Purdue Fall 2018

CS59000-CSC: Crowdsourcing and Social Computing

http://mingyin.org/CS590/Fall2018/index.html

Class 25: Crowd-Powered Systems

2018.11.19

Ming Yin

What have we learned so far in the course?

Participants, Platforms, Process

Incentive Design

Task Assignment

Quality Control

Engagement Control

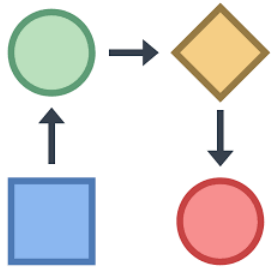

Workflow Design

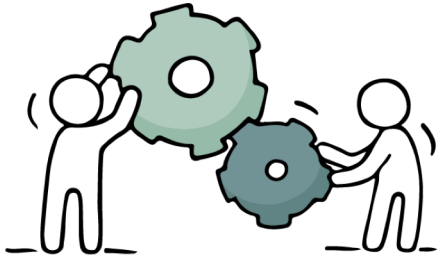

Cooperative Work

Next few classes

Nov 19	Crowd-Powered Systems *Lecture*	**Required** Lasecki et al. Real-Time Captioning by Groups of Non-Experts. UIST'12 Kokkalis et al. MyriadHub: Efficiently Scaling Personalized Email Conversations with Valet Crowdsourcing. CHI'17 **Optional** Vashistha et al. Respeak: A Voice-based, Crowd-powered Speech Transcription System. CHI'17 Nguyen et al. An Interpretable Joint Graphical Model for Fact-Checking from Crowds. AAAI'18
Nov 21	*No class (Thanksgiving)*	
Nov 26	Crowdsourcing: Future Ideas *Lecture*	**Required** Whiting et al. Crowd Guilds: Worker-led Reputation and Feedback on Crowdsourcing Platforms. CSCW'17 **Optional** Morris et al. Subcontracting Microwork. CHI'17 Vaish et al. Crowd Research: Open and Scalable University Laboratories. UIST'17
Nov 28	**Final project presentation (Session 1)**	
Dec 3	**Final project presentation (Session 2)**	

Final Project Timeline

- Final project presentation (Nov 28 & Dec 3)
  - In one week and a half!!
  - Nov 28: 4 presentations; Dec 3: 3 presentations
  - Presentation order

- Final project report (due 11:59pm, Dec 8)
  - This is a HARD deadline. No extension!
  - 10-page maximum, ACM SIG proceedings style
  - Check final project guideline for more details

The Power of Crowds

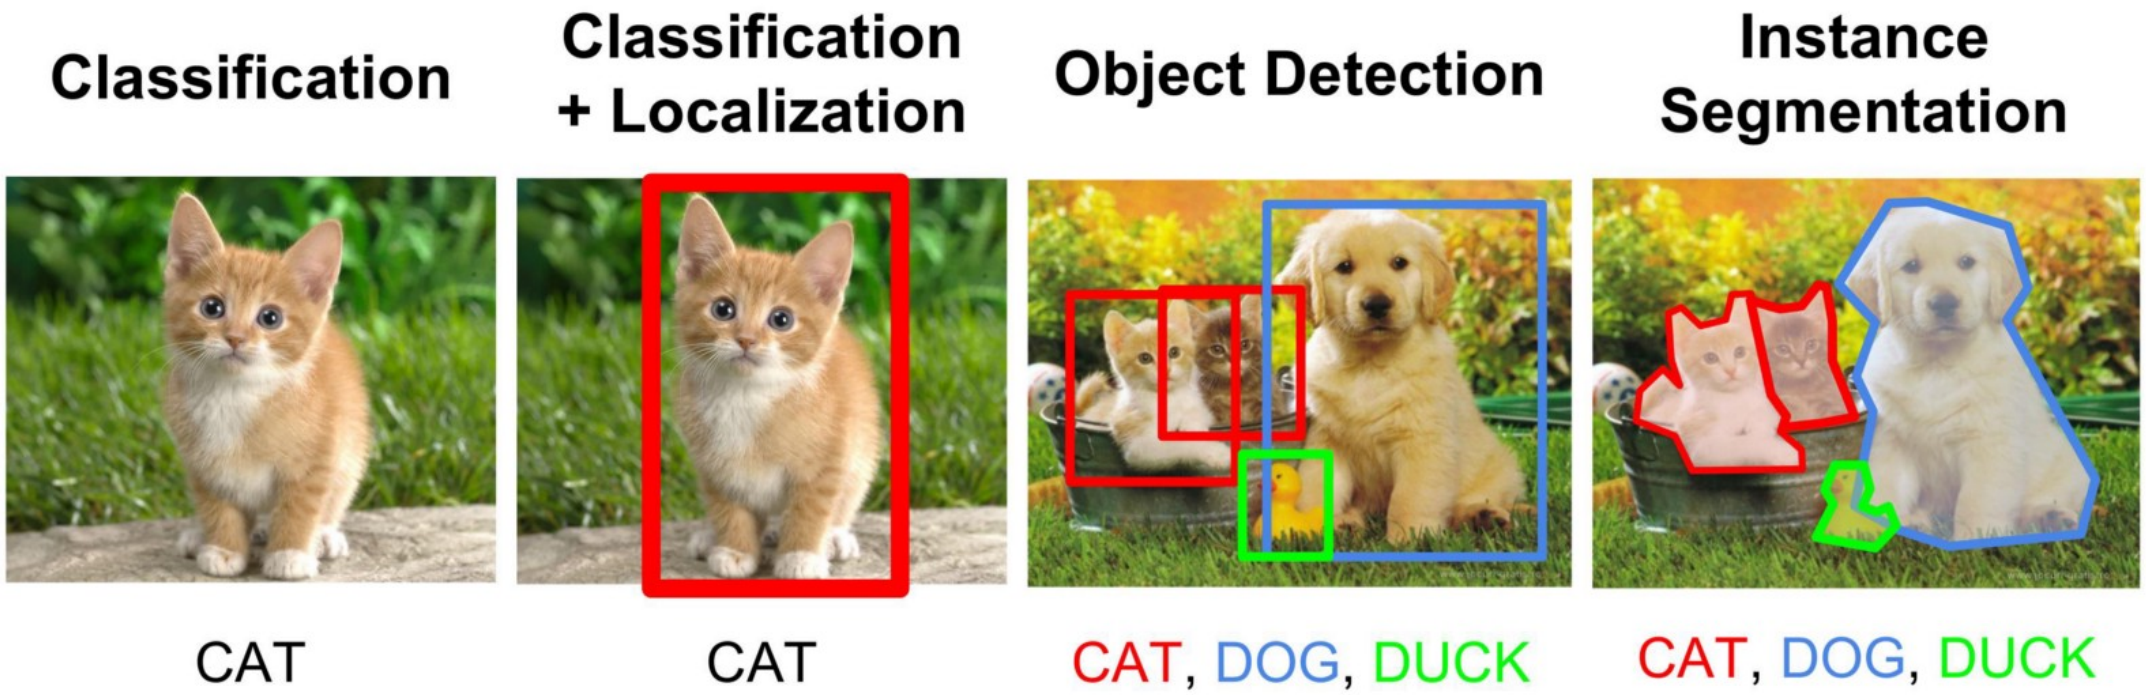

Simple tasks using *basic human intelligence*

The Power of Crowds

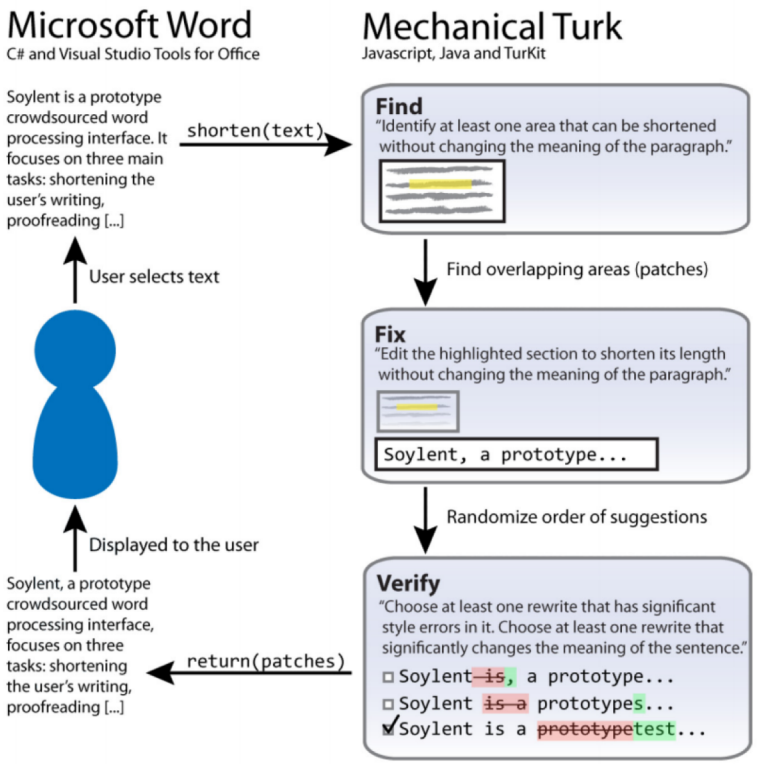

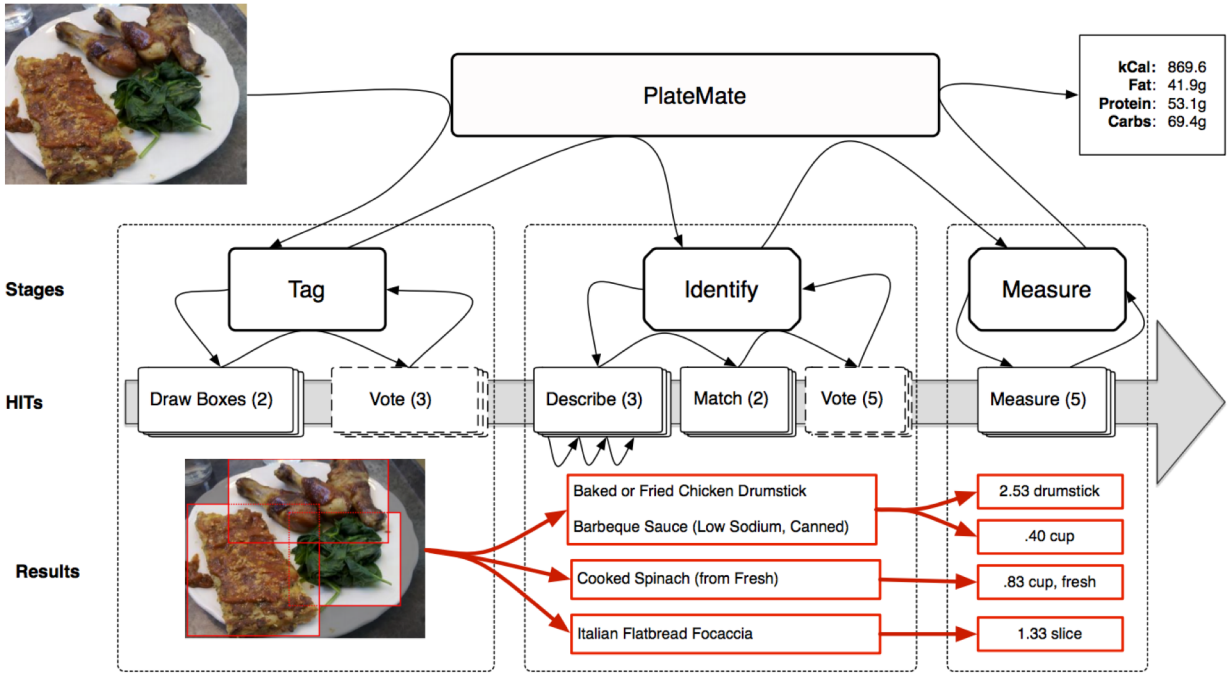

Complex (but dividable) tasks leveraging *the division of labor*

The Power of Crowds

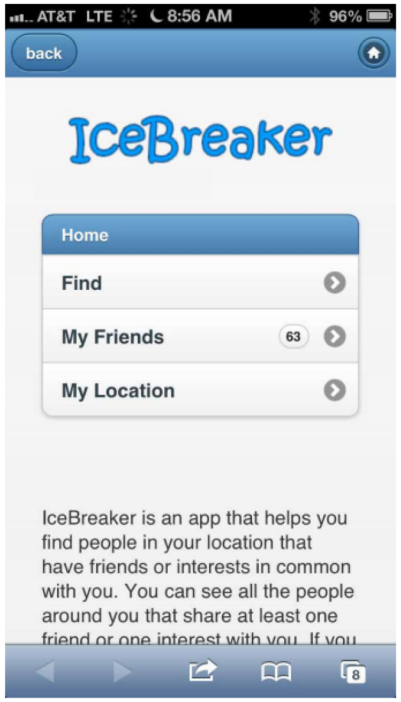

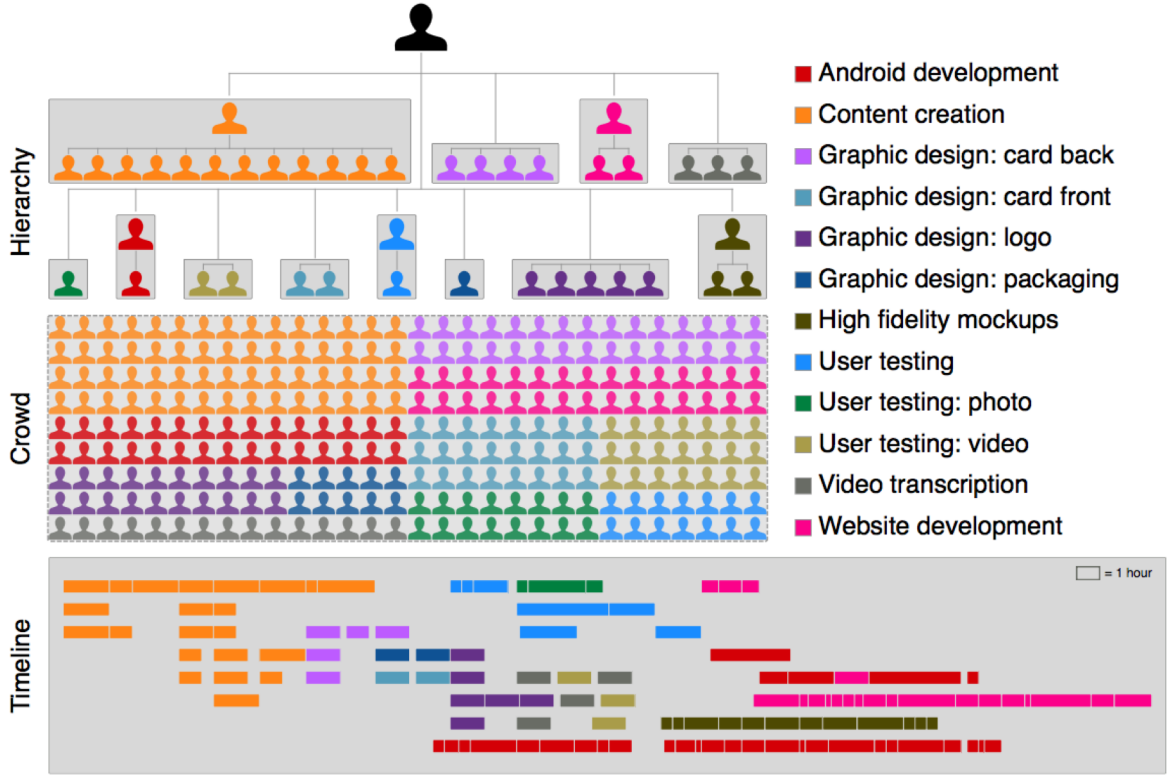

Complex (but undividable) tasks harnessing *teamwork*

Some More Advantages of the Crowd

- On-demand: Get help from the crowd whenever you want
- Scalability: Large number of crowd workers
- **Today:** Explore cool end-to-end systems with crowds in the loop that leverage the crowd's on-demand nature and scalability

Scribe

- *Real-time captioning by Groups of Non-Experts*, Lasecki et. al, UIST 2012
- On-demand: “Real-time”
- Scalability: “Groups of Non-Experts”
- Cool application: useful for deaf and hard of hearing people in classrooms, meetings, etc.

Real-Time Captioning: The State of the Art

- Communications Access Real-Time Translation (CART)
  - Professional stenographers (2-3 years of training)
  - Most expensive
  - Most reliable: average 141 words per minute
- Non-Verbatim Systems (e.g., C-Print)
  - $60/hour
  - Slower, and also not verbatim
- Automated Speech Recognition (ASR)
  - Low performance in real-world settings

System Overview

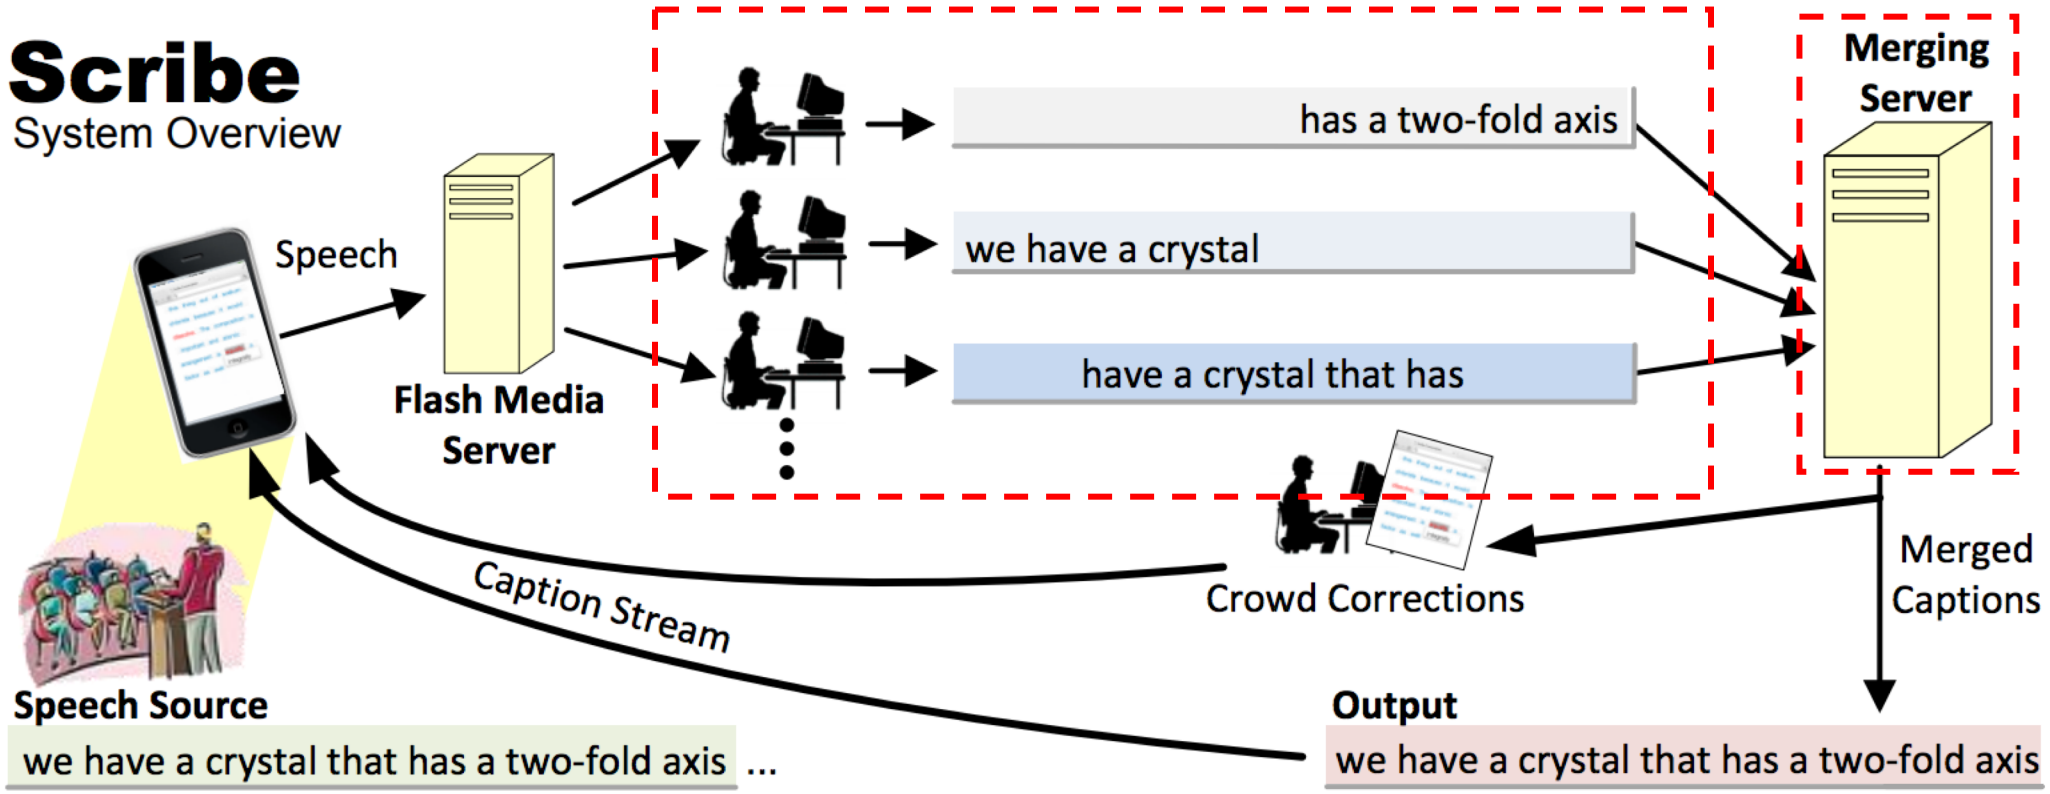

Worker Interface

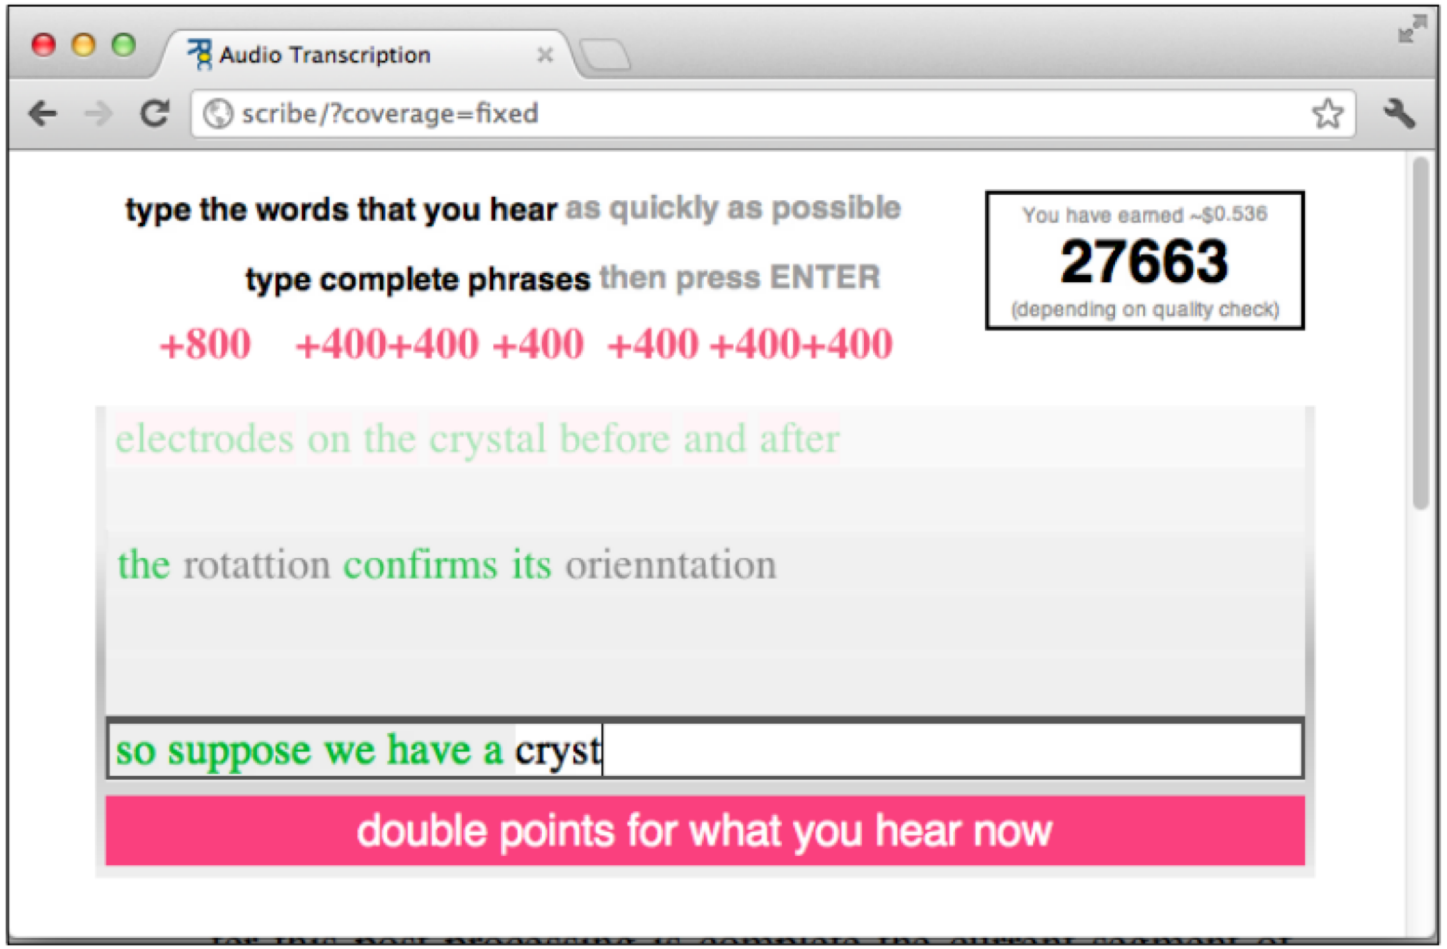

- Separate contiguous sequences of words by the enter key
- "Lock in" a word after 800 milliseconds it is typed
- Performance based incentives!

Key Idea: Adjusting Saliency

- "Guide" workers to work on different parts of the audio by systematically injecting saliency artificially
  - Through varying the volume of the audio signal

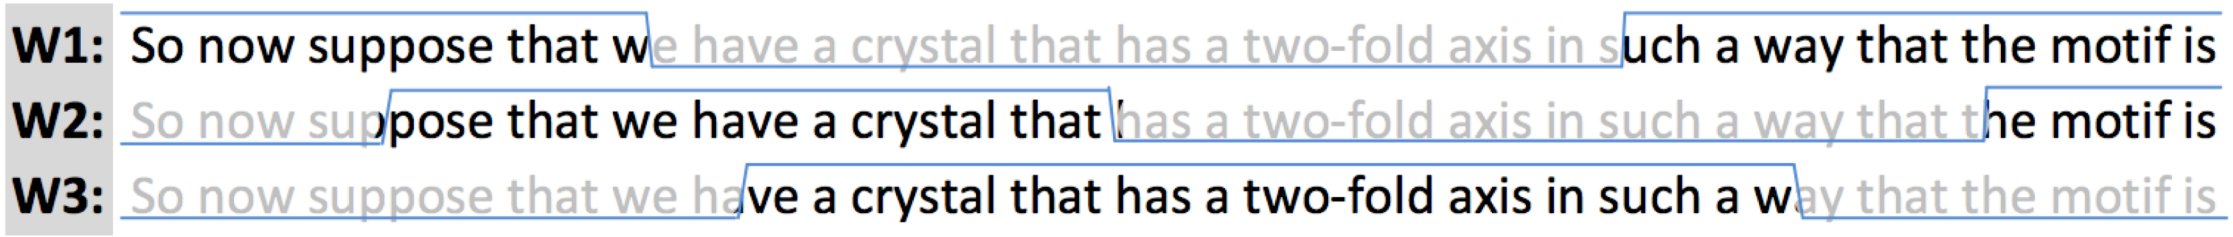

- Better than dividing audio into different segments
  - Contexts help fast transcription
  - Have a sense of "real-time" work

Combining the Inputs

- Multiple Sequence Alignment (MSA):
  - Bioinformatics techniques to align multiple (gene) sequences so as to achieve maximal matching between them
  - Adapting scoring function: incorporating spelling error model

Sequence1	**–TCAGGA–TGAAC–––––**
Sequence2	**ATCACGA–TGAACC–––**
Sequence3	**ATCAGGAATGAATCC––**
Sequence4	**–TCACGATTGAATCGC–**
Sequence5	**–TCAGGAATGAATCGCM**

Online Sequence Alignment

- Maintaining the longest self-consistent path to avoid unnecessary branching.

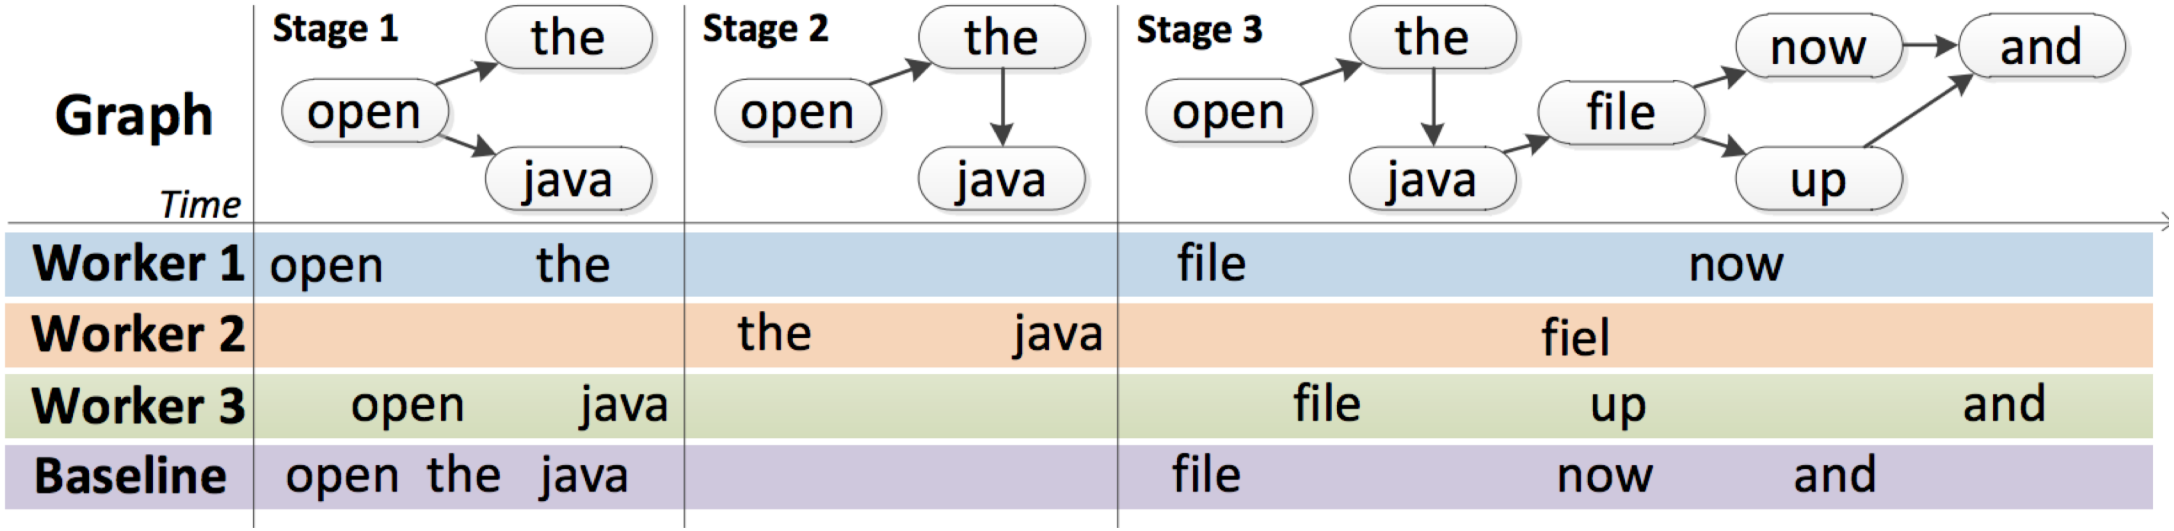

System Evaluation

- **Coverage:** the number of words in the ground truth that has been correctly transcribed (by the group) *within 10 seconds* after the words appear
- **Precision:** the fraction of words in the transcription that are in the ground truth *within 10 seconds* after they appear
- A trade-off

System Evaluation

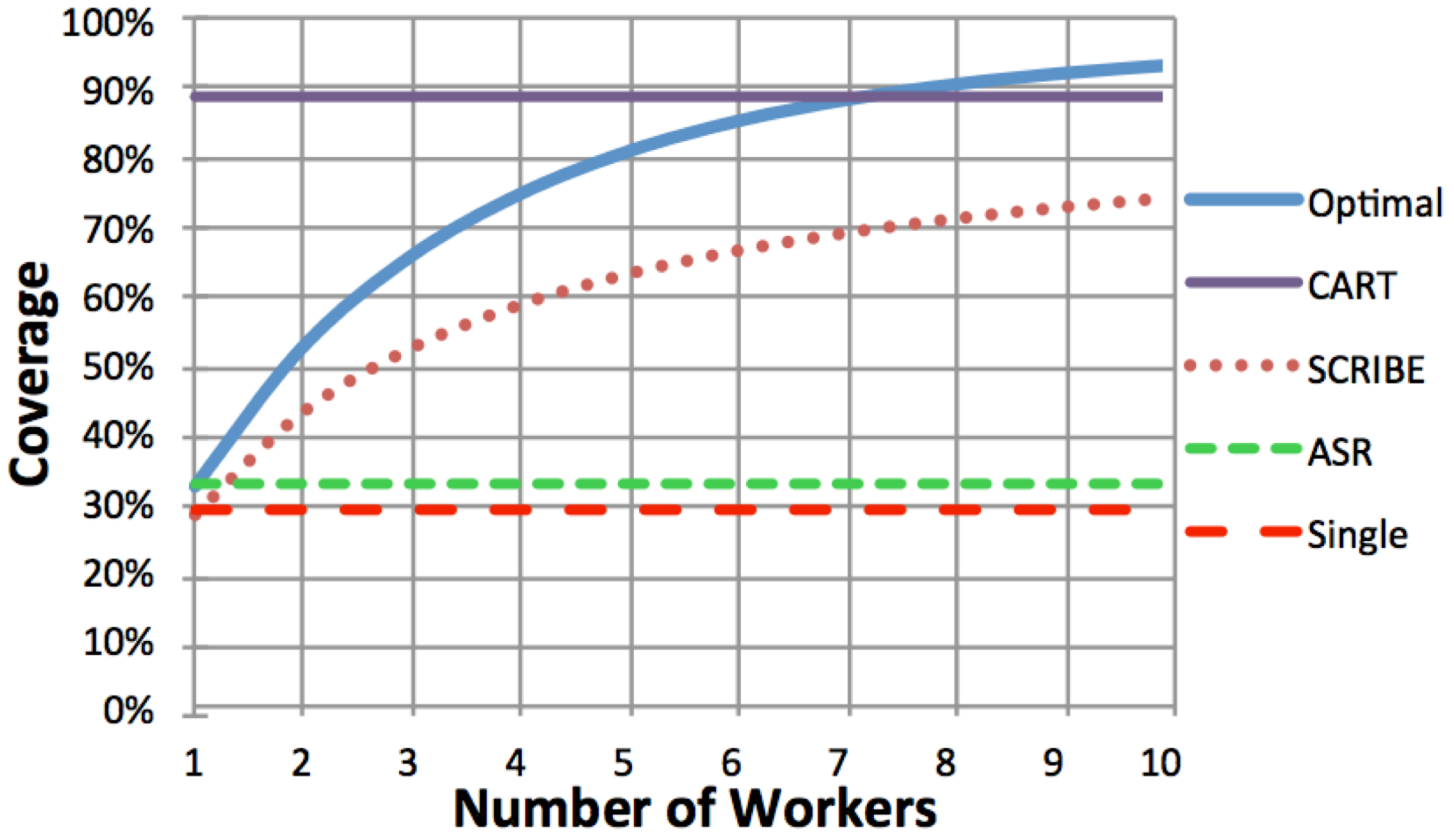

- Scribe outperforms ASR and single worker on coverage
- Average latency of 2.89 seconds (better than CART)!
- Workers also transcribe more words in the salient periods (periods with high volume)

System Evaluation

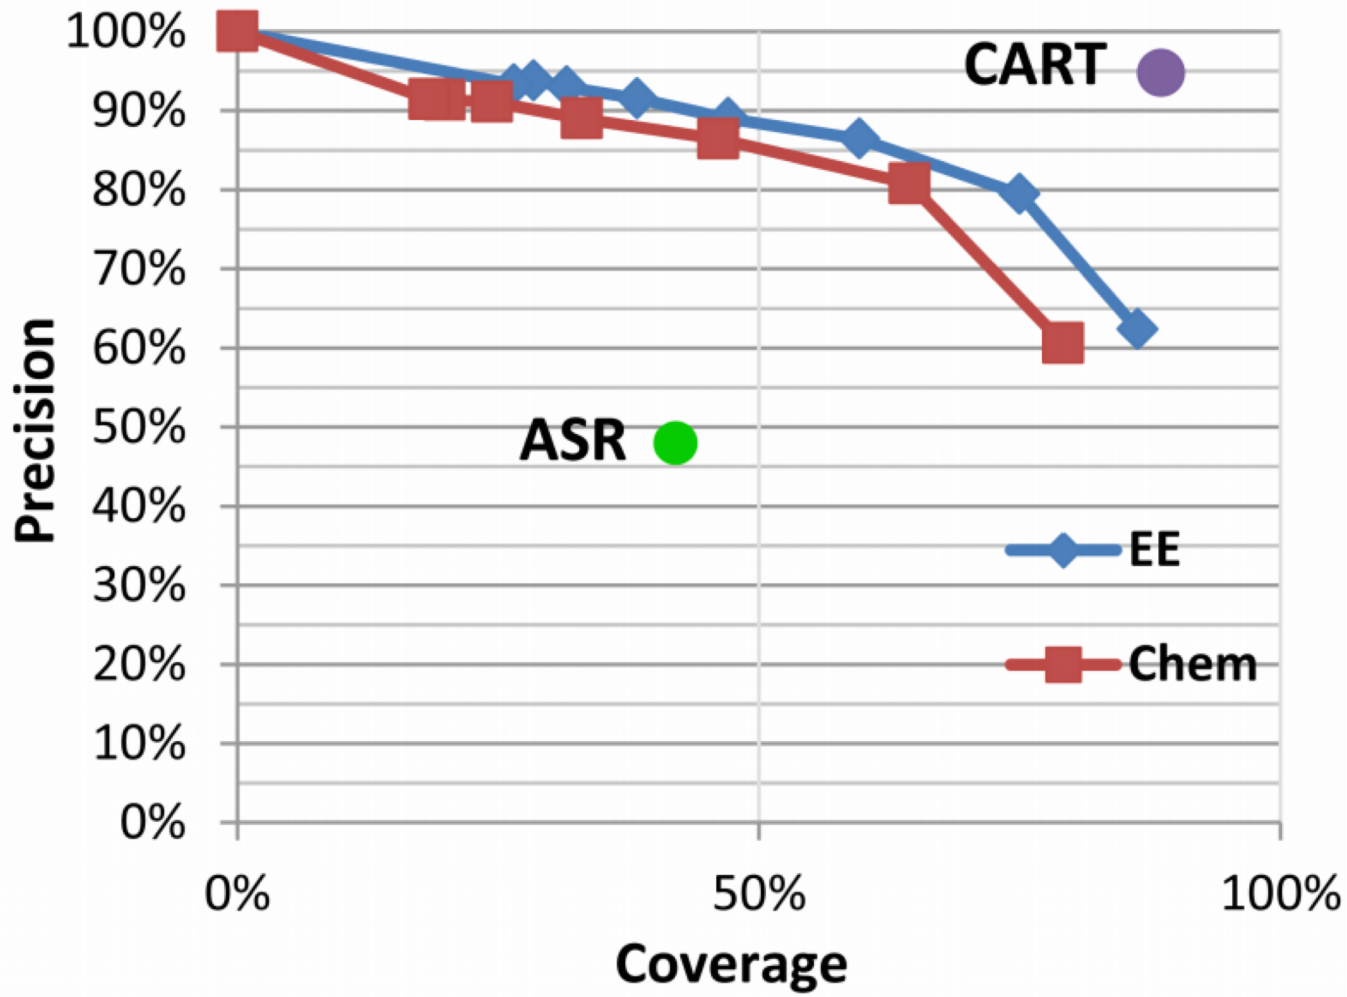

Discussions

- How do you like this system?
- Any opportunities for improvement?

More Crowd-Powered Systems

- *Respeak: A Voice-based, Crowd-powered Speech Transcription System*, Vashistha et. al, CHI 2017
- Using automatic speech transcription tool to transcribe
- But ask crowd workers to re-speak what they heard in a quite environment
- Value for workers: improved vocabulary, pronunciation and oral skills

Utilizing the Crowd for Email Personalization

- *MyriadHub: Efficiently Scaling Personalized Email Conversations with Valet Crowdsourcing*, Kokkalis et. al, CHI 2017

- Key idea: Leverage crowds as valets to analyze email patterns

Use crowds to extract metadata of emails

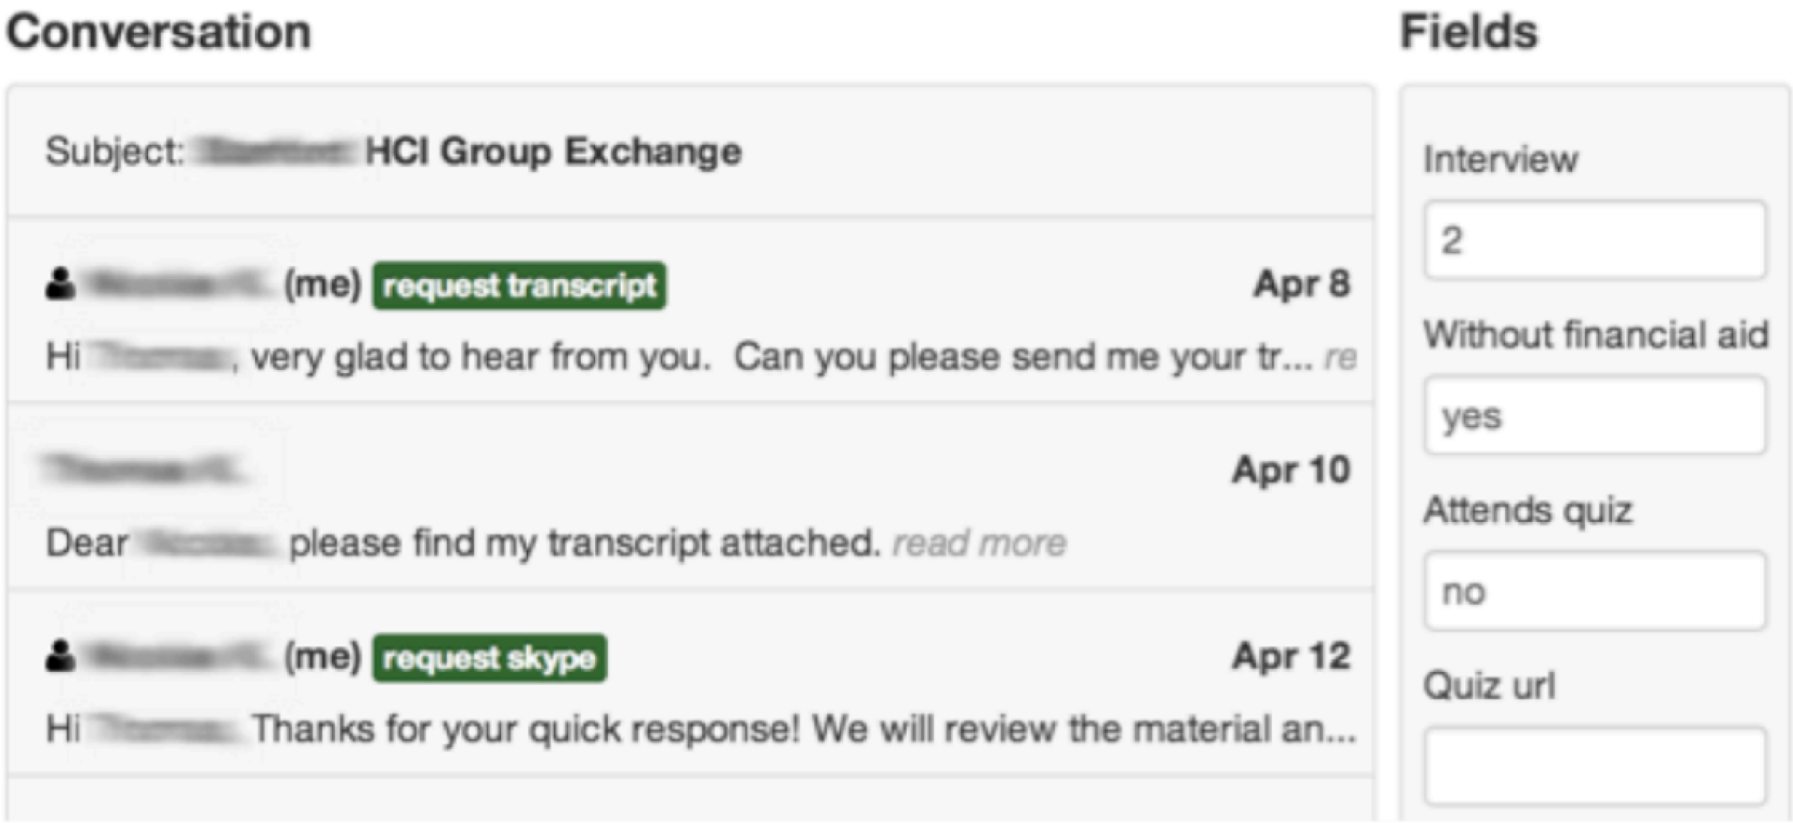

Extracting fields and values from emails

Use crowds to generate templates and rules

Filter Applied	Send Template	New Matches	Automatically?
'university_coded' = *not blank*	request transcript	*none*	ON
Last sent message = request transcript Status = You need to reply. 'university_coded' = *not blank* 'skype_handle' = *blank*	request skype	*none*	OFF
Last sent message = request transcript Status = They need to reply. 'transcripts' = *blank*	transcript - reminder	5 Send	OFF
'send skills survey' = 'TRUE'	skills survey	3 Send	OFF

Other Features

- Visualizations of conversation state (for each receipt)

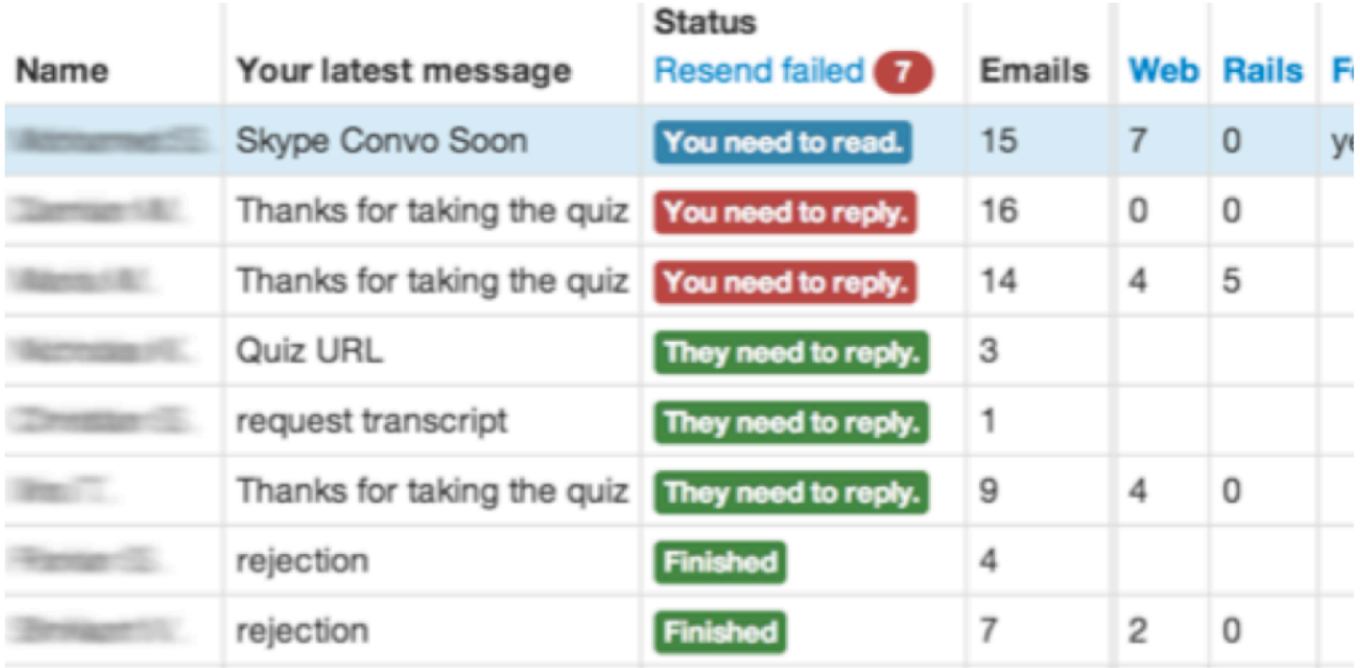

Name	Your latest message	Status Resend failed 7	Emails	Web	Rails	F
[redacted]	Skype Convo Soon	You need to read.	15	7	0	y
[redacted]	Thanks for taking the quiz	You need to reply.	16	0	0	
[redacted]	Thanks for taking the quiz	You need to reply.	14	4	5	
[redacted]	Quiz URL	They need to reply.	3			
[redacted]	request transcript	They need to reply.	1			
[redacted]	Thanks for taking the quiz	They need to reply.	9	4	0	
[redacted]	rejection	Finished	4			
[redacted]	rejection	Finished	7	2	0	

- Integration with existing tools

System Evaluation

- A within-subject field experiment with 12 participants
- Organize a potluck party with 10 invited guests
- Simulated email responses
- Simulated valet crowd workers (only for meta-data extractions)
- Results: 32% saving on time; no mistakes!

System Evaluation

- A between-subject experiment with 172 participants
- Control: redirect participants to fill out a survey
- Experimental: Ask questions directly in the email

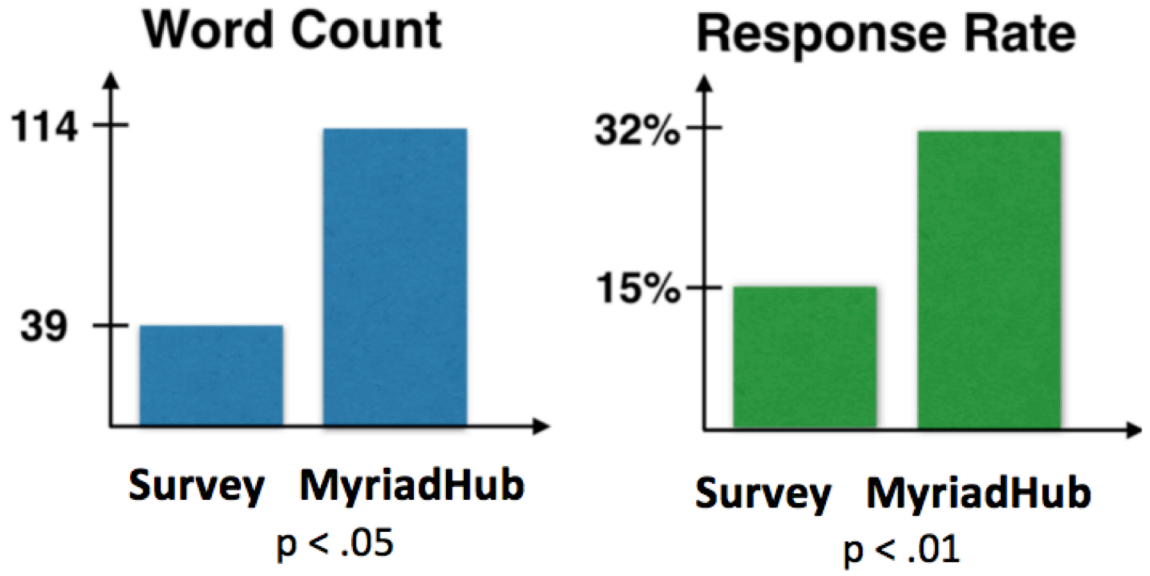

Thoughts?

Discussions

- How do you like this system?
- Any opportunities for improvement?
- If you have a crowd of valet workers, what will you use them for?

More Crowd-Powered Systems

- *VizWiz: Nearly Real-time Answers to Visual Questions*, Bigham et. al, UIST 2010

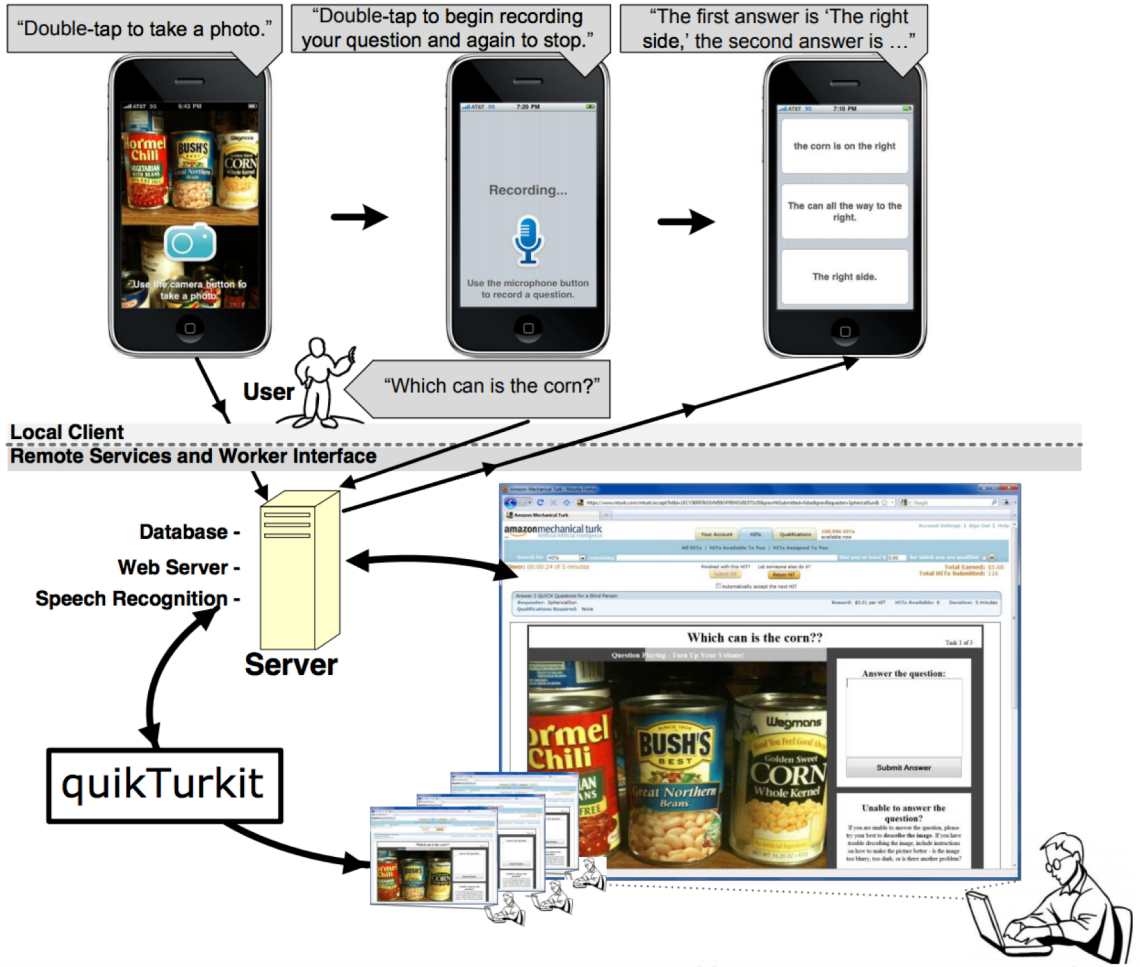

Next Class

- Nov 21: Happy thanksgiving! (No class)
- Nov 26: The last lecture!
  - Crowdsourcing: Future Ideas
  - Required:
    - Whiting et al. Crowd Guilds: Worker-led Reputation and Feedback on Crowdsourcing Platforms. CSCW'17
  - Optional:
    - Morris et al. Subcontracting Microwork. CHI'17
    - Vaish et al. Crowd Research: Open and Scalable University Laboratories. UIST'17