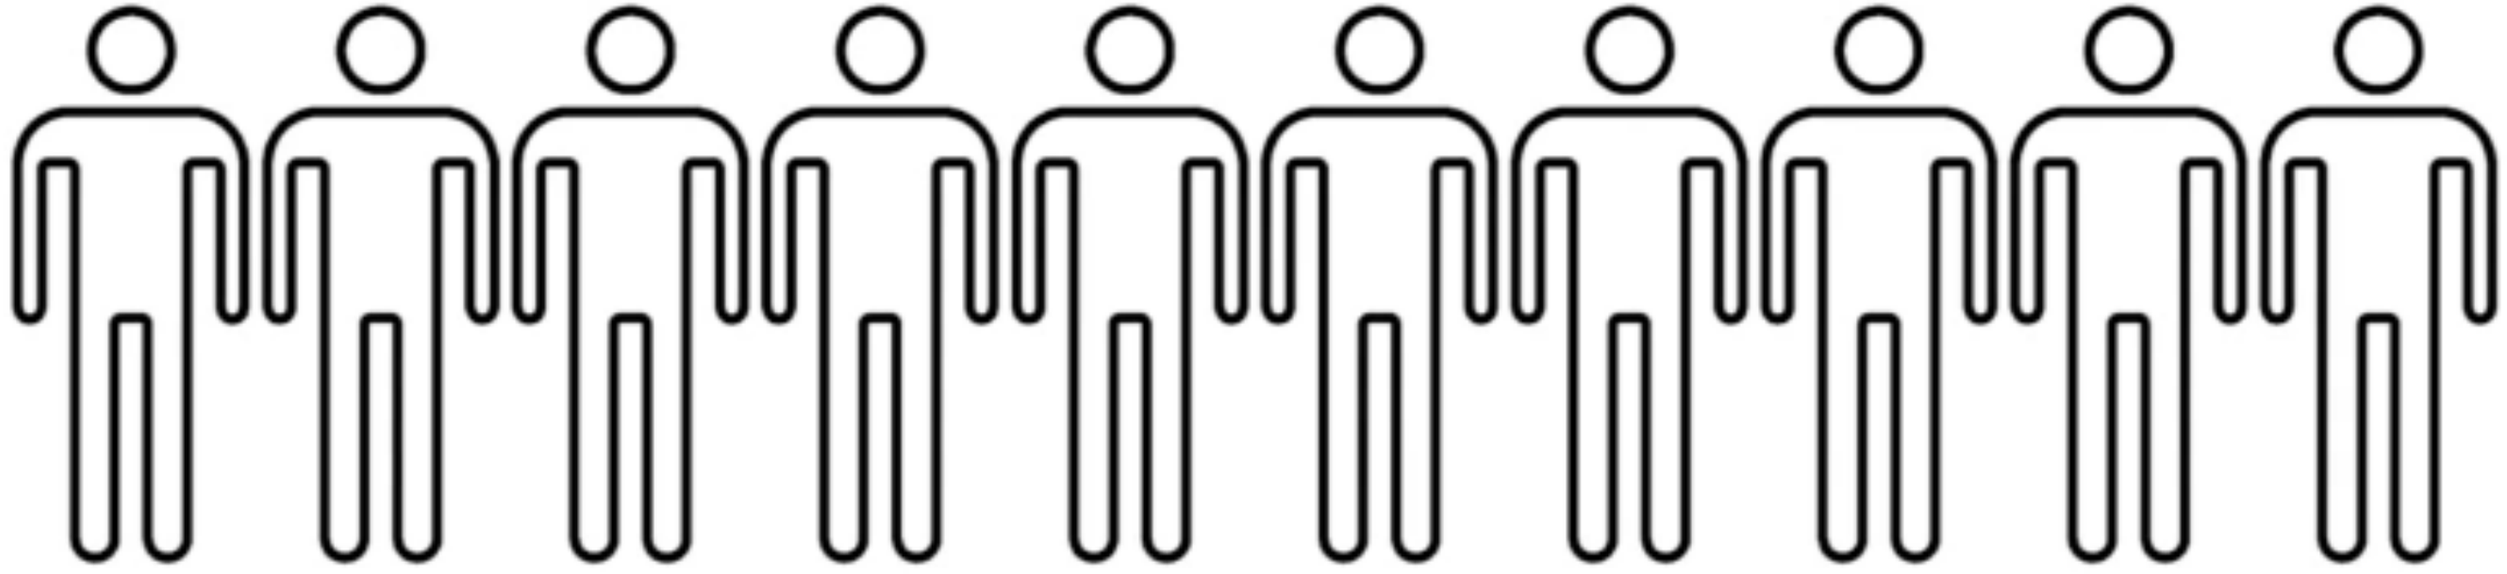

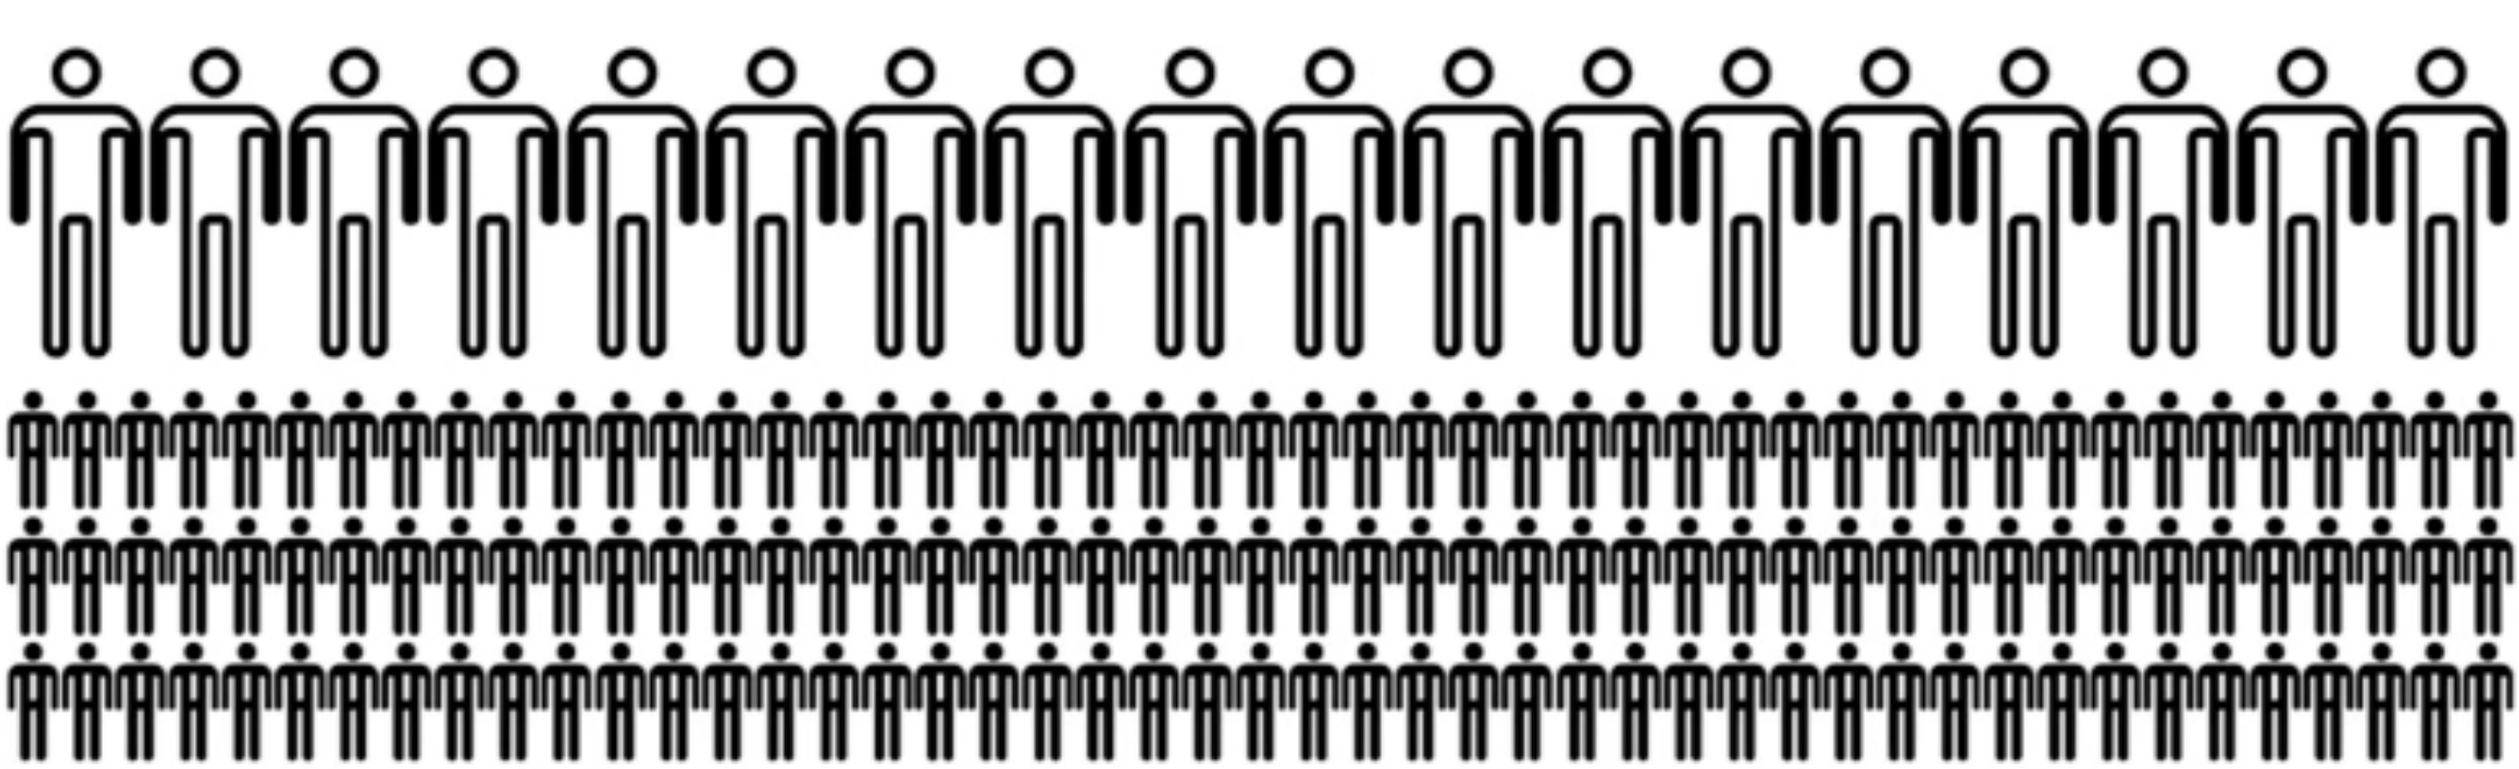

Mechanical Turk

6.810 Engineering Interaction Technologies

Prof. Stefanie Mueller | HCI Engineering Group

amazonmechanical turk
Artificial Artificial Intelligence

Your Account | HITs | Qualifications

Already have an account
Sign in as a Worker | Requeste

Introduction | **Dashboard** | **Status** | **Account Settings**

Mechanical Turk is a marketplace for work.

We give businesses and developers access to an on-demand, scalable workforce.
Workers select from thousands of tasks and work whenever it's convenient.

76,187 HITs available. View them now.

Make Money
by working on HITs

HITs - *Human Intelligence Tasks* - are individual tasks that you work on. Find HITs now.

As a Mechanical Turk Worker you:

- Can work from home
- Choose your own work hours
- Get paid for doing good work

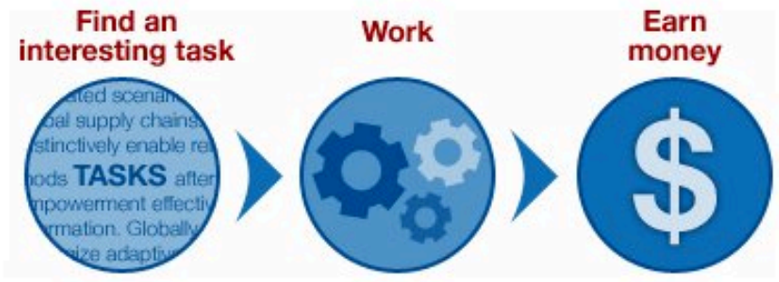

Find HITs Now

or learn more about being a **Worker**

Get Results
from Mechanical Turk Workers

Ask workers to complete HITs - *Human Intelligence Tasks* - and get results using Mechanical Turk. Register Now

As a Mechanical Turk Requester you:

- Have access to a global, on-demand, 24 x 7 workforce
- Get thousands of HITs completed in minutes
- Pay only when you're satisfied with the results

Get Started

FAQ | Contact Us | Careers at Amazon | Developers | Press | Policies

©2005-2010 Amazon.com, Inc. or its Affiliates

An amazon.com. compar

amazon mechanical turk
pay people for completing microtasks (2005)

mechanical turk::

- large pool of workers (turkers)
- sitting at home eager to work and make money
- you post a request e.g., “find spelling mistakes”
- specify payment
- wait until someone picks it up and submits an answer

All HITs | HITs Available To You | HITs

Find HITs containing

All HITs

1-10 of 2249 Results

Sort by: HITs Available (most first) GO!

Show all details | Hide all details

Determine the industry of this business

Requester: Matt Park

HIT Expiration Date: May 13, 2013 (2 days 12 hou

Time Allotted: 3 minutes

Classify Arabic Tweets Dialects SEE REVISED HIT

Requester: Chris Callison-Burch

HIT Expiration Date: Aug 4, 2013 (12 weeks 1 day

Time Allotted: 60 minutes

Classify Arabic Tweets Dialects (No Qualification)

Requester: Chris Callison-Burch

HIT Expiration Date: Aug 22, 2013 (14 weeks 4 da

Time Allotted: 60 minutes

Find the lowest price for the items searched on the website shown

Requester: IntSup

HIT Expiration Date: May 14, 2013 (3 days 3 hours

Time Allotted: 60 minutes

HIT = Human Intelligence Task

often for **micro-pay**

for what tasks would you use humans on mechanical turk vs. an algorithm?

<30 sec brainstorming>

for what tasks would you use humans vs. an algorithm?

- problems computers cannot yet solve
  - image classification & labeling
  - transcription from audio
  - translation
  - content generation for websites
- rating of things human's will perceive (rate logos)

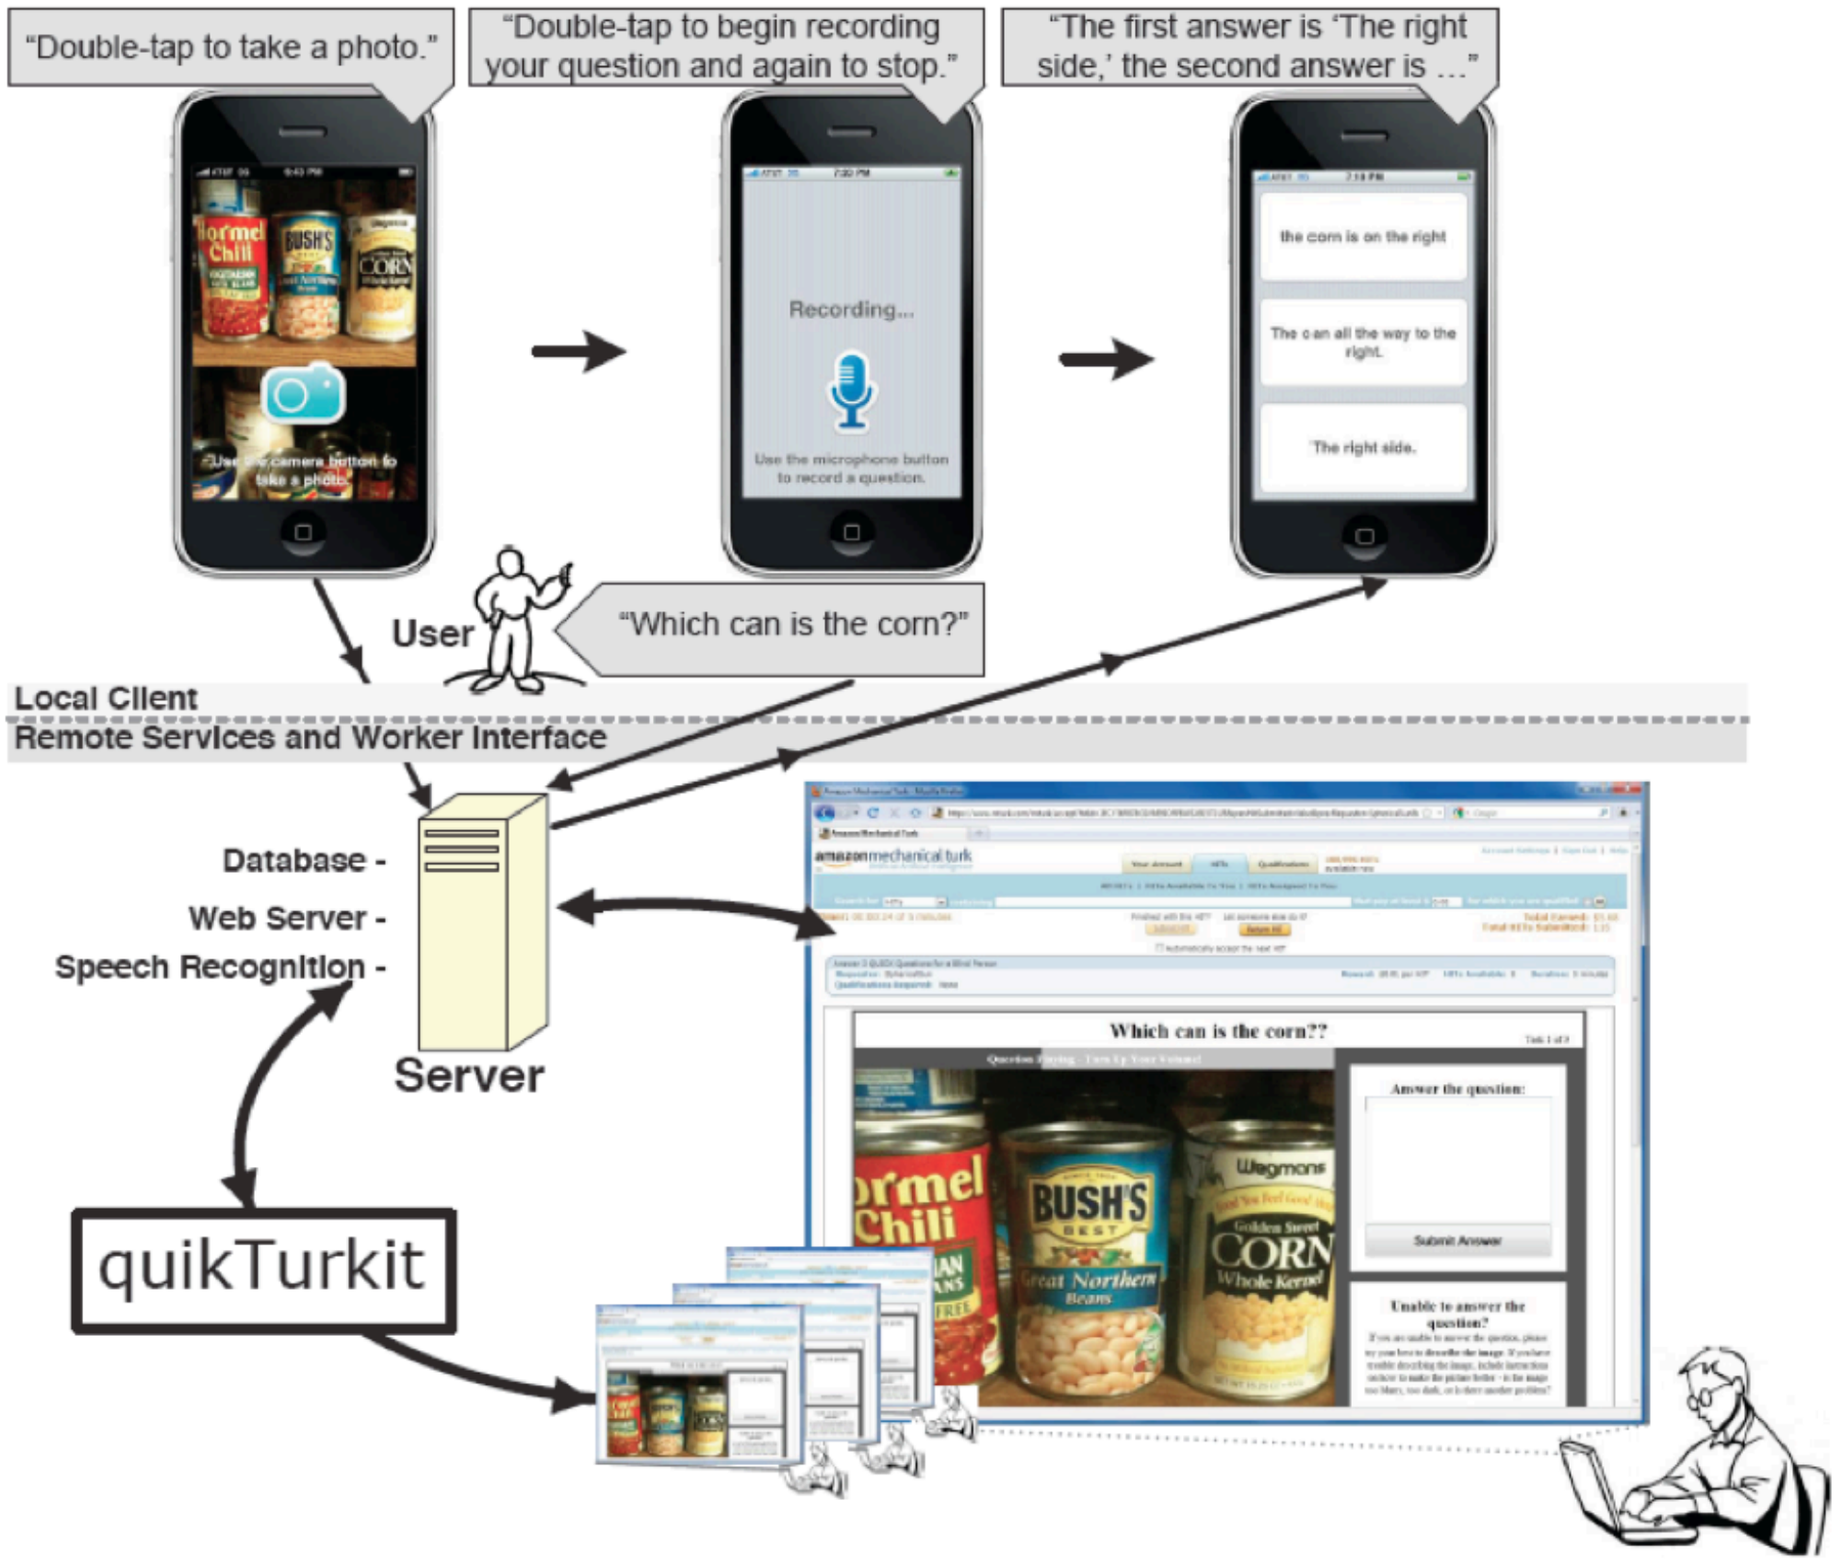

2010 VizWiz: using mechanical turk to help blind users

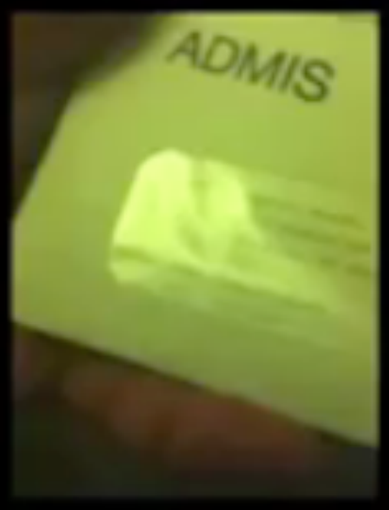

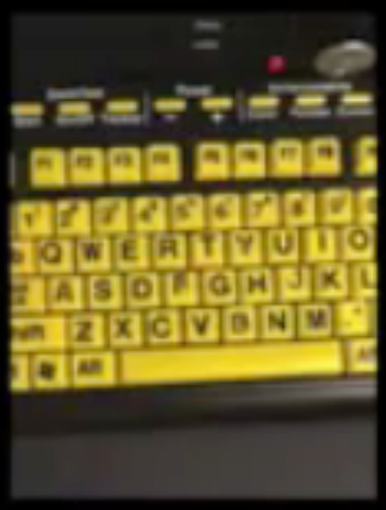


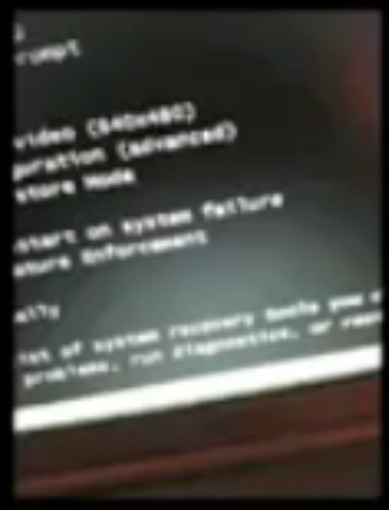
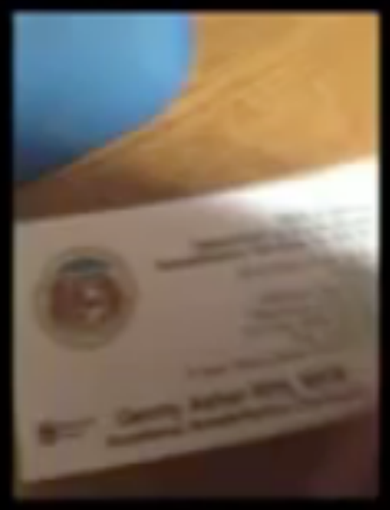

From May 31st, 2011 to May 31st, 2012,

5,329 users asked 40,748 questions

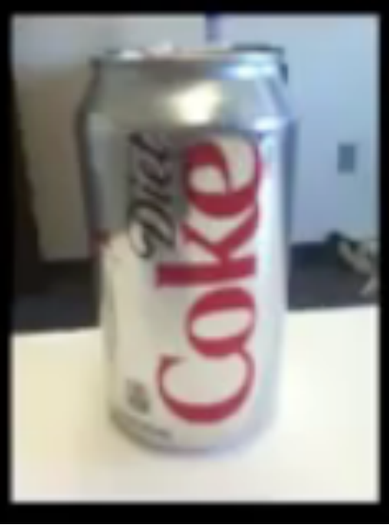



"How much is this?"

Question Type:
Identification – Currency

Primary Subject:
Object - Miscellaneous Objects

Urgency:
Within a few minutes (the question asked must be answered in 1 to 10 minutes)

Subjectivity/Objectivity:
Very Objective: The question is asking for only observations or facts

Photograph Quality:
5 out of 5

Reading - Mail

Primary Subject:
Object - Paper

Urgency:
Within an hour (the question asked must be answered in 10 minutes to 1 hour)

Subjectivity/Objectivity:
Very Objective: The question is asking for only observations or facts

Photograph Quality:
5 out of 5

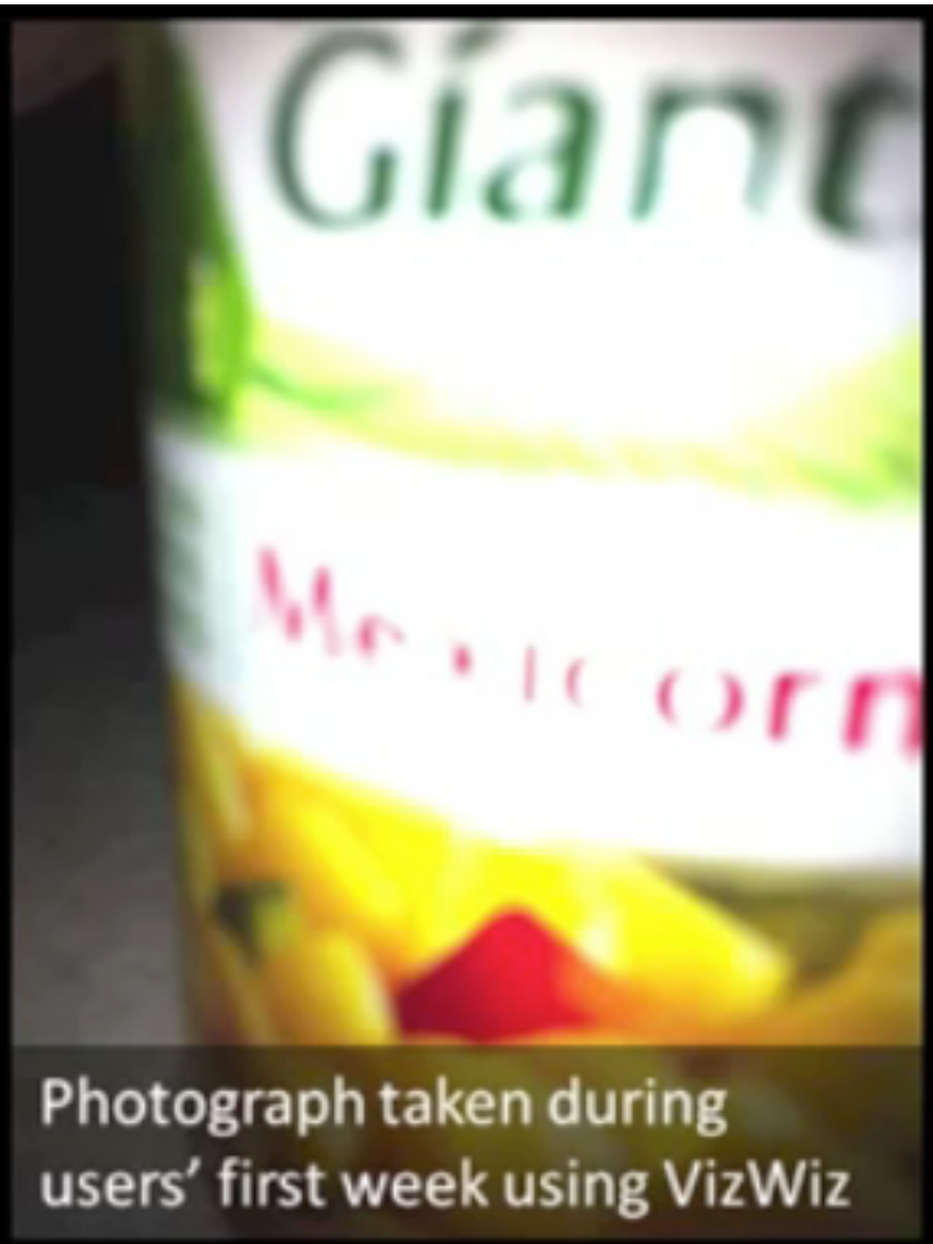

Photograph taken during users' first week using VizWiz

Photograph taken by same user during most recent week using VizWiz (7 months later)

VizWiz: Nearly Real-time Answers to Visual Questions

Jeffrey P. Bigham[†], *Chandrika Jayant*[$], *Hanjie Ji*[†], *Greg Little*[§], *Andrew Miller*[γ],
Robert C. Miller[§], *Robin Miller*[†], *Aubrey Tatarowicz*[§], *Brandyn White*[‡], *Samuel White*[†], *and Tom Yeh*[‡]

[†]University of Rochester Computer Science
Rochester, NY 14627 USA
{jbigham, hji, rmill13, swhite24}@cs.rochester.edu

[‡]University of Maryland Computer Science
College Park, MD 20742 USA
{bwhite, tomyeh}@umiacs.umd.edu

[§]MIT CSAIL
Cambridge, MA 02139 USA
{altat, glittle, rcm}@mit.edu

[γ]University of Central Florida CS
Orlando, FL 32816 USA
amiller@ucf.edu

[$]University of Washington CSE
Seattle, WA 98195 USA
cjayant@cs.washington.edu

ABSTRACT

The lack of access to visual information like text labels, icons, and colors can cause frustration and decrease independence for blind people. Current access technology uses automatic approaches to address some problems in this space, but the technology is error-prone, limited in scope, and quite expensive. In this paper, we introduce *VizWiz*, a talking application for mobile phones that offers a new alternative to answering visual questions in nearly real-time—asking multiple people on the web. To support answering questions quickly, we introduce a general approach for intelligently recruiting human workers in advance called *quikTurkit* so that workers are available when new questions arrive. A field deployment with 11 blind participants illustrates that blind people can effectively use VizWiz to cheaply answer questions in their everyday lives, highlighting issues that automatic approaches will need to address to be useful. Finally, we illustrate the potential of using VizWiz as part of the participatory design of advanced tools by using it to build and evaluate VizWiz::LocateIt, an interactive mobile tool that helps blind people solve general visual search problems.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces

General terms: Human Factors, Design, Experimentation

frustrations—a sighted person can tell in a glance if their clothes match before an important job interview, spot an empty picnic table at the park, or locate the restroom at the other end of the room without having to ask. Blind people often have effective, albeit inefficient, work-arounds that render individual problems into mere nuisances. Collectively, however, small problems can lead to decreased independence.

Talking mobile devices from both research and industry have been designed to help blind people solve visual problems in their everyday lives, but current automatic approaches are not yet up to the task. Products designed for blind people are specialized for a few functions, are prone to errors, and are usually quite expensive. As an example, the popular Kurzweil knfbReader software ($1000 USD) uses optical character recognition (OCR) to convert text to speech in pictures taken by users on their mobile devices [16]. When it works, this product offers the independence of reading printed material anywhere, but unfortunately, OCR cannot yet reliably identify the text in many real-world situations, such as the graphic labels on many products, a hand-written menu in a coffee shop, or even the street name on a street sign. Other popular products identify colors and read barcodes with similar performance (and prices). Filling the remaining void are a large number of human workers, volunteers, and friends

Visual Challenges in the Everyday Lives of Blind People

Erin Brady[1], Meredith Ringel Morris[2], Yu Zhong[1], Samuel White[1], and Jeffrey P. Bigham[1]
University of Rochester, Computer Science[1]
Rochester, NY 14627 USA
{brady, zyu, jbigham}@cs.rochester.edu
samuel.white@rochester.edu

Microsoft Research[2]
Redmond, WA 98052 USA
merrie@microsoft.com

ABSTRACT
The challenges faced by blind people in their everyday lives are not well understood. In this paper, we report on the findings of a large-scale study of the visual questions that blind people would like to have answered. As part of this year-long study, 5,329 blind users asked 40,748 questions about photographs that they took from their iPhones using an application called VizWiz Social. We present a taxonomy of the types of questions asked, report on a number of features of the questions and accompanying photographs, and discuss how individuals changed how they used VizWiz Social over time. These results improve our understanding of the problems blind people face, and may help motivate new projects more accurately targeted to help blind people live more independently in their everyday lives.

Author Keywords
Blind Users; Q&A; Accessibility; Crowdsourcing; Mobile

ACM Classification Keywords
H.5.m Information Interfaces and Presentation: Misc.

General Terms
Human Factors; Experimentation.

INTRODUCTION
Blind people confront a number of visual challenges every day – from reading the label on a frozen dinner to figuring out

a blind person take a picture, speak a question they'd like to know about the picture, and then get an answer back within a minute or so from "the crowd" [5]. VizWiz Social has been released "in the wild" since May 2011, and blind users have asked over 40,000 questions since then. Today's technology is targeted at answering some of them, e.g. "What color is this shirt?" and "What does this letter say?" But, others it cannot, for instance, "How many lines are on this pregnancy test?", "What does the sky look like right now?", and "Is my girlfriend hot?"

To help make sense of this diversity, we developed a taxonomy of the questions asked. By outlining the types of questions asked frequently, we hope to improve understanding of the challenges blind people face and help to motivate research into new technology to answer those questions automatically, which would be cheaper and faster. VizWiz Social also provides a rare look into the adoption of an assistive technology over the long term, and how a human-powered access technology [6] affects the user. For instance, do blind people become better photographers as they use VizWiz Social?

VizWiz Social provides insight into a specific but important subset of challenges faced by blind users, i.e., those that can be represented with a still photograph and brief audio description and that can be answered quickly but asynchronously. Other types of challenges, such as those where a user needs help in a situation requiring conveying and/or receiving continuous information, are beyond the bounds of the current

massive data set
allowed to answer many follow up research questions

asking humans vs. calling functions in your code?
what are some **differences?**

<30 sec brainstorming>

asking humans vs. calling functions in your code? what are the differences?

- **raw results noisy**
  - different workers, different effort (lazy or too eager)
  - 30% are poor quality
- **no guaranteed response time**
  - can be no response at all
- **moral implications**
  - what if blind user gets an intentionally wrong answer?

Soylent
is people.

2010 Soylent: collaborative editing

mhh…
can I really trust sb else to shorten my text?

solution:
get many people involved!
let them rate each other!

iterative process to reduce noise and allow for parallelism

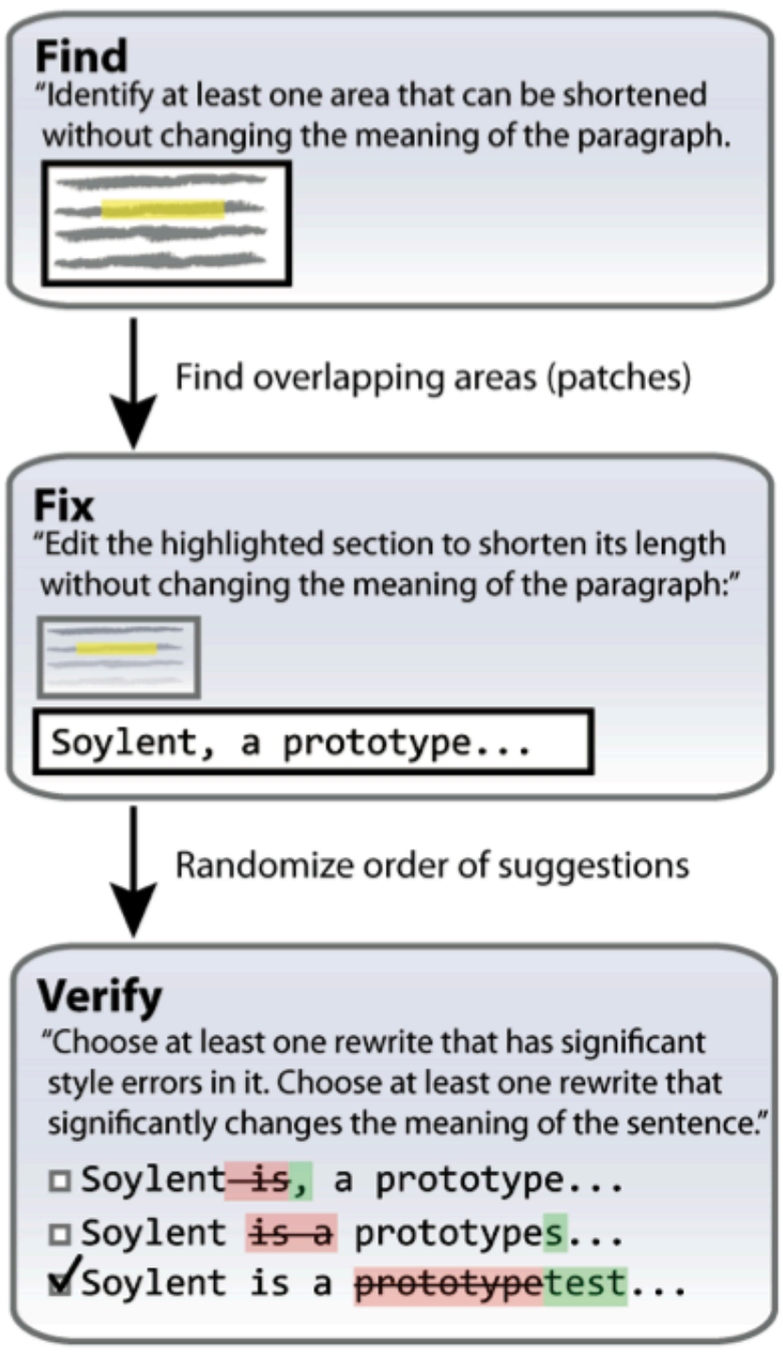

find-fix-verify pattern:

#1 find sections that need work (20% agreement)

#2 fix a single section (produce 3-5 alternatives)

#3 verify by voting on best result – or – flag poor suggestions

Soylent: A Word Processor with a Crowd Inside

Michael S. Bernstein[1], *Greg Little*[1], *Robert C. Miller*[1],
Björn Hartmann[2], *Mark S. Ackerman*[3], *David R. Karger*[1], *David Crowell*[1], *Katrina Panovich*[1]

[1] MIT CSAIL
Cambridge, MA
{msbernst, glittle, rcm,
karger, dcrowell, kp}@csail.mit.edu

[2] Computer Science Division
University of California, Berkeley
Berkeley, CA
bjoern@cs.berkeley.edu

[3] Computer Science & Engineering
University of Michigan
Ann Arbor, MI
ackerm@umich.edu

ABSTRACT

This paper introduces architectural and interaction patterns for integrating crowdsourced human contributions directly into user interfaces. We focus on writing and editing, complex endeavors that span many levels of conceptual and pragmatic activity. Authoring tools offer help with pragmatics, but for higher-level help, writers commonly turn to other people. We thus present Soylent, a word processing interface that enables writers to call on Mechanical Turk workers to shorten, proofread, and otherwise edit parts of their documents on demand. To improve worker quality, we introduce the Find-Fix-Verify crowd programming pattern, which splits tasks into a series of generation and review stages. Evaluation studies demonstrate the feasibility of crowdsourced editing and investigate questions of reliability, cost, wait time, and work time for edits.

ACM Classification: H5.2 [Information interfaces and presentation]: User Interfaces. - Graphical user interfaces.

General terms: Design, Human Factors

Keywords: Outsourcing, Mechanical Turk, Crowdsourcing

INTRODUCTION

Word processing is a complex task that touches on many goals of human-computer interaction. It supports a deep cognitive activity – writing – and requires complicated manipulations. Writing is difficult: even experts routinely make style, grammar and spelling mistakes. Then, when a writer makes high-level decisions like changing a passage answer ourselves [8]; masses of volunteer editors flag spam edits on Wikipedia [13]. Writing is no exception [7]: we commonly recruit friends and colleagues to help us shape and polish our writing. But we cannot always rely on them: colleagues do not want to proofread every sentence we write, cut a few lines from every paragraph in a ten-page paper, or help us format thirty ACM-style references.

As a step toward integrating this human expertise permanently into our writing tools, we present *Soylent*, a word processing interface that utilizes crowd contributions to aid complex writing tasks ranging from error prevention and paragraph shortening to automation of tasks like citation searches and tense changes. We hypothesize that crowd workers with a basic knowledge of written English can support both novice and expert writers. These workers perform tasks that the writer might not, such as scrupulously scanning for text to cut, or updating a list of addresses to include a zip code. They can also solve problems that artificial intelligence cannot, for example flagging writing errors that the word processor does not catch.

Soylent aids the writing process by integrating paid crowd workers from Amazon's Mechanical Turk platform[1] into Microsoft Word. *Soylent is people*: its core algorithms involve calls to Mechanical Turk workers (Turkers). Soylent is comprised of three main components:

1) *Shortn*, a text shortening service that cuts selected text down to 85% of its original length typically without

history

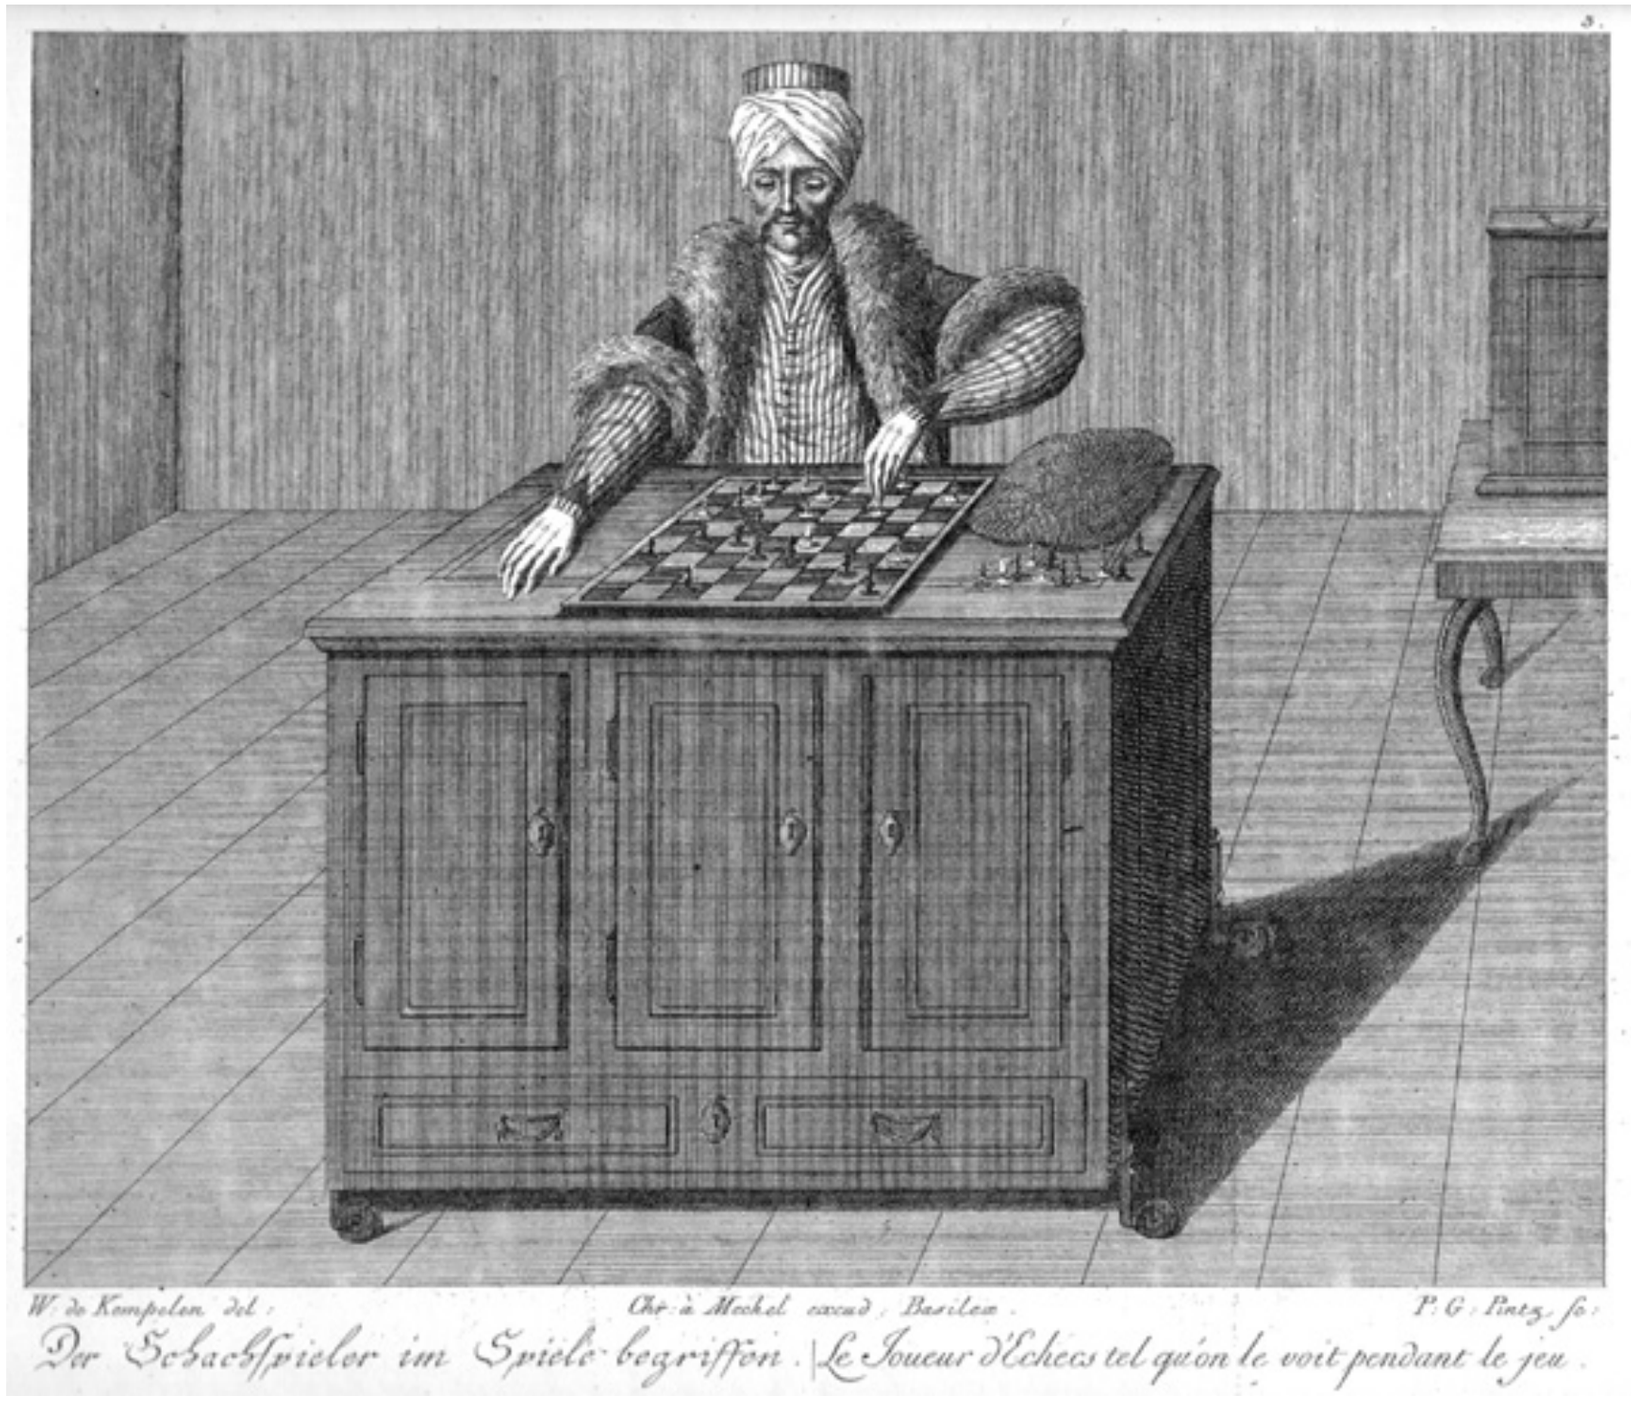

Mechanical Turk

= Automaton Chess Player => Hungarian: A Török

- a fake chess-playing machine in the late 18th century

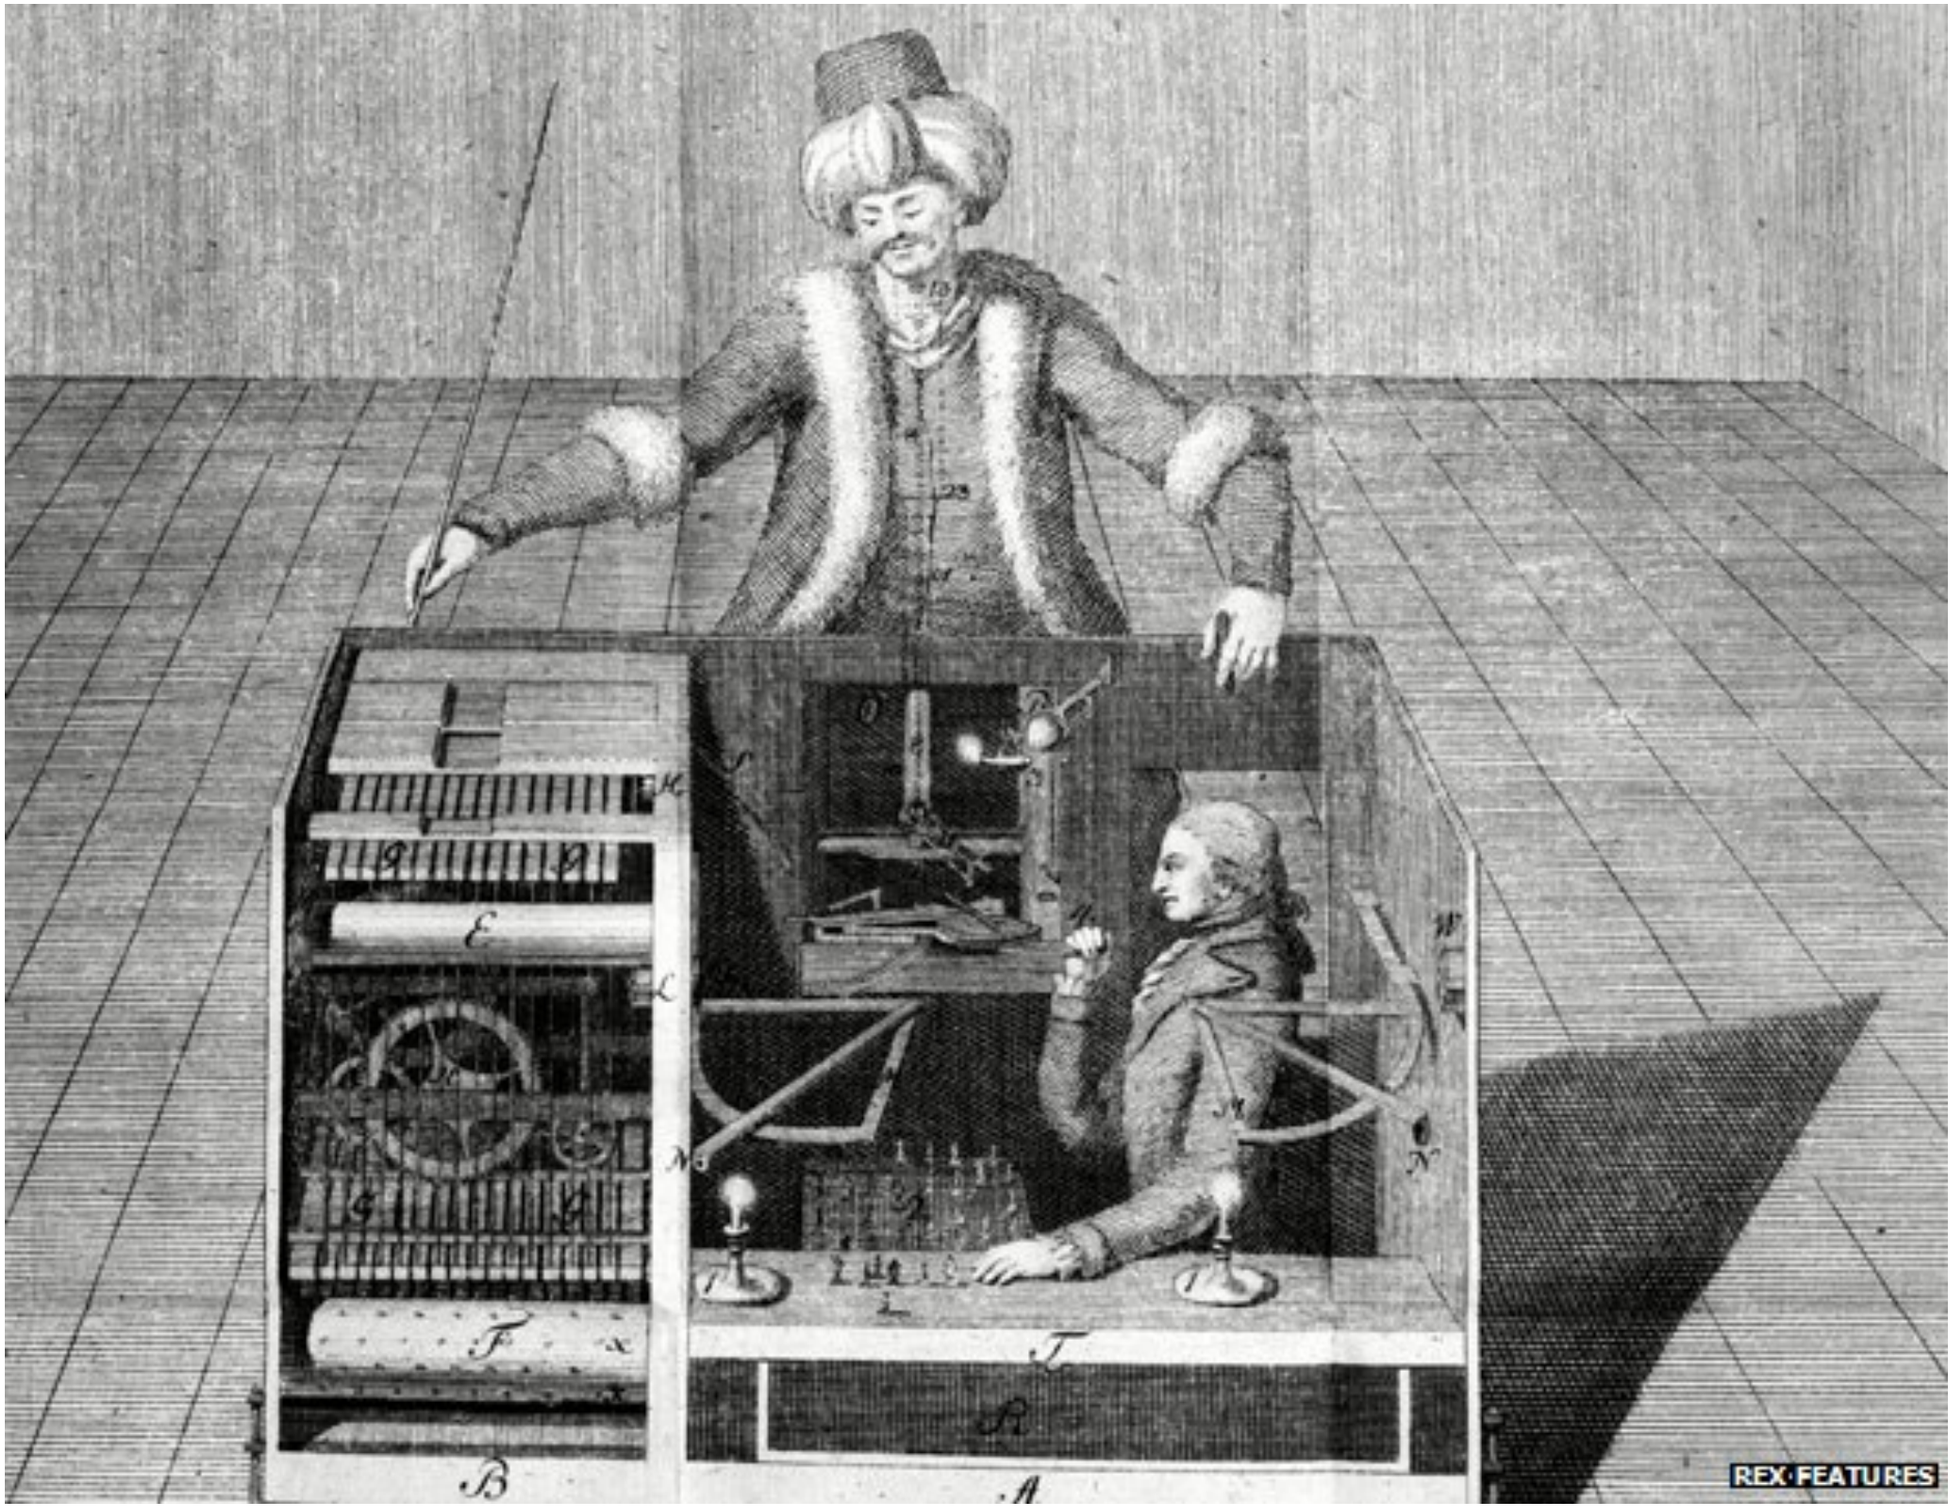

- toured Europe
- took 50 years for anyone to discover the hoax

before computers were digital machines,
they were actually **humans…**

Hidden Figures takes us back to a time when computers were people, women, and black

February 14, 2017 4.10am EST

world wars: maps, artillery tables…

all done by human computing (by hand… wow)

human computation
borrowed strategies from **industrial assembly lines:**
each worker performs one operation that is highly specialized

[Ford 1913]

amazonmechanical turk
Artificial Artificial Intelligence

Your Account | HITs | Qualifications

287,164 HITs available now

All HITs | **HITs Available To You** | **HITs Assigned To You**

Find HITs **containing** **that pay at least $** 0.00 for which you are qual require Master Qualific

All HITs

1-10 of 2249 Results

Sort by: HITs Available (most first) GO!

Show all details | Hide all details

Determine the industry of this business

Requester:	Matt Park	**HIT Expiration Date:**	May 13, 2013 (2 days 12 hours)	**Reward:**	$0.05
		Time Allotted:	3 minutes	**HITs Available:**	27016

Classify Arabic Tweets Dialects SEE REVISED HIT

Requester:	Chris Callison-Burch	**HIT Expiration Date:**	Aug 4, 2013 (12 weeks 1 day)	**Reward:**	$0.00
		Time Allotted:	60 minutes	**HITs Available:**	23532

Classify Arabic Tweets Dialects (No Qualification)

Requester:	Chris Callison-Burch	**HIT Expiration Date:**	Aug 22, 2013 (14 weeks 4 days)	**Reward:**	$0.05
		Time Allotted:	60 minutes	**HITs Available:**	18067

Find the lowest price for the items searched on the website shown

Requester:	IntSup	**HIT Expiration Date:**	May 14, 2013 (3 days 3 hours)	**Reward:**	$0.01
		Time Allotted:	60 minutes	**HITs Available:**	17672

today the assembly line looks like this…

human computers were eventually **replaced by machinery**
… but we still use them today for really hard computing tasks

having humans perform computational tasks

leveraging large crowds of people over the internet

to perform a task typically performed by a single person

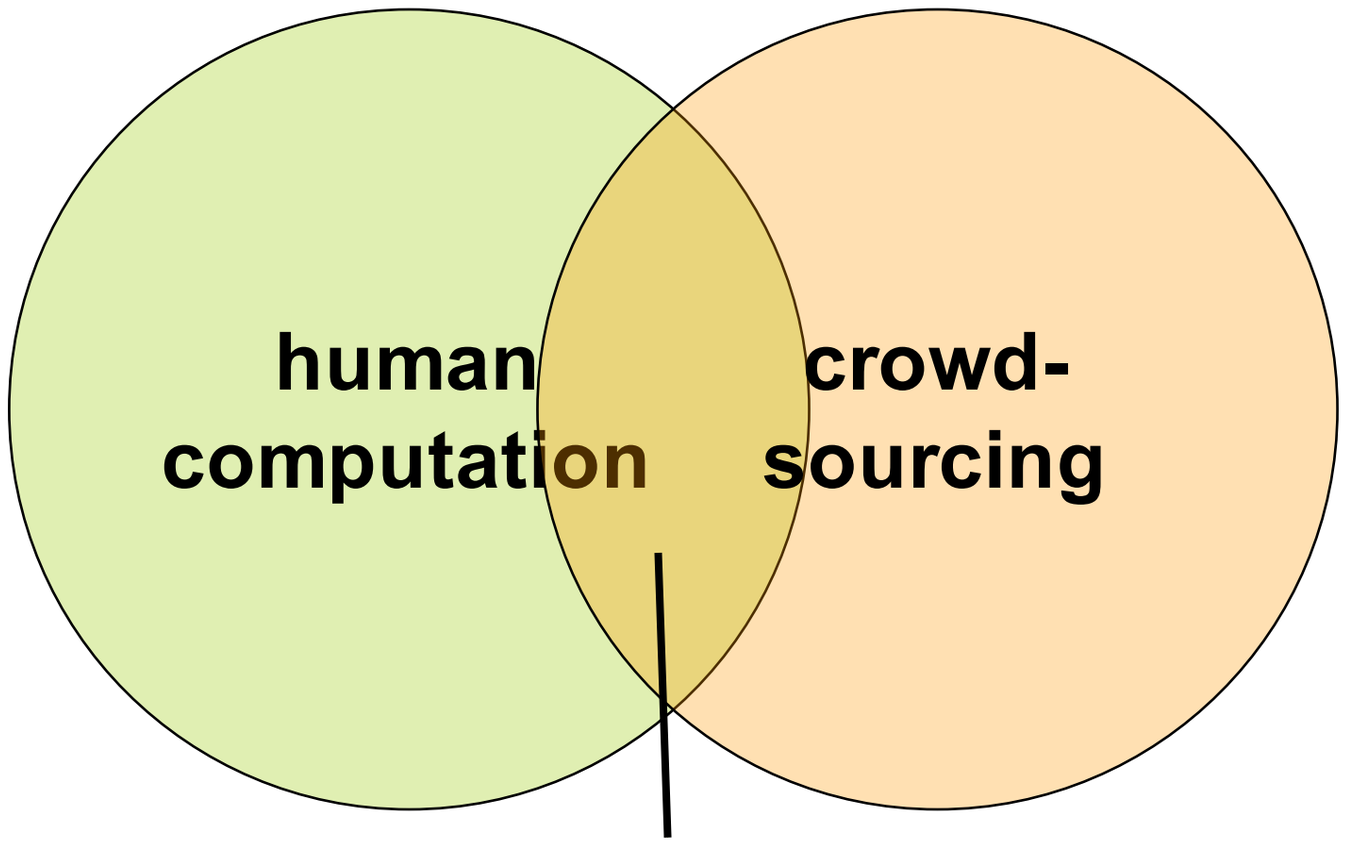

crowdsourced assembly lines

(e.g. mechanical turk)

leveraging large crowds of people over the internet

to perform a task typically performed by a single person

examples?

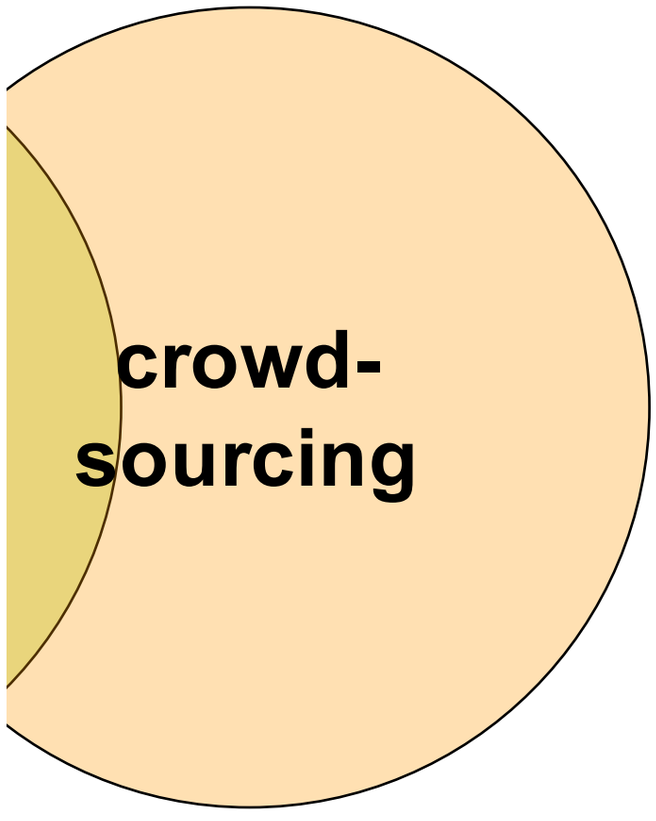

<30 sec brainstorming>

Not logged in Talk Contributions Create accoun

WIKIPEDIA
The Free Encyclopedia

Main page
Contents
Featured content
Current events
Random article
Donate to Wikipedia
Wikipedia store

Interaction
Help
About Wikipedia
Community portal
Recent changes
Contact page

Tools
What links here
Related changes
Upload file
Special pages
Permanent link
Page information
Wikidata item
Cite this page

Print/export
Create a book
Download as PDF

Massachusetts Institute of Technology

From Wikipedia, the free encyclopedia

Coordinates: 42.35982°N 71.0

"MIT" redirects here. For other uses, see MIT (disambiguation).

The **Massachusetts Institute of Technology** (**MIT**) is a private research university in Cambridge, Massachusetts. It is often ranked as one of the world's most prestigious universities.[10][11][12]

Founded in 1861 in response to the increasing industrialization of the United States, MIT adopted a European polytechnic university model and stressed laboratory instruction in applied science and engineering. Researchers worked on computers, radar, and inertial guidance during World War II and the Cold War. Post-war defense research contributed to the rapid expansion of the faculty and campus under James Killian. The current 168-acre (68.0 ha) campus opened in 1916 and extends over 1 mile (1.6 km) along the northern bank of the Charles River basin.

The Institute is traditionally known for its research and education in the physical sciences and engineering, but more recently in biology, economics, linguistics, and management as well. MIT is a member of the Association of American Universities (AAU). For several years, MIT's School of Engineering has been ranked first in various international and national university rankings, while MIT is also often ranked among the world's top universities overall.[10][11][13][12][14] The MIT Engineers compete in 31 sports, most teams of which compete in the NCAA Division III's New England Women's and Men's Athletic Conference, whereas the Division I rowing programs compete as part of the EARC and EAWRC.

As of 2017, 88 Nobel laureates, 52 National Medal of Science recipients, 65 Marshall Scholars, 45 Rhodes Scholars, 38 MacArthur Fellows, 34 astronauts, 21 Turing award winners, 16 Chief Scientists of the U.S. Air Force, and 6 Fields Medalists have been affiliated with MIT. The school has a strong entrepreneurial culture and the aggregated revenues of companies founded by MIT alumni would rank as the eleventh-largest economy in the world.[15][16]

Contents [hide]

Massachusetts Institute of Techno

Motto	*Mens et Manus* (Latin)
Motto in English	Mind and Hand[1]
Type	Private Land grant Space-grant
Established	April 10, 1861
Academic affiliations	AAU AICUM AITU APLU COFHE NAICU[2] URA 568 Group
Endowment	$14.8 billion (2017)[3]
Chancellor	Cynthia Barnhart

Wikipedia (anyone ever tried to put up some fake information?)

KICKSTARTER

68,929
backers

$10,266,845
pledged of $100,000 goal

crowdfunding (Kickstarter takes 5%)

crowd-science (citizen-science)

GLOBE at Night: “asks people to count the number of stars that they can see from their location” to determine global light pollution

motivation of the crowd

what motivates people to contribute?

<30 sec brainstorming>

what motivates people to contribute?

#1 pay
#2 implicit work (you don’t know you are doing it)
#3 enjoyment
#4 altruism
…

#1 pay

amazonmechanical turk
Artificial Artificial Intelligence

Your Account | HITs | Qualifications

Already have an account
Sign in as a Worker | Requeste

Introduction | **Dashboard** | **Status** | **Account Settings**

Mechanical Turk is a marketplace for work.

We give businesses and developers access to an on-demand, scalable workforce.
Workers select from thousands of tasks and work whenever it's convenient.

76,187 HITs available. View them now.

Make Money
by working on HITs

HITs - *Human Intelligence Tasks* - are individual tasks that you work on. Find HITs now.

As a Mechanical Turk Worker you:

- Can work from home
- Choose your own work hours
- Get paid for doing good work

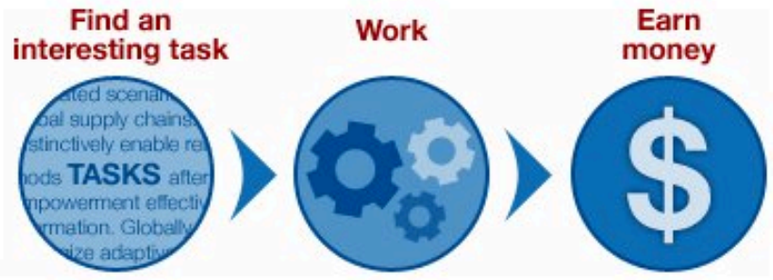

Find HITs Now

or learn more about being a **Worker**

Get Results
from Mechanical Turk Workers

Ask workers to complete HITs - *Human Intelligence Tasks* - and get results using Mechanical Turk. Register Now

As a Mechanical Turk Requester you:

- Have access to a global, on-demand, 24 x 7 workforce
- Get thousands of HITs completed in minutes
- Pay only when you're satisfied with the results

Get Started

FAQ | Contact Us | Careers at Amazon | Developers | Press | Policies
©2005-2010 Amazon.com, Inc. or its Affiliates
An amazon.com. compar

amazon mechanical turk

90% of tasks pay less than $0.10

(**anyone** ever used mechanical turk to make money?)

2012 Umati: paid in snacks

Communitysourcing: Engaging Local Crowds to Perform Expert Work Via Physical Kiosks

Kurtis Heimerl[1], **Brian Gawalt**[1], **Kuang Chen**[1], **Tapan S. Parikh**[2], **Björn Hartmann**[1]
University of California, Berkeley – Computer Science Division[1], School of Information[2]
{kheimerl,gawalt,kuangc,bjoern}@cs.berkeley.edu, parikh@ischool.berkeley.edu

ABSTRACT

Online labor markets, such as Amazon's Mechanical Turk, have been used to crowdsource simple, short tasks like image labeling and transcription. However, expert knowledge is often lacking in such markets, making it impossible to complete certain classes of tasks. In this work we introduce an alternative mechanism for crowdsourcing tasks that require specialized knowledge or skill: *communitysourcing* — the use of physical kiosks to elicit work from specific populations. We investigate the potential of communitysourcing by designing, implementing and evaluating *Umati: the communitysourcing vending machine*. Umati allows users to earn credits by performing tasks using a touchscreen attached to the machine. Physical rewards (in this case, snacks) are dispensed through traditional vending mechanics. We evaluated whether communitysourcing can accomplish expert work by using Umati to grade Computer Science exams. We placed Umati in a university Computer Science building, targeting students with grading tasks for snacks. Over one week, 328 unique users (302 of whom were students) completed 7771 tasks (7240 by students). 80% of users had never participated in a crowdsourcing market before. We found that Umati was able to grade exams with 2% higher accuracy (at the same price) or at 33% lower cost (at equivalent accuracy) than traditional single-expert grading. Mechanical Turk workers had no success grading the same exams. These results indicate that communitysourcing can successfully elicit high-quality expert work from specific communities.

Figure 1. With Umati, the communitysourcing vending machine, users complete tasks on a touchscreen and receive non-monetary rewards.

INTRODUCTION

Crowdsourcing, the division and assignment of tasks to large, distributed groups of online users, has the potential to create new jobs, improve the efficiency of labor markets, and enable a wide variety of new applications. Researchers have demonstrated compelling new systems enabled by crowdsourcing, including applications that assist the blind with visual tasks [8] and that help writers to copy-edit prose [7]. Many crowdsourcing efforts leverage *microtask* markets, which provide platforms for posting and finding short tasks – frequently seconds to minutes long. One of the best-known markets, Amazon Mechanical Turk (MTurk),

#2 implicit work

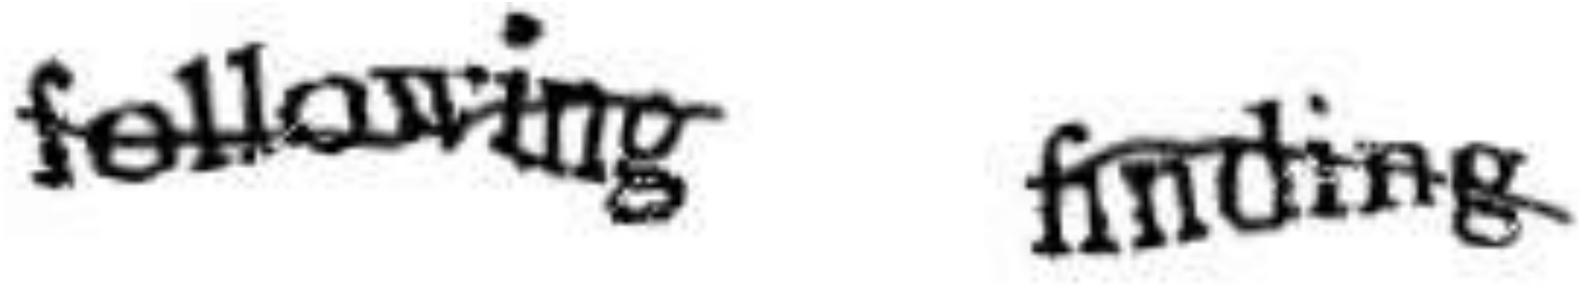

CAPTCHAs…

how to use this for solving **real-world problems?**

<30 sec brainstorming>

The New-York State Yacht Squadron, on its annual cruise to Newport, came into the harbor yesterday afternoon. The following are the names of the boats that came to anchor here: *Jessie*, *Geraldine*, *Evelyn*, *Annie*, *Mannering*, *Julia*, *Bonita*, *Magic*, *Widgeon*, *Rambler*, *Fleur-de-Lis*, *Henrietta*, *Sea-Drift* and *Maria*, with the steamer *America* as a tender. On anchoring, each boat fired a gun, according to custom. The reports were heard distinctly in the city, causing considerable inquiry as to "what was up," and quite a number of sanguine individuals came into our office to inquire if the guns were not annunciatory signals of the successful laying of the Atlantic Cable. We invariably replied in the negative. The squadron will leave to-day for Newport. The yachts *Washington* and *Rambler*, of this city, start with it, with parties of New-Haven people.

"Our apparatus is deployed in more than 40,000 Web sites and has transcribed over 440 million words. "

The Norwich line steamboat train, from New-London for Boston, this morning ran off the track seven miles north of New-London.

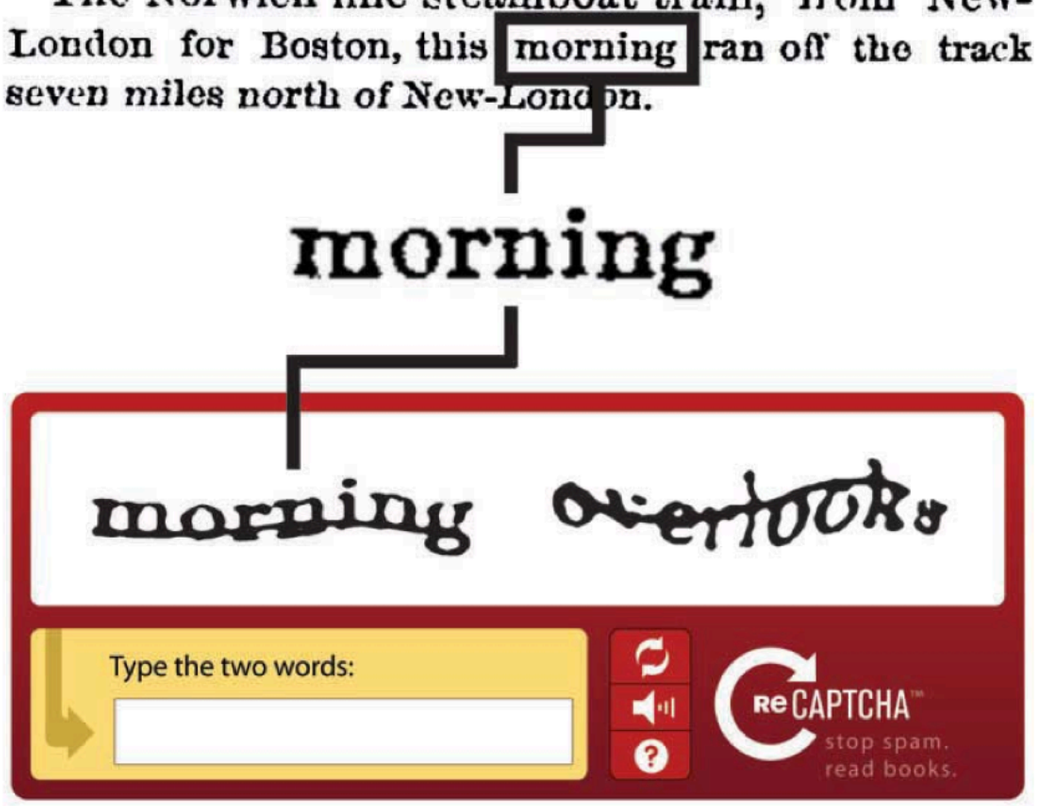

for hard to read text!

reCAPTCHA: Human-Based Character Recognition via Web Security Measures

Luis von Ahn,* Benjamin Maurer, Colin McMillen, David Abraham, Manuel Blum

CAPTCHAs (Completely Automated Public Turing test to tell Computers and Humans Apart) are widespread security measures on the World Wide Web that prevent automated programs from abusing online services. They do so by asking humans to perform a task that computers cannot yet perform, such as deciphering distorted characters. Our research explored whether such human effort can be channeled into a useful purpose: helping to digitize old printed material by asking users to decipher scanned words from books that computerized optical character recognition failed to recognize. We showed that this method can transcribe text with a word accuracy exceeding 99%, matching the guarantee of professional human transcribers. Our apparatus is deployed in more than 40,000 Web sites and has transcribed over 440 million words.

A CAPTCHA (*1*, *2*) is a challenge response test used on the World Wide Web to determine whether a user is a human or a computer. The acronym stands for Completely Automated Public Turing test to tell Computers and Humans Apart. A typical CAPTCHA is an image containing several distorted characters that appears at the bottom of Web registration forms. Users are asked to type the wavy characters to "prove" they are human. Current computer programs cannot read distorted text as well as humans can (*3*), so CAPTCHAs act as sentries against automated programs that attempt to abuse online services. Owing to their effectiveness as a security measure, CAPTCHAs are used to protect many types of Web sites, including free e mail providers, ticket sellers, social networks, wikis, and blogs. For example, CAPTCHAs prevent ticket scalpers from using computer programs to buy large numbers of concert tickets, only to re sell them at an inflated price. Sites such as Gmail and Yahoo Mail use CAPTCHAs to stop spammers from obtaining millions of free e mail accounts, which they would use to send spam e mail.

According to our estimates, humans around the world type more than 100 million CAPTCHAs every day (see supporting online text), in each case spending a few seconds typing the distorted characters. In aggregate, this amounts to hundreds of thousands of human hours per day. We report on an experiment that attempts to make positive use of the time spent by humans solving CAPTCHAs. Although CAPTCHAs are effective at preventing large scale abuse of online services, the mental effort each person spends solving them is otherwise wasted. This mental effort is invaluable, because deciphering CAPTCHAs requires people to perform a task that computers cannot.

We show how it is possible to use CAPTCHAs to help digitize typeset texts in nondigital form by enlisting humans to decipher the words that computers cannot recognize. Physical books and other texts written before the computer age are currently being digitized en masse (e.g., by the Google Books Project and the nonprofit Internet

Computer Science Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA.

*To whom correspondence should be addressed. E mail: biglou@cs.cmu.edu

#3 enjoyment

Player 1 guesses: purse
Player 1 guesses: bag
Player 1 guesses: brown

Success! Agreement on “purse”

Player 2 guesses: handbag

Player 2 guesses: purse
Success! Agreement on “purse”

2004: games with a purpose for **image labeling**

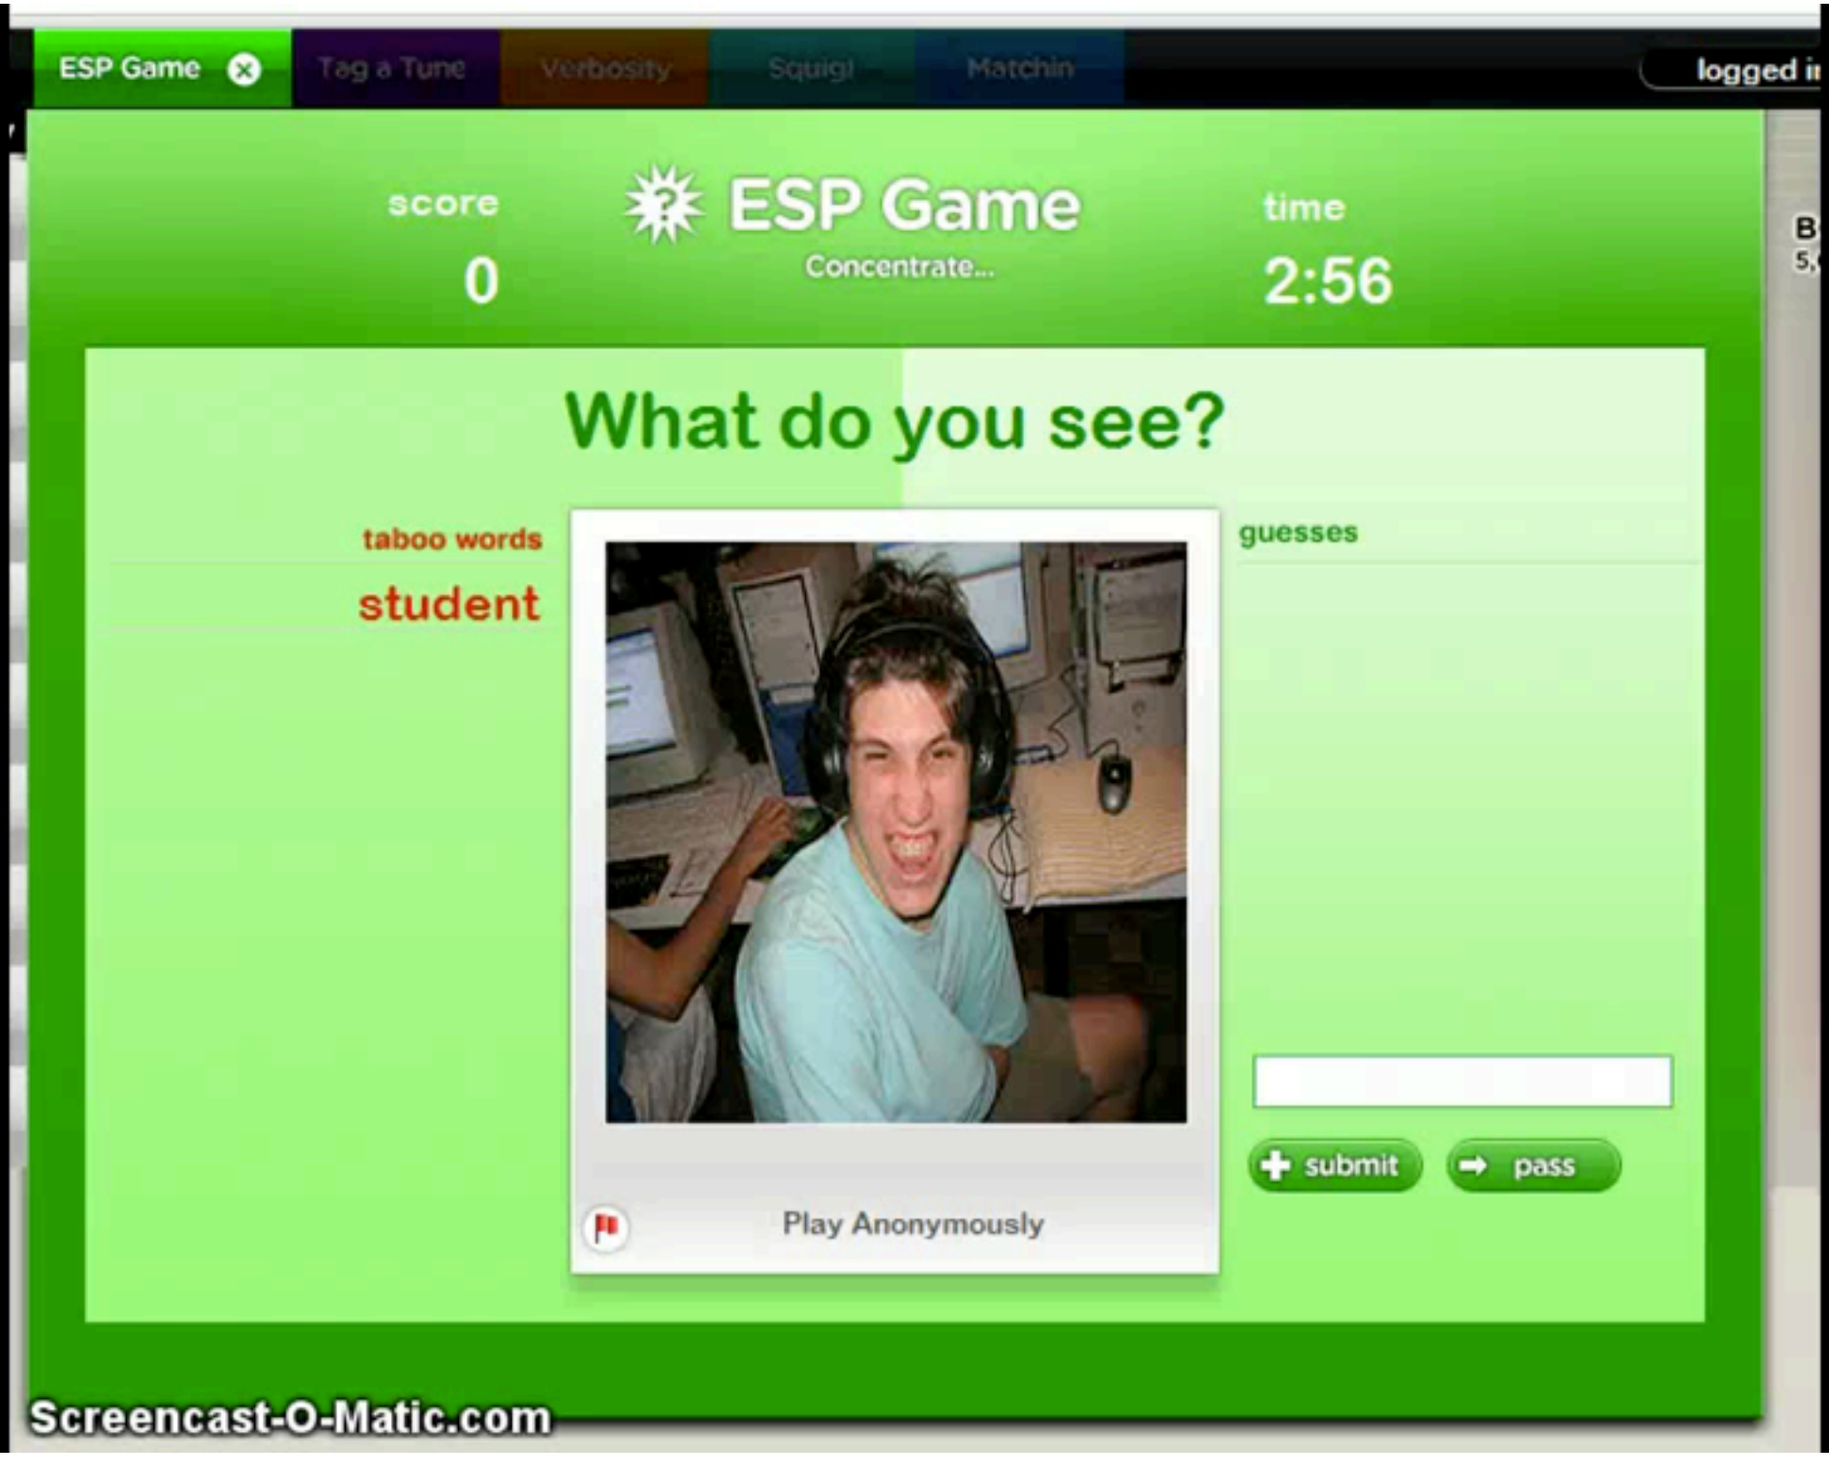

2004: games with a purpose for **image labeling**

Labeling Images with a Computer Game

Luis von Ahn and Laura Dabbish
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA, USA
{biglou,dabbish}@cs.cmu.edu

Abstract

We introduce a new interactive system: a game that is fun and can be used to create valuable output. When people play the game they help determine the contents of images by providing meaningful labels for them. If the game is played as much as popular online games, we estimate that most images on the Web can be labeled in a few months. Having proper labels associated with each image on the Web would allow for more accurate image search, improve the accessibility of sites (by providing descriptions of images to visually impaired individuals), and help users block inappropriate images. Our system makes a significant contribution because of its valuable output and because of the way it addresses the image-labeling problem. Rather than using computer vision techniques, which don't work well enough, we encourage people to do the work by taking advantage of their desire to be entertained.

Categories & Subject Descriptors: I.2.6 [**Learning**]: Knowledge acquisition. H.3.m [**Information Retrieval**]: miscellaneous. H.5.3 [**HCI**]: Web-based interaction.

General Terms: Design, Human Factors, Languages

Keywords: Distributed knowledge acquisition, image labeling, online games, World Wide Web.

INTRODUCTION

Images on the Web present a major technological challenge. There are millions of them, there are no guidelines about providing appropriate textual descriptions for them, and computer vision hasn't yet produced a

The only method currently available for obtaining precise image descriptions is manual labeling, which is tedious and thus extremely costly. But, what if people labeled images without realizing they were doing so? What if the experience was enjoyable? In this paper we introduce a new interactive system in the form of a game with a unique property: the people who play the game label images for us.

The labels generated by our game can be useful for a variety of applications. For accessibility purposes, visually impaired individuals surfing the Web need textual descriptions of images to be read aloud. For computer vision research, large databases of labeled images are needed as training sets for machine learning algorithms. For image search over the Web and inappropriate (e.g., pornographic) content filtering, proper labels could dramatically increase the accuracy of current systems.

We believe our system makes a significant contribution, not only because of its valuable output, but also because of the way it addresses the image-labeling problem. Rather than making use of computer vision techniques, we take advantage of people's existing perceptual abilities and desire to be entertained.

Our goal is ambitious: to label the majority of images on the World Wide Web. If our game is deployed at a popular gaming site like Yahoo! Games and if people play it as much as other online games, we estimate that most images on the Web can be properly labeled in a matter of weeks. As we show below, 5,000 people continuously playing the

YOU AND A RANDOM PARTNER TAKE TURNS **PEEKING** AND **BOOMING**

PEEK : GUESS WHAT YOUR PARTNER IS REVEALING

HINTS HELP YOU GUESS

PASS FOR DIFFICULT IMAGES

BOOM : REVEAL PARTS OF THE IMAGE TO YOUR PARTNER

GIVE HINTS IF NECESSARY

TELL YOUR PARTNER IF A GUESS IS **HOT** OR **COLD**

2006: **local** image labeling

Peekaboom: A Game for Locating Objects in Images

Luis von Ahn, Ruoran Liu and Manuel Blum
Computer Science Department, Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh PA 15213
{biglou, royliu, mblum}@cs.cmu.edu

ABSTRACT
We introduce Peekaboom, an entertaining web-based game that can help computers locate objects in images. People play the game because of its entertainment value, and as a side effect of them playing, we collect valuable image metadata, such as which pixels belong to which object in the image. The collected data could be applied towards constructing more accurate computer vision algorithms, which require massive amounts of training and testing data not currently available. Peekaboom has been played by thousands of people, some of whom have spent over 12 hours a day playing, and thus far has generated millions of data points. In addition to its purely utilitarian aspect, Peekaboom is an example of a new, emerging class of games, which not only bring people together for leisure purposes, but also exist to improve artificial intelligence. Such games appeal to a general audience, while providing answers to problems that computers cannot yet solve.

Author Keywords
Distributed knowledge acquisition, object segmentation, object recognition, computer vision, Web-based games.

ACM Classification Keywords:
I.2.6 [**Learning**]: Knowledge acquisition. H.5.3 [**HCI**]: Web-based interaction.

INTRODUCTION

training an algorithm for testing whether an image contains a dog would involve presenting it with multiple images of dogs, each annotated with the precise location of the dog in the image. After processing enough images, the algorithm learns to find dogs in arbitrary images. A major problem with this approach, however, is the lack of training data, which, obviously, must be prepared by hand. Databases for training computer vision algorithms currently have hundreds or at best a few thousand images [13] — orders of magnitude less than what is required.

In this paper we address the problem of constructing a massively large database for training computer vision algorithms. The target database will contain millions of images, all fully annotated with information about what objects are in the image, where each object is located, and how much of the image is necessary to recognize it. Our database will be similar to those previously shown to be useful for training computer vision algorithms (e.g. [13]).

To construct such a database, we follow the approach taken by the ESP Game [1] and introduce a new game called Peekaboom. Peekaboom is an extremely enjoyable networked game in which, simply by playing, people help construct a database for training computer vision algorithms. We guarantee the database's correctness even if the people playing the game don't intend it. As we will

vision:

‘I want to label all the images on the internet’

Translate this text

Woher hast du das?

Type In English

2009 DuoLingo:

translate the web while learning a language

Duolingo: Learn a Language for Free while Helping to Translate the Web

Luis von Ahn
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213, USA
biglou@cs.cmu.edu

ABSTRACT

I want to translate the Web into every major language: every webpage, every video, and, yes, even Justin Bieber's tweets.

With its content split up into hundreds of languages — and with over 50% of it in English — most of the Web is inaccessible to most people in the world. This problem is pressing, now more than ever, with millions of people from China, Russia, Latin America and other quickly developing regions entering the Web. In this talk, I introduce my new project, called Duolingo, which aims at breaking this language barrier, and thus making the Web truly "world wide."

We have all seen how systems such as Google Translate are improving every day at translating the gist of things written in other languages. Unfortunately, they are not yet accurate enough for my purpose: Even when what they spit out is intelligible, it's so badly written that I can't read more than a few lines before getting a headache.

With Duolingo, our goal is to encourage people, like you and me, to translate the Web into their native languages.

BIO

Luis von Ahn is the A. Nico Habermann Associate Professor of Computer Science at Carnegie Mellon University. He is working to develop a new area of computer science that he calls Human Computation, which aims to build systems that combine the intelligence of humans and computers to solve large-scale problems that neither can solve alone.

An example of his work is reCAPTCHA, in which over one billion people — 15% of humanity — have helped digitize books and newspapers.

Among his many honors are a MacArthur Fellowship, a Packard Fellowship, a Sloan Research Fellowship, a Microsoft New Faculty Fellowship, the ACM Grace Hopper Award, and CMU's Herbert A. Simon Award for Teaching Excellence and Alan J. Perlis Teaching Award.

vision:

‘I want to translate the entire Internet.’

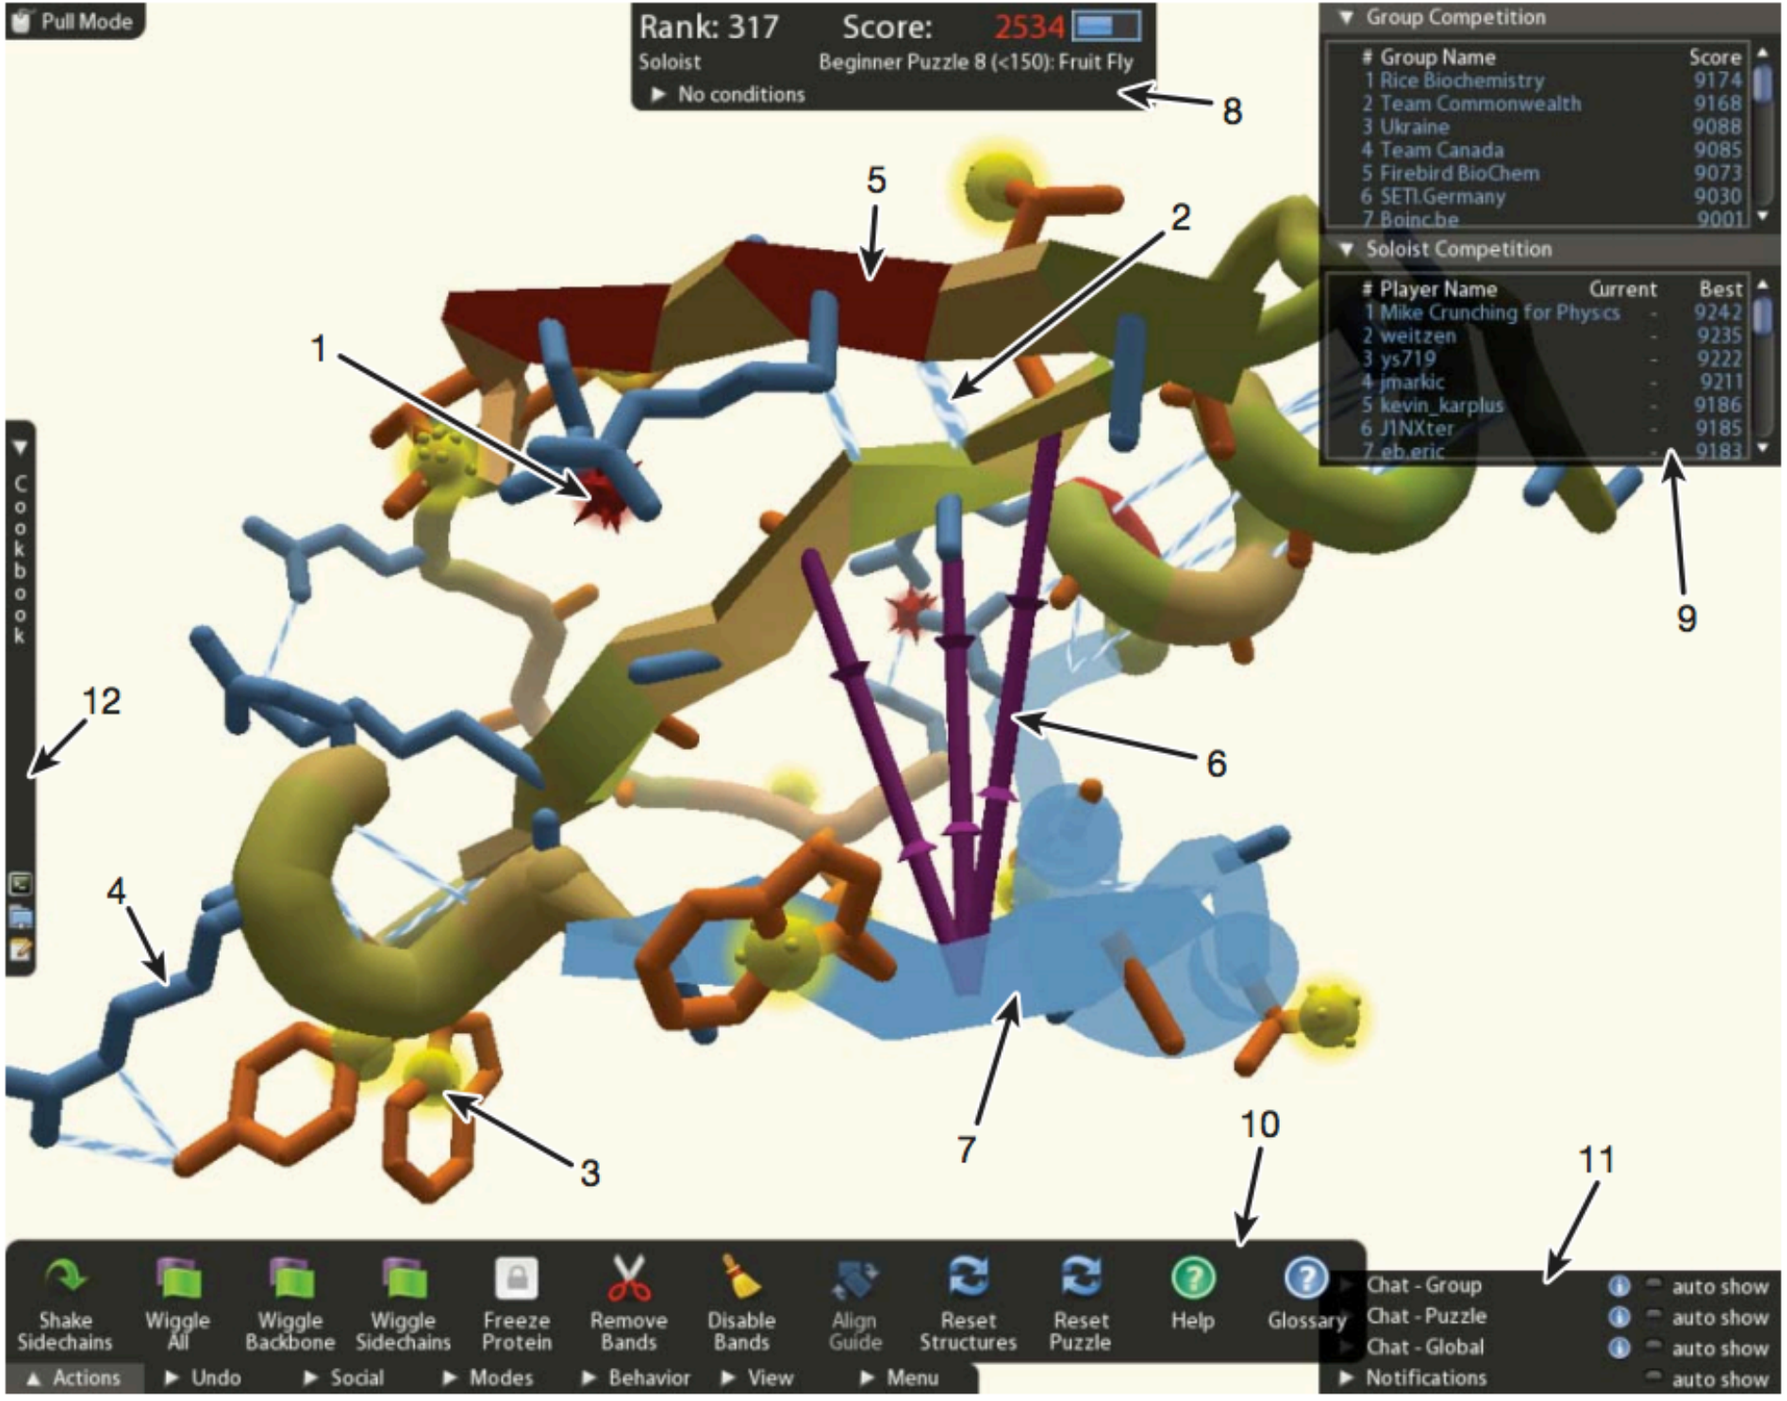

many more examples…

e.g. Foldit: fold protein structures as perfectly as possible highest scoring results are analyzed by researchers to eradicate diseases

#4 altruism

Jim Gray:

- went missing while on a boat trip
- satellite took **560,000 images** of the area
- mechanical turkers helped find his boat in the images

www.helpfindjim.com

#5 fear?

BlueServo™

About Us | Virtual BorderWatch℠ | Information | BorderWatch℠ Archives | Don

Virtual Stake Outs - Live Border Cameras

CLICK ON IMAGE OR LINK TO VIEW VIDEO

Camera 1

This is a known drug traffic area. If you see people walking along this trail carrying backpacks or packages please report this activity.

Camera 2

Look for individuals on foot carrying backpacks. If you see this activity please report it immediately.

homeland security:
watching the Texas-Mexico border

concerns & limitations

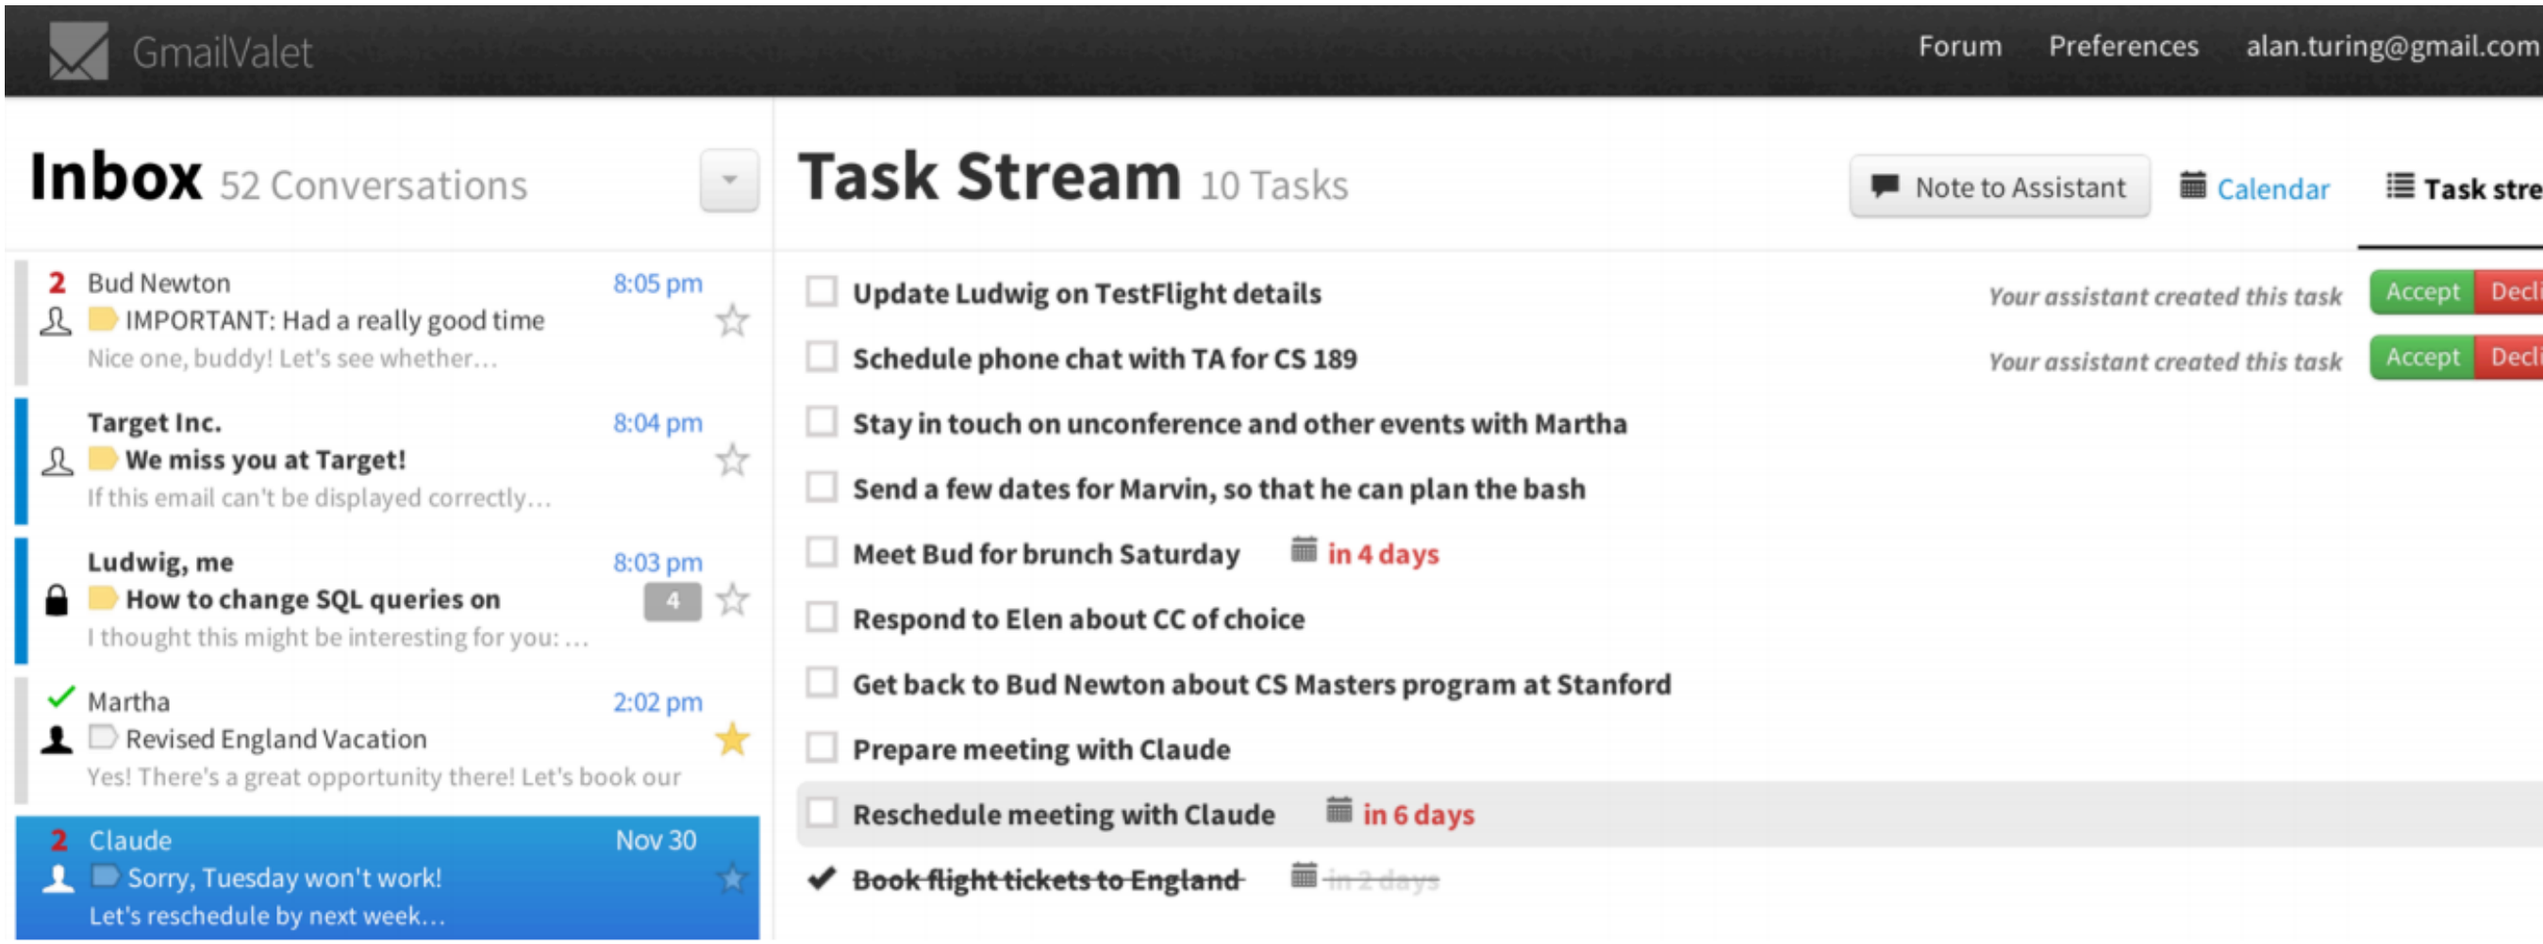

privacy, sensitive data:

2013 managing email overload

privacy, sensitive data:
2013 managing email overload

EmailValet: Managing Email Overload through Private, Accountable Crowdsourcing

Nicolas Kokkalis, Thomas Köhn, Carl Pfeiffer, Dima Chornyi, Michael S. Bernstein, Scott R. Klemmer
Stanford University HCI Group
Computer Science Department, 353 Serra Mall, Stanford, CA
{nicolas, thomasko, carlp87, chornyi, msb, srk}@cs.stanford.edu

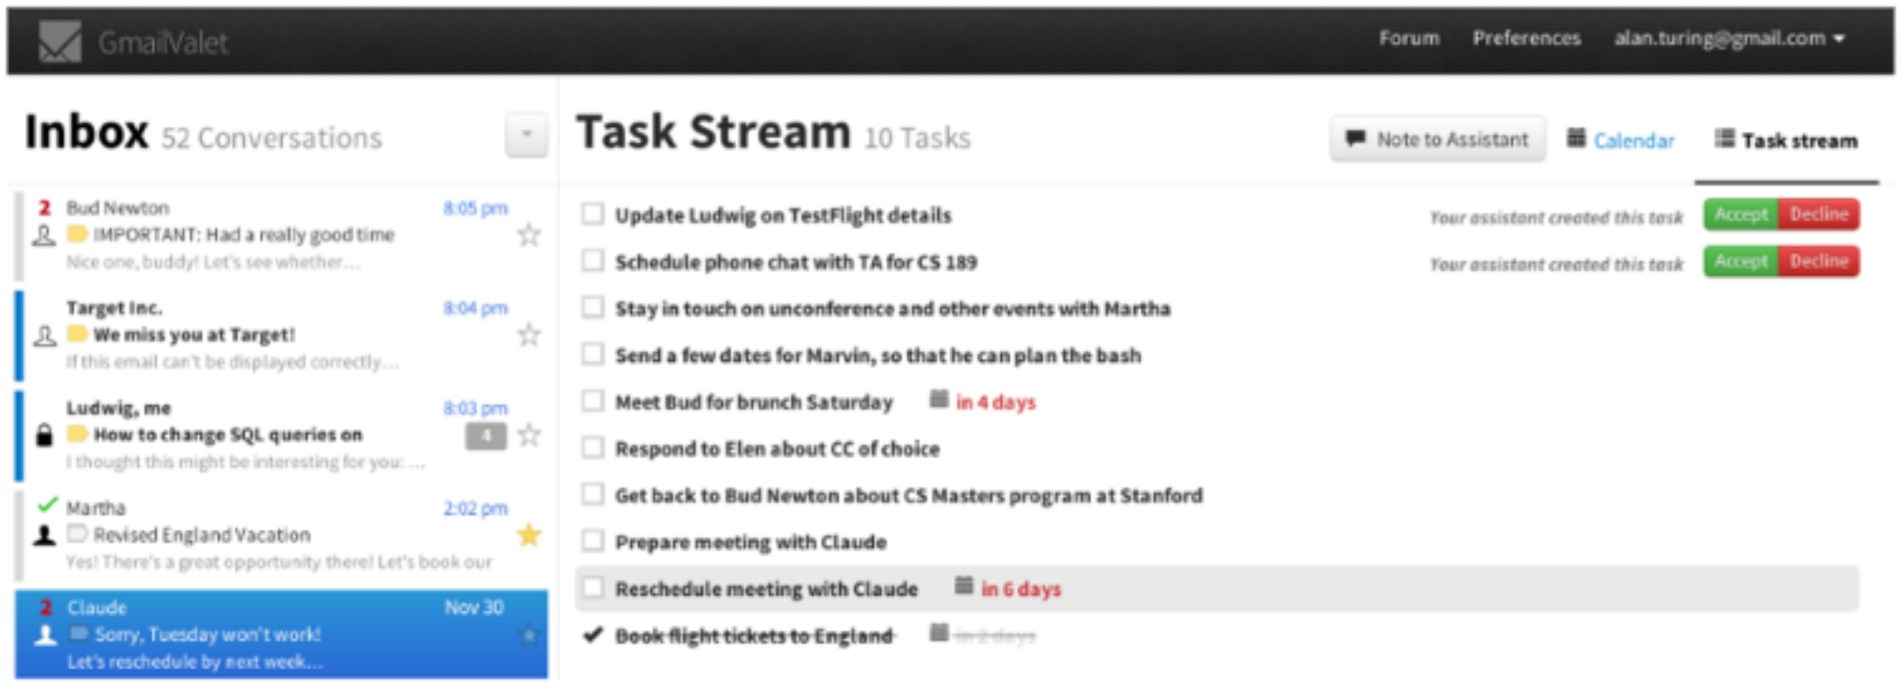

Figure 1. The EmailValet email client draws on crowdsourced expert assistants to transform a cluttered inbox into an organized task stream. Assistants are given limited, accountable access to the user's inbox so that they may extract tasks from each email.

ABSTRACT

This paper introduces privacy and accountability techniques for crowd-powered systems. We focus on email task management: tasks are an implicit part of every inbox, but the overwhelming volume of incoming email can bury important requests. We present *EmailValet*, an email client that recruits remote assistants from an expert crowdsourcing marketplace. By annotating each email with its implicit

Author Keywords

Crowdsourcing; Email Overload; Human Assistants; Task Management; Access Control.

ACM Classification Keywords

K.4.3 [Organizational Impacts]: Computer-supported collaborative work.

General Terms

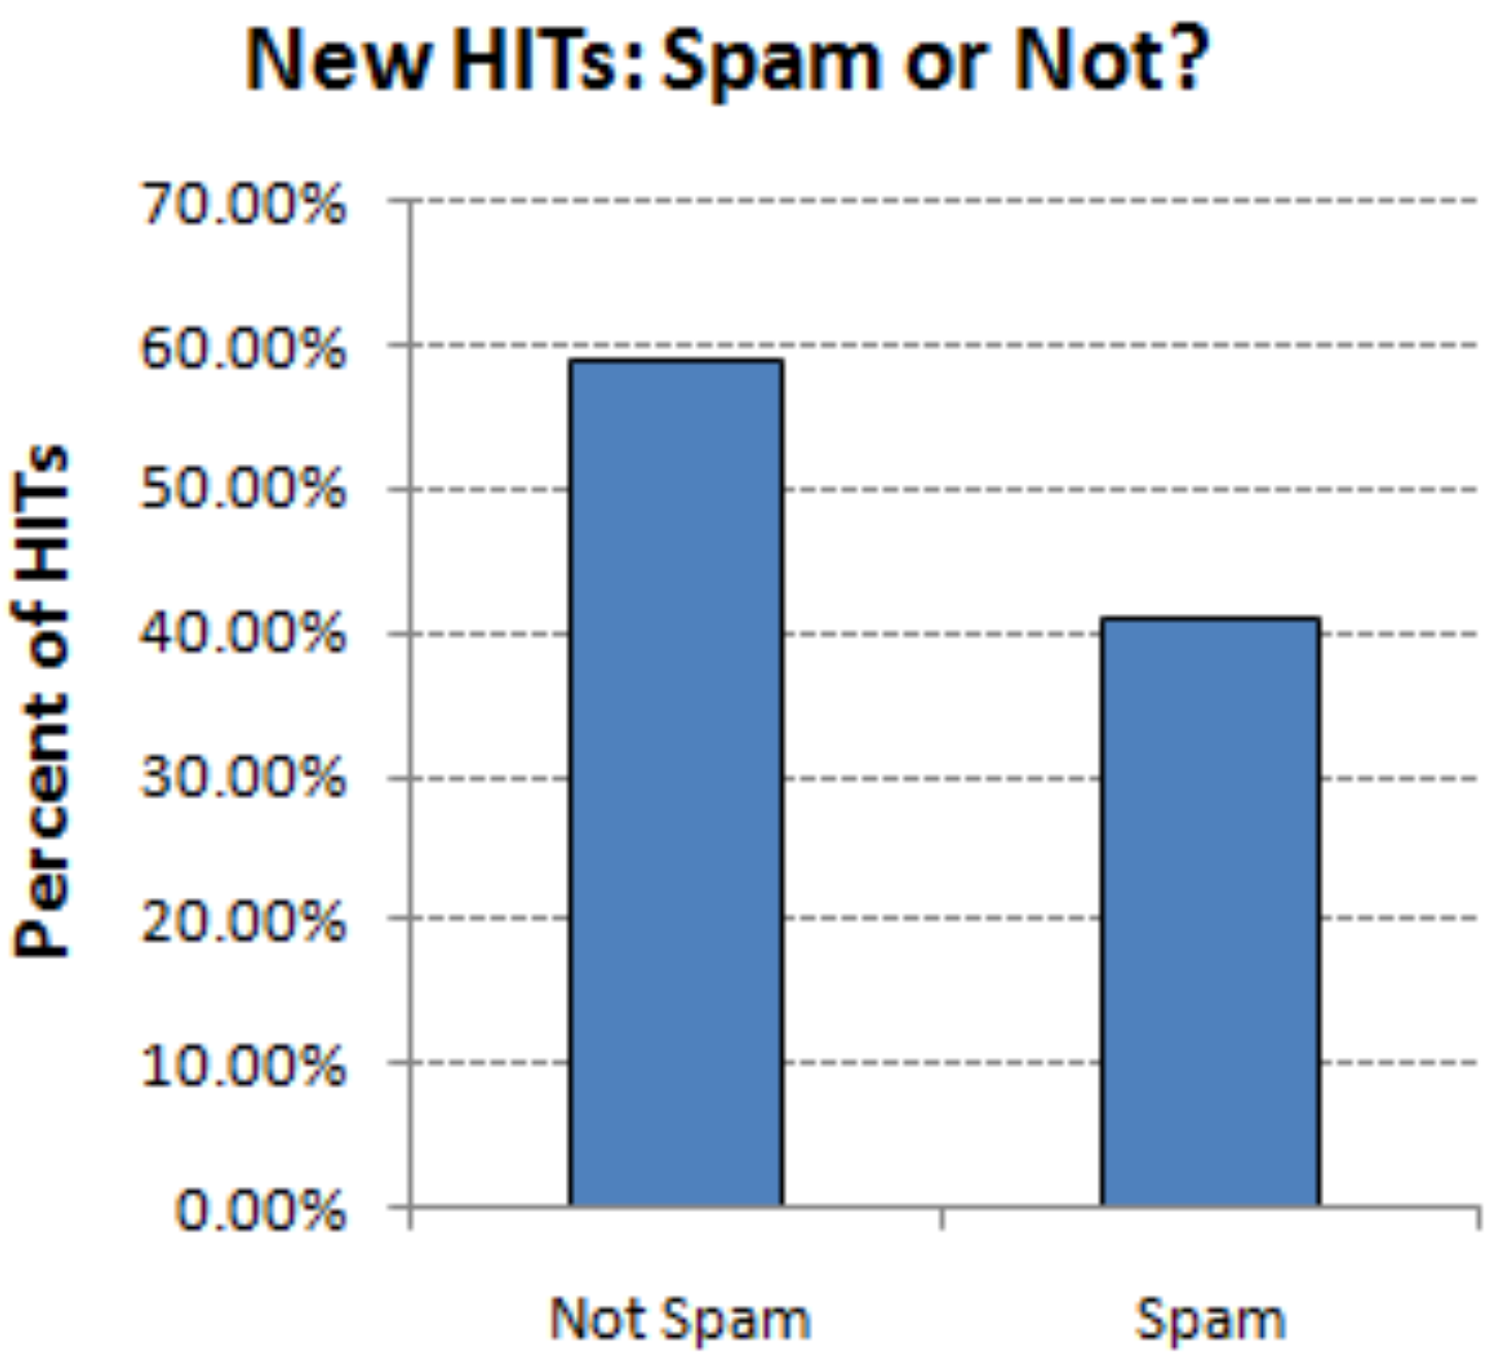

issues for task workers:
ca. 40.92% is spam (2010)

spam examples:
fake accounts and **fake clicks** for better SEO

“Create a Twitter account and follow me”

“Write a positive review on Yelp”

“Like my YouTube video”

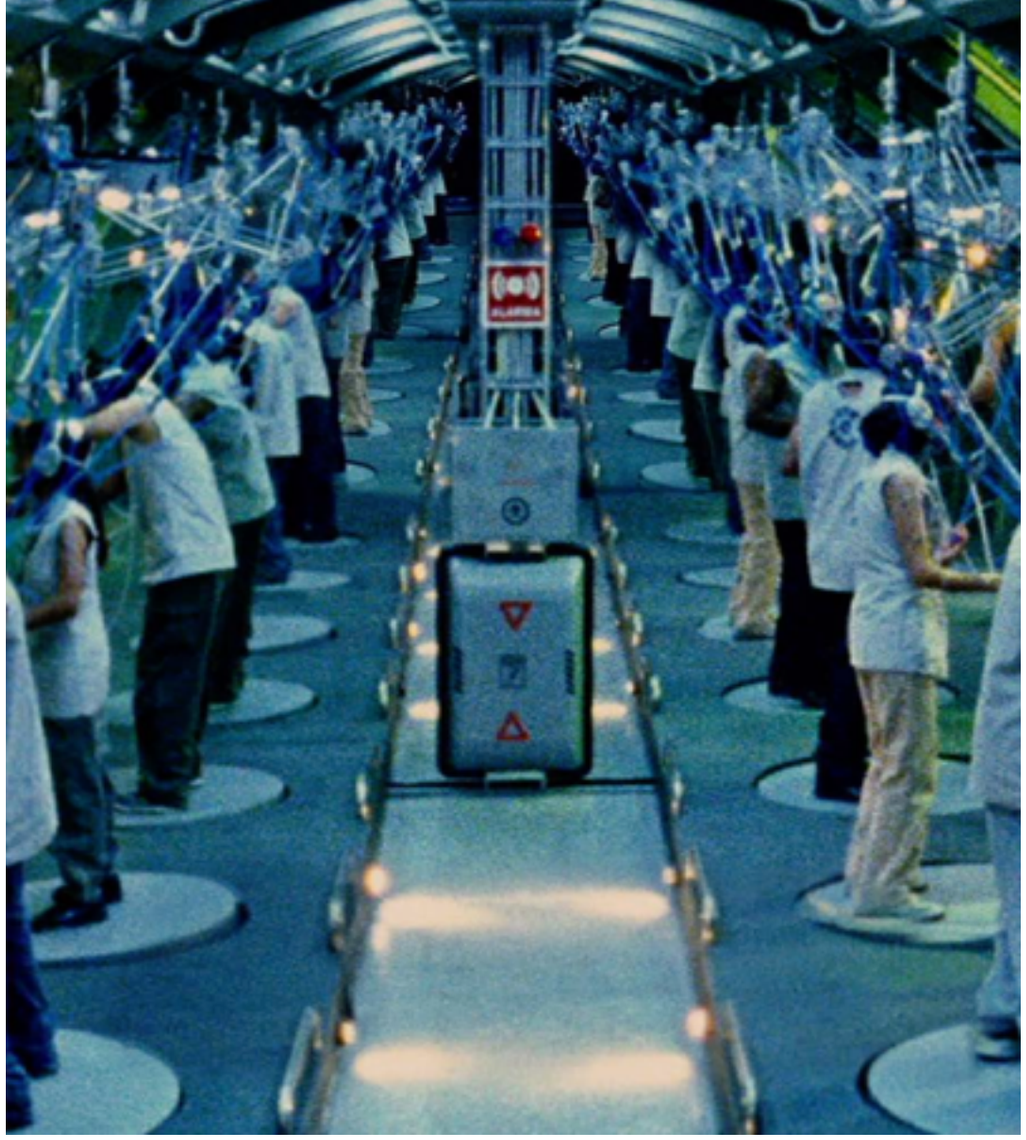

moral issues: digital sweatshops?

- unregulated sector – no labor laws
- 500,000 turkers, for 18% this is their main employment
- every mouse click is monitored

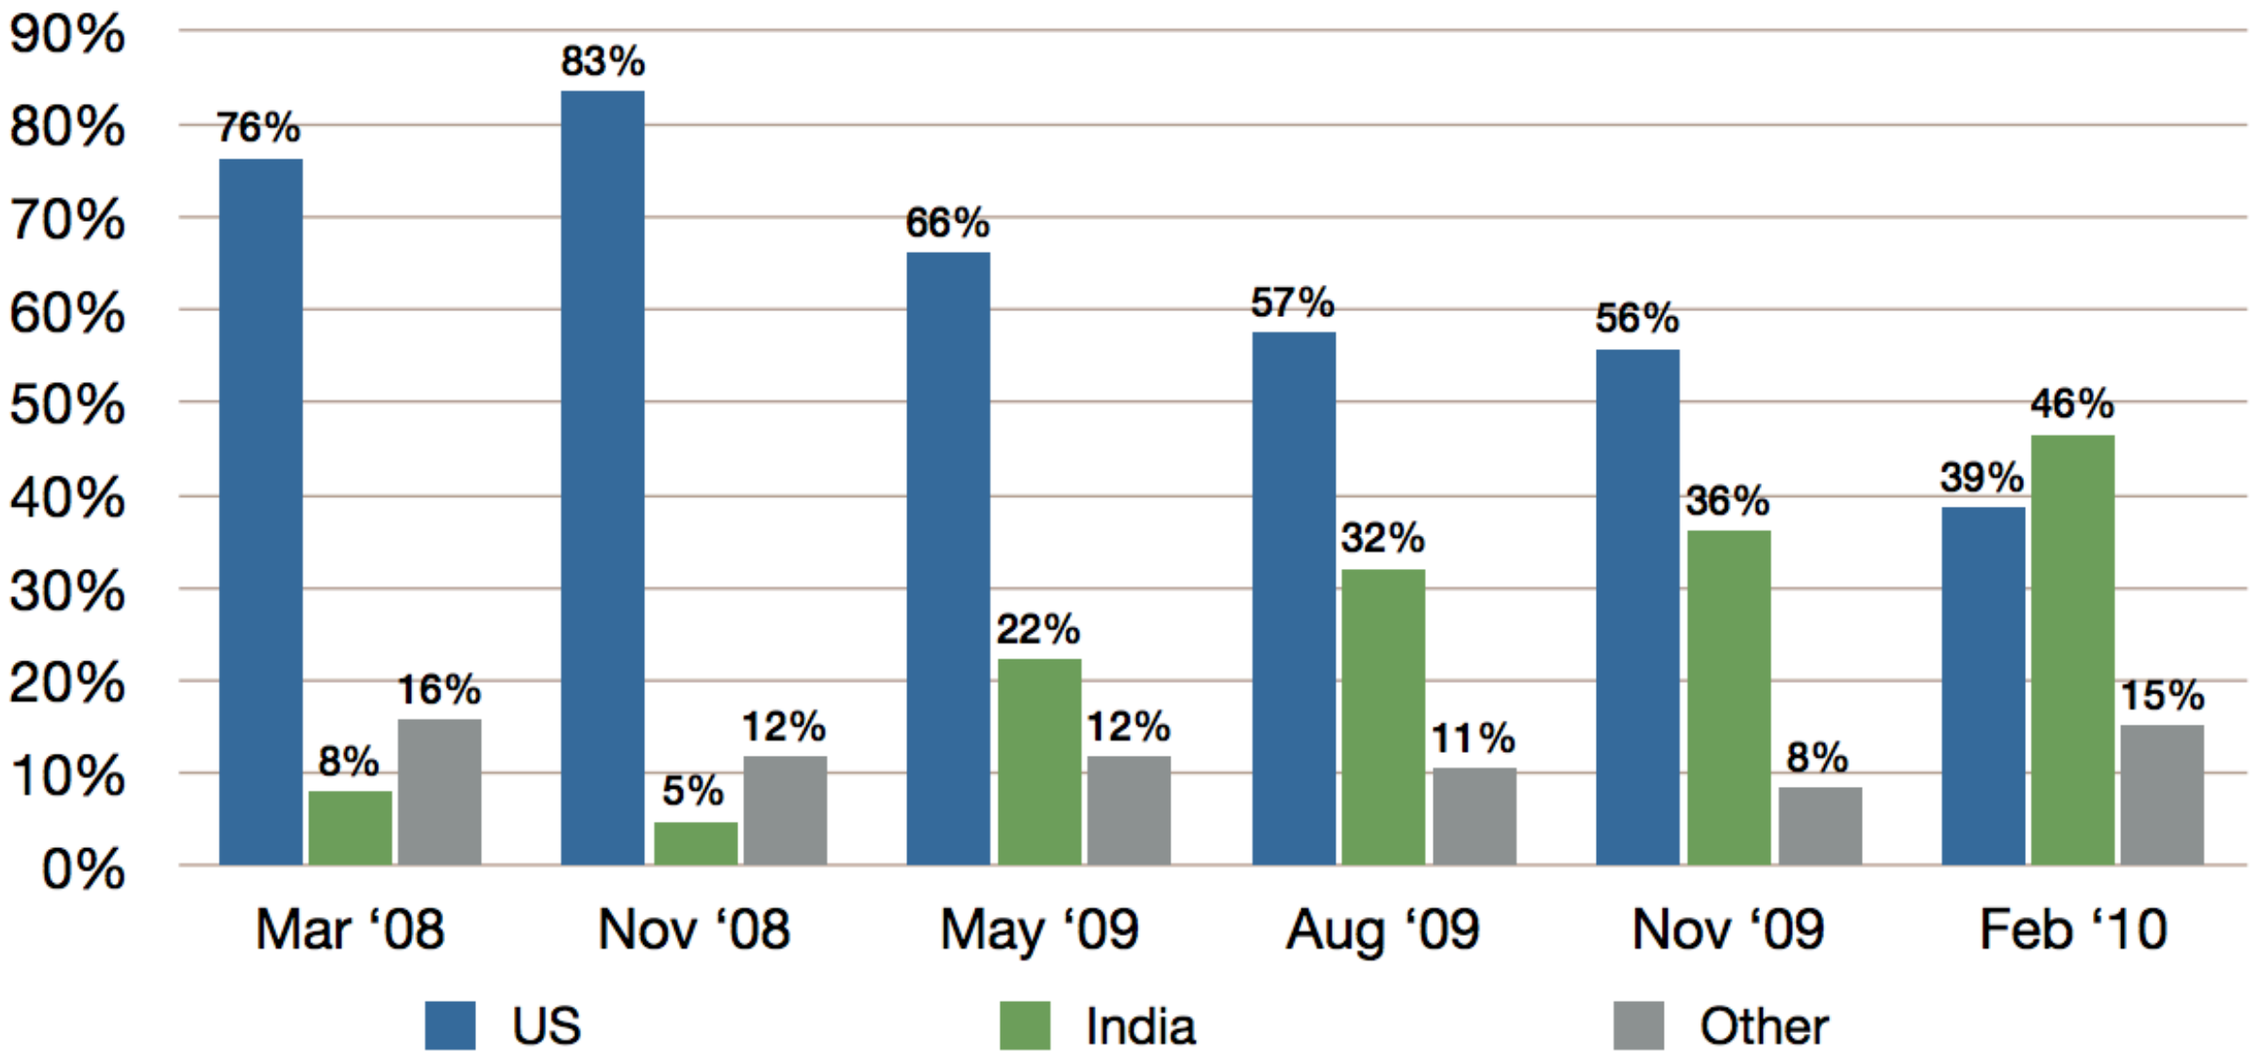

who are the turkers?

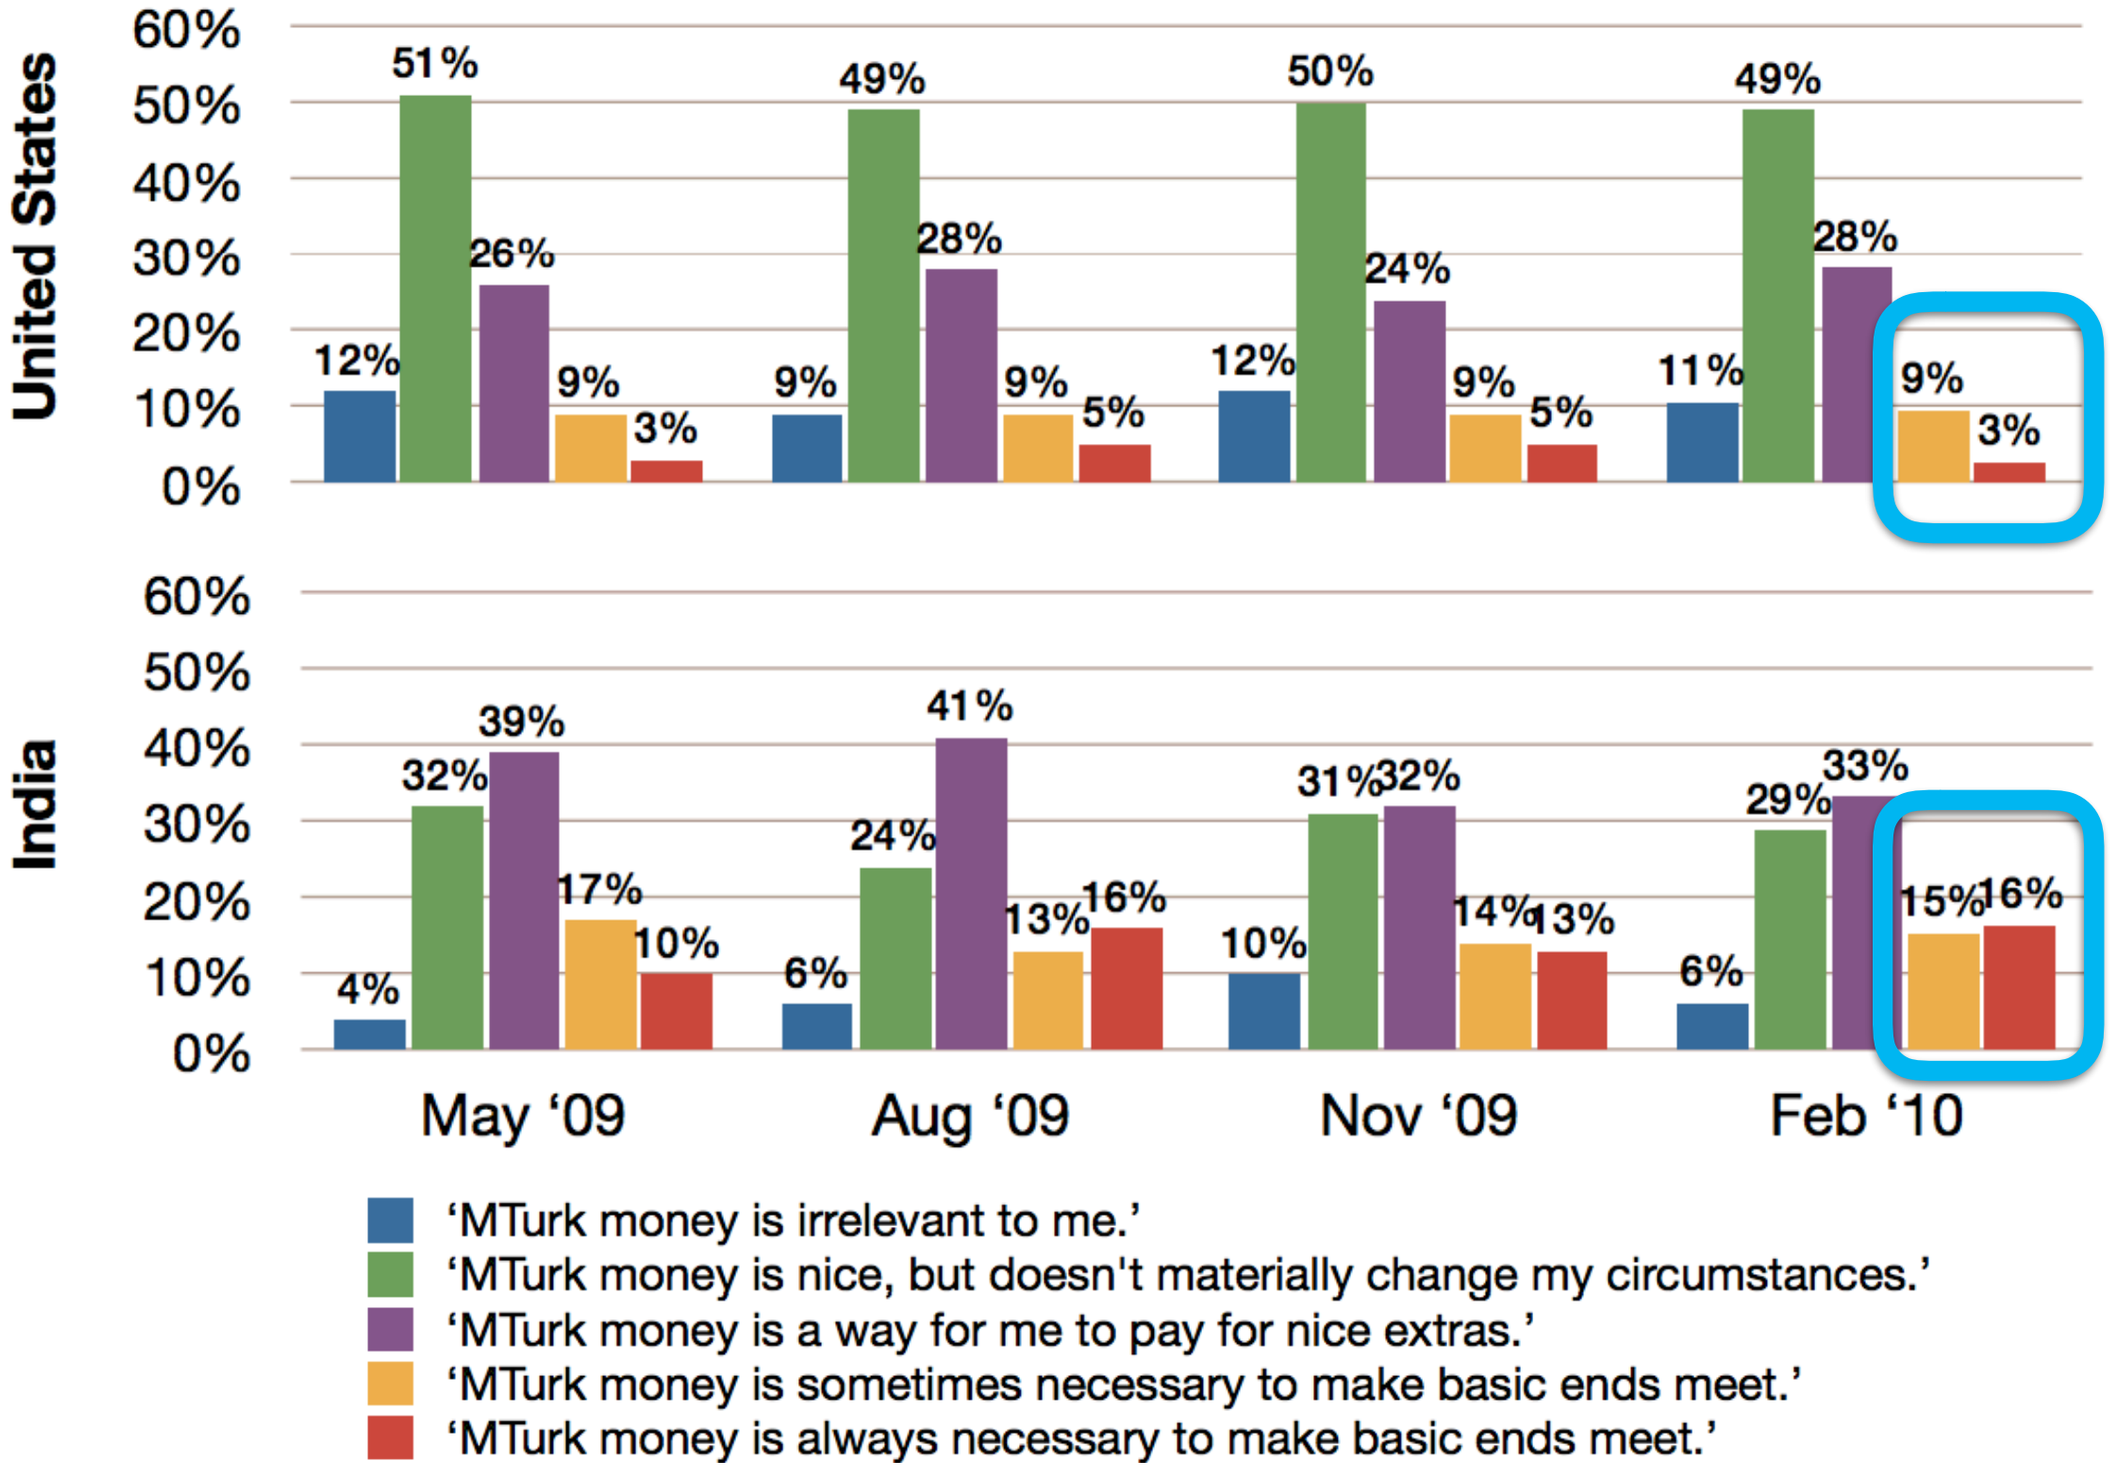

reliance on mechanical turk income
(average turker makes **below $2 per hour**)

Who are the Crowdworkers? Shifting Demographics in Mechanical Turk

Joel Ross
Department of Informatics
University of California, Irvine
Irvine, CA 92697-3440 USA
jwross@uci.edu

Andrew Zaldivar
Department of Informatics
University of California, Irvine
Irvine, CA 92697-3440 USA
azaldiva@uci.com

Lilly Irani
Department of Informatics
University of California, Irvine
Irvine, CA 92697-3440 USA
lirani@ics.uci.edu

Bill Tomlinson
Department of Informatics
University of California, Irvine
Irvine, CA 92697-3440 USA
wmt@uci.edu

M. Six Silberman
Department of Informatics
University of California, Irvine
Irvine, CA 92697-3440 USA
msilberm@uci.edu

Copyright is held by the author/owner(s).
CHI 2010, April 10–15, 2010, Atlanta, Georgia, USA.
ACM 978-1-60558-930-5/10/04.

Abstract
Amazon Mechanical Turk (MTurk) is a crowdsourcing system in which tasks are distributed to a population of thousands of anonymous workers for completion. This system is increasingly popular with researchers and developers. Here we extend previous studies of the demographics and usage behaviors of MTurk workers. We describe how the worker population has changed over time, shifting from a primarily moderate-income, U.S.-based workforce towards an increasingly international group with a significant population of young, well-educated Indian workers. This change in population points to how workers may treat Turking as a full-time job, which they rely on to make ends meet.

Keywords
Mechanical Turk, demographics, user surveys, crowdsourcing, human computation

ACM Classification Keywords
H5.3. Group and Organization Interfaces: Computer-supported cooperative work; H5.3 Group and Organization Interfaces: Web-based interaction.

General Terms
Human factors.

conclusions

human computation::

- using human processing power to **solve problems that computers cannot yet solve**
- humans are directed by a computer
- enabled by the platform Mechanical Turk
  - a paid crowdsourcing platform for microtasks

What's the craziest thing your cat has ever done?

Requester:	Bjoern Hartmann	**Assignments Pending Review:**	0	
HIT Expiration Date:	Jul 27 2008, 07:33 AM PDT	**Reviewed Assignments:**	30	Dow
Reward:	$0.08	**Remaining Assignments:**	0	
Assignments Requested:	30	**Remaining Time:**	Expired	

Sketch a cat

Requester:	Bjoern Hartmann	**Assignments Pending Review:**	0	
HIT Expiration Date:	Aug 14 2008, 08:24 AM PDT	**Reviewed Assignments:**	50	Dov
Reward:	$0.05	**Remaining Assignments:**	0	
Assignments Requested:	50	**Remaining Time:**	Expired	

My cat has only one eye and he can open anything. He is a master at undoing the lid to his cat food container. Anything that screws on or latches he can get open. I have tried every kind of container and within 1 day he has it figured out. He can open any door. He loves to open the refrigerator and let the dog eat all she wants. He has even opened the front door for a visitor and they were shocked when the door opened and he was the only one sitting there. He can open the sliding door and any cabinet he feels the need to open.

create your own crowd-sourced project some time!

end.