

Crowdsourcing

September 24, 2018

Fall 2018

COMP 7920

1

Today

Introduction to crowdsourcing

What is it?

Platforms

Applications

Challenges and Issues

3rd Case Study from Wed (if time)

Fall 2018

COMP 7920

2

Crowdsourcing

.. the act of a company or institution taking a function once performed by employees and outsourcing it to an undefined (and generally large) network of people in the form of an open call.

- Howe, 2006

Fall 2018

COMP 7920

3

Crowdsourcing

the process of obtaining needed services, ideas, or content by soliciting contributions from a large group of people, and especially from an online community, rather than from traditional employees or suppliers.

-- Merriam Webster

Fall 2018

COMP 7920

4

Crowdsourcing Platforms

Aim to encourage and/or support crowdsourced work

Differ according to their:

- Level of organization
- Size of community
- Incentives



Example Platform : Amazon Mechanical Turk

Name origin: 18th century “automatic” Chess-playing machine

Requestors (aka employers) post tasks known as **Human-Intelligence Tasks (HITs)**

Workers (aka online users) complete tasks for small financial reward

The screenshot shows the Amazon Mechanical Turk interface. At the top, it says 'amazonmechanicalturk' and 'Your Account'. Below that, there are tabs for 'HITs' and 'Qualifications'. A search bar is visible with the text 'Search for HITs' and a filter 'that pay at least \$ 0.00'. The main content area is titled 'All HITs' and shows '1-10 of 1895 Results'. The results are sorted by 'HITs Available (most first)'. There are five task listings visible, each with a requester name, HIT Expiration Date, Reward, and Time Allotted. The tasks are: 'Image Tagging', 'Find Restaurant Web Addresses', 'Product Search Relevance', 'Verify Restaurant Websites', and 'Find Business Web Addresses'.

Mechanical Turk

2010 Data:

~ five million tasks completed per year at 1-5¢ each [Ipeirotis 2010]

Population: 40% U.S., 40% India, 20% elsewhere

Gender, education and income are close mirrors of overall population distributions [Ross 2010]

Mechanical Turk

More recent data:

~ 600k people participating

2018 study on wages [Hara 2018]

Median: ~\$2/h

Mean: \$3/h

The following website has live demographics data:

<http://demographics.mturk-tracker.com/>

Fall 2018

COMP 7920

9

Harnessing Crowds

Example application domains

Crowd work

Crowd-powered interfaces

Online experiments

User-generated content

Citizen science

Crowdfunding

Fall 2018

COMP 7920

10

Crowd work

Using crowd workers to complete work that would be difficult to complete otherwise

Why?

Volume of work

Tedious tasks

Expertise / abilities involved

Need for outsider insight

Can be explicit (e.g., through crowdsourcing platforms) or implicit

Fall 2018

COMP 7920

11

Crowd work: Examples

One of the earliest examples is Luis Von Ahn's ReCaptcha




Fall 2018

COMP 7920

12

Crowd-powered interfaces: Examples

WizViz [Bigham et al.]



Fall 2018
COMP 7920
13



Crowd-powered interfaces: Soylent [Berstein et al.]

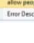
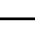
A crowd-powered word processor

Text Shortening

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, usernames, 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 differences in structure aren't relevant to a specific 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 editing. Clustering and selection generalization would also be improved by recognizing common text structure like URLs, usernames, email addresses, dates, times, etc.


While GUIs  computers more intuitive and easier to learn, they didn't let  be able to control computers efficiently.

While GUIs  computers more intuitive and easier to learn, they didn't  allow people to control computers efficiently.


Error Descriptions

Fall 2018
COMP 7920
14

Crowdsourced Experiments



https://www.experiencedynamics.com/sites/default/files/blog_images/ED-Website-Shoot_030414_40.jpg



Fall 2018
COMP 7920
15

Crowdsourced Experiments


Why?

Potential to:

- Increase sample sizes
- Increase participant diversity

Fall 2018
COMP 7920
16

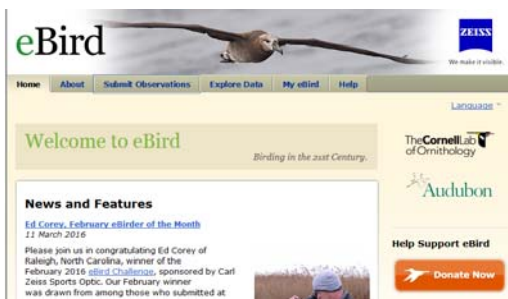
User-Generated Content




Citizen Science

Science conducted in whole, or in part, by non-professional scientists

Citizen Science



Citizen Science



Crowdfunding

The practice of funding a project or venture by raising money from a large number of people, often via web-based platforms

Fall 2018

COMP 7920

21

Crowdfunding: High-Profile Examples



Fall 2018

COMP 7920

22

Crowdsourcing

Challenges / Issues

Motivating contributions

Data / work quality

Workflows and coordination

Performance

Fall 2018

COMP 7920

23

Motivating Contributions

How to get people to participate?

Some strategies

Payment

Gamification

Community building

Ensuring personal benefit

Embedding crowd work into another task

Fall 2018

COMP 7920

24

Data Quality

Data and/or work coming from crowds has the potential to be “noisy”

Components / solutions:

- Ensure clear task descriptions
- Establish checks
- Include redundancy
- Incorporate voting
- Have intrinsic motivations align with task goals

Fall 2018

COMP 7920

25

Workflows

How can crowd work benefit from Divide & Conquer?

Considerations:

- Task decomposition
- Sequencing
- Ensuring task coverage
- Providing sufficient context
- Worker coordination?

Fall 2018

COMP 7920

26

Performance

Performance requirements will vary dramatically base on application

Solutions explored

- Over-recruiting
- Having workers on hand
- Different pricing models

Fall 2018

COMP 7920

27

Other Considerations

What are the ethical and/or legal considerations of crowdsourcing?

E.g.,

- Demographics of crowd workers
- Imbalance of benefit
- Data storage

We will cover these in more detail later...

Fall 2018

COMP 7920

28

Relevance to HCI

Crowdsourcing HCI work

E.g., interface evaluation, user studies

HCI of crowdsourcing

Interface/interaction design

Lots of interesting socio-technical questions

Fall 2018

COMP 7920

29

Intro to crowdsourcing: Recap

Now you are:

Familiar with the term crowdsourcing and example platforms

Are aware of a number of HCI-relevant application domains

Are familiar with some of the challenges and issues

Fall 2018

COMP 7920

30

Case Study 3: Switter

General research methodology

Exploratory study

Prototype design / implementation
addressing a subset of challenges / issues
raised in study

Limited field deployment

Fall 2018

COMP 7920

31

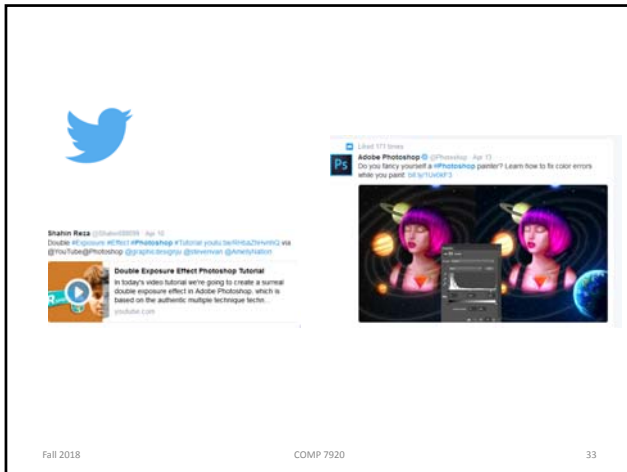
expert designers continually seek to learn
new things

even after 10 years of experience

Fall 2018

COMP 7920

32



Challenges

Large volume of content

Assessing utility is difficult

Fall 2018 COMP 7920 34

Switter

Supporting Exploration of Software Learning Materials on Social Media

Fall 2018 COMP 7920 35

Field Study

Deployed Switter to 9 design practitioners

Used at least once per day over 7 days

"Wizard of Oz" content population (~30 tweets/day)

Data collection

- Logs, journal entries, semi-structured interviews

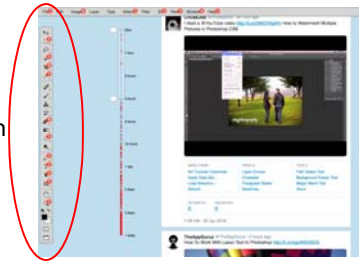
Fall 2018 COMP 7920 36

Findings

Exploration based on popularity

Addressing known weaknesses

Command comparisons



Fall 2018

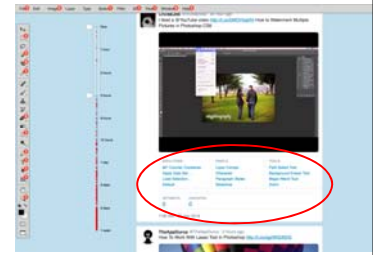
COMP 7920

37

Findings

Quickly assess the content

unexpected/
unknown
use-cases



Fall 2018

COMP 7920

38

Discussion

Wizard-of-oz components

Missing features

Choice of evaluation method

Fall 2018

COMP 7920

39