

CS 5306  
INFO 5306:  
Crowdsourcing and  
Human Computation

Lecture 2  
Haym Hirsh



CLASSIFY

STORY

SCIENCE



DISCUSS

PROFILE

LANGUAGE



## Classify



SDSS



Favourite



Invert

Examples

Restart

*Note: Please always classify the galaxy in the centre of the image.*


### SPIRAL


How many spiral arms are there?

 1	 2	 3	 4
 More than 4	 Can't tell		


Killer whales are beautiful animals . I remember seeing  
these huge , smooth , black and white creatures jumping  
high into the air at Sea World , as a kid .

Greg Little, Lydia B. Chilton, Max Goldman, and Robert C. Miller.  
"Exploring iterative and parallel human computation processes."  
*In Proceedings of the ACM SIGKDD Workshop on Human Computation, 2010*






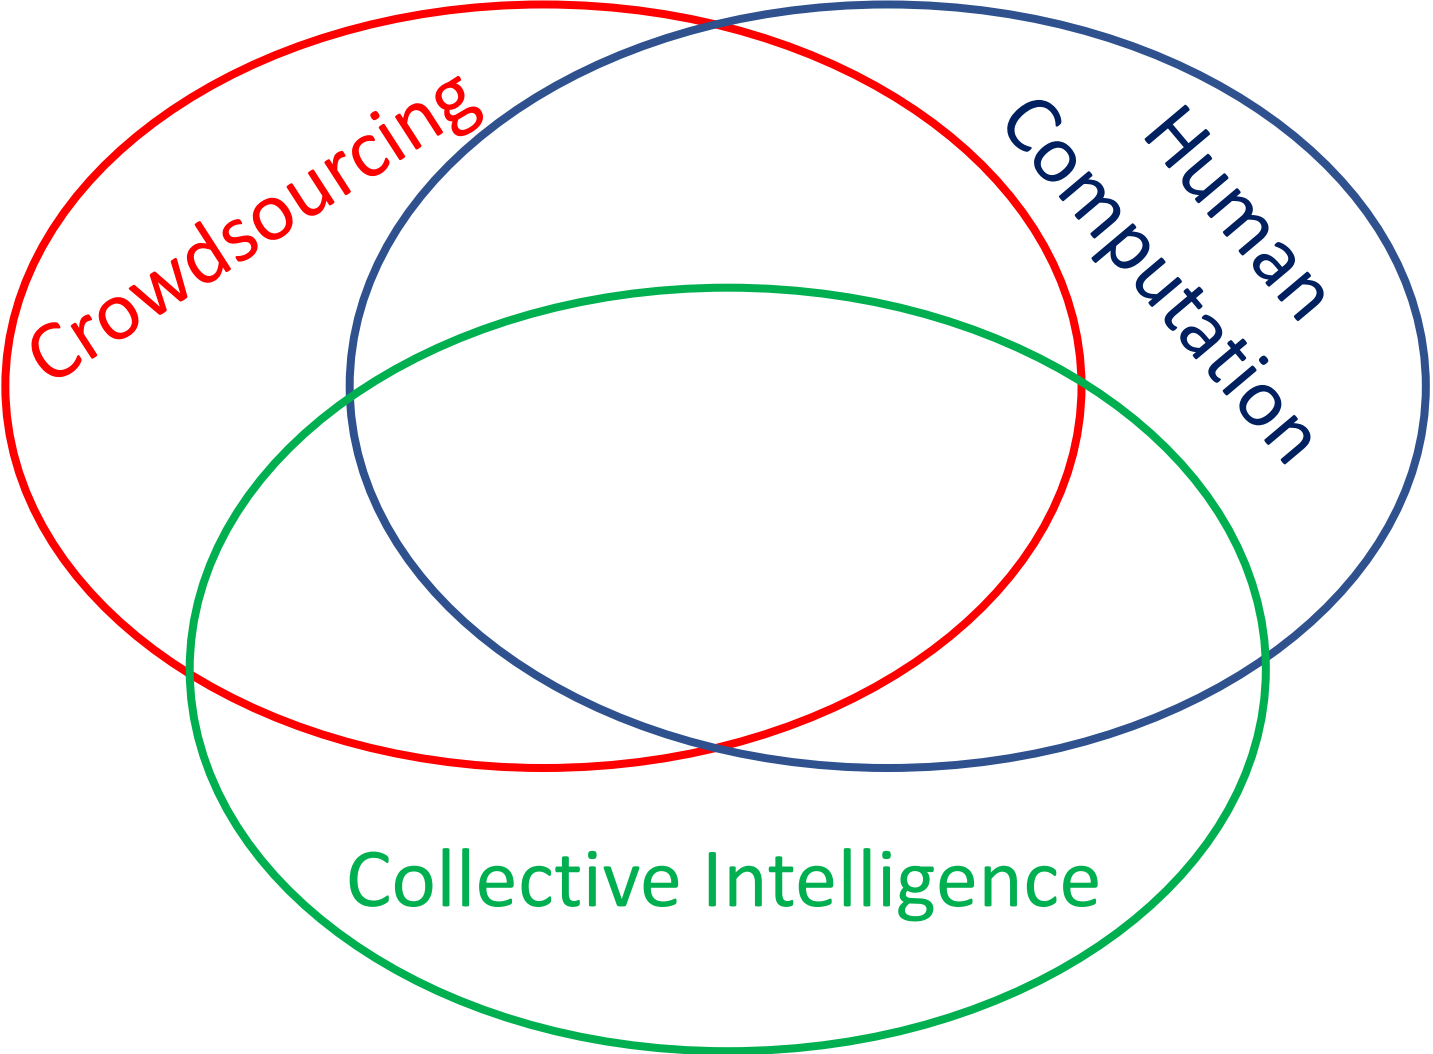
[mcaf.ee/vlqsga](https://mcaf.ee/vlqsga)




[tinyurl.com/y8fut9tn](https://tinyurl.com/y8fut9tn)









Social  
Computing

# Crowdsourcing

“Crowdsourcing is the act of taking a job traditionally performed by a designated agent (usually an employee) and outsourcing it to an undefined, generally large group of people in the form of an open call.”

“The Rise of Crowdsourcing,” Jeff Howe, *Wired*, Jun 2006

From: torvalds@klaava.Helsinki.FI (Linus Benedict Torvalds)  
Newsgroups: comp.os.minix  
Subject: What would you like to see most in minix?  
Summary: small poll for my new operating system  
Message-ID: <1991Aug25.205708.9541@klaava.Helsinki.FI>  
Date: 25 Aug 91 20:57:08 GMT  
Organization: University of Helsinki

Hello everybody out there using minix -

I'm doing a (free) operating system (just a hobby, won't be big and professional like gnu) for 386(486) AT clones. This has been brewing since april, and is starting to get ready. I'd like any feedback on things people like/dislike in minix, as my OS resembles it somewhat (same physical layout of the file-system (due to practical reasons) among other things).

I've currently ported bash(1.08) and gcc(1.40), and things seem to work. This implies that I'll get something practical within a few months, and I'd like to know what features most people would want. Any suggestions are welcome, but I won't promise I'll implement them

Linus (torvalds@kruuna.helsinki.fi)

PS. Yes - it's free of any minix code, and it has a multi-threaded fs. It is NOT protable (uses 386 task switching etc), and it probably never will support anything other than AT-harddisks, as that's all I have :-).

# Crowdsourcing

“We say that a system is a [crowdsourcing] system if it enlists a crowd of humans to help solve a problem defined by the system owners”

“Crowdsourcing systems on the World-Wide Web,” Anhai Doan, Raghu Ramakrishnan, Alon Y. Halevy. *Communications of the ACM*, Vol. 54 No. 4, Pages 86-96, April 2011

# Human Computation

“a paradigm for utilizing human processing power to solve problems that computers cannot yet solve.”

*Human Computation*, Luis van Ahn. Doctoral Dissertation, Department of Computer Science, Carnegie Mellon University (2005).

# Human Computation

- Human computation:
  - The problems fit the general paradigm of computation, and as such might someday be solvable by computers.
  - The human participation is directed by the computational system or process

“Human Computation: A Survey and Taxonomy of a Growing Field”

Quinn and Bederson CHI 2011

# Human Computation

Computer algorithms that call on human effort  
in order to solve the problems they're given



# Human Computation

Computer algorithms that call on human effort  
in order to solve the problems they're given

Programs that call people as subroutines

# Collective Intelligence


- “Collective intelligence is a shared or group intelligence that emerges from the collaboration and competition of many individuals and appears in consensus decision making in bacteria, animals, humans and computer networks”

Wikipedia, “Collective Intelligence”

# Complementary

- Crowdsourcing: Focuses on bringing people together
- Human computation: Focuses on the algorithms that utilize people
- Collective intelligence: Focuses on the emergent behavior

Behavioral  
Science



Computer  
Science

# Can Still Be a Computer Scientist...

... even though you're building computer programs whose functionality depends on the work of people

- Reliability
- Efficiency
- ....

# Common Application Domain: Labeling Data

- “Get Another Label? Improving Data Quality and Data Mining Using Multiple, Noisy Labelers”, Victor S. Sheng, Foster Provost, Panagiotis G. Ipeirotis, *KDD 2008*
- “Cheap and Fast — But is it Good? Evaluating Non-Expert Annotations for Natural Language Tasks”, Rion Snow, Brendan O’Connor, Daniel Jurafsky, and Andrew Y. Ng, *EMNLP 2008*
- “Utility Data Annotation with Amazon Mechanical Turk”, Alexander Sorokin and David Forsyth, *IEEE Workshop on Internet Vision, CVPR 2008*

# Common Application Domain: Labeling Data

Major workhorse: Majority vote

Give it to  $n$  workers

Take the majority vote

# OpenSurfaces

Sean Bell, Paul Upchurch, Noah Snavely, Kavita Bala  
Cornell University

New! Our SIGGRAPH 2014 paper is now online.

[Intrinsic Images in the Wild](#)

New! The code for our full crowdsourcing pipeline is online.

[Code \(Github repository\)](#)

[Documentation](#)

## Statistics

	Good	All
Labeled Scenes	22,214	91,912
Whitebalanced Photos	18,199	25,262
Segmentations	105,082	150,316
Planar Segmentations	52,395	105,082
Named Materials	87,091	104,505
Named Objects	71,460	86,816
Rectified Textures	27,823	35,192
Reflectances	51,719	63,175
Reflectance Judgements	N/A	3,113,257
MTurk Assignments	382,642	385,362
MTurk Users	3,174	3,555
Hours of Work	N/A	13,478
Submitted Items	N/A	12,426,283

## Materials



## Reflectances



## Textures




OpenSurfaces is a large database of annotated surfaces created from real-world consumer photographs. Our annotation framework draws on crowdsourcing to segment surfaces from photos, and then annotate them with rich surface properties, including material, texture and contextual information.



# Computer Science Question

Given an unlimited supply of unlabeled data, and you want to label each item with probability of error less than  $\epsilon$ , how many items can you label within a given budget  $B$ ?

“Get Another Label? Improving Data Quality and Data Mining Using Multiple, Noisy Labelers”  
Victor S. Sheng, Foster Provost, Panagiotis G. Ipeirotis



# Can Still Be a Computer Scientist...

... even though you're building computer programs whose functionality depends on the work of people

- Reliability
- Efficiency
- ....

Can Still Be a Computer Scientist...


... and think about people

# Financial Incentives and the “Performance of Crowds”

Winter Mason and Duncan J. Watts  
*2009 KDD Human Computation Workshop*


**Instructions**

At the beginning of a task, you will be presented with a list of images taken from traffic cameras. An example list is shown below.



1. 2. 3.

Your goal is to reorder the list chronologically from left to right and top to bottom. The sorted list is shown below.



1. 2. 3.

Notice that in the sorted row, the truck on the right moves away from the camera, and the blue cab on the left approaches the camera. To correctly sort the photos, you need to determine the flow of the traffic.

To reorder a list, click and drag a photo to the position it belongs. The other photos will move accordingly. Once you believe the list is in the correct order, click on the "Submit" button at the top of the page.

If you do not want to complete any more tasks, click on the "Finished" button at the bottom of the page.

*(This button will not be available in the next 3 practice examples)*

Click here to practice: [Practice](#)

Images provided by Trafficland.com

To select a word, first click on the first letter of the word, then click on the last letter of the word. If you are correct, it will turn red and the word will appear to the right of the puzzle.

For each puzzle you will see a set of *possible* words and their category. **Not all of the words listed are in the puzzle!** In addition, the number of words in each puzzle changes. The list of *possible* words follows:  
ACHIEVE, ATTAIN, BUILDING, CHAIR, COMPETE, GREEN, LAMP, MASTER, MUSIC, PLANT, STAPLE, STEREO, STRIVE, SUCCEED, TURTLE

For this practice puzzle, you will have to find at least 8 words to continue.

### RANDOM WORDS



SUCCEED
BUILDING
ACHIEVE

Submit Puzzle

# Financial Incentives and the “Performance of Crowds”

Winter Mason and Duncan J. Watts

*KDD Hcomp Workshop 2009*

- More pay for a task
  - Means you get more workers
  - Does *not* mean you get more accurate answers
- Pay rate sets the human computation algorithm’s “clock rate” – how fast you want things done

# Financial Incentives and the “Performance of Crowds”

Winter Mason and Duncan J. Watts

*KDD Hcomp Workshop 2009*


- More pay for a task
  - Means you get more workers
  - Does *not* mean you get more accurate answers
- Pay rate sets the human computation algorithm’s “clock rate” – how fast you want things done
- If you pay on a quota system, performance improves



# Honesty in an Online Labor Market

Siddharth Suri, Daniel G. Goldstein, and Winter A. Mason


In *Proceedings of the Third Human Computation Workshop, 2011*



# Honesty in an Online Labor Market

Siddharth Suri, Daniel G. Goldstein, and Winter A. Mason

*In Proceedings of the Third Human Computation Workshop, 2011*



# Class Mechanics

- Instructor: Prof. Haym Hirsh, [cs.cornell.edu/~hirsh](http://cs.cornell.edu/~hirsh)
- TA: Eric Wang, [ericwang0701@gmail.com](mailto:ericwang0701@gmail.com), [ericwang.info](http://ericwang.info)
- Course website: [cs.cornell.edu/courses/cs5306](http://cs.cornell.edu/courses/cs5306)
- Discussion: [piazza.com/cornell/fall2017/cs5306](http://piazza.com/cornell/fall2017/cs5306)
- Coursework submission: [gradescope.com](http://gradescope.com) (entry code 9ZRR3V)
- Exams:
  - Prelim (tentative): 10/17/2017, in class
  - Final: 12/12/2017, 9:00AM - 11:30AM

# Required Background

“knowledge of basic computer science principles and skills (such as CS 1110, CS 1114, CS 2110, CS 3110, or equivalent)”

You should feel comfortable doing some programming:

- Project 1: Data analysis
- Project 2: Human computation system or analysis

# Coursework

- Prelim: 15-20%
- Final: 15-20%
- Assignments (~5): 10-20%
- Projects (2): 45-55%
- Extra credit: Used if you are near the boundary between grades
  
- READINGS!

# Readings

- *Infotopia: How Many Minds Produce Knowledge*, Cass Sunstein
  - <http://library.books24x7.com.proxy.library.cornell.edu/toc.aspx?site=KD708&bookid=42073>
- Technical papers 2-4 per lecture

# Assignments

- Readings for Tuesday:

- “Judgment under Uncertainty: Heuristics and Biases”, Amos Tversky and Daniel Kahneman, *Science* 27 Sep 1974

<https://www.csc2.ncsu.edu/faculty/mpsingh/local/Social/f16/wrap/readings/Tversky+Kahneman-heuristics-biases-1974.pdf>

- "They saw a game; a case study." Albert H. Hastorf and Hadley Cantril, *The Journal of Abnormal and Social Psychology* 49, no. 1, 1954

<http://www2.psych.ubc.ca/~schaller/Psyc590Readings/Hastorf1954.pdf>

- "Opinions and social pressure", Solomon E. Asch, *Scientific American* Nov 1955

<http://www.uvm.edu/pdodds/teaching/courses/2009-08UVM-300/docs/others/everything/asch1955a.pdf>

# Assignment 0

Setting up accounts on Amazon Mechanical Turk and Crowdfunder  
(to be posted on the course website)