Lecture 2:
CS 6306 / INFO 6306:
Advanced Human Computation

Today: Big Picture Landscape

- Required readings:
  - Quinn, Alexander J., and Benjamin B. Bederson. "Human computation: a survey and taxonomy of a growing field." *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 2011.
- Additional readings:
  - Doan, Anhai, Raghu Ramakrishnan, and Alon Y. Halevy. "Crowdsourcing systems on the world-wide web." *Communications of the ACM* 54.4 (2011): 86-96.
  - Howe, Jeff. "The rise of crowdsourcing." *Wired magazine* 14.6 (2006): 1-4.
  - Kalil, Thomas A. "Leveraging cyberspace." *IEEE Communications Magazine* 34.7 (1996): 82-86.
  - Law, Edith, and Luis von Ahn. "Human computation." *Synthesis Lectures on Artificial Intelligence and Machine Learning* 5.3 (2011): 1-121. Chapters 1 and 2.
  - Malone, Thomas W., Robert Laubacher, and Chrysanthos Dellarocas. "The collective intelligence genome." *MIT Sloan Management Review* 51.3 (2010): 21.
  - Surowiecki, James. *The wisdom of crowds*. Anchor, 2005. Chapter 1.

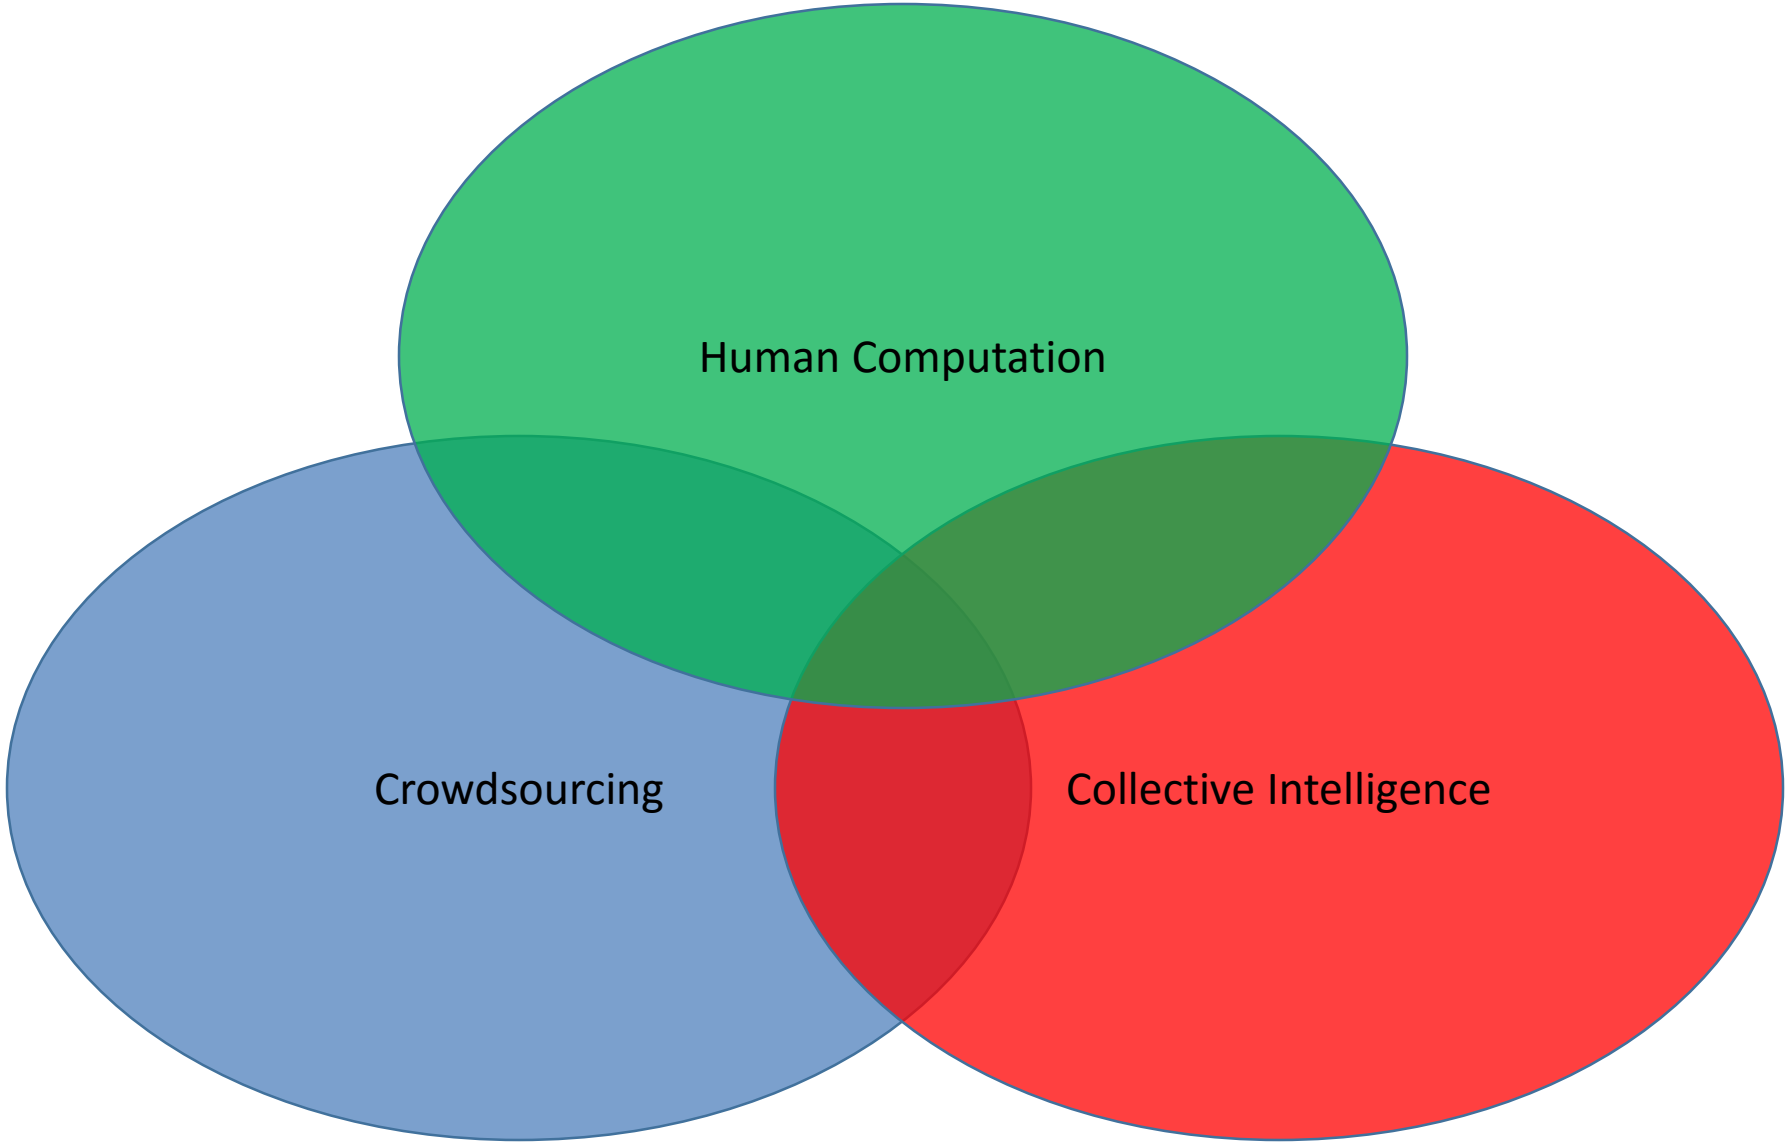
Human Computation
Crowdsourcing
Collective Intelligence

Human Computation

- Writing programs that turn to people as if they were subroutines to do things that we don’t know how to get computers to do (yet).

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

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, and if in doing so, it addresses the following four fundamental challenges:
  - How to recruit and retain users?
  - What contributions can users make?
  - How to combine user contributions to solve the target problem?
  - How to evaluate users and their contributions?"

"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

Collective Intelligence

- "groups of individuals doing things collectively that seem intelligent"

"The Collective Intelligence Genome", Malone, Laubacher, and Dellarocas, *MIT Management Review*, April 1, 2010

- "How can people and computers be connected so that—collectively—they act more intelligently than any individuals, groups, or computers have ever done before?"

http://cci.mit.edu

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”

Quinn and Bederson, CHI 2011

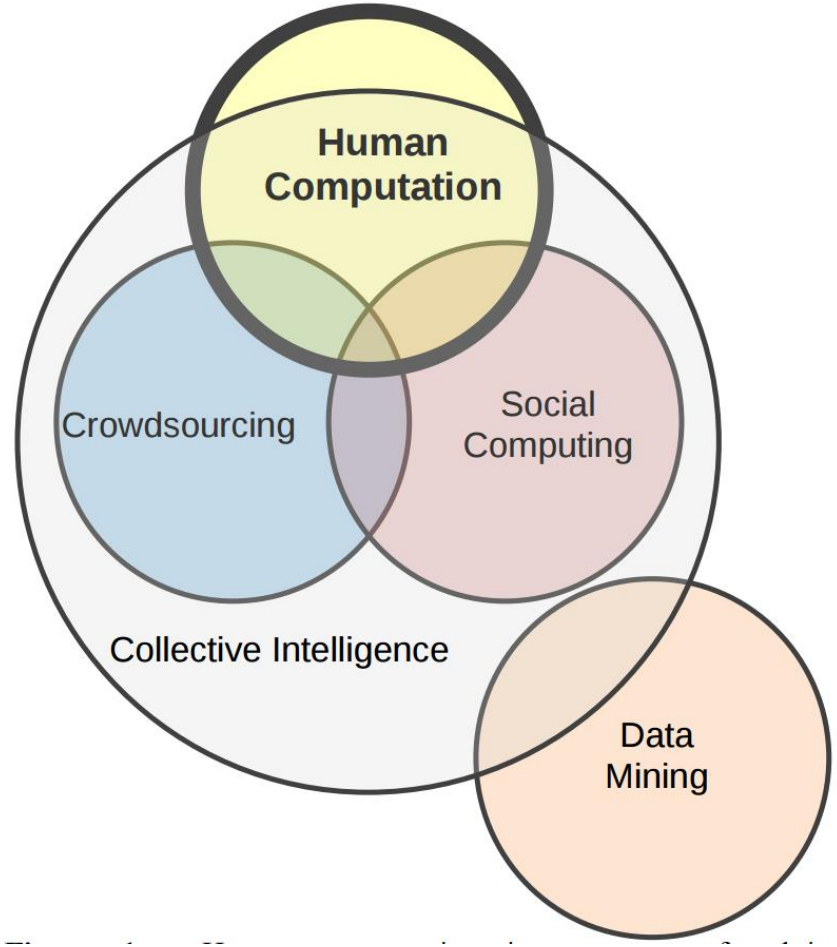

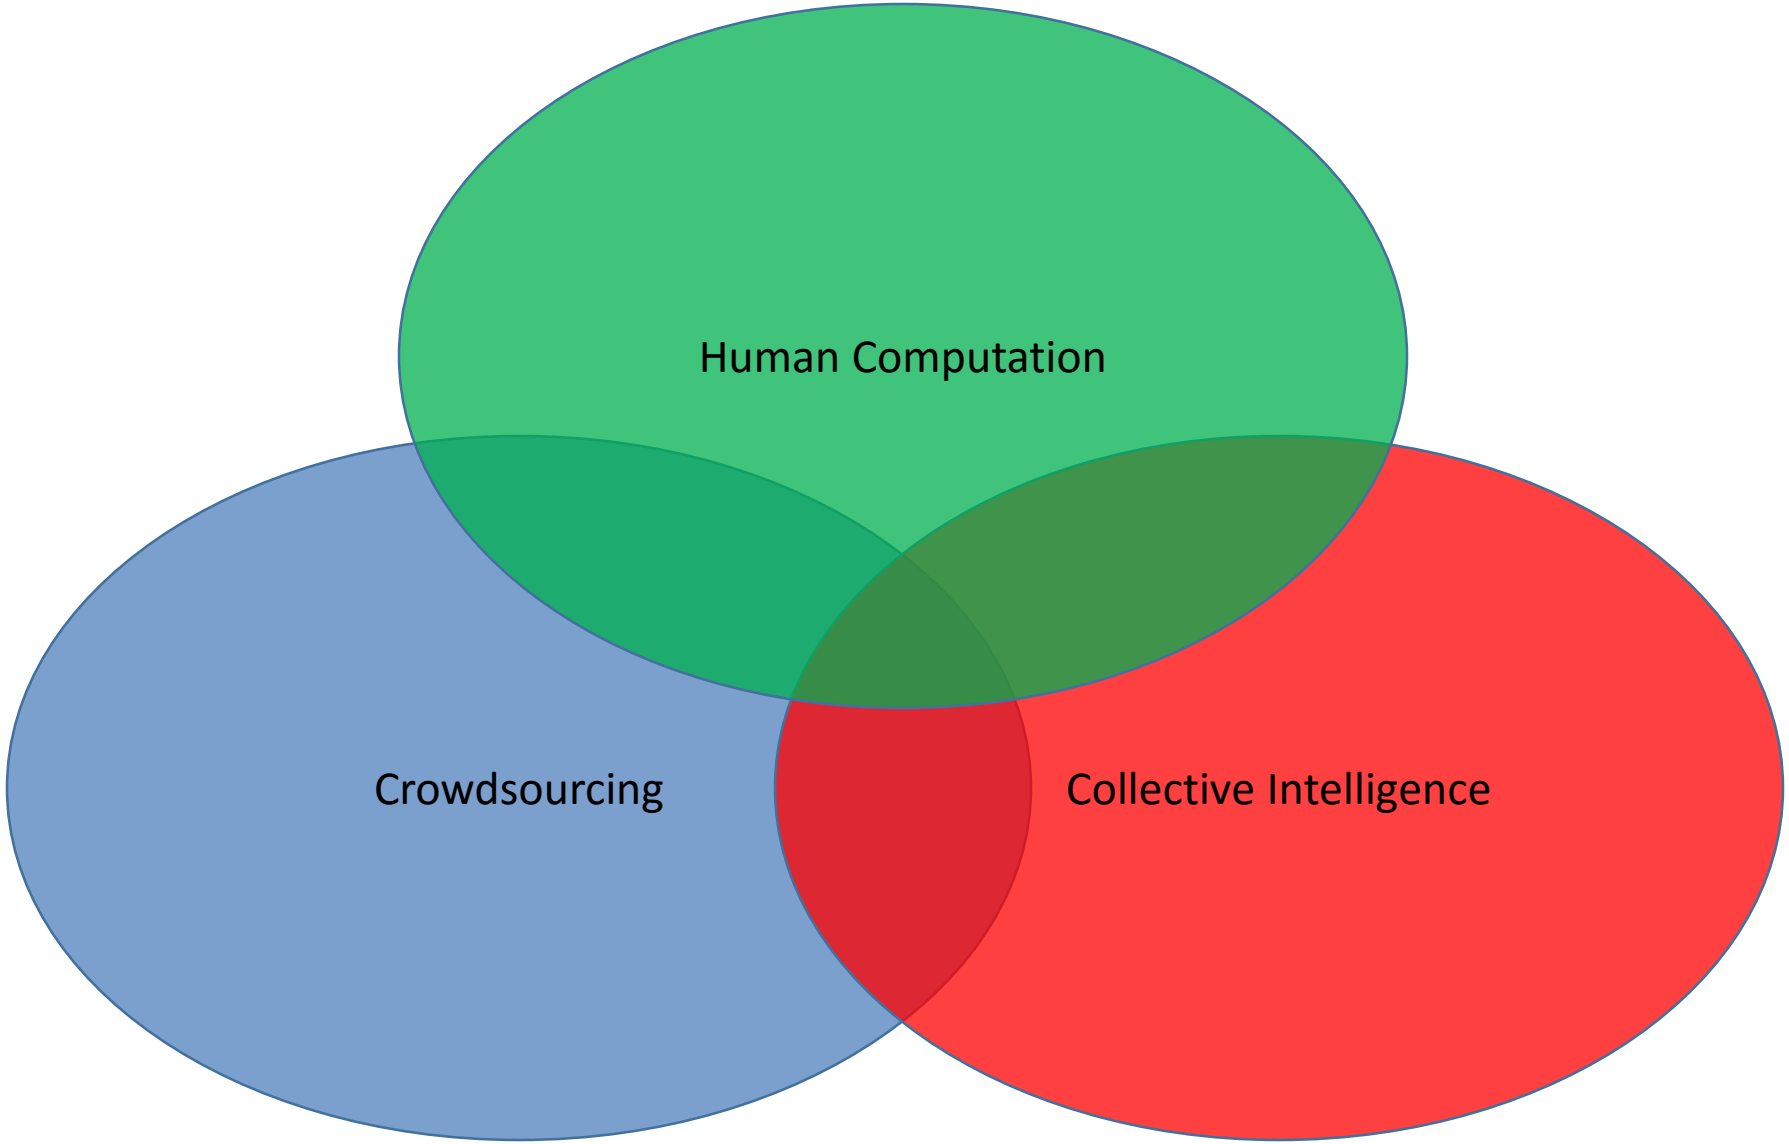
Human Computation
Crowdsourcing
Collective Intelligence

Quinn and Bederson, CHI 2011

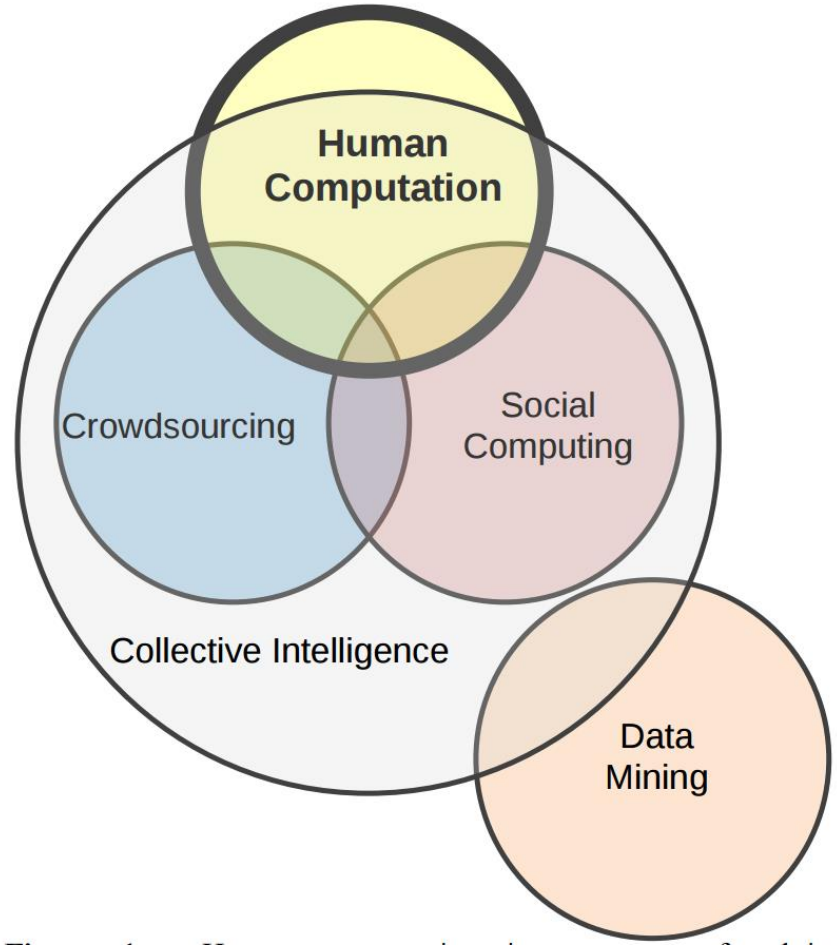

- Collective Intelligence:
  All crowdsourcing is intelligent?
  All social computing intelligent?
  “Emergence” not always present
- Data Mining:
  Used in HC
  Uses HC
- Social Computing:
  People interacting online
  Platform
  Too vague to be useful?

Quinn and Bederson, CHI 2011

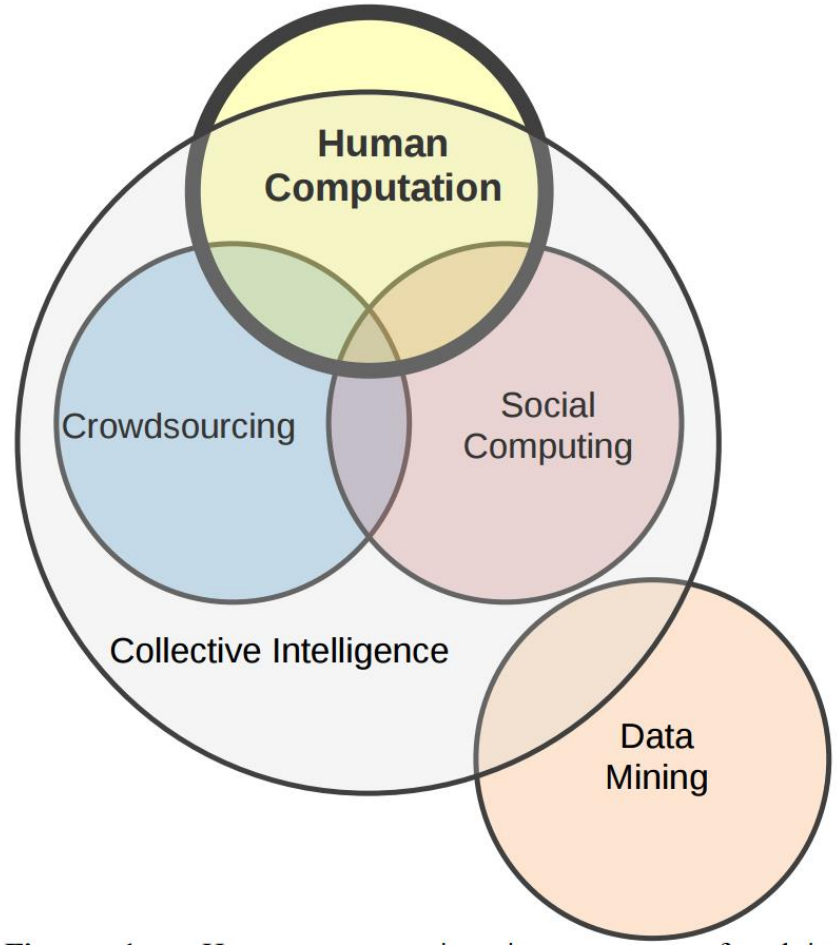

- “human computation does not encompass online discussions or creative projects where the initiative and flow of activity are directed primarily by the participants’ inspiration, as opposed to a predetermined plan designed to solve a computational problem.”
- “Whereas human computation replaces computers with humans, crowdsourcing replaces traditional human workers with members of the public.”

Quinn and Bederson, CHI 2011

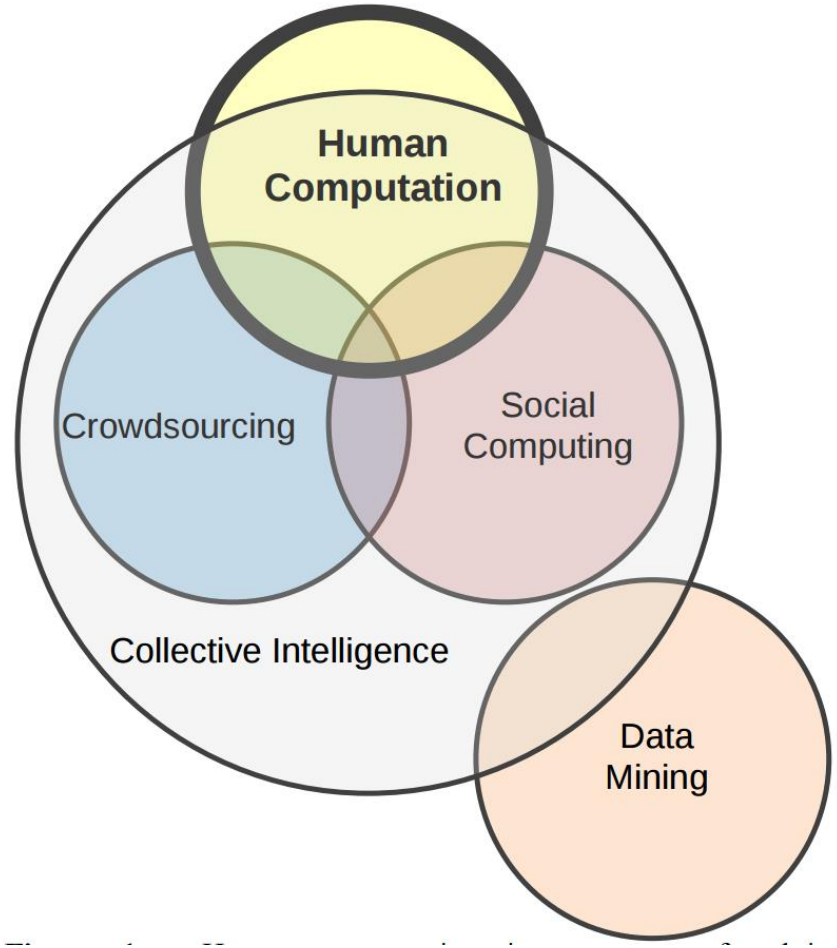

- “Technologies such as blogs, wikis, and online communities are examples of social computing. The scope is broad, but always includes humans in a social role where communication is mediated by technology. The purpose is not usually to perform a computation.”

Quinn and Bederson, CHI 2011

- Motivation:
  - Pay (AMT)
  - Altruism (search for Jim Gray)
  - Fun (ESP Game)
  - Reputation (Threadless)
  - Implicit Work [Coercion] (ReCAPTCHA)

- (Tom Malone (MIT): “Glory, fun, money”)

Quinn and Bederson, CHI 2011

- Reliability
  - Output agreement (ESP Game)
    - Do they independently generate same label
  - Input agreement (Tag-a-tune)
  - Economical models (e.g., game theory)
  - Defensive task design (AMT)
    - Use image of text rather than text itself for a translation task
  - Reputation system (AMT, Wikipedia)
  - Redundancy/voting (AMT)
  - Ground truth seeding (AMT)
  - Statistical filtering/aggregating
  - Multilevel review (Soylent, TurkIt)
    - Peer grading
  - Expert review
  - Automatic checking (fold.it)

Quinn and Bederson, CHI 2011

- Aggregation
  - Collection (Wikipedia, reCaptcha)
  - Statistical processing (e.g., averaging) (IEM, Wisdom of the Crowds, Kasparov against the World)
  - Iterative improvement (TurkIt, Soylent)
  - Active learning
  - Search
  - [Genetic algorithm]
  - None

Quinn and Bederson, CHI 2011

- Human Skill
- Process Order
  - Relative roles of people vs computer
  - (= Different algorithms / Design patterns)
- Task Request Cardinality
  - How many tasks does each worker do?
  - How many workers for each task?

Malone, Laubacher, and Dellarocas 2010

- Who is doing it?
- Why are they doing it?
- How is it being done?
- What is being done?

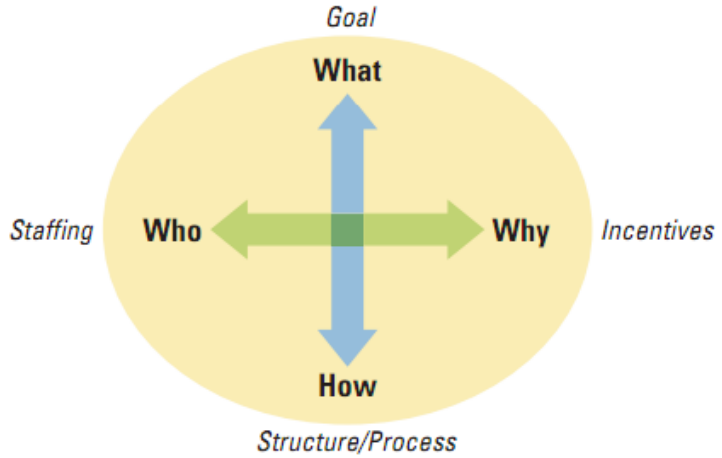

Malone, Laubacher, and Dellarocas 2010

- Who is doing it?
  - Organization / Authority selects
  - Crowd: Individuals select
- Why are they doing it?
  - Money
  - Love (e.g., enjoyment, socializing)
  - Glory
- What is being done?
  - Create
  - Decide

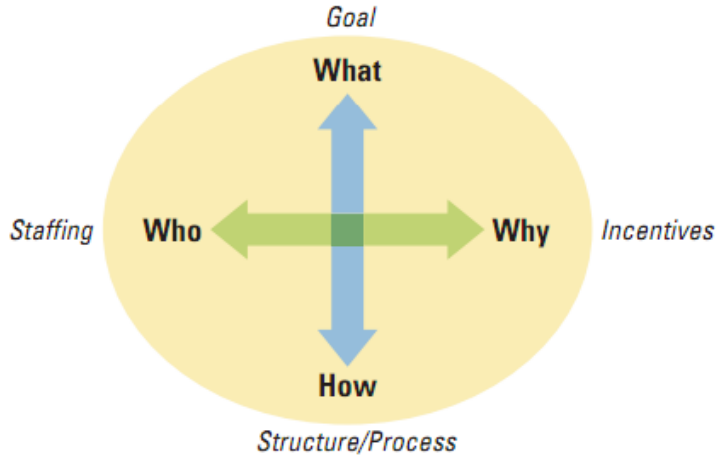

Malone, Laubacher, and Dellarocas 2010

- How is it being done?
  - Create
    - Collection (incl. Contest)
    - Collaboration
  - Decide
    - Group
      - Voting, Consensus, Averaging, Markets
    - Individual
      - Markets
      - Social Network

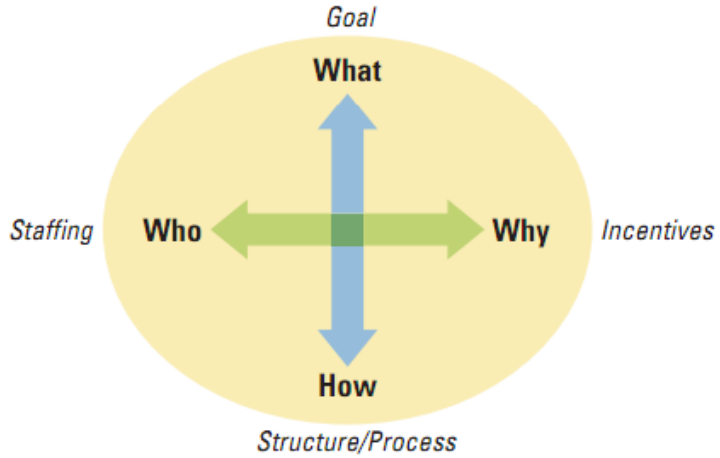

Doan, Ramakrishnan, and Halevy 2011

- “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, and if in doing so, it addresses the following four fundamental challenges:
  - How to recruit and retain users?
  - What contributions can users make?
  - How to combine user contributions to solve the target problem?
  - How to evaluate users and their contributions?”

Law and von Ahn 2011

- Aggregating outputs: Simple vs complex
- Task routing: Push vs pull
- Quality: Correctness, efficiency
- Task design: [Information,] granularity, independence, incentives, quality control
  - Incentives:
    - Extrinsic: money, rewards
    - Intrinsic: social standing, influence/power, curiosity, social contact, idealism, to collect/ownership, to win/competition

Surowiecki, 2004
Wisdom of the Crowds

- Requires participant:
  - Diversity
  - Independence
  - Decentralization
  - Aggregation

ABSTRACT

People with shared interests are using the Internet to solve problems, accomplish tasks, and create resources that would be well beyond the reach of any one person or organization. The Internet is being used to create virtual libraries, factor large numbers, organize massive volunteer efforts, and filter information in a collaborative fashion. The ability to leverage the efforts of large numbers of networked users has important economic, social, and political consequences. This phenomenon is important to policy makers because it can potentially be used toleverage scarce taxpayer dollars and promote applications of the information infrastructure.

Leveraging Cyberspace

Thomas A. Kalil, National Economic Council

"Give me a lever long enough and a place to stand, and I will move the Earth."

– Archimedes

The rapid growth in the ubiquity and functionality of the Internet is amazing. The Internet now connects 10 million computers, tens of millions of users, and more than 100 countries. At its current rate of growth, the Internet will connect 100 million computers by the year 2000. Because anyone with a computer and a connection to the Internet can publish, the global information space is also growing rapidly. Developers of search engines such as Altavista and Lycos believe that the Web currently contains 30–50 million pages of information, or 200 to 330 Gbytes of text. At current growth rates, the Web could surpass the 29 Tbytes of the Library of Congress in two years [1].

In addition to allowing anyone to publish, the open architecture of the Internet also allows anyone to add to its func-

In my view, one of the more important applications of the Net is its ability to enable "communities of interest" to solve problems, accomplish tasks, or create resources that would be well beyond the ability of any one individual or organization to complete. Steven Whitehead refers to this phenomenon as "cyberspace leveraging," which he defines as "using computer networks to harness the power of a large population of networked users" and leveraging the "small efforts of the many" as opposed to the "big efforts of the few" [2]. Think of the Internet as a distributed, massively parallel supercomputer connecting not only microprocessors but people, information repositories, sensors, intelligent agents, and mobile code.

There are many examples of Internet users leveraging cyberspace, some of which will be described in greater detail below. Six hundred volunteers from five continents used 1600 computers to factor RSA-129 in eight months, a mathematical feat that was projected to take 40 quadrillion years. Users of multi-user dungeons (MUDs) collaborate to build elaborate text-based virtual reality environments. Archives of scientific

Kalil, 1996

- Can dependencies between parts of the task be eliminated or managed?
- What will motivate people to participate?
- Is there part of the task that must be centrally administered?
- Does the initiative demonstrate increasing returns?

Types of Human Computation

- Overt / Explicit Work
  - Collecting
  - Smartest in the Crowd
  - Collaborative Decisions
  - Microlabor Markets
  - Games with a Purpose
  - Prediction Markets
- Covert / Implicit Work
  - Games with a Purpose
  - Prediction Markets
  - Trojan Horse systems
  - Mining user behavior

Logistics

- Course meeting time
  - 1/week vs 2/week?
  - October 4?
- Workload:
  - Projects and Assignments: 66%
    - Smaller ones relevant to week's papers
    - Final Project (Report and Presentation)
  - Class presentations and participation: 33%
    - Weekly paper write-ups
    - Leading in-class discussions of papers
  - Completing course evaluation: 1%

Workload

- Class presentations and participation: 33%
  - Weekly readings and paper write-ups
    - Responsible for *Required* readings – prior to class
    - 12 class weeks will be focused on papers:
      - Must do 9 ~1 page “reaction papers”
      - 1 additional reaction paper for any missed class
      - Due by 11pm the night before class – on Piazza
      - Grade: 1 (inadequate), 2 (needs work), 3 (adequate), 4 (great), 5 (perfect)
      - Content:
        1. Briefly summarize the most interesting/important idea or insight in each paper (in your own words)
        2. For *one* of the required readings, suggest how it could have been improved
        3. Suggest one question or topic that you would like to discuss in class concerning the readings

Workload

- Class presentations and participation: 33%
  - Leading in-class discussions of papers
    - Responsible for 2 sessions
    - I will handle required readings
    - You will select (at least) 3 additional readings, either from reading list or others you suggest
      - We will briefly meet for the last portion of class
        - 2 weeks beforehand to discuss paper selections
        - 1 week beforehand to discuss possible exercise/assignment
      - In class:
        - Summarize the selected readings, relating them to the required readings as best as you can
        - Propose an original idea for possible future research in the area
          - What question are you asking?
          - Why is it interesting?
            (Could improve results, address a limitation, lead to a new application, etc.)
          - How would you evaluate it?
      - By later that day: Rank the reaction papers

Workload

- Class presentations and participation: 33%
  - Constructive feedback on your classmates' presentation and report

Workload

- Projects and Assignments: 66%
  - Smaller ones relevant to week's papers
    - Some proposed by weekly session leader
    - Collectively explore a topic?
  - Final Project
    - Flexible, based on your interests
    - Projects that are related to work in other classes or research or other settings are welcome and even encouraged
    - However, the work performed for this course must be separate from that work. You may not submit the same work for both this and another class, and all relevant parties (other instructors, collaborators, etc.) must be informed and agree to your doing so
    - Project proposal, mid-semester class discussion
    - Final conference-like paper
    - In-class presentation

Status Reports

1. Email haym.hirsh@cornell.edu:
   - Name and email
   - Program of study, year
   - Interests and background
2. Read:
   - Alexander J. Quinn and Benjamin B. Bederson. "Human computation: a survey and taxonomy of a growing field."
     Proceedings of the SIGCHI conference on human factors in computing systems. ACM, 2011.

Status Reports

3. Initiate setting up a *worker* account at Amazon Mechanical Turk:
   - Go to mturk.com, click on “Sign in as a Worker” at upper right, enter your email, click on “No, I am a new customer” and proceed from there
4. Complete IRB Training:
   - https://www.oria.cornell.edu/documents/IRB%20Training%20(for%20new%20human%20participant%20researchers).pdf

Readings for Next Time

- Khatib, Firas, Seth Cooper, Michael D. Tyka, Kefan Xu, Ilya Makedon, Zoran Popović, and David Baker. "Algorithm discovery by protein folding game players." *Proceedings of the National Academy of Sciences* 108, no. 47 (2011): 18949-18953.
- Tuite, Kathleen, Noah Snavely, Dun-Yu Hsiao, Adam M. Smith, and Zoran Popović. "Reconstructing the world in 3D: bringing games with a purpose outdoors." In *Proceedings of the Fifth International Conference on the Foundations of Digital Games*, pp. 232-239. ACM, 2010.