

Hybrid Human-Machine Information Systems

Gianluca Demartini
gianlucademartini.net

Gianluca Demartini



- B.Sc., M.Sc. at U. of Udine, Italy
- Ph.D. at U. of Hannover, Germany
 - Entity Retrieval
- Worked at the eXascale Infolab U. Fribourg (Switzerland), UC Berkeley (on Crowdsourcing), Yahoo! (Spain), L3S Research Center (Germany)
- Senior Lecturer in Data Science at the iSchool, U. of Sheffield
- Tutorials on Entity Search at ECIR 2012 and RuSSIR 2015, on Crowdsourcing at ESWC 2013, ISWC 2013, ICWSM 2016, WebSci 2016, Facebook

g.demartini@sheffield.ac.uk

www.gianlucademartini.net

Research Interests

- **Entity-centric Information Access (2005-now)**
 - Structured/Unstruct data (SIGIR 12), TRank (ISWC 13, WSemJ 16)
 - NER in Scientific Docs (WWW 14), Prepositions (CIKM 14)
 - IR Evaluation (ECIR 16 Best Paper Award, IRJ 2015)
- **Hybrid Human-Machine Systems (2012-now)**
 - ZenCrowd (WWW 12, VLDBJ), CrowdQ (CIDR 13)
 - Human Memory based Systems (WWW 14, PVLDB)
 - Hybrid systems overview (COMNET, 2015)
- **Better Crowdsourcing Platforms (2013-now)**
 - Platform Dynamics (WWW 15)
 - Pick-a-Crowd (WWW 13), Malicious Workers (CHI 15)
 - Scale-up Crowdsourcing (HCOMP 14), Scheduling (WWW 16)

Hybrid Human-Machine Info Systems


- Use of **paid micro-task crowdsourcing** (i.e., Amazon MTurk)
- To build information systems that leverage
 - the power of machines to **scale over large amounts of data**
 - the ability of **human intelligence** to keep the quality of data processing high

5-year Analysis of MTurk workload

- Mturk-tracker.com
 - Collects metadata about each visible **batch** (Title, description, rewards, required qualifications, HITs available, requester, etc), that is, set of similar tasks or **HITs**
 - Records batch progress (every ~20 minutes)
 - Covers 130M tasks
 - 2009-2014

Djellel Eddine Difallah, Michele Catasta, Gianluca Demartini, Panagiotis G. Ipeirotis, and Philippe Cudré-Mauroux. **The Dynamics of Micro-Task Crowdsourcing -- The Case of Amazon MTurk**. In: 24th International Conference on World Wide Web (**WWW 2015**), Research Track.

HIT Pricing



5-cents is the new
1-cent

#mturkdynamics

- HIT reward has increased over time
- **Audio transcription** is the most popular task
- Demand for Indian workers has decreased
- **Surveys** are most popular for US workers
- 1000 new requesters per month join
- 10K new HITs arrive and 7.5K HITs get completed every hour

- Check #mturkdynamics for the main findings

Facebook Buys Instagram for \$1 Billion

BY EVELYN M. RUSLI

2:02 p.m. | Updated

Facebook is not waiting for its initial public offering to make its first big purchase.

In its largest acquisition to date, the social network has purchased Instagram the popular photo-sharing application, for about \$1 billion in cash and stock, the company said Monday.



<http://dbpedia.org/resource/Facebook>

<http://dbpedia.org/resource/Instagram>

owl:sameAs

fbase:Instagram

HTML:

<p>Facebook is not waiting for its initial public offering to make its first big purchase.</p><p>In its largest acquisition to date, the social network has purchased Instagram, the popular photo-sharing application, for about \$1 billion in cash and stock, the company said Monday.</p>

RDFa enrichment

<p><cite e property="rdfs:label">Facebook</cite> is not waiting for its initial public offering to make its first big purchase.</p><p>In its largest acquisition to date, the social network has purchased <cite property="rdfs:label">Instagram</cite>, the popular photo-sharing application, for about \$1 billion in cash and stock, the company said Monday.</p>


CNET > News > Mobile

Instagram for Android is now available

At long last, Instagram finally releases the Android version of its app.

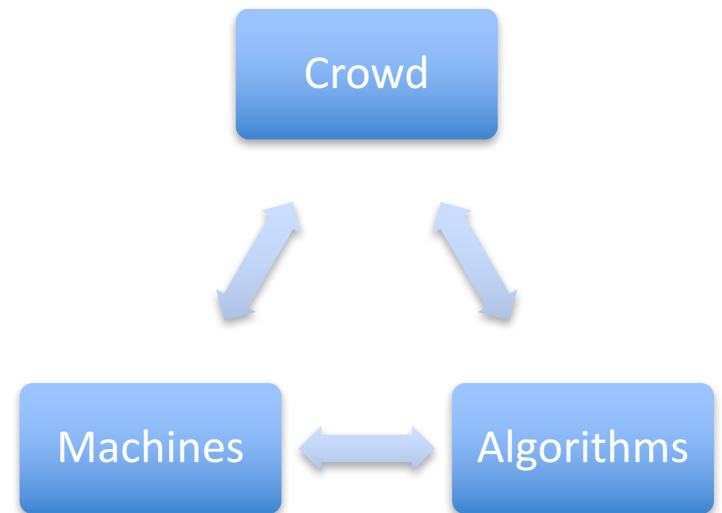
by Jason Cipriani | April 3, 2012 10:07 AM PDT

Instagram has been around since 2010, available only to iOS devices. Android users have been waiting patiently, with repeated promises of an Android version arriving soon.

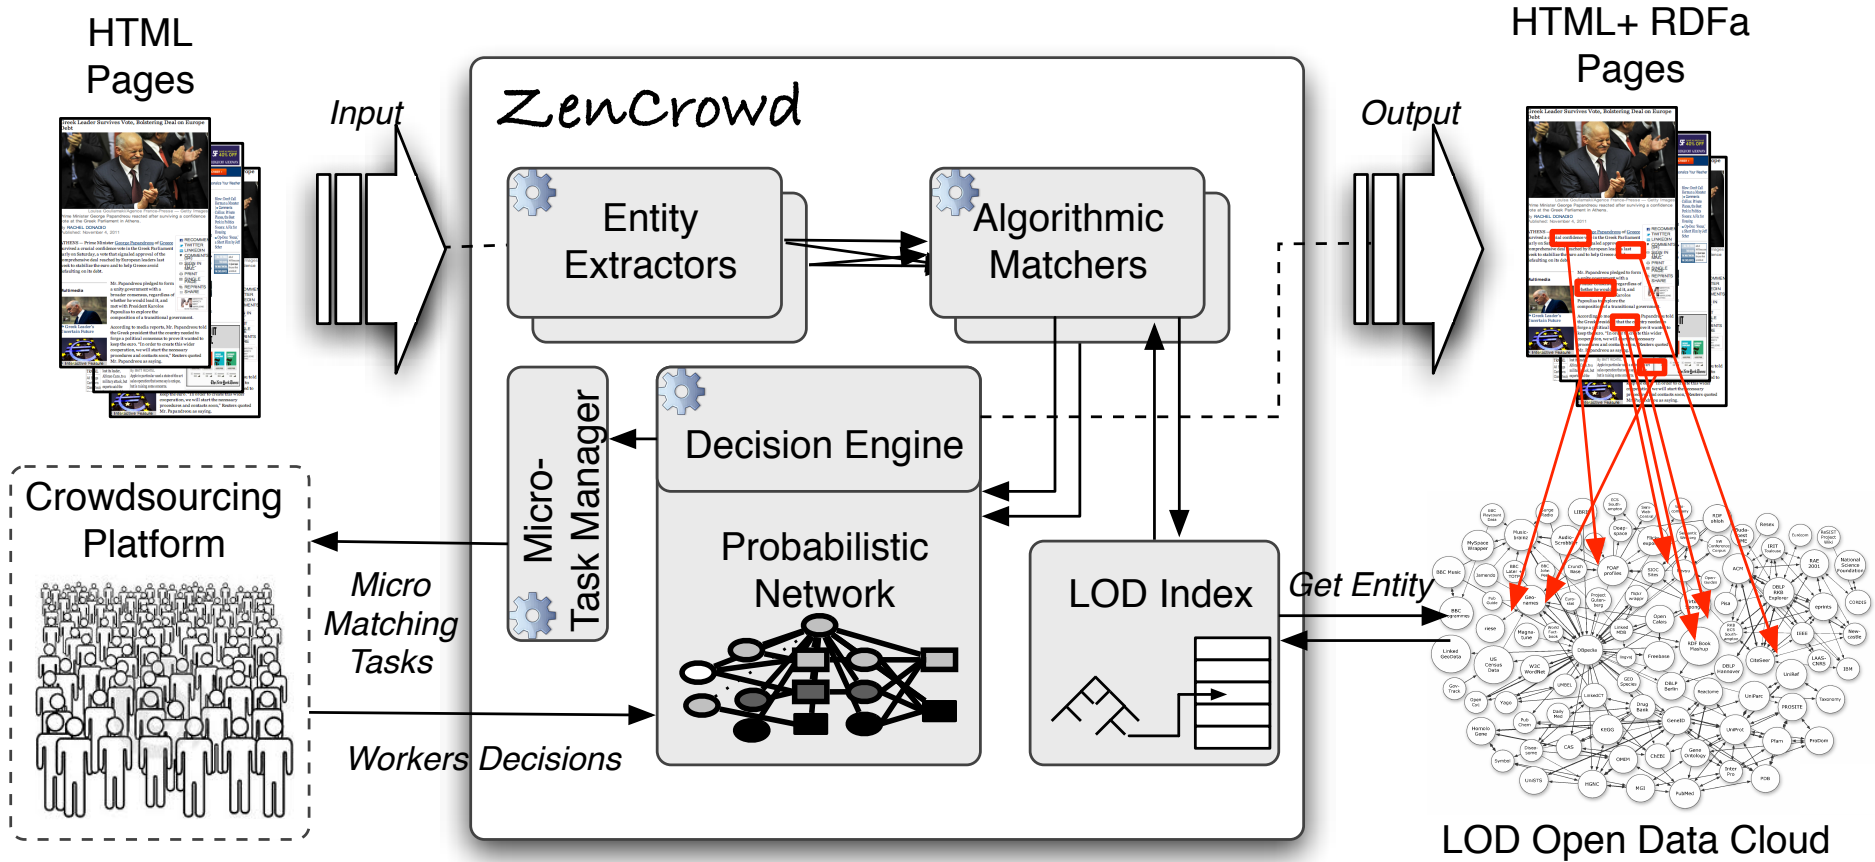


ZenCrowd

- Combine both algorithmic and manual linking
- Automate manual linking via crowdsourcing
- Dynamically assess human workers with a probabilistic reasoning framework



ZenCrowd Architecture




Gianluca Demartini, Djellel Eddine Difallah, and Philippe Cudré-Mauroux. ZenCrowd: Leveraging Probabilistic Reasoning and Crowdsourcing Techniques for Large-Scale Entity Linking. In: 21st International Conference on World Wide Web (**WWW 2012**).

Lessons Learnt

- Crowdsourcing + Prob reasoning works!
- But
 - Different worker communities perform differently
 - Many low quality workers
 - Completion time may vary (based on reward)
- Need to **find the right workers** for your task (see WWW2013 and CHI2015 papers)
- Need to make sure **high priority tasks** are completed fast (see WWW2016 paper)

My customized list of batches:

Batch description	Challenge	Number of tasks	Reward
Football players identifications	Recommend	5	Completed \$0.25
What movie is this scene from?	Recommend	9	31 available \$0.25
Comics, mangas and characters	Recommend	5	41 available For Fun



Number of tasks	Reward
10 available	\$0.25
31 available	\$0.25
18 available	\$0.25
11 available	\$0.25

Behavioral Patterns of Malicious Workers

Ineligible
Workers (IW)

Instruction: Please attempt this microtask ONLY IF you have successfully completed 5 microtasks previously.

Response: *'this is my first task'*

Fast Deceivers
(FD)

eg: Copy-pasting same text in response to multiple questions, entering gibberish, etc.

Response: *'What's your task?' , 'adasd', 'fgfgf gsd ljlkj'*

Rule Breakers
(RB)

Instruction: Identify 5 keywords that represent this task (separated by commas).

Response: *'survey, tasks, history' , 'previous task yellow'*

Smart Deceivers
(SD)

Instruction: Identify 5 keywords that represent this task (separated by commas).


Response: *'one, two, three, four, five'*

Gold Standard
Preys (GSP)

These workers abide by the instructions and provide valid responses, but stumble at the gold-standard questions!

Scheduling HITs

- Fair Scheduling
 - Priority of HITs but avoid starvation
 - Assign HITs of the same type (no context switch)



Happy to help

- Take away: “Crowdsourcing works but...”
- I have used crowdsourcing (i.e., **Amazon MTurk**) over the last 6 years
- And worked on improving its effectiveness/efficiency over the last 4 years
- If you have tried / want to try it but
 - Find that you get very bad quality data back
 - Don’t know how to best design the task / reward / instructions
- Let me know!

gianlucademartini.net