Algorithmic Crowdsourcing:

Current State and Future Perspective

Jie Wu

Dept. of Computer and Info. Sciences

Temple University

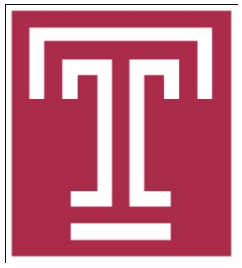

Road Map

- Introduction and Motivation
- Mechanical Turk
- Applications
- Paradigms
- Challenges and Opportunities
- Social Crowdsourcing
- Conclusion

What

Why

Basic components

Motivation examples

INTRODUCTION

Big Data is Everywhere!

- Lots of data is being collected:

  Volume, Variety, Velocity

  - Web data, e-commerce
  - Purchases
  - Bank/credit card transactions
  - Video and images
  - Social networks

How Much Data?

- Google processes 100 PB a day
- Wayback Machine has 3 PB + 100 TB/month (3/2009)
- WeChat has 600 M users and 20 B message per day
- Facebook has 2.5 PB of user data + 15 TB/day (4/2009)
- eBay has 6.5 PB of user data + 50 TB/day (5/2009)

Big Data Era

- "In information technology, big data consists of datasets that grow so large that they become awkward to work with using on-hand database management tools."

- Computers are not efficient in processing certain data (e.g., image processing)

What is Crowdsourcing?

- Coordinating a crowd (a large group of people online) to do micro-work (small jobs) that solves problems (that software or one user cannot easily do)

The Benefits of Crowdsourcing

- Performance
  - Inexpensive
  - Fast
- Human Processing Unit (HPU)
  - More effective than CPU (for some apps)
    - Image labeling
    - Language translation
    - Social network survey
    - …

Basic Components

- Requester
  - People submit jobs
  - Human Intelligence Tasks (HITs)
- Worker
  - People work on jobs
- Platform
  - Job management
  - Amazon Mechanical Turk (MTurk)

Requester

amazon mechanical turk beta

Worker Pool

Malaysia Airlines Flight MH 370

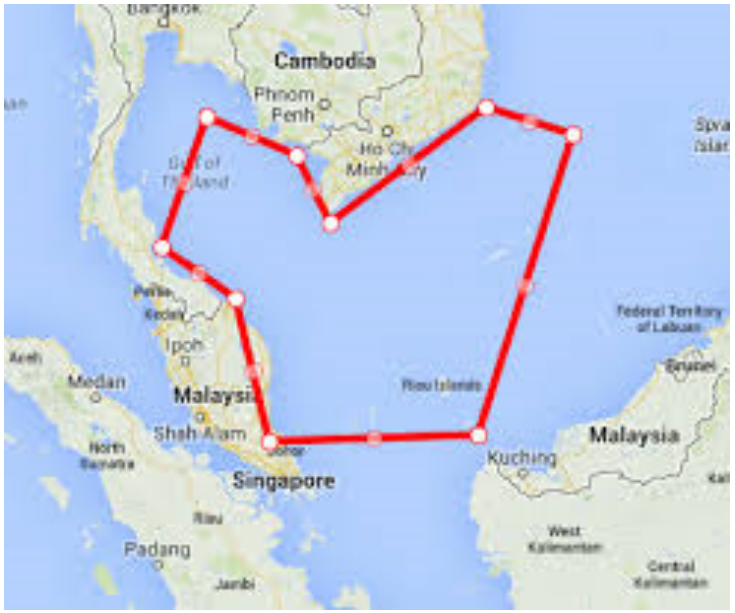

- DigitalGlobe
  - Crowdsourcing volunteers comb satellite photos for Malaysia Airlines jet

- March 11 (from a CSU prof. email)

I just saw on our local Denver Fox news (KDVR.com) that a local company, DigitalGlobe, has reoriented their satellites to take high-res images in the area where the plane may have crashed. Crowdsourcing efforts are on to have people scan these images and find signs of debris. I was reminded of Jie Wu's talk earlier this month.

Help Find Jim Gray

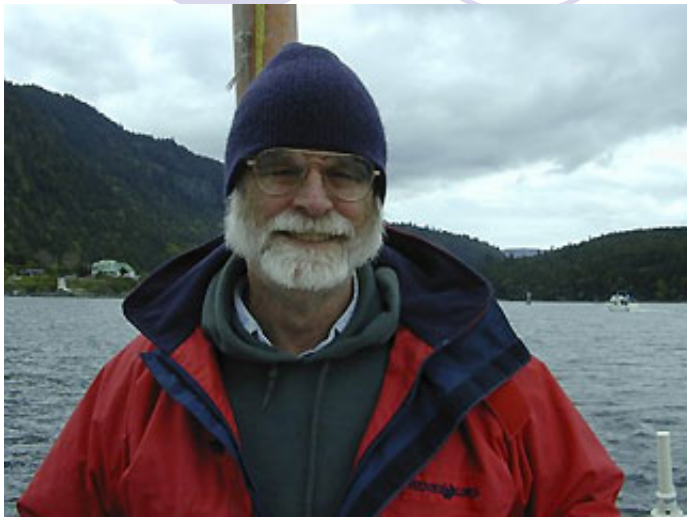

- Jim Gray, Turing Award winner, went missing with his sailboat outside San Francisco Bay in January 2007.

- Use satellite image to search for his sailboat.

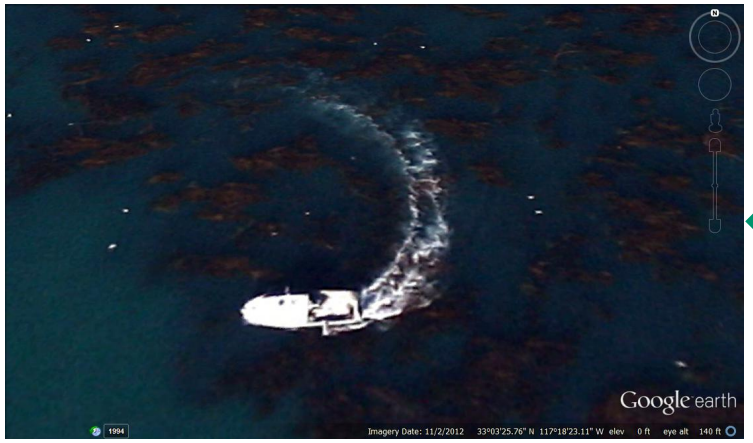

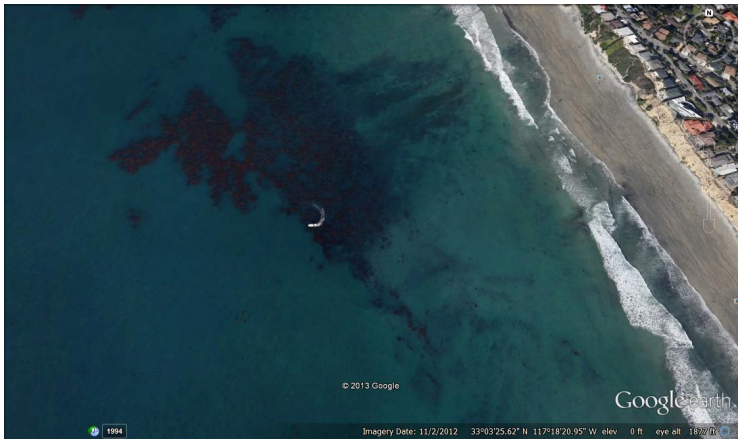

DARPA Network Challenges

- Problem (2009): $40,000 challenge award for the first team to find 10 balloons.
- MIT team won under 9 hours.
- Winning strategy
  - $2,000 per balloon to the first person to send the correct location
  - $1,000 to the person who invited the winner
  - $500 to whoever invited the inviter
  - … (or to charity) …

Tag Challenges

- Problem (March 31, 2012): Find five suspects in Washington, D.C., New York, London, Stockholm, and Bratislava.
- Winner from UCSD CrowdScanner: located 3 of the 5 suspects.
- Winner strategy: same as MIT. Also, recruiters of the first 2,000 get $1.

Washington DC

New York City

Bratislava

AI Could End Human Race (Stephen Hawking)

- Stephen Hawking
  - "Humans, who are limited by slow biological evolution, couldn't compete, and would be superseded."
- Recent movies
  - Her (2014) & Ex-Machina (2015)

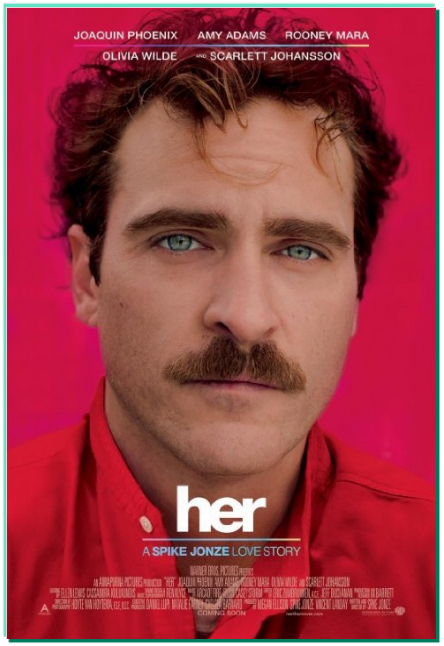

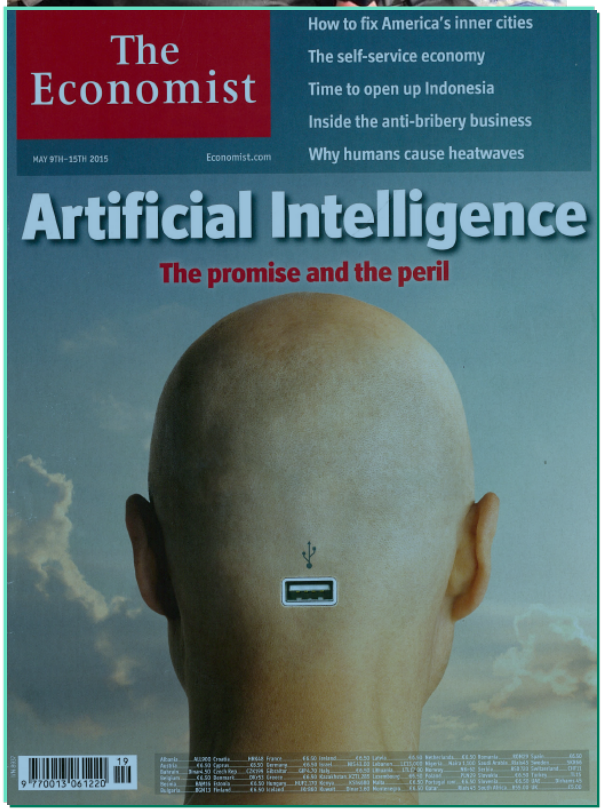

Smarter Than You Think

- Who is smarter
  - Human or computer?
- AI will redefine
  - What it means to be human

Our Machine Masters

NY Times, Oct. 31, 2014

- 1997 (Chess)
  - Kasparov vs. Deep Blue
- 1998
  - Kasparov vs. Topalov: 4:0
  - Kasparov + machine vs. Topalov + machine: 3:3
- 2005 (freestyle tournament)
  - Grand-master (>2,500)
  - Machine (Hydra)
  - Grand-master + machine
  - Amateurs (>1,500) + machine *
- 2016 (Go game)
  - AlphaGo vs. Lee Sedol: 4:1

Worker

HIT

Dashboard

Requester

MECHANICAL TURK

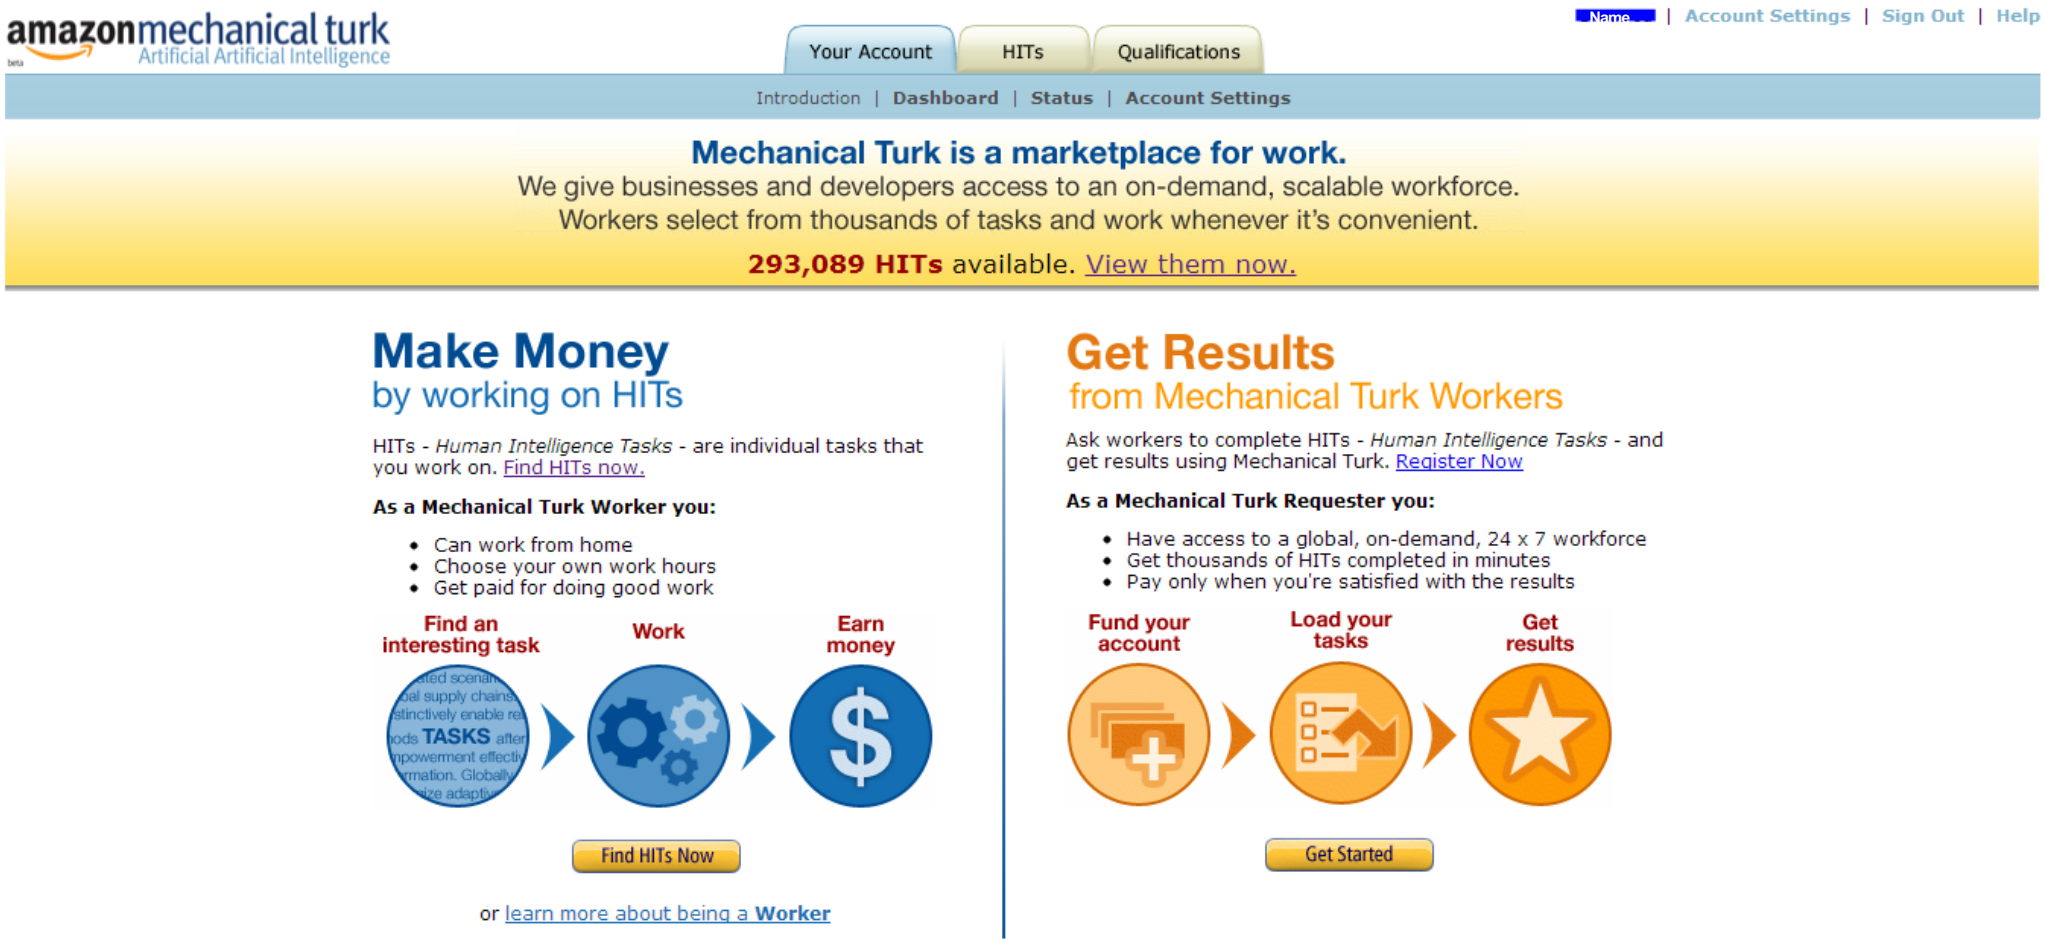

- As a worker, make an average of $0.03 per task
- Paid directly to Amazon account
- As requester, set up simple tasks for workers to complete
- Quality control is possible through MTurk services

Worker: Contract for a HIT

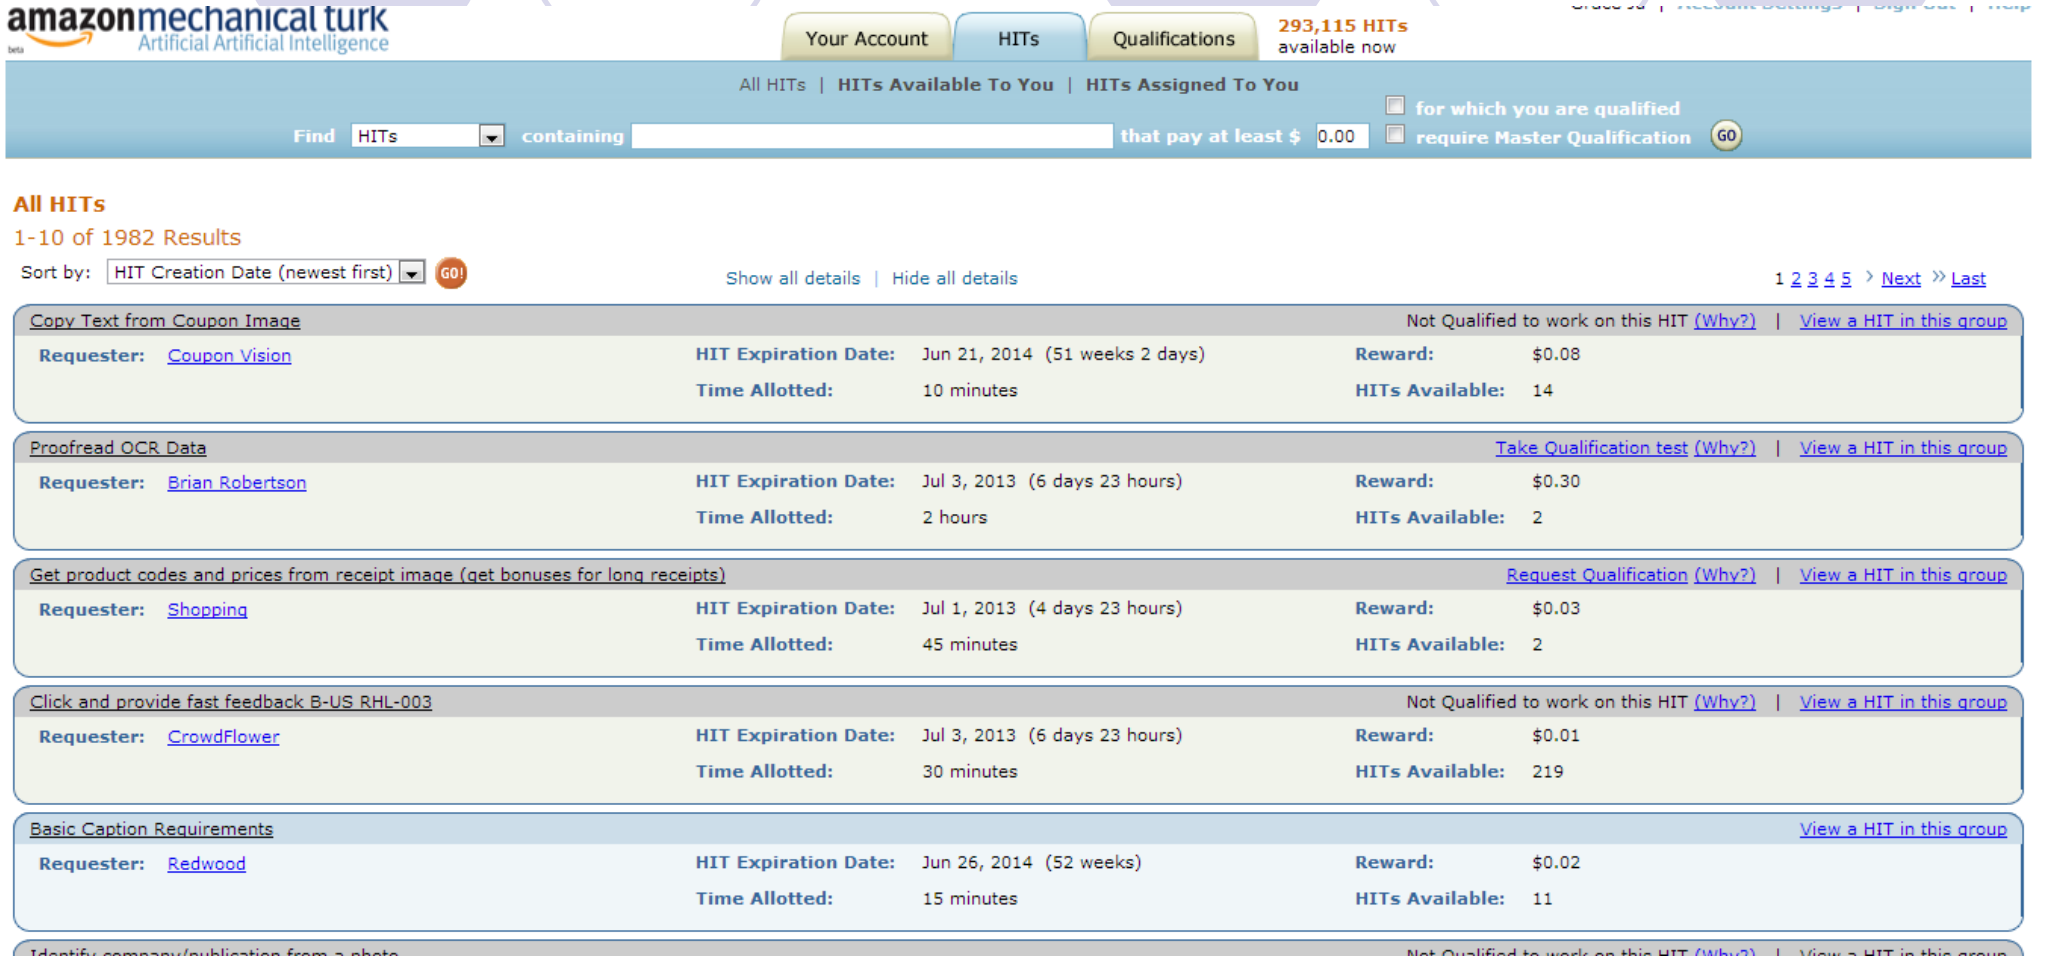

- Select a HIT
  - By creation date, payment amount, time allotment

Worker: Reviewing a HIT

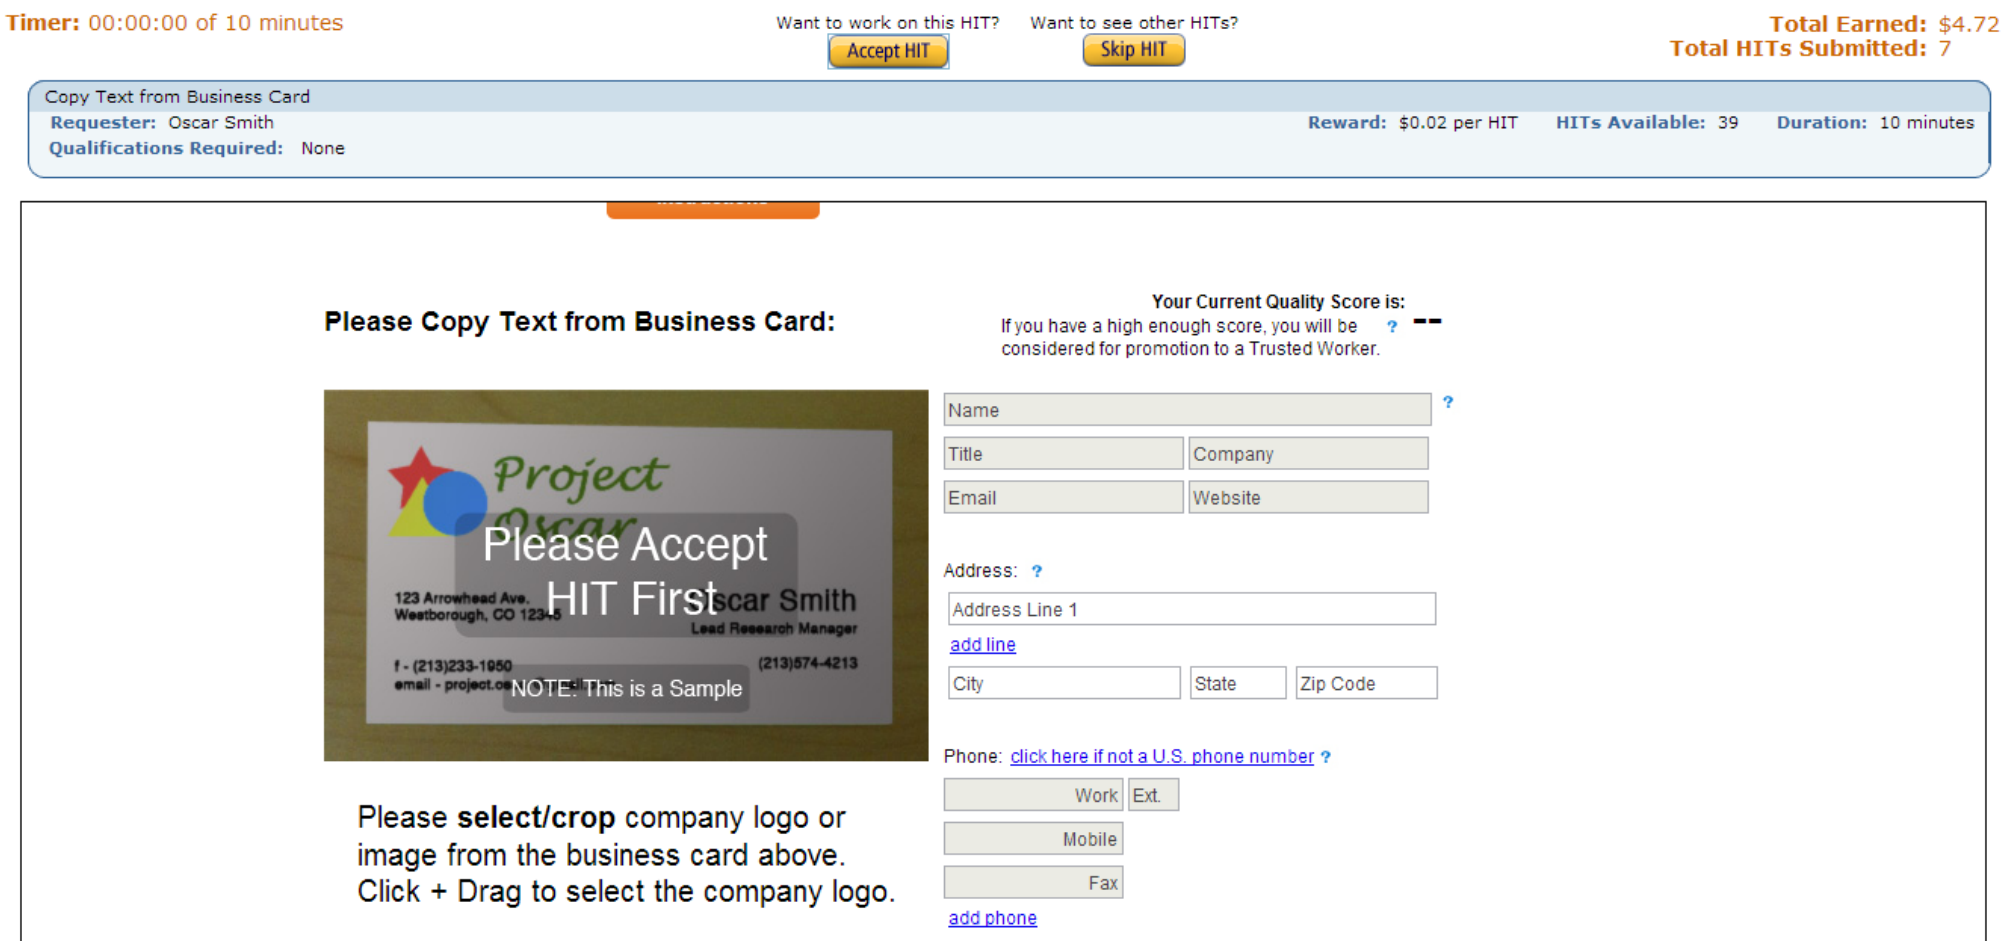

- Review the HIT before accepting
  - Shown full task, allotted time (10 minutes), reward amount ($0.02)

Worker: Completing a HIT

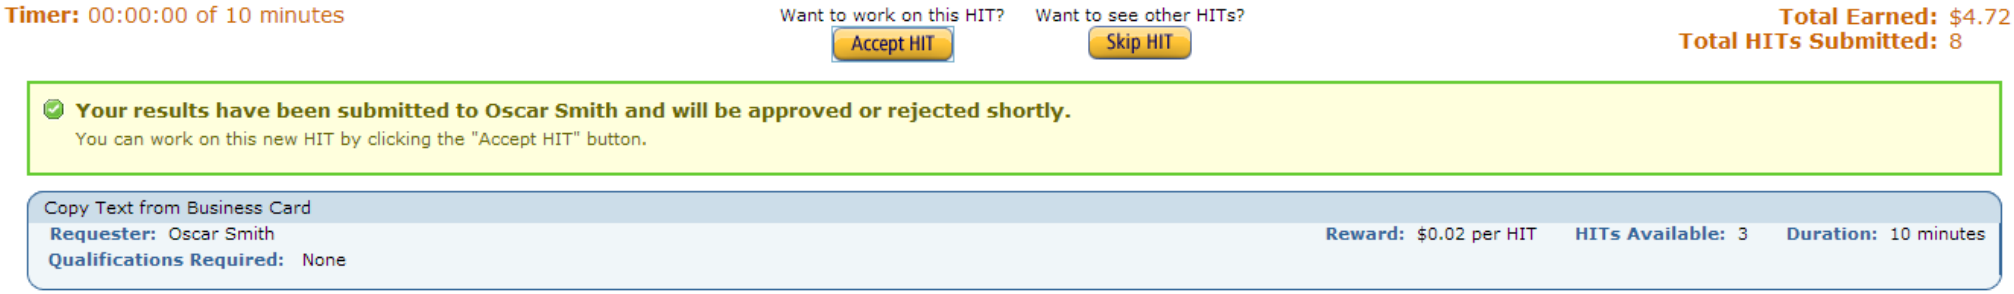

- Confirmation message in green
- Automatically shows the next HIT submitted by the same requester
- Check Dashboard to see if HIT is accepted

Worker: Sample Dashboard

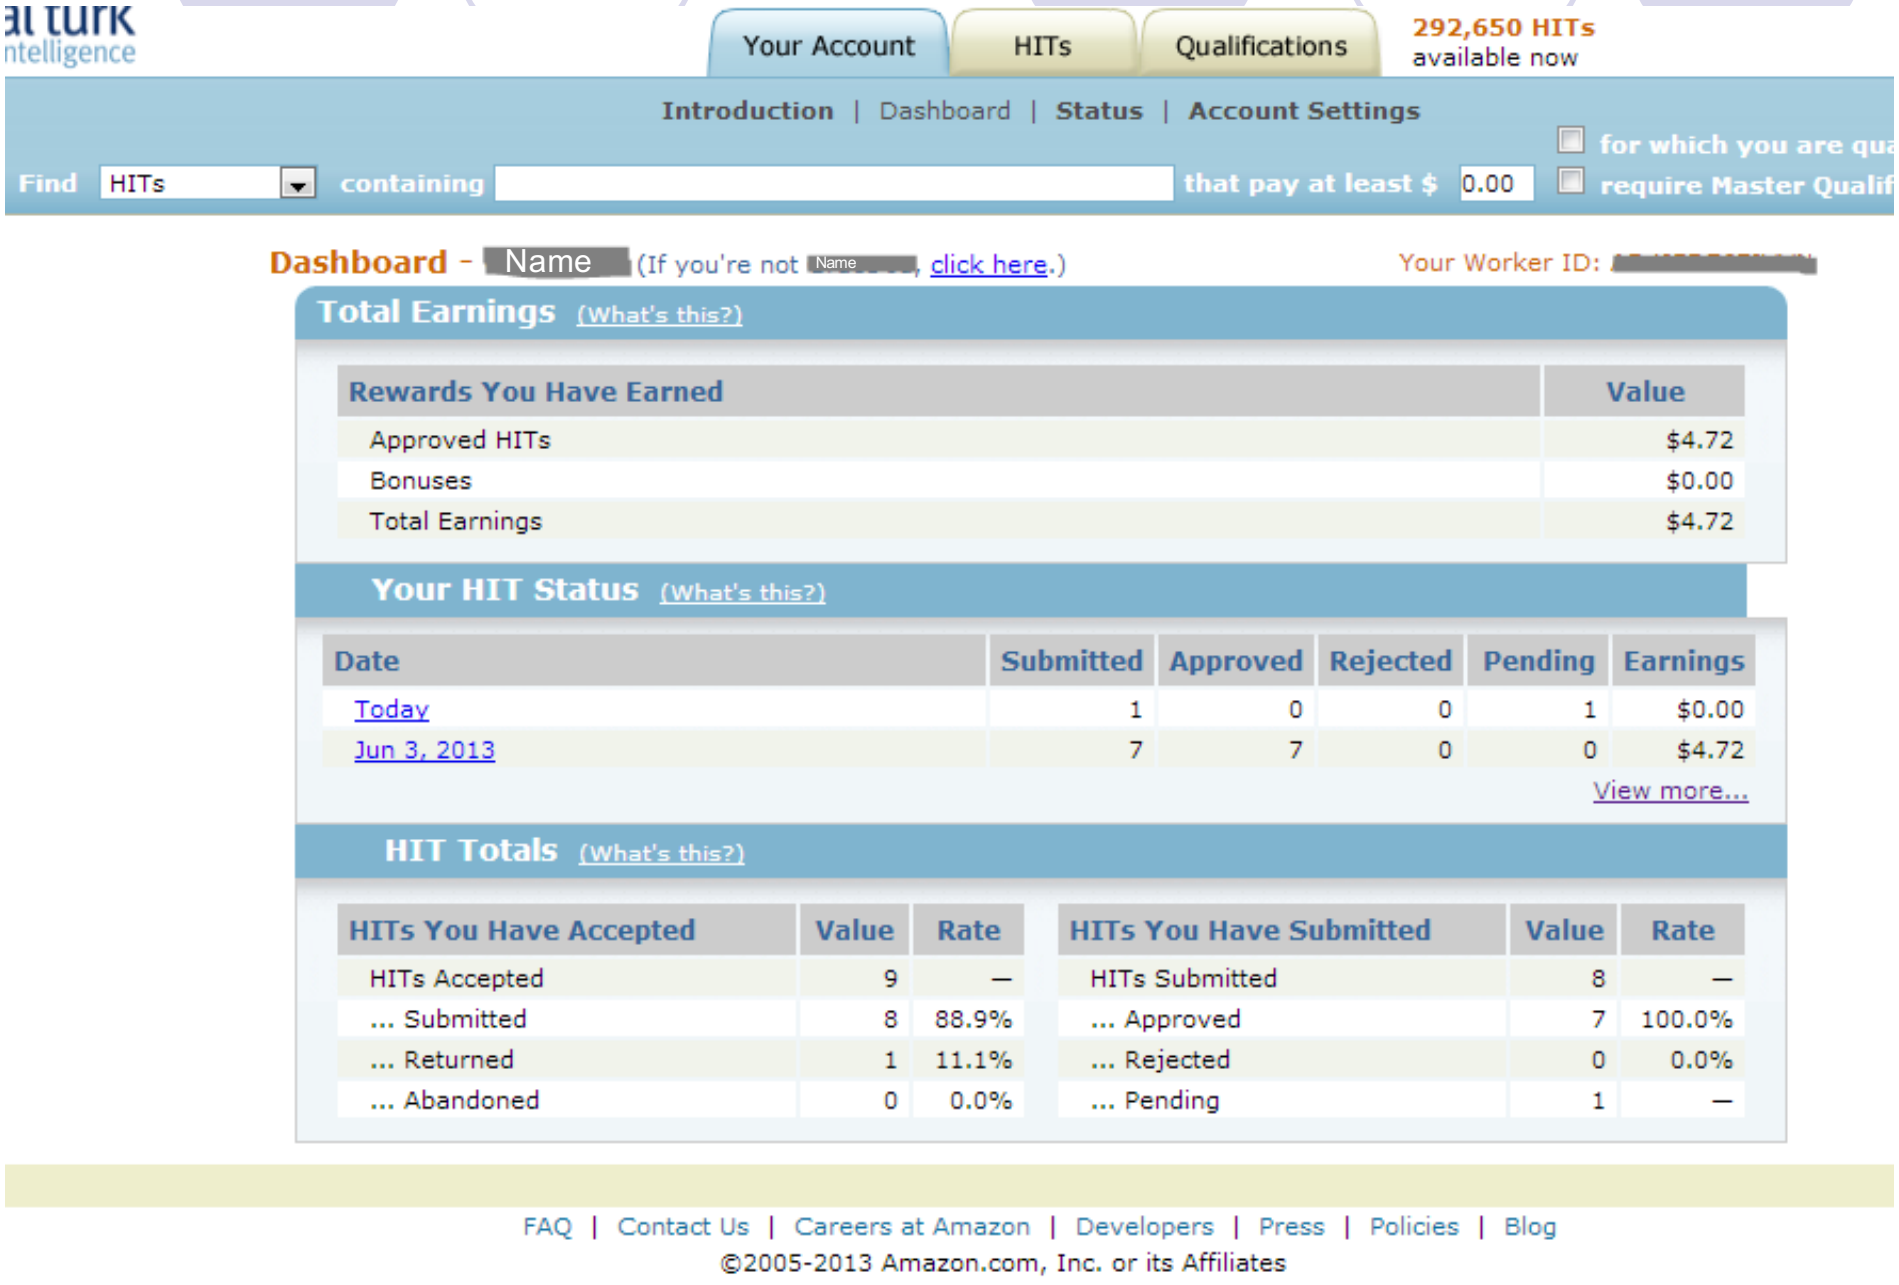

al turk
ntelligence

Your Account | HITs | Qualifications

292,650 HITs available now

Introduction | Dashboard | Status | Account Settings

Find HITs containing that pay at least $ 0.00

for which you are qua

require Master Qualif

Dashboard - Name (If you're not Name, click here.)

Your Worker ID: [redacted]

Total Earnings (What's this?)

Rewards You Have Earned	Value
Approved HITs	$4.72
Bonuses	$0.00
Total Earnings	$4.72

Your HIT Status (What's this?)

Date	Submitted	Approved	Rejected	Pending	Earnings
Today	1	0	0	1	$0.00
Jun 3, 2013	7	7	0	0	$4.72

View more...

HIT Totals (What's this?)

HITs You Have Accepted	Value	Rate
HITs Accepted	9	—
... Submitted	8	88.9%
... Returned	1	11.1%
... Abandoned	0	0.0%

HITs You Have Submitted	Value	Rate
HITs Submitted	8	—
... Approved	7	100.0%
... Rejected	0	0.0%
... Pending	1	—

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

Avoid Shady Requester

How Turkopticon works:

Turkopticon adds functionality to Amazon Mechanical Turk as you browse for HITs and review status of work you've done. As you browse HITs, Turkopticon places a button next to each requester and highlights requesters for whom there are reviews from other workers. Bad reviews let you avoid shady employers and good reviews help you find fair ones. You can view reports made against requesters with a quick click.

As you review HITs you've completed, are there HITs you weren't fairly paid for? Turkopticon adds a button that lets you review requesters from your "Status Detail" page.

EteRNA

Galaxy Zoo

Fine-grained Recognition

APPS: IMAGE PROCESSING

Biology: EteRNA: CMU, Stanford

- Aim: to gain mastery over the way RNA molecules folds.

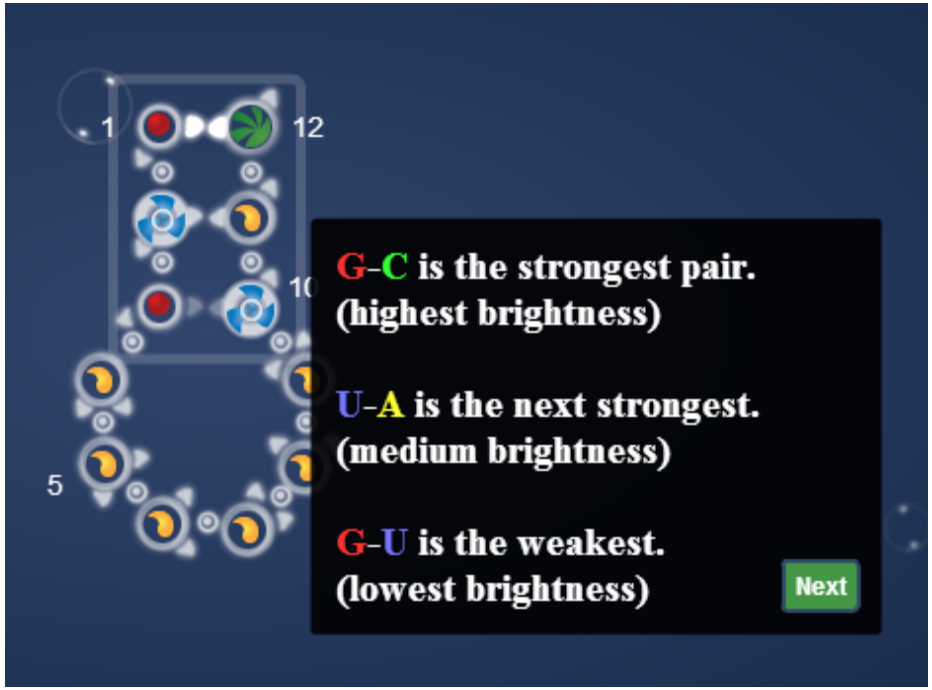

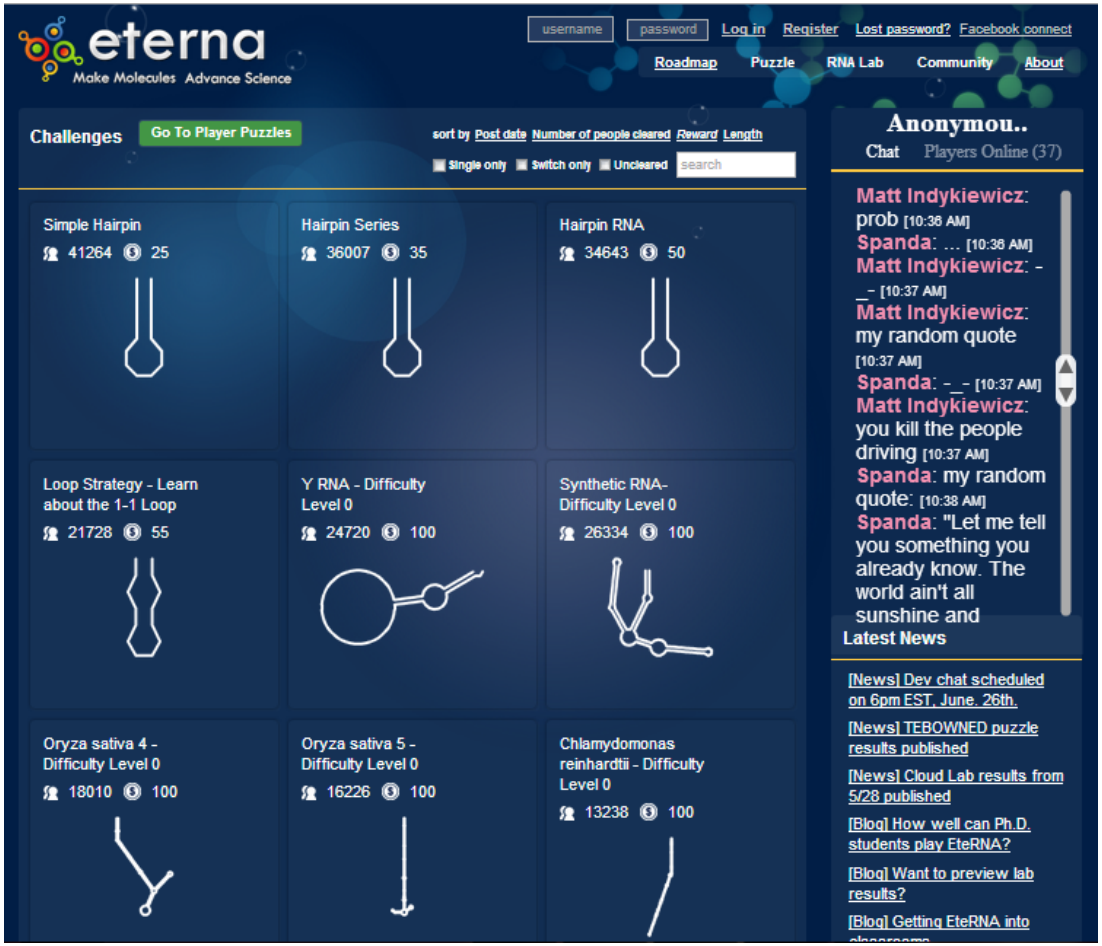

EteRNA: CMU, Stanford

- By assigning different colors (RNA nucleotides), a RNA chain will fold into different structure

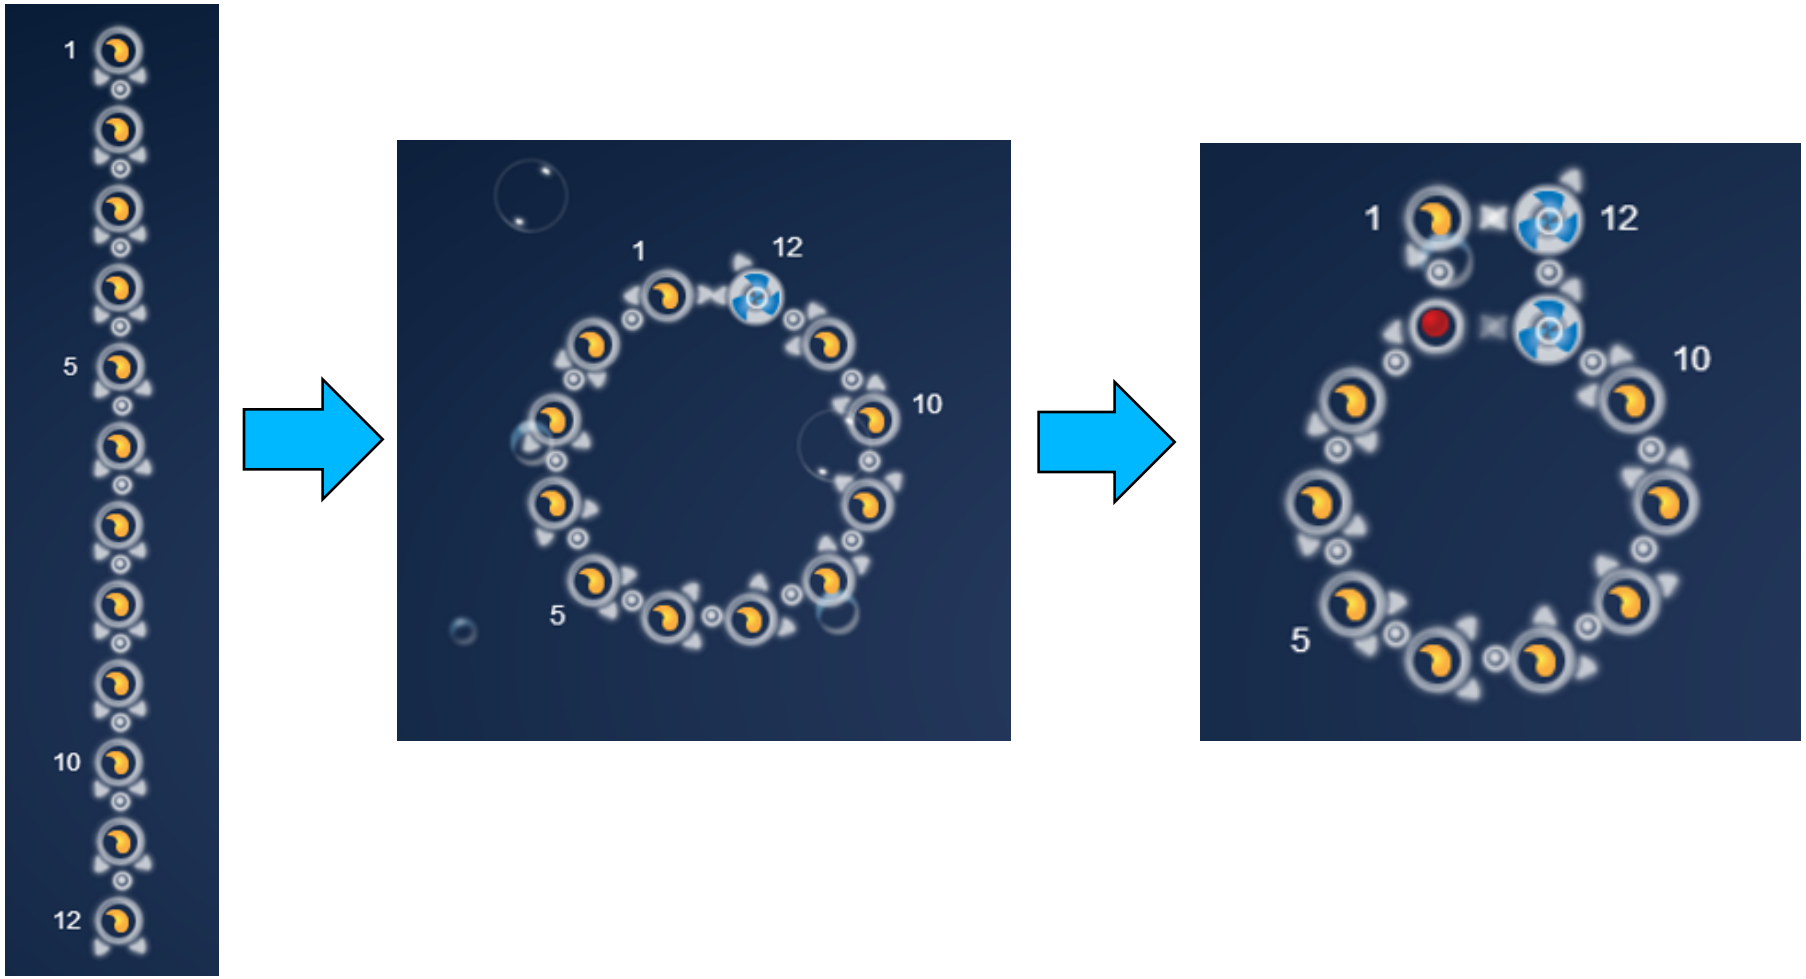

GalaxyZoo: Zooniverse

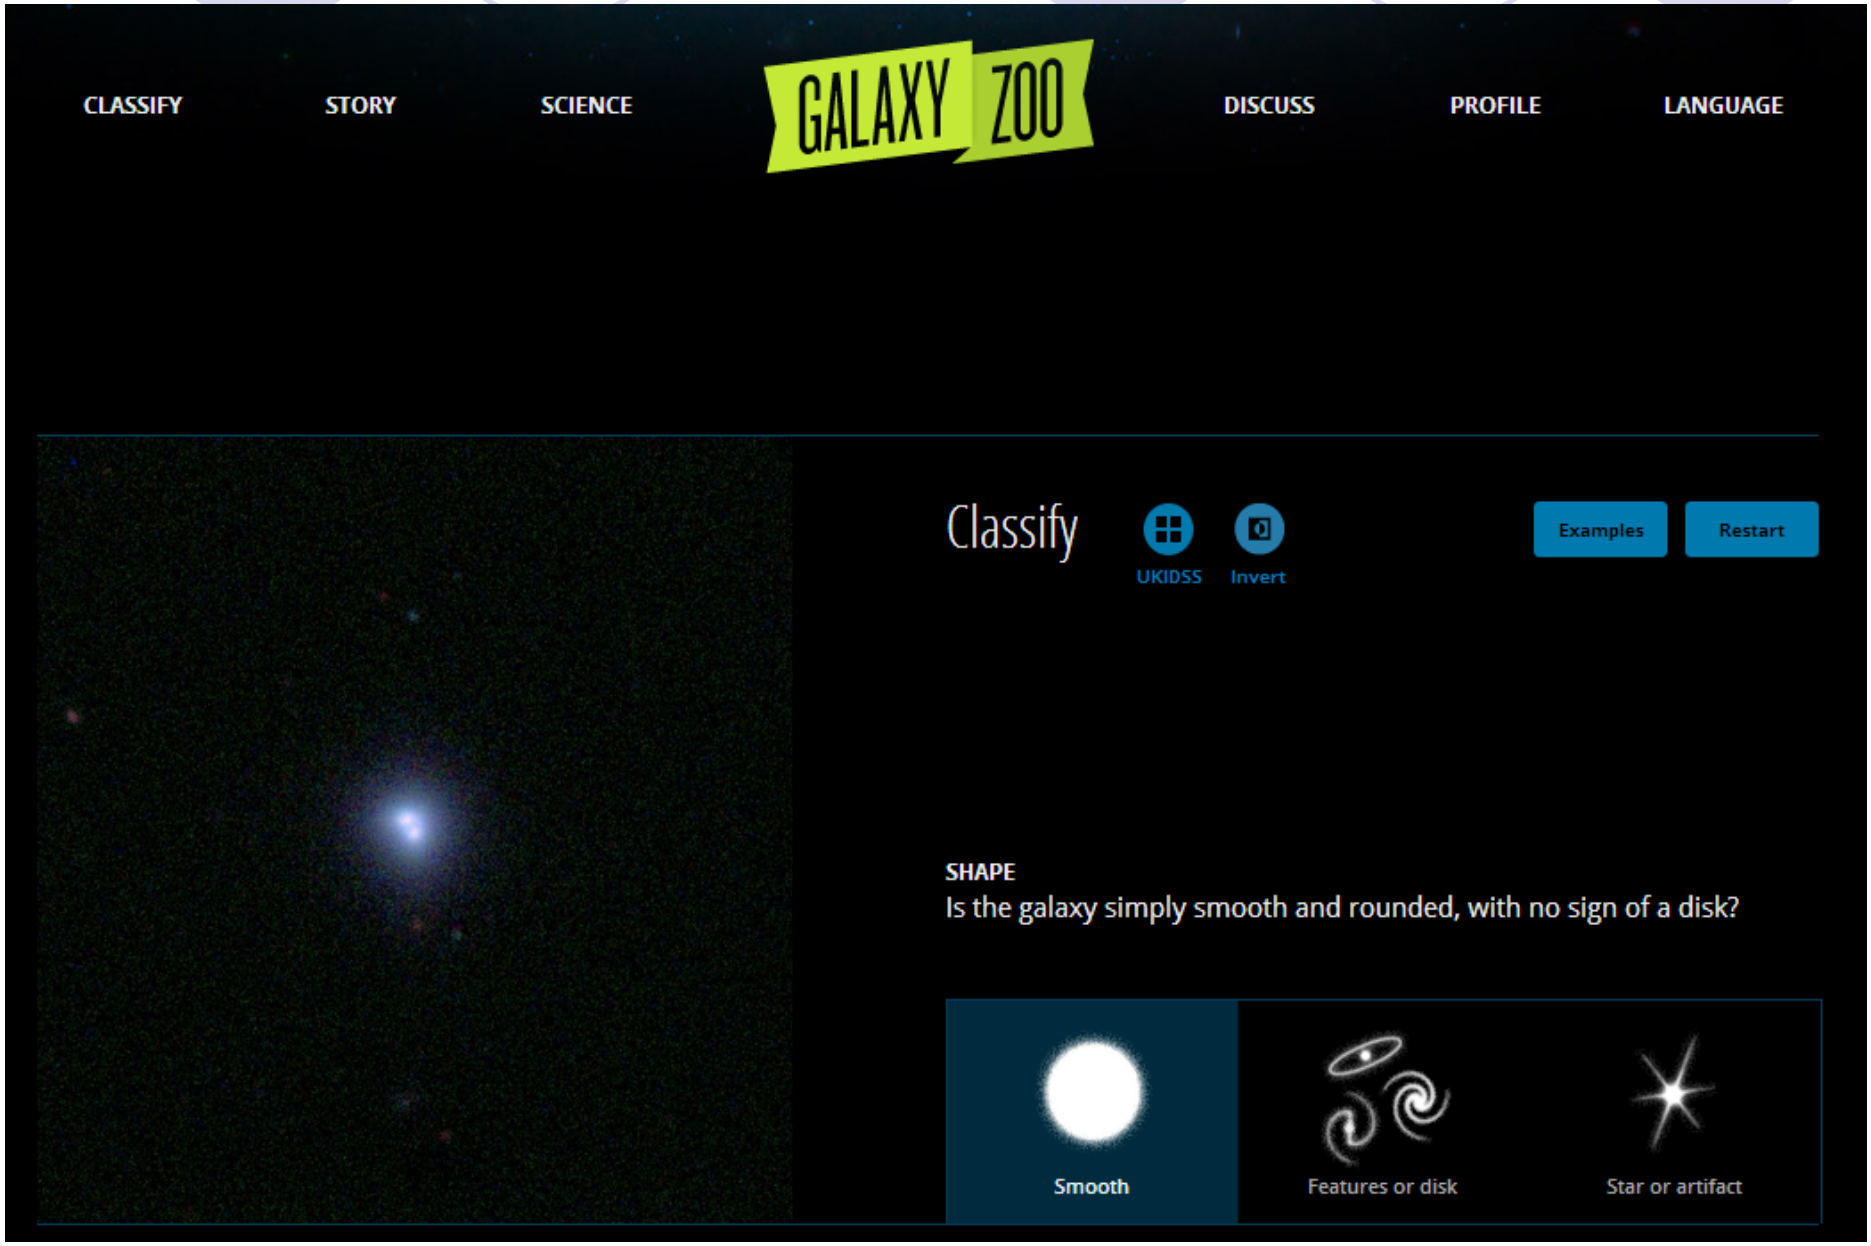

GalaxyZoo: Zooniverse

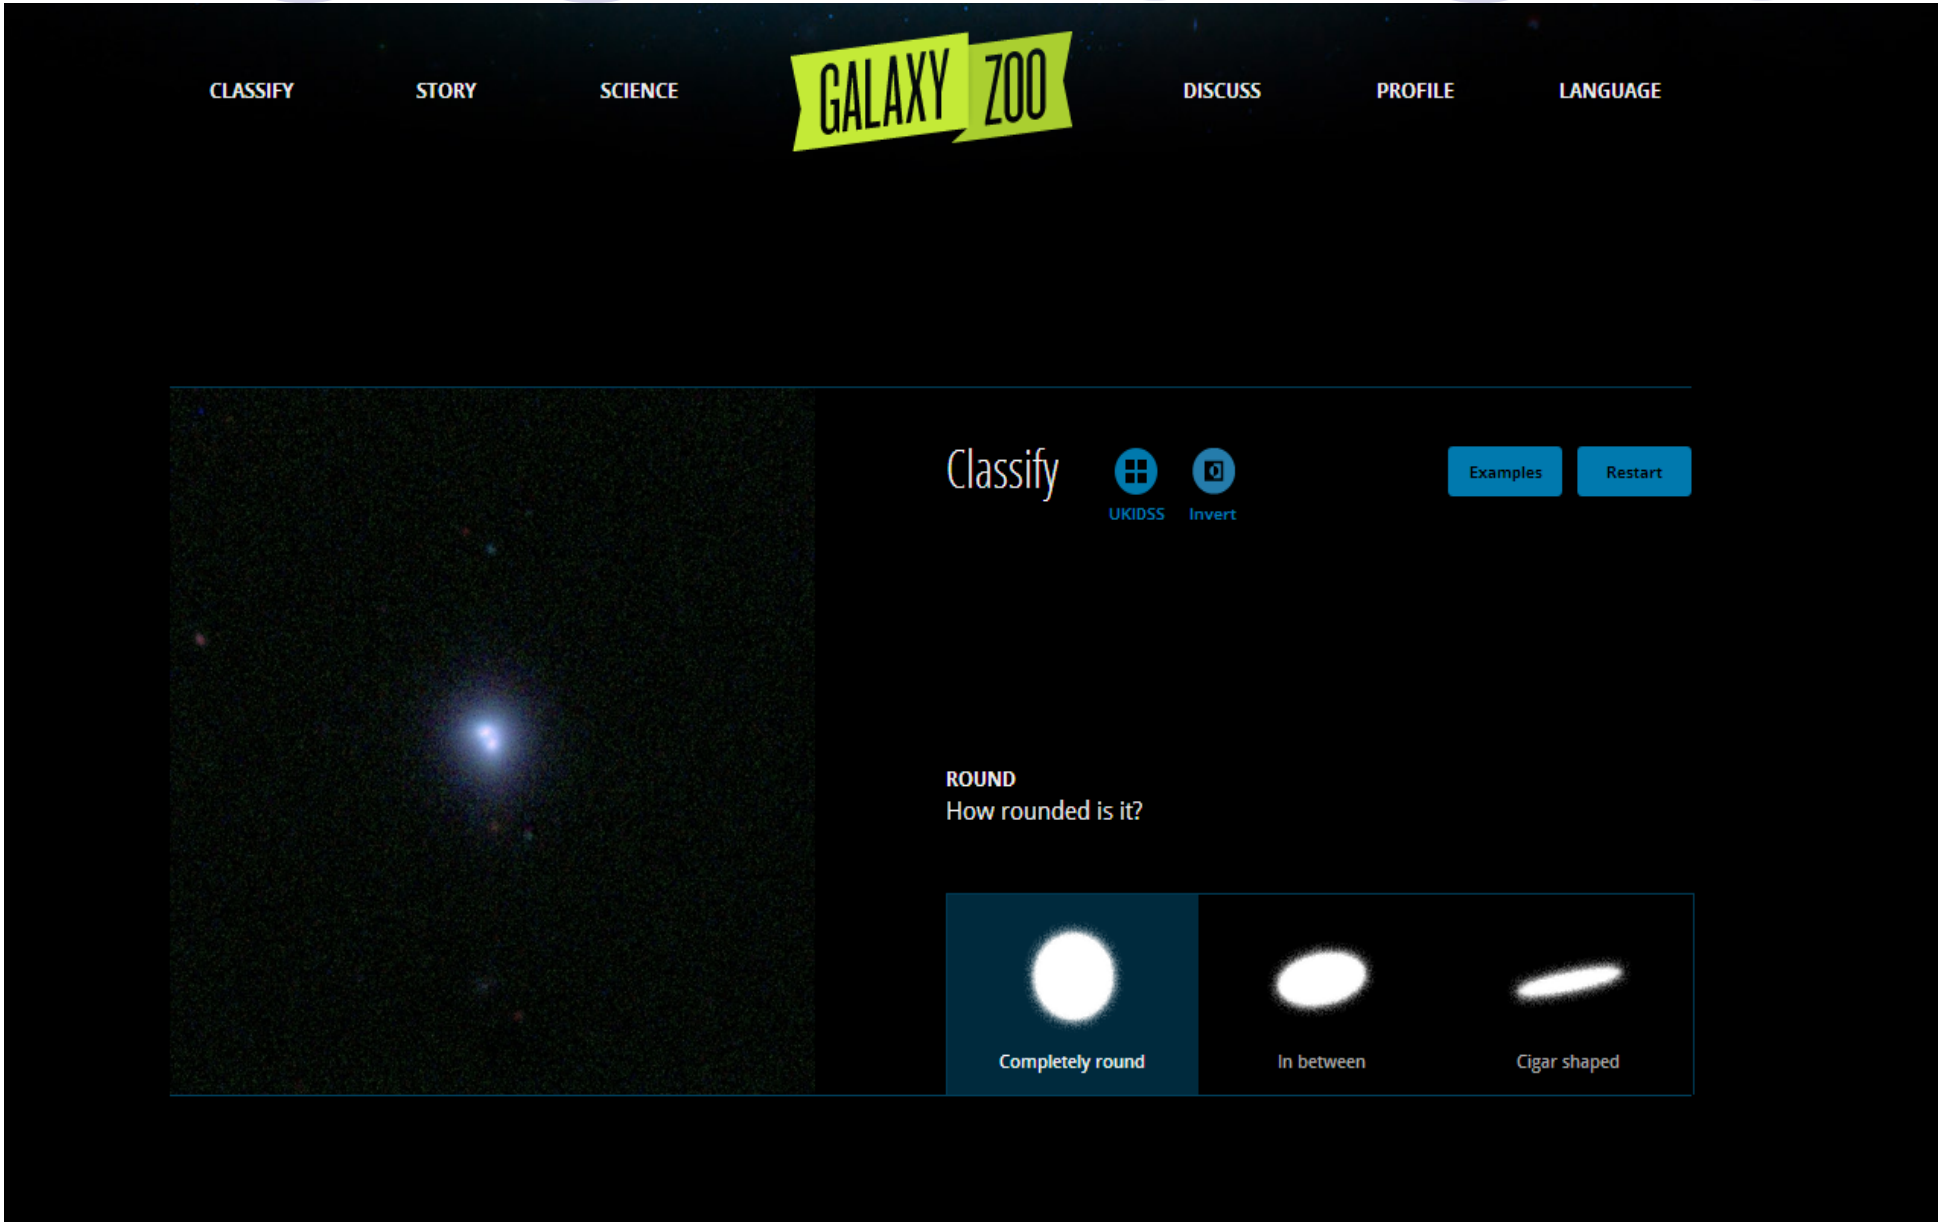

Fine-Grained Recognition: Tohme

K. Hara et al, "Tohme: Detecting Curb Ramps in Google Street View Using Crowdsourcing, Computer Vision, and Machine Learning," UIST 2014

GWAP.com

OnToGalaxy

reCAPTCHA

ChaCha

Crowrdsourcing

APPS: COMMONSENSE KNOWLEDGE

GWAP.com: CMU

ESP Game

- Labeling images

Tag a Tune

- Labeling tunes

OnToGalaxy: University of Bremen

- Given a keyword
  - e.g., "tourism"
- Collect pods with words related to keyword
  - e.g., "voyage"
- Shoot down pods with unrelated words
  - e.g., "resist"
- An experimental game platform

reCAPTCHA: CMU

reCAPTCHA™

- WHAT IS reCAPTCHA
- GET reCAPTCHA
- PROTECT YOUR EMAIL
- MY ACCOUNT
- RESOURCES: DOCS & PLUGINS

reCAPTCHA IS A FREE ANTI-BOT SERVICE THAT HELPS DIGITIZE BOOKS.

steamboat train, from New
this morning ran off the track
New-London. Four cars plunged

- LEARN HOW reCAPTCHA WORKS

USE reCAPTCHA ON YOUR SITE

- STRONG SECURITY
- ACCESSIBLE TO BLIND USERS
- 30+ MILLION SERVED DAILY

NEW See how accurate reCAPTCHA is at digitizing content!

ChaCha Search Engine: powered by people

- ChaCha is a human-guided search engine.
- Questions are answered real-time in an "ask a smart friend" format.
- One can access ChaCha via its website, text message, or mobile apps.

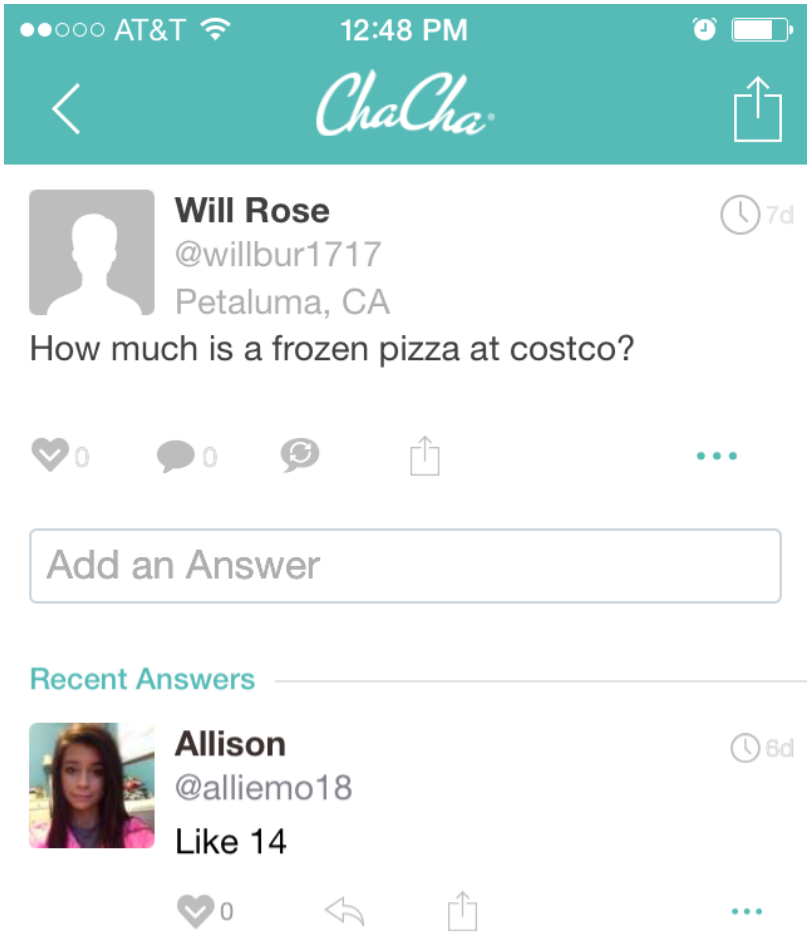

Crowdsensing: Smart City

Traveling safety

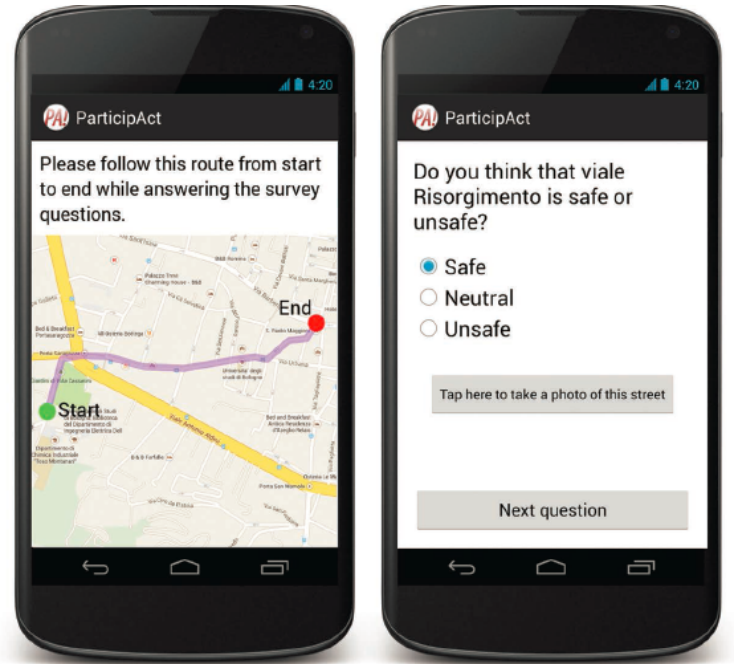

G. Cardone et al, "The ParticipAct Mobile Crowd Sensing Living lab: The Testbed for Smart Cities," IEEE Communications Magazine 2014

Activity census

R. Vaish et al, "Twitch Crowdsourcing: Crowd Contributions in Short Bursts of Time," CHI 2014

Waze Social GPS, Maps & Traffic

- Waze is the world's largest community-based traffic and navigation app.

Smartphone-based Crowdsensing

- Smart city
  - Personalized recommendation

P. Organisciak et al, "A Crowd of Your Own: Crowdsourcing for On-demand Personalization," AAAI 2014

  - Public transportation

Z. He et al, "High Quality Participant Recruitment in Vehicle-based Crowdsourcing using Predicatable Mobility," INFOCOM 2015

  - Indoor map construction

X. Guo et al, "ShopProfiler: Profiling Shops with Crowdsourcing Data," INFOCOM 2014

  - Speaker counting

C. Xu et al, "Crowdsensing the Speaker Count in The Wild: Implications and Applications," IEEE communications magazine 2014

- Some challenges
  - GPS-less (energy efficient)

X. Sheng et al, "Leveraging GPS-less sensing Scheduling for Green Mobile Crowd Sensing," JIoT 2014

  - Trustfulness & Game

Z. Feng et al, "TRAC: Truthful Auction for Location-aware Collaborative Sensing in Mobile Crowdsourcing," INFOCOM 2014

T. Luo et al, "Crowdsourcing with Tullock Contests: a New Perspective," INFOCOM 2015

  - Coverage

M. Karliopoulos et al, "User Recruitment for Mobile Crowdsensing over Opportunistic Networks," INFOCOM 2015

  - Privacy

L. Pournajaf et al, "Spatial Task Assignment for Crowd Sensing with Cloaked Locations," MDM 2014

Sequential

Iterative and Parallel

Divide-and-Conquer and Aggregate

Map and Reduce: a Special Case

Publish/Subscribe

PARADIGMS

Sequential: Collaborative Workflow

- Lexical translation

  (weak bilinguals or machine)
- Assistive translation

  (strong bilinguals)
- Refine sentence

  (monolinguals)

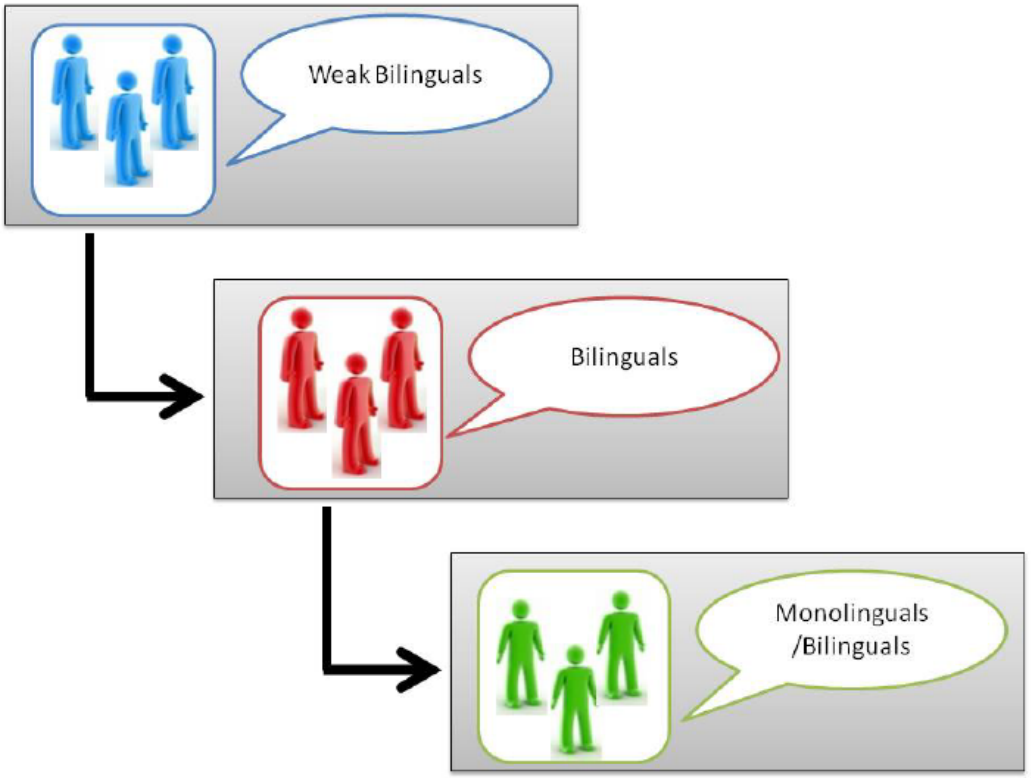

V. Ambati et al, "Collaborative Workflow for Crowdsourcing Translation," CSCW 2012

Iterative and Parallel

- Iterative improve and vote

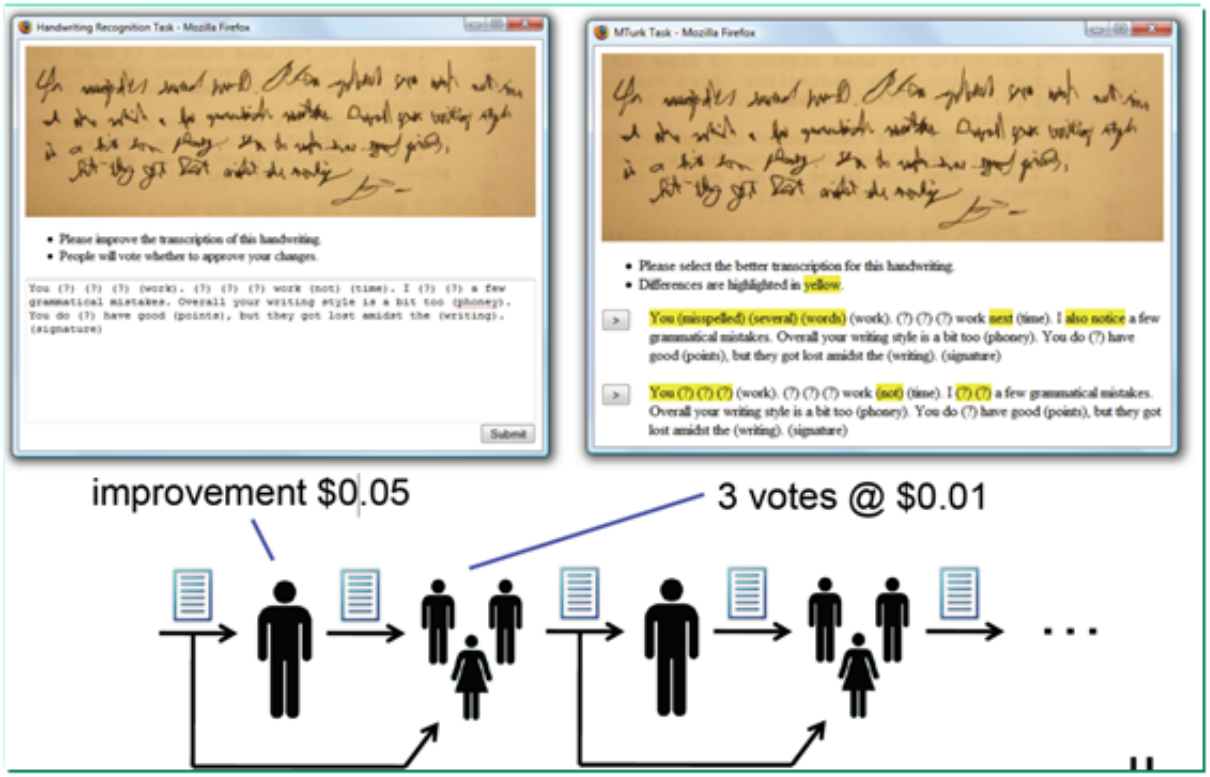

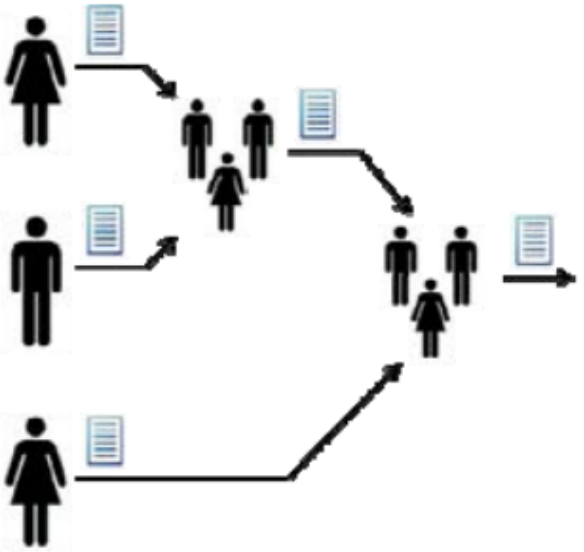

G. Little et al, "Exploring Iterative and Parallel Human Computation Processes," HCOMP 2010

Divide-and-Conquer and Aggregate

- Divide-and-Conquer and Aggregate
  - Decompose a problem statement and aggregate the results

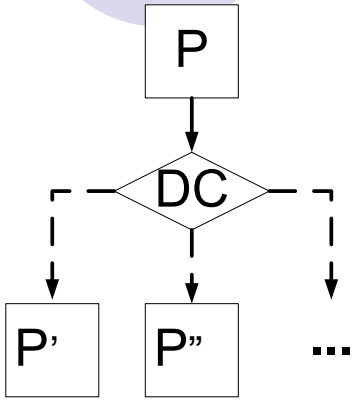

Divide and Conquer

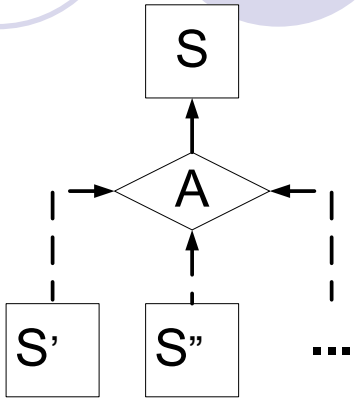

Aggregate

- Two special aggregates
  - Merge
  - Reduce

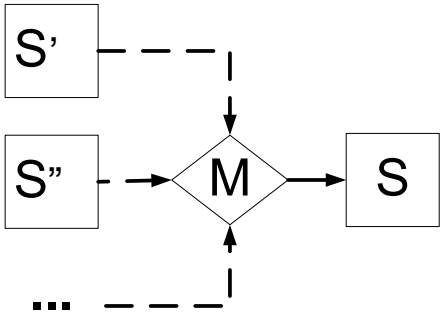

Merge

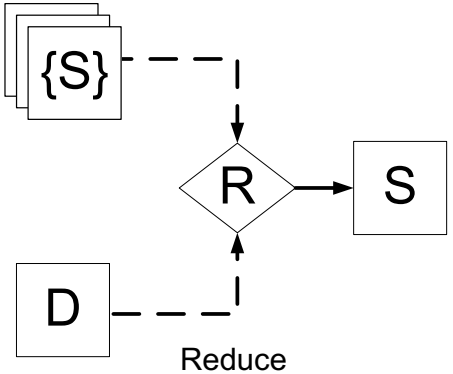

Reduce

P. Minder et al, "Crowdlang - First Steps Towards Programmable Human Computers for General Computation," *AAAI 2011*.

Map and Reduce: A Special Case

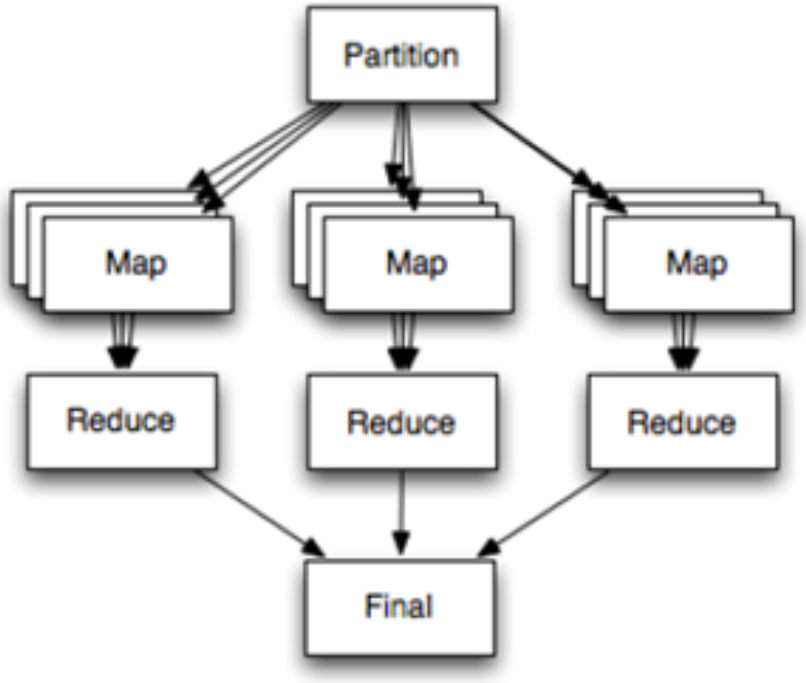

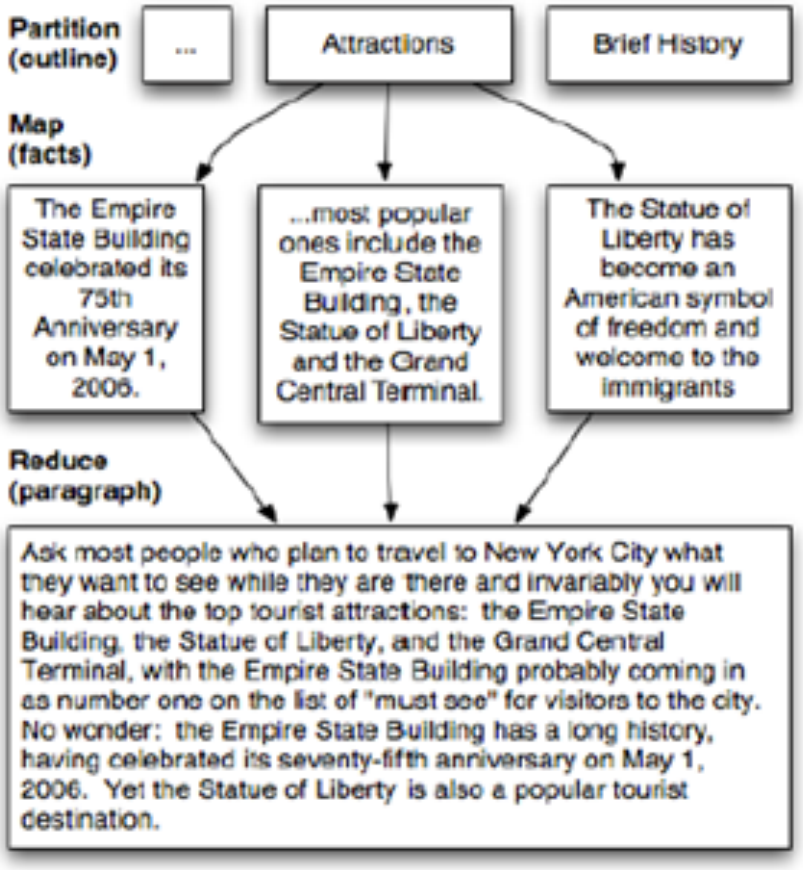

A. Kittur et al, "Crowdforge: Crowdsourcing Complex Work," UIST 2011

Publish and Subscribe

- Pub/Sub middleware-based task assignments

I. Zarko et al, "IoT Data Management Methods and Optimization Algorithms for Mobile Publish/Subscribe Services in Cloud Environments," EU FP 7 OpenIoT Project 2013

Challenges

Opportunities

CHALLENGES AND OPPORTUNITIES

Challenges

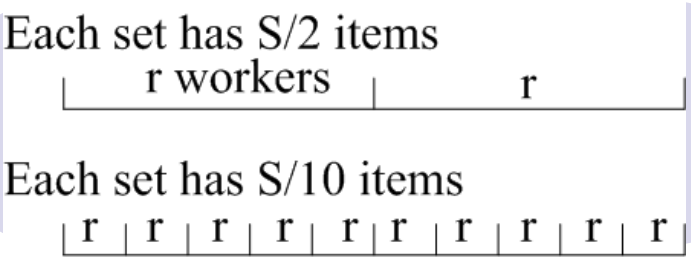

- Trade-offs: time, cost, and quality
  - Max algorithm with human error (with a probability)
  - Maximize quality (via redundancy) subject to cost and time

P. Venetis et al, "Max Algorithms in Crowdsourcing Environments," WWW 2012
R. Kawajiri et al. "Steered Crowdsensing: Incentive Design Towards Quality-oriented Place-centric Crowdsensing", UBICOMP 2014

- Incentive: money, glory, and love
  - Platform-centric: a Stackelberg game
  - User-centric: auction-based incentive mechanism

D. Yang et al, "Crowdsourcing to Smartphones: Incentive Mechanism Design for Mobile Phone Sensing," MobiCom 2012.

Challenges: HPU + CPU

- CrowdDB:

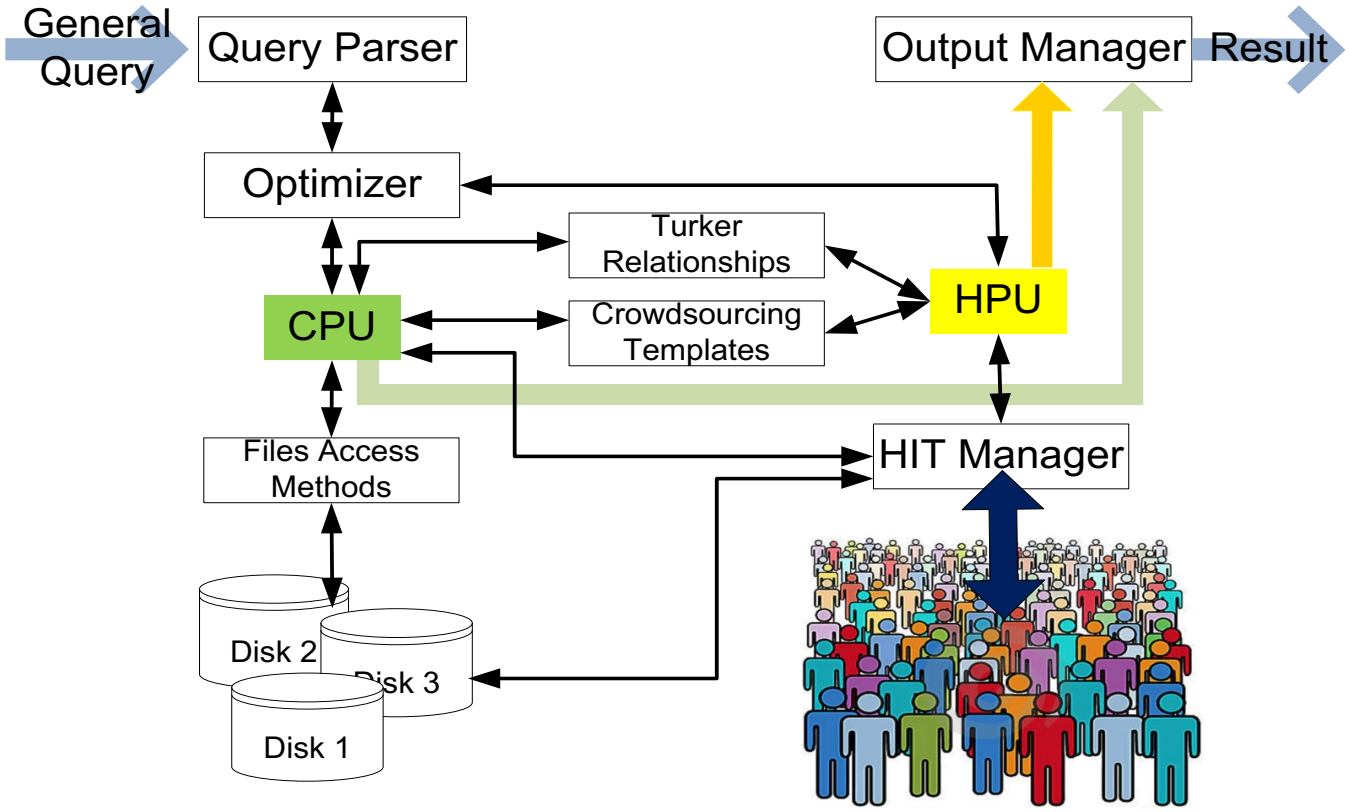

M. Franklin et al, "CrowdDB: Answering Queries with Crowdsourcing," SIGMOD 2011

CPU-assisted HPU

- Entity Matching

  e.g. JHU matches John Hopkins Univ.

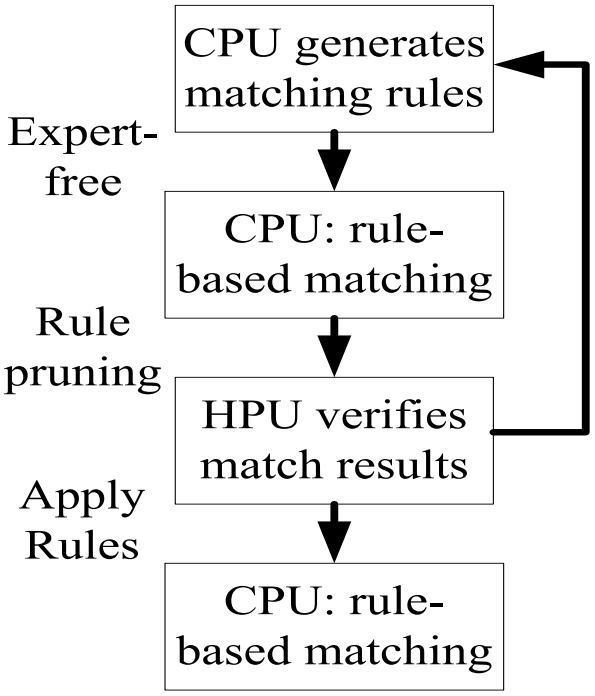

- Radiology
  - Filter

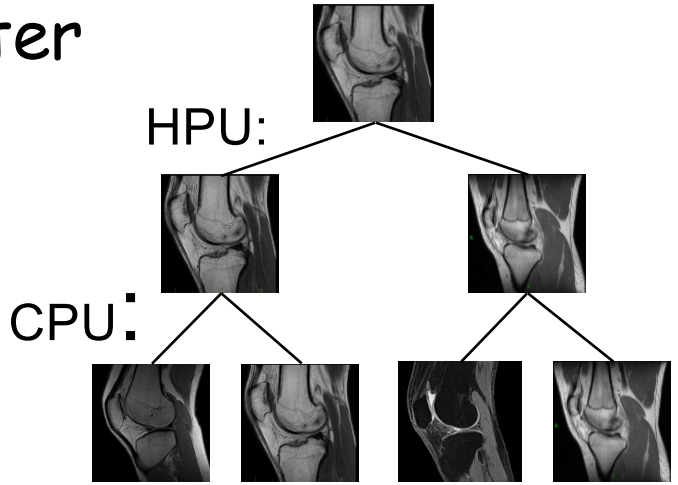

  - Diagnose

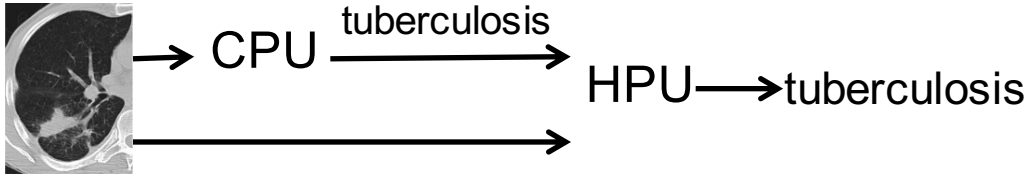

C. Gokhale etal, "Corleone: Hands-off Crowdsourcing for Entity Matching," SIGMOD 2014

Challenges: Collaborative Workflows

Turkomatic

- Complex works require careful and accurate design workflow
- Problems:
  - Loop subtasks
  - Task starvation
  - Multi-stage task with limited budget

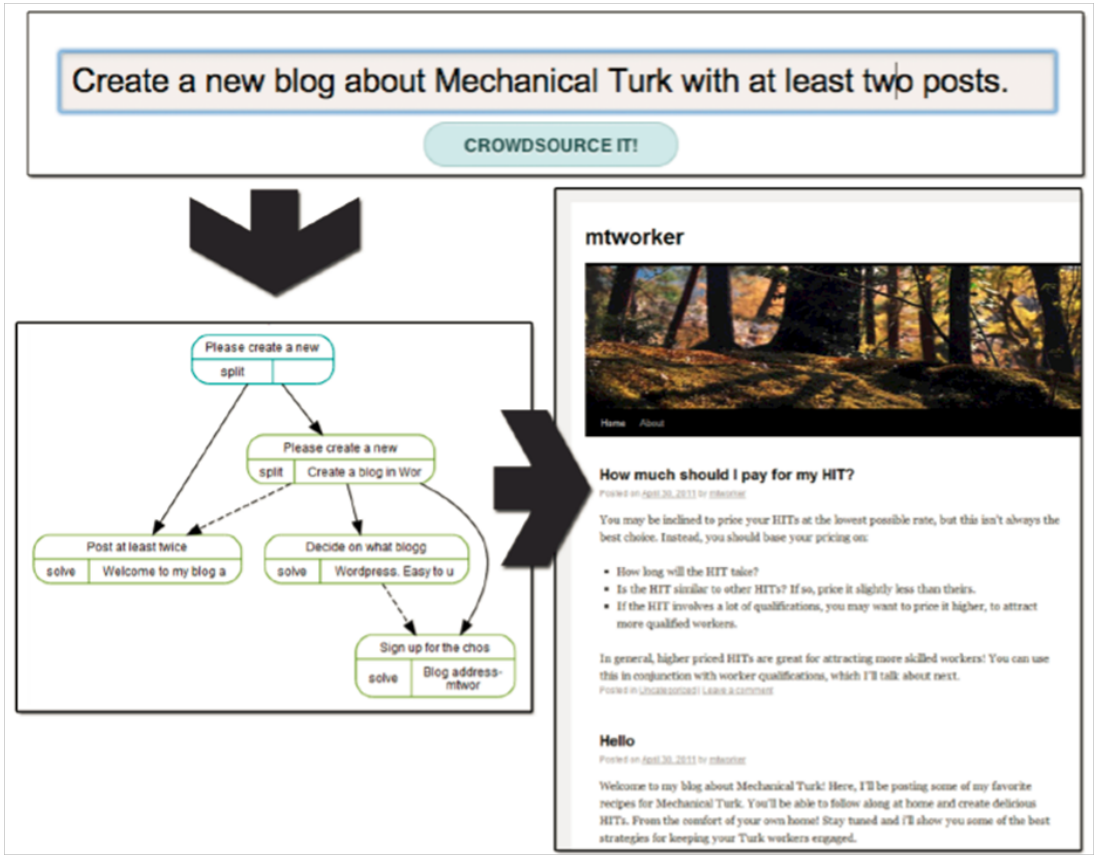

Kulkarni et al, "Collaboratively Crowdsourcing Workflows with Turkomatic," CSCW 2012

C. Fofi et al, "Design Patterns for Hybrid Algorithmic-Crowdsourcing Workflows," CBI 2014

Challenges: Multi-dimensional Data

Multi-dimension

- Personal activity data:
  - Eating hobby
  - Shopping preferences
  - Incomes
  - Emotional state
- Social data:
  - Close friends
  - Similar users
- Environmental data:
  - Locations
  - Climates

Integrate all context

X. Hu et al, "Multidimensional Context-aware Social Network Architecture for Mobile Crowdsensing," IEEE Communications magazine 2014

Opportunities

- Beyond simple workflows
  - Graph search
  - Graph match
- Beyond simple worker selection
  - Dynamic procurement
- Beyond independent workers
  - Social networks

Beyond Simple Workflows

- Graph search
  - Human-assisted graph search
  - Best sequence of questions with simple Y/N answers

  A. Parameswaran et al, "Human-Assisted Graph Search: It's Okay to Ask Questions," VLDB 2010

- Graph match
  - People graph (who knows and/or communicates with whom)
  - Puzzle graph (ideas are compatible and can merge)
  - Natural dynamic for people to merge their compatible ideas

  C. Brummitt et al, "Jigsaw Percolation: What Social Networks Can Collaboratively Solve a Puzzle," 2012

Beyond Simple Worker Selection

Dynamic Procurement (multi-armed bandit)

- A gambler facing a row of slot machines
- Which one to play, how many times, and in which order
- Each machine having a random reward from a fixed distribution
- Objective: maximizing the sum of rewards earned through a sequence of lever pulls

A. Badanidiyuru et al, "Bandits with Knapsacks: Dynamic Procurement for Crowdsourcing," 2013

Beyond Independent Workers

- Social network of workers
- Iterative recruitment of workers through social ties
- Challenges
  - Graph searching
  - Timeliness of responses
  - Stoppage condition

Computational Surplus Around

QQ Example

SOCIAL CROWDSOURCING

Computational Surplus Around

- Friends help friends
  - Fixed individual capability
  - Probabilistic friends' capability
- Makes dissemination decisions
  - Based on the estimations of the fixed and potential computational capacities

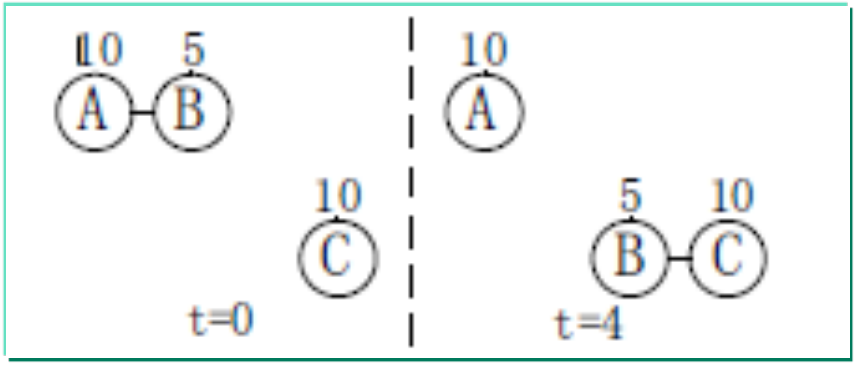

S. Zhang et al, "Minimum Makespan Workload Dissemination in DTNs: Making Full Utilization of Computational Surplus Around ," MobiHoc 2013

Water Filling Schedule

- Response delay
- Computation (by a friend)
- Reply delay

- d_i: response + reply

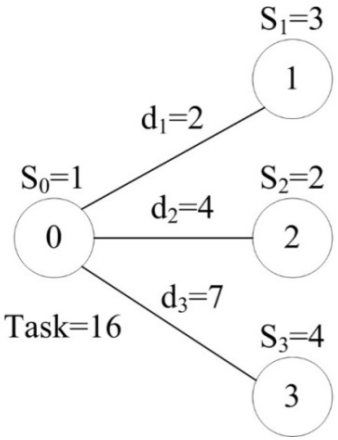

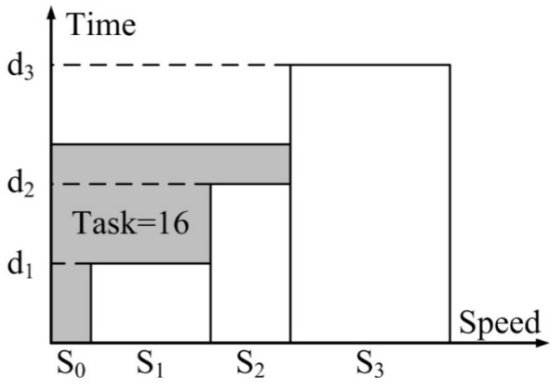

M. Xiao et al, "Multi-task Assignment for Crowdsensing in Mobile Social Networks," INFOCOM 2015

- Scheduling across time: assign jobs to workers

J. Bragg et al, "Parallel Task Routing for Crowdsourcing," AAAI 2014

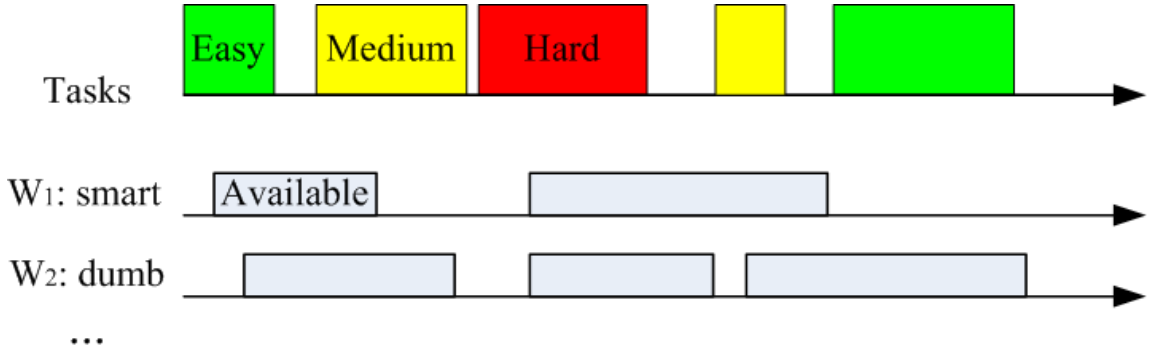

QQ Example

- Tencent QQ, or QQ
  - Instant messaging
- As of March 2013
  - 798.2 million active QQ accounts
  - Peak of 176.4 million simultaneous online users
- QQ experiment
  - Exploring social status of QQ users by responses

Recursive Doubling (reduce)

- Initial label is L = "2" (subtract L by 1 when forwarding this request to QQ friends)

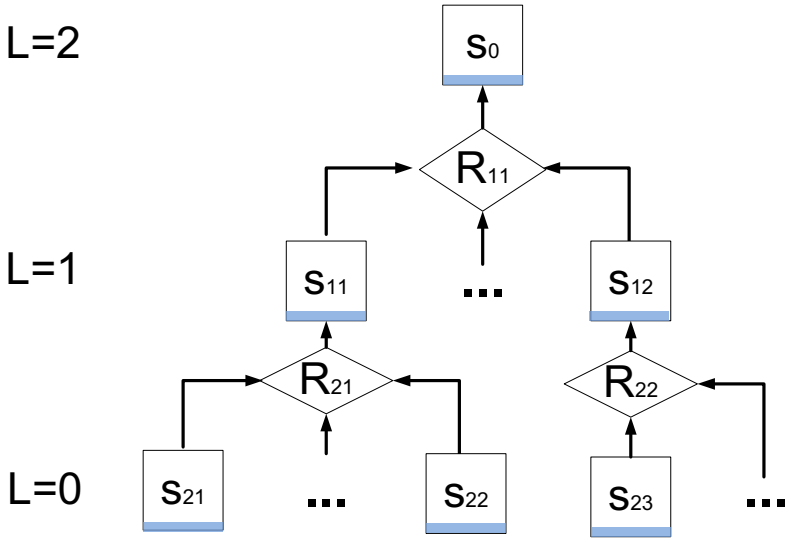

- When L = 0, return the total number of QQ friends
- When L > 0, do the following:
  - Forward this request to all QQ friends
  - After receiving the first 10 replies, compute the average number of friends, and send them back to me

Recursive Doubling (merge)

- Initial label is L = "2" (subtract L by 1 when forwarding this request to QQ friends)
- When L = 0, return the following:
  - Basic information (B)
  - Number of friends (N)
  - Timestamps (T)
- When L > 0, do the following:
  - Forward this request to all QQ friends
  - Pack the first 10 replies, together with your own information (B, N, T), and send them back to me

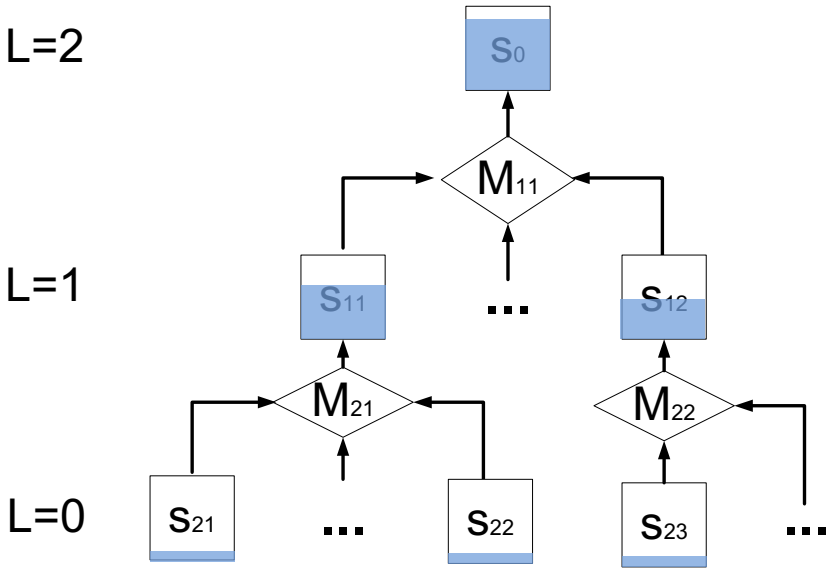

Summary

Acknowledgements

CONCLUSION

Summary

- HPU as a new paradigm to compliment the traditional CPU-based computing for big data
- Many unexplored algorithmic problems
  - Worker selection
  - Social connections of workers
  - Workflow design
  - Cost-time-quality trade-offs
  - Incentive mechanisms

Acknowledgements

- Wei Chang

  Temple University

- Grace Ju

  Carnegie Mellon University

- Wenjun Jiang and Ying Dai

  Hunan University and LinkedIn