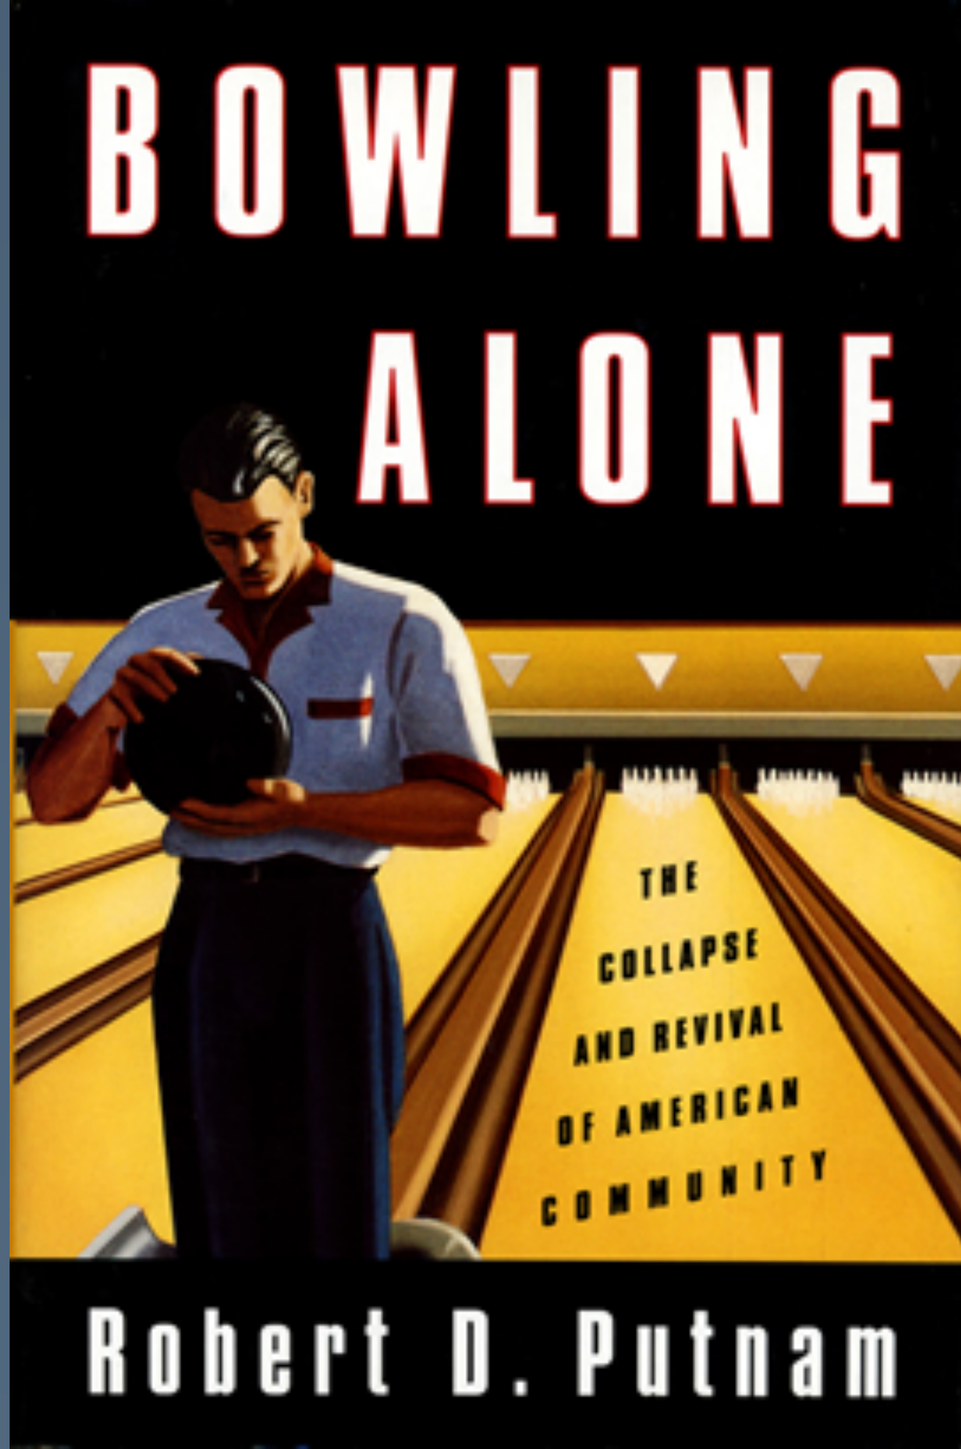

Social capital in social network sites (SNSes)

- Facebook usage increases all types of social capital, especially bridging social capital

[Ellison, Steinfeld and Lampe, JCMC '07]

Regression predicting bridging capital scale →

	Scaled Beta	p
Intercept	3.80	****
Gender: male	−0.02	
Gender: female	0.02	
Ethnicity: white	0.08	*
Ethnicity: nonwhite	−0.08	*
Income	0.04	
Year in school	0.00	
State residence: in-state	−0.05	
State residence: out-of-state	0.05	
Local residence: on campus	−0.04	
Local residence: off campus	0.04	
Fraternity/sorority member	−0.01	
Not member of fraternity/sorority	0.01	
Hours of Internet use per day	−0.03	
Self-esteem	0.20	***
Satisfaction with life at MSU	0.66	****
Facebook (FB) intensity	0.34	****

Conflict and coordination

- **What happens to collaboration costs as Wikipedia grows?** [Kittur, Suh, Pendleton, and Chi, CHI '07]

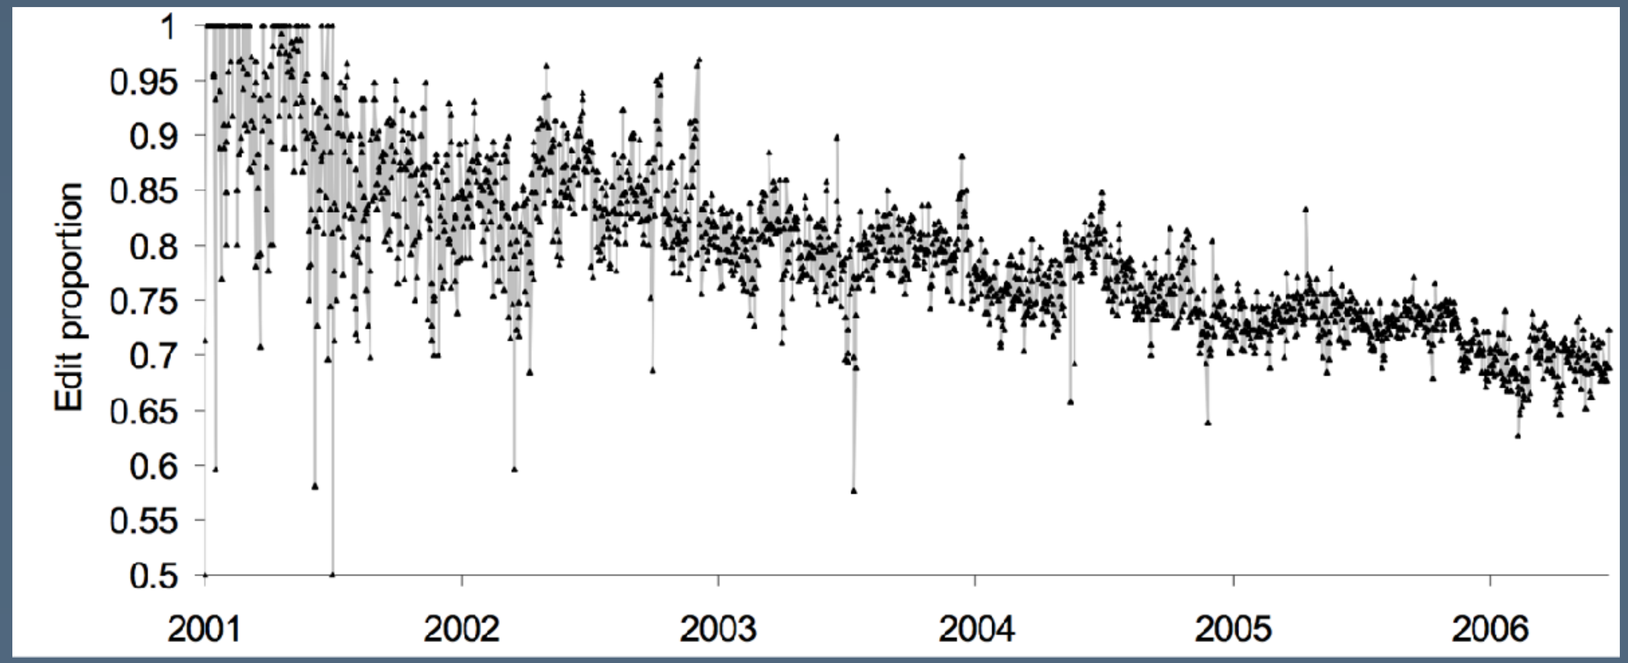

Amount of direct work on articles goes down, and activity on coordination pages goes up

Conflict and coordination

- **As more editors join, which kinds of coordination techniques succeed?** [Kittur and Kraut, CSCW '08]
  - **Explicit: participation in talk pages**
  - **Implicit: set direction by making edits**

More editors only improves article quality only with implicit coordination — a few take on a disproportionate amount of work.

Predicting Tie Strength

- The Strength of Weak Ties [Granovetter, Am. Jour. of Soc. '73]

  Cited by 23692

  - Strong ties: a small number of people you know very well
  - Weak ties: your large number of acquaintances
  - Theory: your weak ties are bridges to other parts of the network; they can help you find jobs and information
- How well can you predict tie strength observationally using social media?

Predicting tie strength

[Gilbert and Karahalios, CHI '09]

- Can we observationally model tie strength?
- Most predictive:
  - Days since last communication
  - Days since first communication
  - Wall words exchanged
  - Mean strength of mutual friends

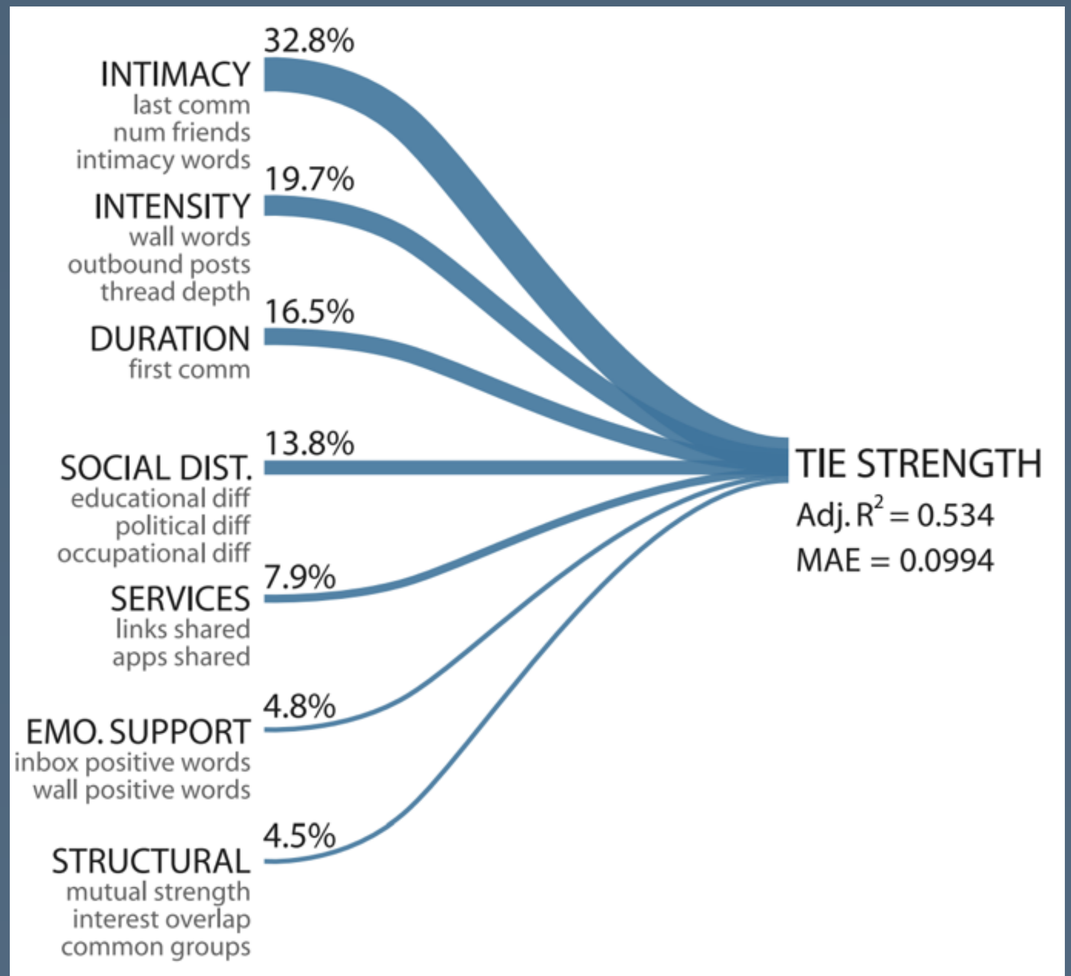

Emotional contagion

[Kramer et al., PNAS '14]

- If you see positive or negative status updates via social media, does it put you in a more positive or negative mood?

- Method: selectively hide positive or negative status updates, and measure how many positive and negative status updates were posted

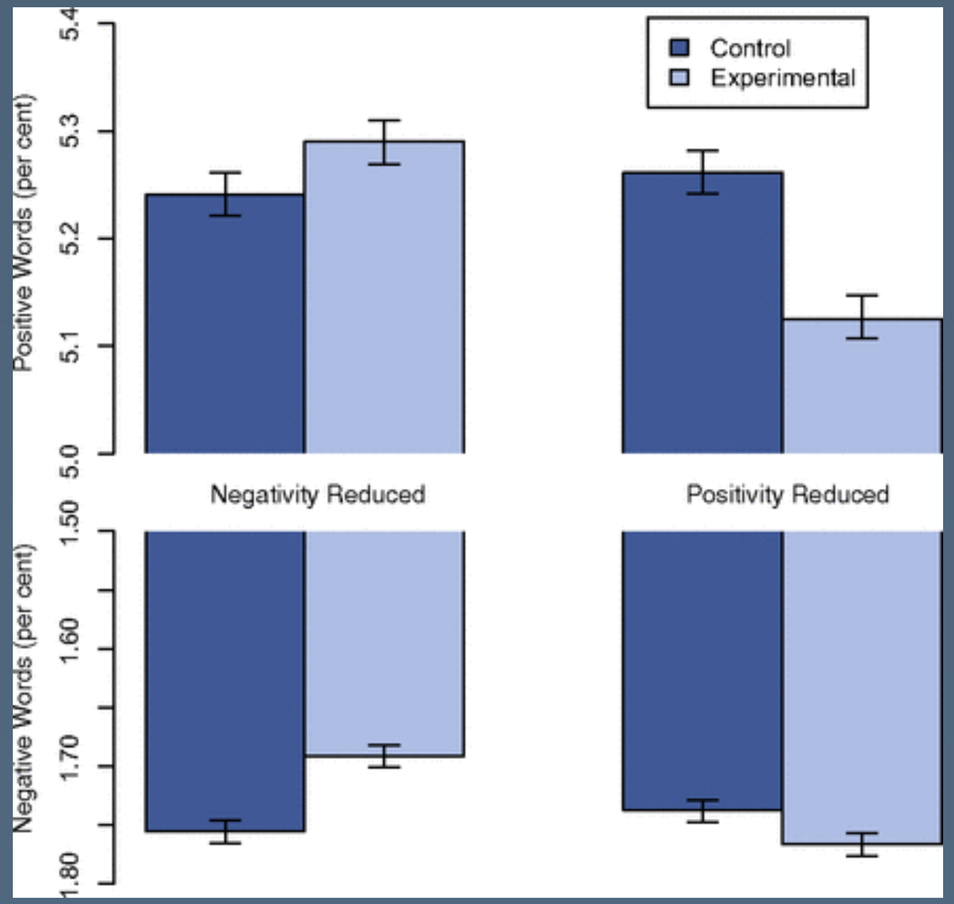

Crowdsourcing

Open call participation

- **Data collection, machine learning training, user studies, social science experiments**
  [Ipeirotis 2010, Heer et al. 2010, Kittur et al. 2008]

- **Games with a purpose**
  [von Ahn and Dabbish 2004, Cooper et al. 2011]

- **Collective action**
  [Wikipedia, Polymath Project, Search for Jim Gray]

Games with a Purpose

Label every image on the internet using a game

[von Ahn and Dabbish, CHI '06]

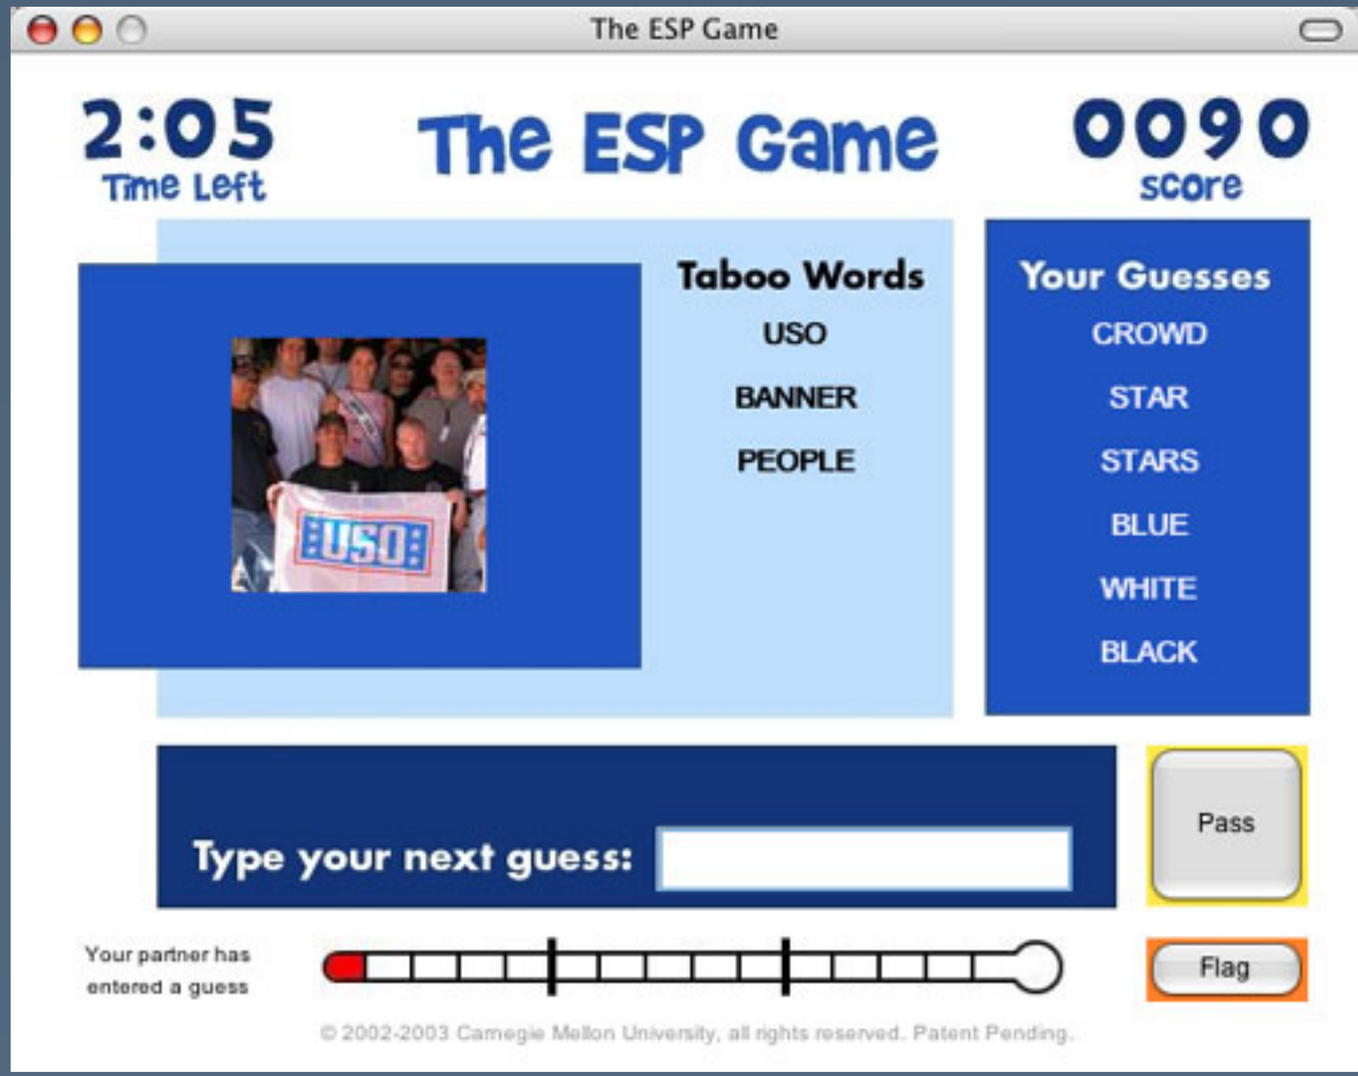

Scientific Collaboration

- FoldIt: protein-folding game
  - Amateur scientists have found protein configurations that eluded scientists for years

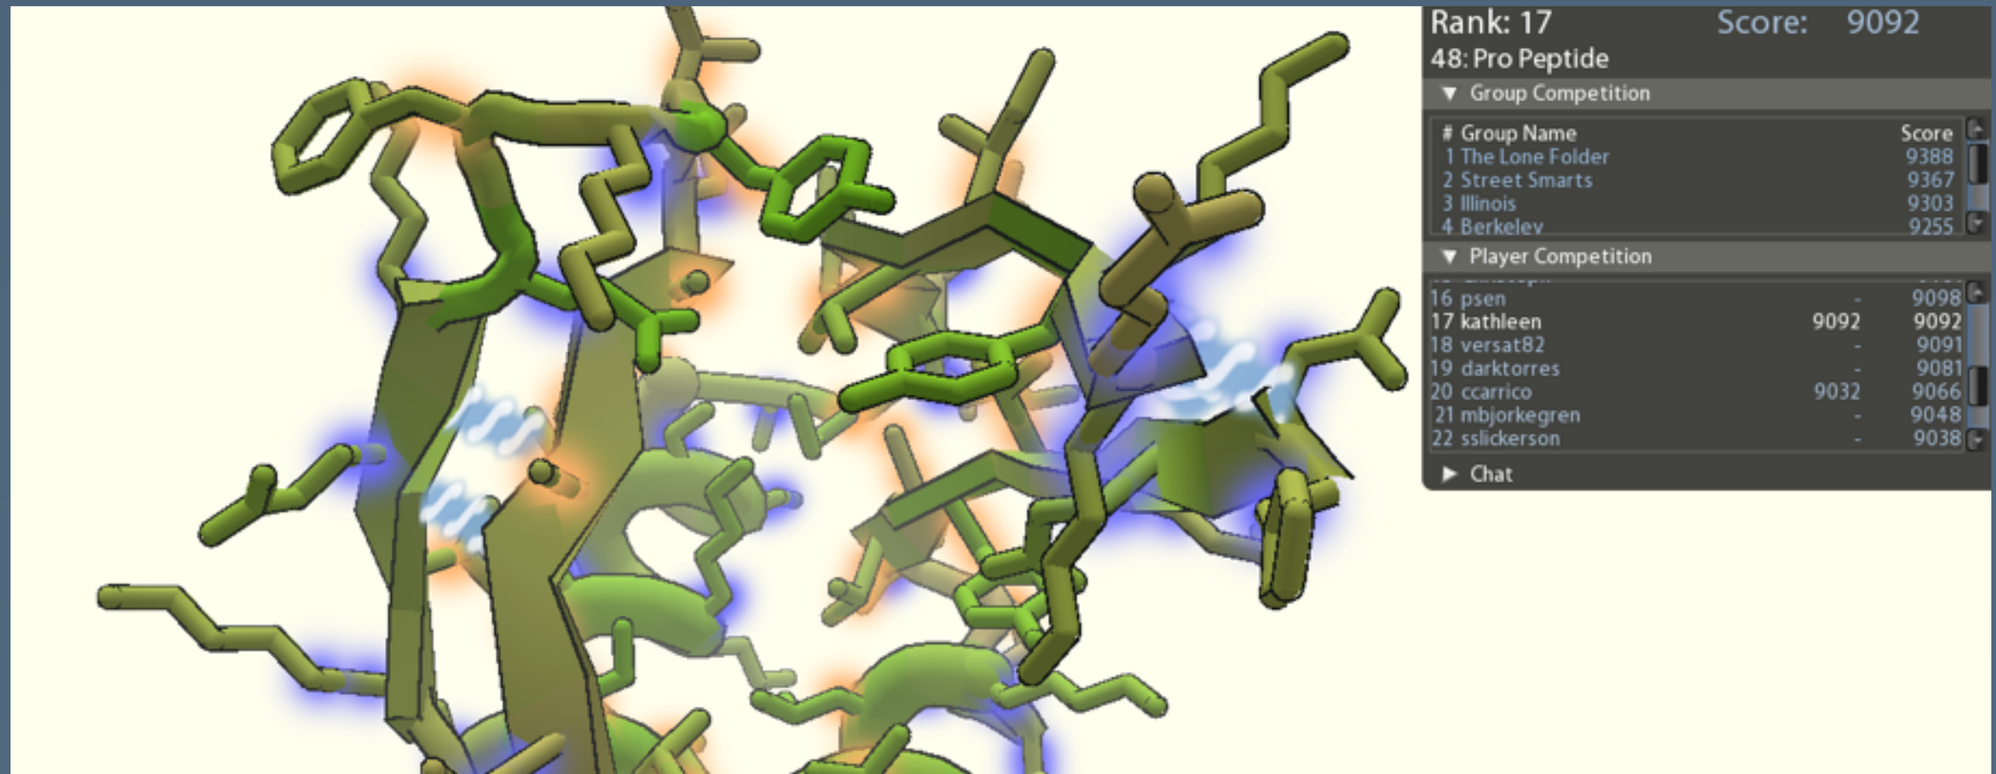

LabInTheWild

- "Buzzfeed-ifying" online studies through narcissism

What is your website aesthetic?

Compare your visual preferences to people around the world. This experiment takes

How good is your implicit memory?

Have you ever had a gut feeling about something? Your implicit memory might have

What is your thinking style?

Find out how your thinking style compares to others. This experiment takes around 5

Quantifying Visual Preferences Around the World

Katharina Reinecke
University of Michigan
Ann Arbor, MI 48109

Krzysztof Z. Gajos
Harvard University
33 Oxford St., Cambridge, MA

Paid Crowdsourcing

- Pay small amounts of money for short tasks
- Amazon Mechanical Turk: Roughly five million tasks completed per year at 1-5¢ each [Ipeirotis 2010]

Label an image

Reward: $0.02

Transcribe audio clip

Reward: $0.05

- Rough population (needs to be updated):
  40% U.S., 40% India, 20% elsewhere
- Gender, education and income are close mirrors of overall population distributions

Soylent

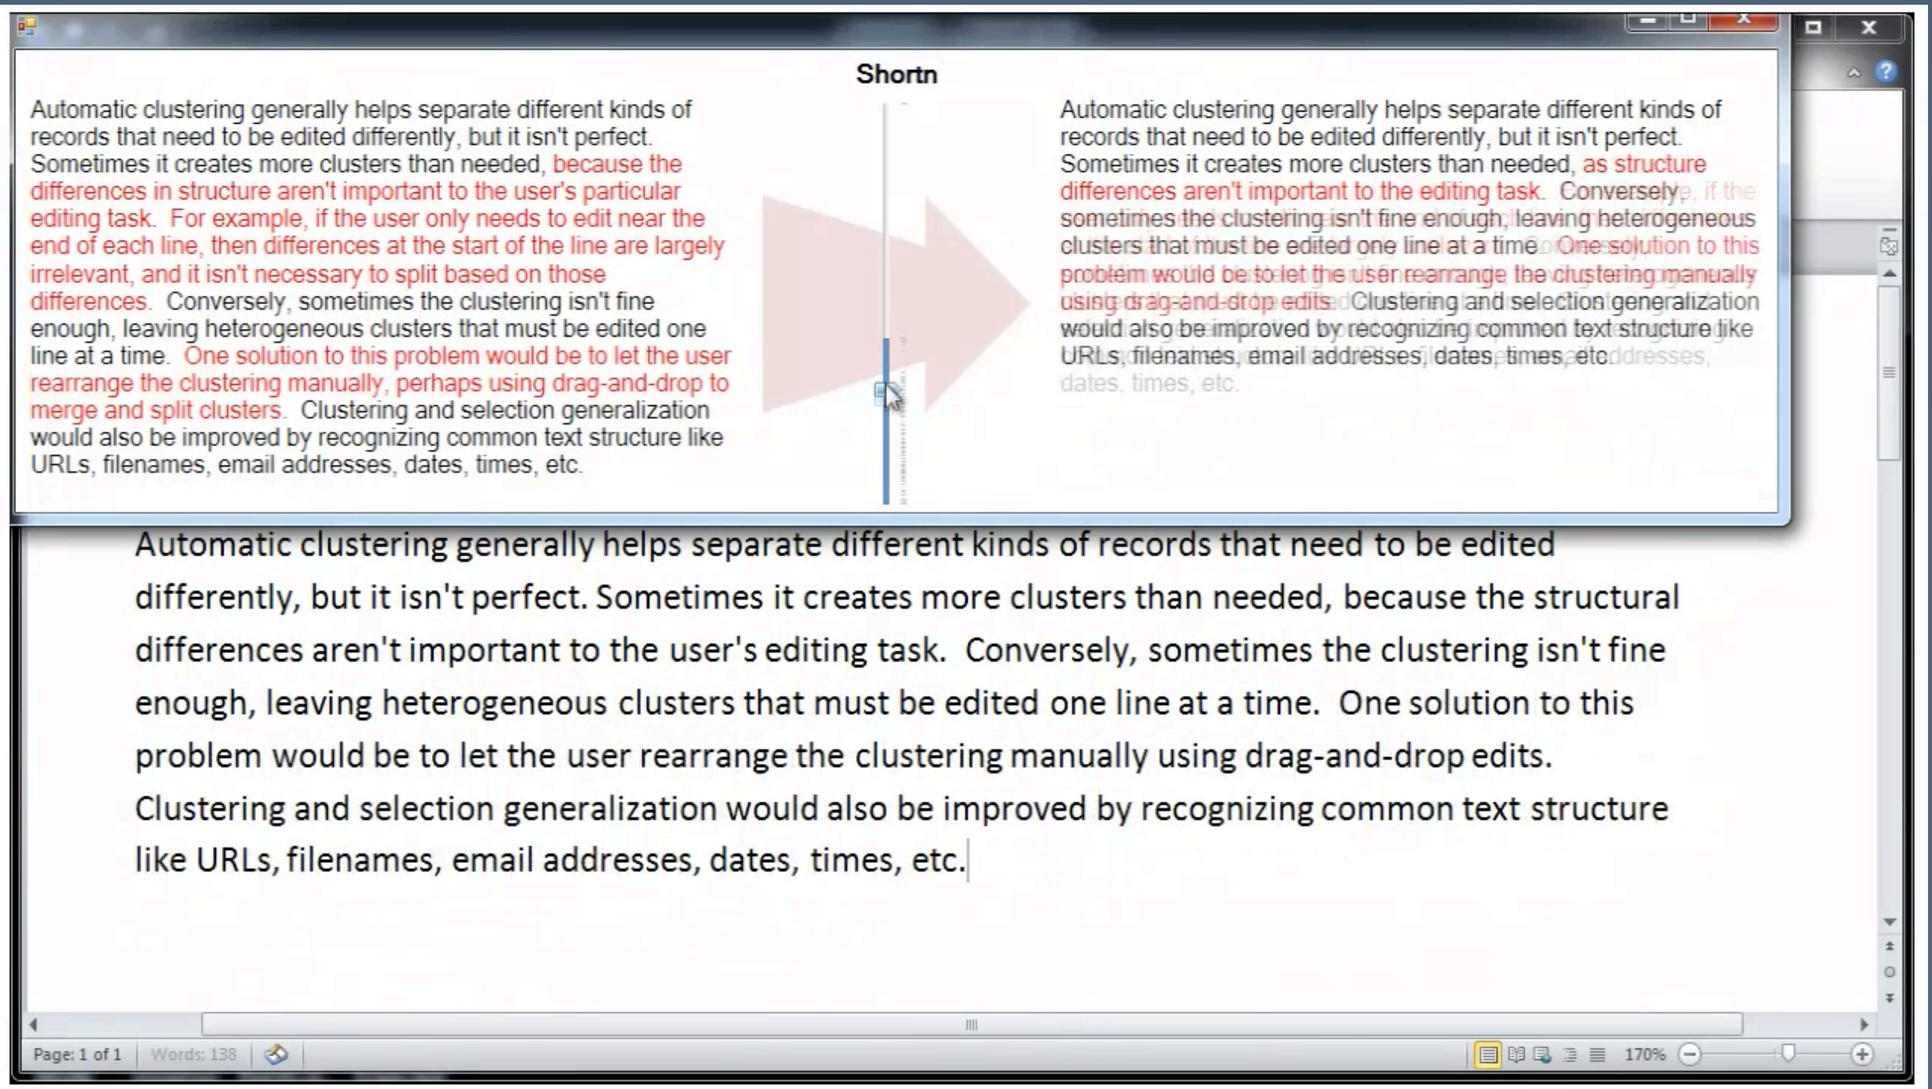
Shortn
Automatic clustering generally helps separate different kinds of records that need to be edited differently, but it isn't perfect. Sometimes it creates more clusters than needed, because the differences in structure aren't important to the user's particular editing task. For example, if the user only needs to edit near the end of each line, then differences at the start of the line are largely irrelevant, and it isn't necessary to split based on those differences. Conversely, sometimes the clustering isn't fine enough, leaving heterogeneous clusters that must be edited one line at a time. One solution to this problem would be to let the user rearrange the clustering manually, perhaps using drag-and-drop to merge and split clusters. Clustering and selection generalization would also be improved by recognizing common text structure like URLs, filenames, email addresses, dates, times, etc.
Automatic clustering generally helps separate different kinds of records that need to be edited differently, but it isn't perfect. Sometimes it creates more clusters than needed, as structure differences aren't important to the editing task. Conversely, sometimes the clustering isn't fine enough, leaving heterogeneous clusters that must be edited one line at a time. One solution to this problem would be to let the user rearrange the clustering manually using drag-and-drop edits. Clustering and selection generalization would also be improved by recognizing common text structure like URLs, filenames, email addresses, dates, times, etc.
Automatic clustering generally helps separate different kinds of records that need to be edited differently, but it isn't perfect. Sometimes it creates more clusters than needed, because the structural differences aren't important to the user's editing task. Conversely, sometimes the clustering isn't fine enough, leaving heterogeneous clusters that must be edited one line at a time. One solution to this problem would be to let the user rearrange the clustering manually using drag-and-drop edits. Clustering and selection generalization would also be improved by recognizing common text structure like URLs, filenames, email addresses, dates, times, etc.
Page: 1 of 1
Words: 138
170%

Paid Crowdsourcing: Goals

- Design and create crowd-powered systems (e.g., Soylent)
- Design algorithms and design patterns for complex tasks
- Understand worker motivation
- Quality control
- Coming up in a future class...

Social computing contributions

- Using sociotechnical systems as a lens to better understand human social behavior
  - e.g., How do we grow friendships? What role do they play as we undergo major life changes?
- Creating sociotechnical systems that demonstrate new kinds of social or collective behavior
  - e.g., How might the internet come together to write the Great American Novel?