
A Fun and Engaging Interface for Crowdsourcing Named Entities

Kara Greenfield, Kelsey Chan, Joseph P. Campbell

4 Oct 2016



This work is sponsored by the Department of the [Air Force and/or other appropriate department(s)] under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the author and are not necessarily endorsed by the United States Government.




Named Entity Recognition (NER)

- **Automatically extract named mentions of entities from natural language text**
- **Ontology of named entity types**
 - **Person**
 - **Organization**
 - **Location**



Don't NER Corpora Already Exist?

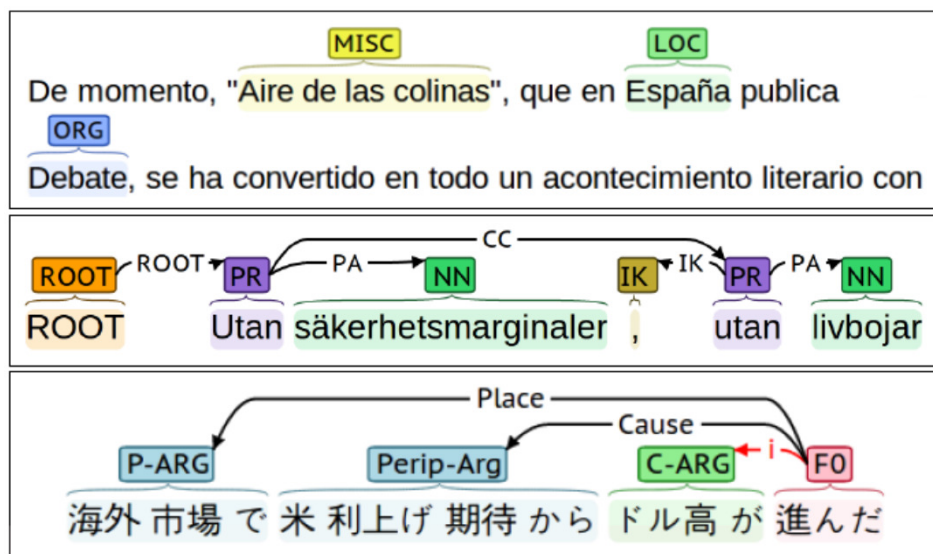
- Domain adaptation is still a huge problem
 - New genre
 - New language
 - New ontology



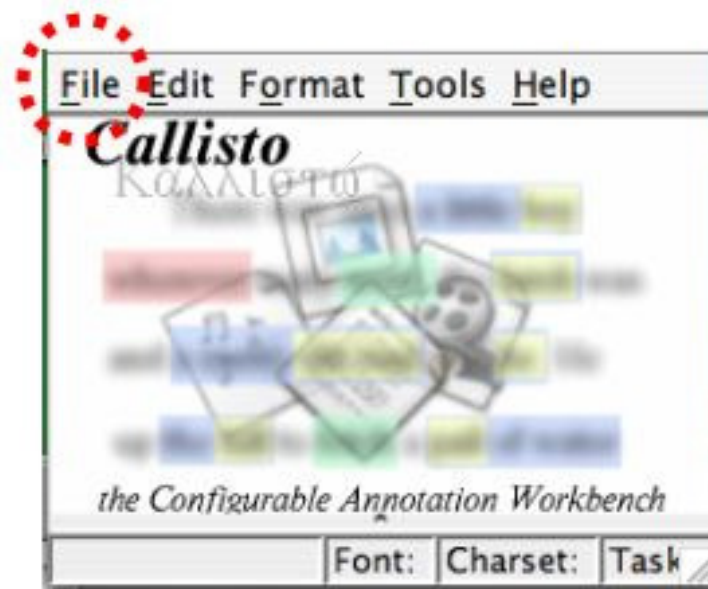


Standard NER Annotation Tools

- Applicable to many tasks, not just NER
- Require linguistic knowledge
- Designed for experts



BRAT Rapid Annotation Tool
(Stenetorp, et al., 2012)



Callisto
(MITRE, 2013)



Early NER Crowdsourcing Interfaces

- Easy for the researcher to create human intelligence tasks
- Difficult for the worker to read the passage
- Limited to short spans of text
- Requires the worker to annotate multiple entity types simultaneously

Twitter NER Annotation System

Timer: 00:00:00 of 10 minutes

Want to work on this HIT? Want to see other HITs?

Label named entities in Twitter data

Requester: [REDACTED] Reward: \$1.00 per HIT HITs Available: 445 Duration: 10 minutes

Qualifications Required: HIT approval rate (%) is not less than 95

on the way to Tomales Bay for a BBQ w/ friends, discussing politics!

Word	Person	Place	Organization	None	???
on	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
the	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
way	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
to	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
Tomales	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
Bay	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
for	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
a	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
BBQ	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
w/	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
friends	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
discussing	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
politics	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>
and	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>	<input checked="" type="radio"/>	<input type="checkbox"/>

Help

An entity is a object in the world like a place or person and a **named entity** is a phrase that uniquely refers to an object by its proper name (Hillary Clinton), acronym (IBM), nickname (Opra) or abbreviation (Minn.). Here are some more examples of named entities for each of the types we are interested in.

PER: Barack Obama, the Palsen, John;
ORG: IBM, Coca-Cola Bottling Co., the Yankees, U.S.;
PLACE: Baltimore, MD; Washington; Mt. Everest; the Hoover dam;

When tagging named entities remember to:

- Tag words according to their **meaning** in the context of the tweet
- Only tag **names**, i.e. words that directly and uniquely refer to entities
- **Do not** tag names of the tweet: PER, ORG, and LOC

(Finin et al, 2010)



Early NER Crowdsourcing Interfaces

- Allows worker to annotate an entire entity mention at once instead of word by word
- Presents document in a more natural way
- Interface helps workers to distinguish named mentions from nominals at the expense of slightly more work

Span-based NER Annotation

• Choose Category ▾
Joe Sutton (remove)

• Choose Category ▾
Kevin Sweeney (remove)

Currently selected:
Fernley

Approve Cancel

EMAIL TEXT:

<Subject>
Helsinki
</Subject>

<Body>

Here is the abridged draft of the document for the presentation to Joe Sutton.

The 'expected loss' calculation in the previous version was intended to show the outcome if we assume some recovery of amounts owing from Vneshtopprom, but until we re-engage in the negotiation we can only speculate as to what this may be. Hence, we have restricted the analysis to the maximum exposure.

Please e-mail Kevin Sweeney and me with any comments/amendments.

Fernley

</Body>

(Lawson et al, 2010)



MITLL NER Annotation Interface

- **Make it easy for the worker to accurately complete the task**
- **Recognize when the worker didn't accurately complete the task**
- **Make the worker want to complete additional HITs**



MITLL NER Annotation Interface

- **Make it easy for the worker to accurately complete the task**
- Recognize when the worker didn't accurately complete the task
- Make the worker want to complete additional HITs



Help the Workers

- Minimize the effort that a worker must go through to annotate a document
- Minimize the mental burden on a worker

Please select **all PERSON entity mentions** in the following excerpt.

Instructions - click to show

Auto-selection Help - click to show

Person Entities

Elon Reeve Musk
Musk
Maye
Errol Musk
Musk

Remove from list

Currently Selected

Add to list

Text

Elon Reeve Musk (born June 28, 1971) is a South African-born, Canadian-American business magnate, engineer, inventor and investor. He is the CEO and CTO of SpaceX, CEO and product architect of Tesla Motors, and chairman of SolarCity. He is the founder of SpaceX and a cofounder of Zip2, PayPal, and Tesla Motors. He has also envisioned a conceptual high-speed transportation system known as the Hyperloop and has proposed a VTOL supersonic jet aircraft with electric fan propulsion. **Musk** was born June 28, 1971, in Pretoria, Gauteng, South Africa, the son of **Maye**, a Canadian-born model, and **Errol Musk**, a South African-born electromechanical engineer (with whom he has a troubled relationship). His paternal grandmother was British, and he also has Pennsylvania Dutch ancestry. After his parents divorced in 1980, **Musk** lived mostly with his father in locations in South Africa.



Clear Instructions Matter

- **The use of a small set of pilot experiments to tune the examples included in the instructions greatly eliminated requests for clarification in further runs**
- **Quickly responding to worker requests for clarification resulted in workers completing tasks accurately and completing more HITs**



MITLL NER Annotation Interface

- Make it easy for the worker to accurately complete the task
- **Recognize when the worker didn't accurately complete the task**
- Make the worker want to complete additional HITs

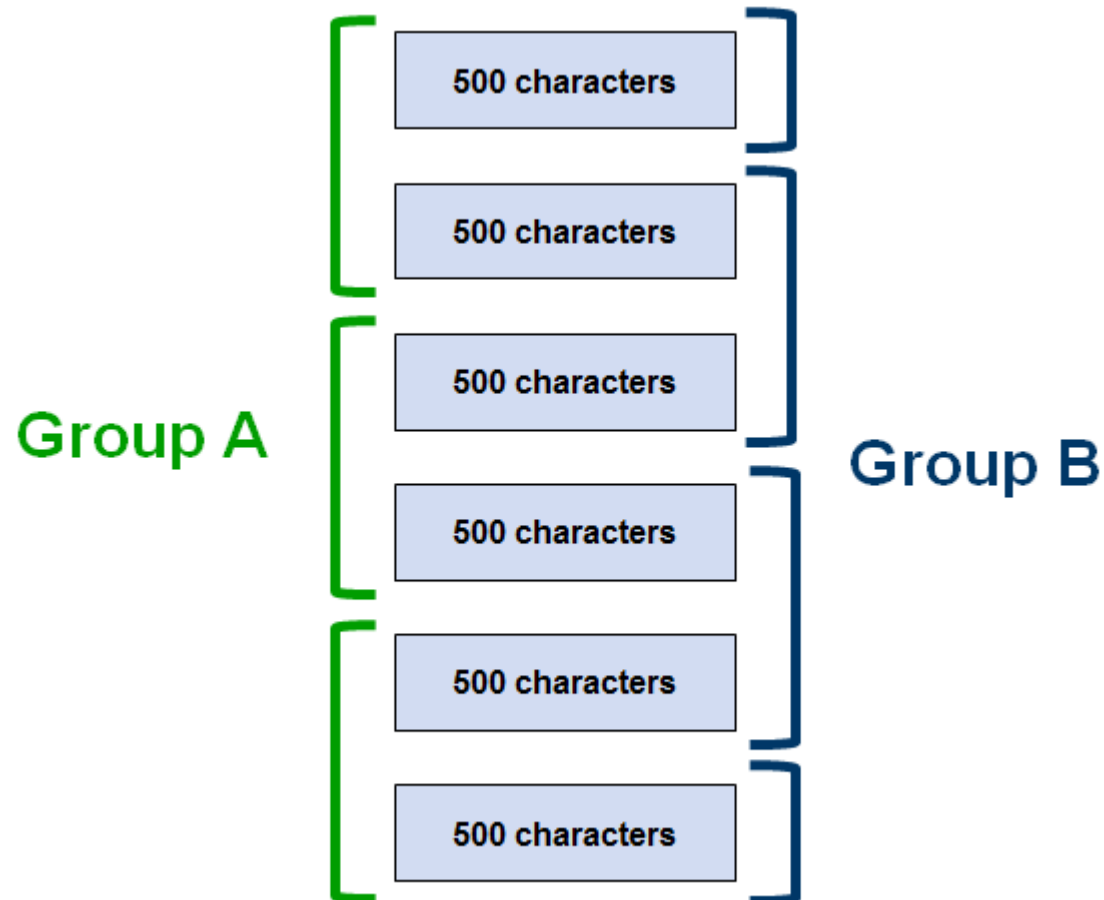


De-incentivizing Spam Workers

- Workers who annotate too little
 - **Financial:** incentive based pay structure where the monetary reward increases with the number of annotations (Lawson et. Al. 2010)
 - **Psychological:**
 - Asking workers who didn't annotate anything on a HIT if they were sure there were no entity mentions
 - Threats of lowered approval ratings
- Workers who annotate too much
 - **Psychological:**
 - Threats of lowered approval ratings



Identifying Worker Fatigue



- Compared the responses of a worker who was less likely to be fatigued with one who was more likely to be fatigued
- Document sizes based on preliminary experimentation with time workers took to complete a HIT and not ending documents mid-sentence



Identifying Worker Fatigue

- **Compare workers who may be fatigued to automated system output using MITIE (MIT Information Extraction system)**
- **Used a third worker to verify that system false positives (as compared to the first 2 workers) were in fact false positives**



MITLL NER Annotation Interface

- Make it easy for the worker to accurately complete the task
- Recognize when the worker didn't accurately complete the task
- **Make the worker want to complete additional HITs**



Easy to Use Interface

- **An easy to use interface can overpower a lack of strong financial incentives**

Morning :) Just some friendly advice :)
I have done about 140 of your hits. I really like the names ones.
I am guessing your account is a new, based on the # of reviews it has on the workers Turkopticon sight. I also noticed that it seems like your batches are not really being worked as fast as you likely hope, and I wanted to offer some advice on that. Though I really enjoy your hits (and the interface I must say is really fantastic! Kudos!), the pay does leave something to be desired.



Conclusions

- **Strongest motivator seems to be an implicit threat of lowered approval ratings**
- **Most workers with high approval ratings didn't suddenly switch to deliberately performing poorly on our task, but did sometimes make honest mistakes due to fatigue or lack of comprehension**
- **An easy to use interface can be a stronger motivational factor than financial incentives**



Future Work

- **More rigorously determine the relationship between financial incentives and ease of use incentives**
- **Develop similar easy to use interfaces for more complicated annotation tasks**



References

- Finin, T. et al., 2010. *Annotating Named Entities in Twitter Data with Crowdsourcing*. s.l., s.n., pp. 80-88.
- Geyer, K., Greenfield, K., Mensch, A. & Simek, O., 2016. *Named Entity Recognition in 140 Characters or Less*. s.l., s.n.
- King, D., n.d. *MITLL/MITIE*. [Online] Available at: <https://github.com/mit-nlp/MITIE>
- Lawson, N., Eustice, K., Perkowitz, M. & Yetisgen-Yildiz, M., 2010. *Annotating Large Email Datasets for Named Entity Recognition with Mechanical Turk*. s.l., s.n., pp. 71-79.
- MITRE, 2013. *Callisto*. [Online] Available at: <http://mitre.github.io/callisto/index.html>
- Nadeau, D. & Sekine, S., 2007. A survey of named entity recognition and classification. *Linguisticae Investigationes*, 30(1), pp. 3-26.
- Poibeau, T. & Kosseim, L., 2001. *Proper Name Extraction from Non-Journalistic Texts*. s.l., s.n.
- Stenetorp, P. et al., 2012. *BRAT: a Web-based Tool for NLP-Assisted Text Annotation*. s.l., s.n., pp. 102-107.