uPick: Crowdsourcing Based Approach to Extract Relations Among Named Entities

Deepti Aggarwal, Rohit Ashok Khot, Vasudeva Varma, Venkatesh Choppella
IIIT Hyderabad, Hyderabad, INDIA 500032
{deepti.aggarwal, rohit_a}@research.iiit.ac.in, {vv, venkatesh.choppella}@iiit.ac.in

ABSTRACT
Despite the advancement in the information extraction area, the task of identifying associated relations among named entities within a text document remains a significant challenge. Existing automated approaches lack human precision and they also struggle to handle erroneous documents. In this paper, we propose a crowdsourcing-based approach to improve the accuracy of the generated relations from the existing extraction techniques. Our idea is to gather judgments on the extracted relations of an article from the interested users. By contributing, the users in return remember the facts related to a document. This paper presents the complete design of the approach along with a user study done with twelve participants. Results show that the users rated the proposed system positively and were willing to contribute their time and energy for the task.

Categories and Subject Descriptors
I.2.6 [Learning]: Knowledge acquisition. H.5.3 [HCI]: Web-based interaction.

General Terms
Design, Experimentation, Human Factors.

Keywords
Crowdsourcing, named entities, information extraction.

1. INTRODUCTION
The World Wide Web (WWW) is rapidly moving away from the search domain and entering into the domain of discovery wherein users are acquiring, analyzing and adapting to the knowledge that exist over the web. However, the rapid growth in the user-generated content (UGC) [27] has complicated the matter and users must now crawl multiple websites on the web to get precise and up-to-date information on the topic they are interested in. The area of Information Extraction [4] strives to minimize the effort by presenting a holistic view of the available information to the user on his selected topic. Named Entity Recognition [21] is one such subtask of the information extraction domain useful in gathering the factual information from a large body of the text. Named Entity is an atomic element in a body of text, associated with a particular type (domain). Common examples of the named entity types are person, location, organization etc. When named entities are linked together, they form a relation. This relation generates the factual information required to grasp the meaning of the sentence (text). Every relation is made up of three parts: Subject-relation-Object, where Subject and Object are the named entities belonging to the same or different types and relationship between them is defined by verb, adjective, adverb etc. The relations thus formulated can then be used to present facts about a particular topic. Moreover, the extracted relations are also useful in building improved Question Answering systems [20].

India HCI'12, April 18–21, 2012, Pune, Maharashtra, India.

To understand named entity and the associated relations, consider the following body of text:

"Sachin Tendulkar was born in Bombay. His mother Rajni worked in the insurance industry, and his father Ramesh Tendulkar, a Marathi novelist, named Tendulkar after his favourite music director, Sachin Dev Burman."

Now, in the first line of the above text, *Sachin Tendulkar* is a named entity of type *Person*, while *Bombay* is a named entity of type *Location.* While the relation '*born in*' associates '*Sachin Tendulkar'* and '*Bombay*' together to form a relation: '*Sachin Tendulkar was born in Bombay*'. Here, 'Sachin' is the subject while 'Mumbai' is the object. All extracted relations from the above body of text are shown in Table 1.

Table 1: All possible relations from the above body of text

Subject (Named Entity)	Relation	Object (Named Entity)
Sachin Tendulkar	born in	Bombay
Sachin Tendulkar	mother	Rajni
Rajni	worked in	Insurance company
Sachin Tendulkar	father	Ramesh Tendulkar
Ramesh Tendulkar	named	Tendulkar
Ramesh Tendulkar	favourite music director	Sachin Dev Burman
Tendulkar	named after	Sachin Dev Burman

Various successful methods exist in the field of information extraction to correctly identify named entities from the corpus for different domains, e.g., medical, newswire domains. [7, 12, 15, 34]. A significant challenge, however, is in identifying associated relations among the identified named entities along with their co-references at the document level. To understand the co-references consider the above body of text, where Sachin Tendulkar is referred various times by different expressions. For example, 'His mother Rajni', refers to Sachin's mother by using a pronoun.

'Ramesh Tendulkar, named Tendulkar' refers to Sachin Tendulkar by using his last name. The reason behind this is that finding named entities and the associated co-references require extraction of the semantic knowledge from a large body of text. Various techniques and tools [12, 15, 34] are developed to address this problem using Natural Language Processing based approaches and other heuristics. However, at present, these techniques are insufficient in handling scenarios where spelling errors and ambiguous acronyms introduce ambiguity in the sentences [18, 24]. An example of ambiguous acronym in the medical domain is, a term RA has the following eight senses: "rheumatoid arthritis", "renal artery", "right atrium", "right atrial", "refractory anemia", "radioactive", "right arm", "rheumatic arthritis".

Techniques to tackle with the problems of spelling errors and ambiguous acronyms seek knowledge from other vocabularies such as WordNet [32], which help in the disambiguation of such words. But such resources introduce words of different contexts (noun, verb, adjectives, adverbs) and make the task tedious to tackle with newly generated words along with their different contexts. For example, a word 'crane' has different meanings, as a noun, it is a bird, and as a verb, it means lift or move. Therefore, using them alone for extraction will allow unpredictable errors associated with the retrieved data and will lead to inaccurate information being communicated to the users (readers). However, the situation could be improved if we utilize human judgments on the extracted relations among named entities. Humans are generally quite proficient in judging the accuracy of the presented information provided they are already familiar with the presented topic. However, humans find the task of filtering (analyzing the accuracy) cumbersome and not particularly engaging [33]. Therefore, we aim to design an immersive environment that motivates a human to be particularly engaged with tasks of analyzing or filtering.

To improve the accuracy of the extracted relations, we present a simple crowdsourcing environment, called as *uPick,* where we utilize the knowledge and expertise of interested users to perform the task. Interested users are readers, editors or frequent browsers of an article. Any person who spends sufficient amount of time on reading a particular article can be termed as an interested user. Our idea is to ask interested users to solve an accompanying challenge once they have finished reading the given article. Each challenge is composed of the relations among the named entities extracted from the article user is reading or has just read. We believe the users will be happy to contribute and participate in the challenge for two reasons: First it will help them to test their understanding of the article they are reading. And second it could also be a good way to remember the important factual information about the article (repeated learning). For fun, we introduce competition through leaderboard, which is a reputation system that measures and displays top scores of a day, to encourage user participation. The responses of the user for a given challenge are then verified against the responses from other users who have played the same challenge before, to obtain the truth-value for the relations within the challenge. We believe if the system becomes popular among the readers then the same approach can be generalized to other Natural Language Processing or Artificial Intelligence related problems.

To evaluate the viability of the proposed approach, we did a user study with 12 participants. Results of the study show that accuracy of the extracted relations has improved by a considerable factor using our proposed approach. For example, for one of the selected documents for user study, we achieved an accuracy of 84% whereas it was only 65% using the automated system. Users appreciated the task performed and were willing to contribute their time and energy for the given tasks. The major shortcoming of our user study was the limited number of users. In future, we are planning to make our system available online to obtain a significant number of users for getting majority votes for the generated relations.

The rest of the paper is organized as follows. Section 2 states the background and related work, reviewing all the techniques related to our work. Section 3 describes the architecture of our proposed scheme, uPick. The experimental design set up and the results are presented in Section 4. The paper is concluded in Section 5, stating all the possible future work of the proposed design.

2. BACKGROUND AND RELATED WORK

Automated techniques used to extract relations among the named entities can be classified based on the usage of the following methods: Natural Language Processing, Machine Learning, Statistical methods and others. Some of the examples of the existing systems using above classified approaches are reviewed in Table 2.

Table 2: Examples of existing systems to extract named entities relations from various automated techniques

Various methods to extract relations	Examples of existing systems
Natural language Processing	Using CRF [15]
Machine Learning	Using HMM [7], using SVM [34], using supervised learning [2, 24]
Statistical methods	Statsnowball [35], using MEM [12]
Others	Context based clustering [9], vocabulary based [28]

Although the automated techniques to extract relations are able to obtain accuracy of 80-90%, but they suffer from the following two drawbacks:

1. **Dependency:** A majority of such systems relies on external reliable sources or vocabularies, such as a pre-defined set of related Named Entities derived from a corpus like Wikipedia, WordNet [32] and MindNet [19, 28]. These resources provide various semantic relations between different words (such as: hypernym/hyponym, hierarchical relationship etc.) and support automatic text analysis for applications like automatic text classification and automatic text summarization. Mostly vocabularies are maintained manually therefore, updating them is costly and requires expertise. Moreover, they are limited in their size and are not specific to any domain; therefore introduce context related noise issues.
2. **Scalability:** Most of these systems are developed for particular domain (medical, protein) or a corpus (structured like Wikipedia, Encarta) in mind. In some cases they also aim to determine only some specific relations [3, 9] e.g., extracting only DATE-and-PLACE-of-EVENT type of relations. Such systems however do not scale well when applied to other domains or corpuses. Therefore, their use remains limited to the domain or corpus for which they were designed.

Therefore, despite of the promising results and improvements proposed by such systems, there remains a significant gap between the system performance and the human intelligence. Since, humans are exceptionally good at not only extracting or finding the relations (e.g., game of protein folding OntoPronto [26]) among the named entities but also in verifying their accuracy. We believe bringing human intelligence into the picture through the means of crowdsourcing [16] can make a system more accurate and precise.

Crowdsourcing is a distributed way of solving a problem online by outsourcing it to a group of people (crowd) with diverse socio and economic backgrounds, through an open call. However, humans are not like computers, they need a stronger incentive to participate in non-engaging tasks like filtering relations. Therefore, crowdsourcing involves a multi-player experience that relies on collective, thoughtful engagement of many online contributors towards solving a problem. The interest in crowdsourcing has flourished recently with the rapid growth of the online users. Researchers are now looking at interesting ways to channel the energy and intelligence of online users to solve various problems over the Web. For example, many recent works [1, 6, 16, 30, 31] in this area have shown that by providing proper incentives, humans can contribute substantially to a system, at different processing levels.

Crowdsourcing applications are divided into four categories [33] (refer Figure 1): 1) Voting system 2) Information sharing system 3) Creative system and 4) Human computational games.

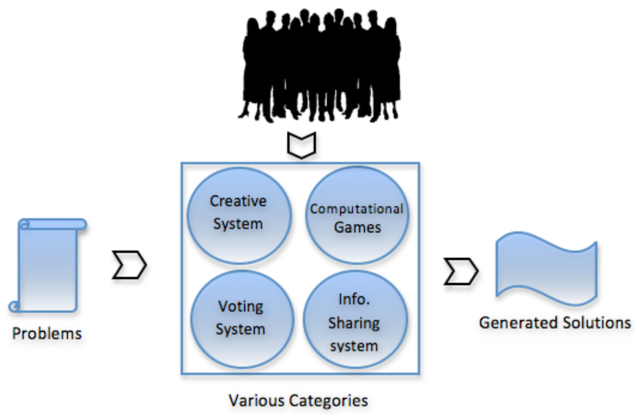

Figure 1:Existing crowdsourcing approaches to generate solutions for an open problem

We briefly discussed Voting system and the human computational games since they are more popular than the other crowdsourcing categories. Detailed survey of all the categories can be found in the survey paper by Yuen *et. al.* [33].

2.1 Voting systems

Voting systems are driven by the monetary incentives where a worker gets paid for his online contribution. The best example of the Voting system is Amazon Mechanical Turk (AMT) [1]. AMT is a financial market place for small tasks such as: labeling images, annotating Named Entities and spam identification etc. [33]. However, studies performed by Mason and Watts [17] concluded that the increased monetary incentives only increased the quantity of the performed work but not its quality. It is believed that with larger financial incentives, there will be more problems associated with validating the task done by the users along with other issues of dissatisfaction [13,14]. We believe that the quality of the work is often better when the human is intrinsically motivated to perform the task.

2.2 Human Computational Games

Human Computational Games [8, 23] provide an interesting way to project a problem in an entertaining game like environment where fun, curiosity and intellectual challenges ensure enriched user experience. Such games are also called as Games with a purpose [8], where the purpose of the game is normally kept hidden from the players. Therefore, people play such games mostly for the sake of entertainment but as a side effect (unknowingly) of playing such games, the underlying tasks get accomplished. The first and the best-known example of human computation game is ESP Game [31], designed to tag images in order to produce better image indexes (hidden task). These indexes are then used to improve the search engines for retrieving the images. Other examples of such games are Peekaboom [23, 30], Verbosity [23], TagATune [23], with the computational purpose to collect database of image related tags, commonsense facts about words and music respectively. The data generated by such games remain at the level of lexical resources, i.e., terms and tags are not connected semantically**.** Recently, Siorpaes *et al.* [26] have reviewed multiple Human Computational Games with the focus on the development of Semantic Web [25], where people contribute to weave the Web with a meaningful linked structure. The idea is to bring human intelligence as part of a game in building of such a formal knowledge structure of Web, which otherwise, cannot be fully automated for certain tasks. Some examples of such tasks include collecting named entities, finding relational hierarchy, phrase detection, finding neighbors in graphs and many others.

However, most of the Human Computational Games are collaborative and competitive in nature [16] and often require minimum of two players to play the game. But multi-player gaming model restricts a single player from playing when there is an absence of other willing partners. To mitigate this problem, Siorpaes et al. [26] has proposed an idea of a single player OntoBay game where the player plays against the previous games' challenges and the past user inputs.

The second challenge with human computational games is clear articulation of the cognitive activities behind the games. For example, some of the human computational games require high cognitive efforts such as: writing descriptive tags, lengthy processing task of watching, listening from the participants, thereby, making the game uninteresting for the users to get involved. We believe the games should involve minimum cognitive efforts from the participants so that the purpose of playing games get exposed to the users (supposed to be hidden) and thus hindering the entertainment value. Moreover, games can also be intellectually stimulating along with being entertaining.

3. OUR SCHEME: UPICK

Our approach is to develop an accurate, scalable and domain-independent system to produce Named Entity relations by introducing humans in the loop that makes a minimal use of other automated resources. Therefore, we design a simple human computational game called *uPick*, where, instead of pairing two or more random users for a game, we propose a single-player environment. We engage the interested users of an article towards filtering of the named entity relations retrieved from the article she is reading. Working of the game is simple. We first extract all the possible relations from the document using an automated technique such as Part-of-Speech (POS) Tagging [22]. The extracted relations are then presented to the users in the form of a challenge for filtering. User then identifies all the valid relations among the presented ones and marks them. Please note that our system is not going through any learning phase because system interference is not required after generation relations. We therefore display all the facts related to a document to every interested. We are willing to take the consensus of each reader for each relation (derived from the document), to define the majority vote corresponding to that relation.

To encourage active participation from the users, we utilize a reputation-based system where contributions of the individual users are listed publically for others to view in terms of the scores. We believe that doing so will initiate competitive nature among the players (interested users) and will motivate them to contribute heavily for the given task. Finally we verify the collected responses by cross checking them with responses from other players. The system architecture is explained in the Figure 2.

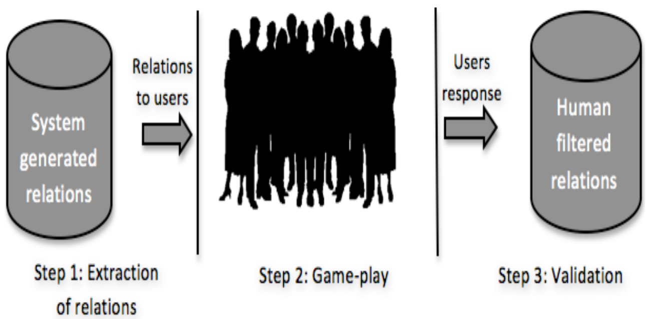

Figure 2: System Architecture of uPick presenting the extracted relations to the users as a challenge

To elaborate the architecture, uPick involves the following steps: 1) Extraction of relations 2) Game Play 3) Validation 4) Scoring. Let us explain them one by one.

3.1 Extraction of relations

We first extract all possible relations among named entities from a document using any of the existing automated technique. In the current prototype, we have used the Stanford Part-of-Speech tagger [27] to determine and tag the basic constructs of English sentences such as noun, verb, adjective, adverb, etc. To extract relations, we utilized eleven relation extraction rules proposed by Chen et al. [5] along with some other heuristics based rules. These rules help to identify a named entity, its relationship type and the corresponding attributes from the tagged constructs. Named entity relations are then generated automatically for a given text. Figure 3 shows the working of the automated techniques used to extract the relations automatically.

Please note that we have made use of corpus independent techniques, namely, POS tagging and rules based on English language structure, which provide very less accuracy by themselves. We believe that our system along with the human iterations will improve the accuracy of the automatically generated relations over the time.

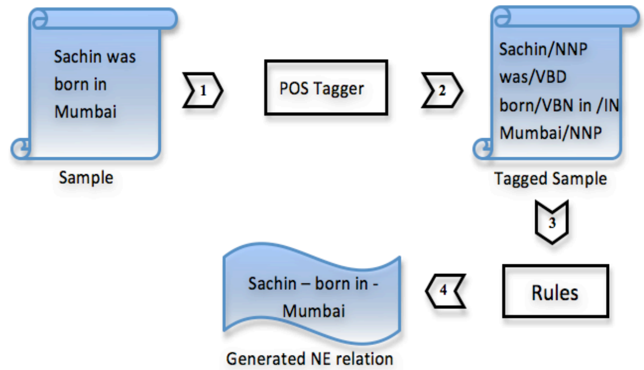

Figure 3: Overview of named entity relations extraction process on a line using the automated techniques

3.2 Game Play

In this game, users are challenged to provide their knowledge of an online document they have just read. The challenge is in the form of a set of questions users are asked to validate. The facts related to the document are presented on a browser interface with the document to read followed by the corresponding named entity relations. The users provide their judgments about the authenticity of the each system-generated relation. The relations that are not selected by the user are considered to be irrelevant and counted as invalid. Detailed steps are explained below.

3.2.1 Steps

Following are the steps involved to play the uPick game:

1. First a user goes to the uPick website and reads a given article.

2. All the relations extracted from the document are presented to the user as a challenge at the bottom of each article. Participation is not compulsory and the user can participate in it if she is interested to play.

3. The user must tick all facts that she thinks are true in relation to the given article. In the current prototype, we present all the extracted facts with check boxes, where the user responds with her judgments.

4. For each judgment, user gets some score based on the majority voting explained in the later sub-section.

Figure 4 shows a snapshot of the uPick game in action where a user is asked to respond with correct facts about Sachin Tendulkar.

Figure 4: The uPick game in action: The player is challenged with a set of questions related to Sachin Tendulkar

In future, we are planning to design better interactive features, such as providing three filters to ask user vote such as True, False, and Don't know. Also, a provision at the end of the game play where the player can compare his performance with his friends and check how well he has performed with respect to them.

3.3 Validation (post processing)

Once the game has been iterated with a significant number of players, we compare the collected responses from each game against the corrected facts stored in the database and filter out erroneous response data. The relation instances having a majority of votes are taken as true facts corresponding to the document. Therefore, retrieving such facts produces the filtered relations associated with the named entities appearing in the document, which are then stored as valid relations in the database.

3.4 Scoring

Our system is based on the majority voting of the interested users on the relations presented to them for a document. For each fact the user judges, a score is awarded to him. This score is based on the correctness of his judgment, which is calculated in two ways:

1. For the first user of uPick, user judgments are compared against the expert corrected set of relations.
2. For all other users, responses of each user are matched with the majority voting, i.e., if majority of the users who have played the game had voted for a particular fact, and the user response matches with majority, then score is awarded. Majority in our game is more than 50%.

4. USER STUDY

We conducted a supervised laboratory study to test the accuracy of our uPick scheme against an automated system. We recruited 12 participants from our university campus by sending invitation e-mails. The average age of the participants was 15 years, the youngest participant being 10 and the oldest being 29. Four of them were male and eight were female. To evaluate the efficiency of our proposed approach, we searched for a sample population who read online. Therefore, we targeted a specific population group of younger people to perform our experiments.

Our usability test consisted of two sessions that span an hour. The first session was dedicated to registration and training. At first, the participants got an introduction to the study then the procedure to play the game was explained.

To perform user testing we selected four articles on Ashok Maurya, Sachin Tendulkar, Shahrukh Khan, and Sonia Gandhi from Wikipedia and named them as D_1, D_2, D_3, and D_4 respectively. All the named entity relations are extracted for these selected documents by using the technique discussed earlier. Table 3 shows the total number of extracted relations for each of the document along with their accuracy verified manually with the help of an expert.

Table 3: Accuracy of the extracted relations from the selected documents set using Automated technique

	D_1	D_2	D_3	D_4
Total number of extracted facts	37	39	40	33
Valid relations among the extracted facts	24	24	23	16
Invalid relations	13	15	17	17
Accuracy (Valid relations / total relations)	65%	61%	57%	49%

Each document gives different number of relations, out of which some are invalid and the remaining are valid. Valid relations are those that are complete, i.e., contain a subject, a relation and an object and convey a correct meaning. We call the relations as invalid, if they are either incomplete (subject or relation or object missing) or do not convey any meaning. Therefore, the accuracy of the relations generated is the ratio of valid relations to the system-generated relations.

The extracted relations from each of the four documents are then used to formulate a challenge (one per document) where we ask the user to verify the authenticity of each relation after she has finished reading the document, i.e., whether the relation holds true or false for the given document.

The second (last) session was dedicated to the actual gameplay. Each participant was given a task of reading two of the four test documents (randomly picked) and then to solve the accompanying challenge. Randomization in document selection was followed for counterbalancing and to minimize the learning effect. Therefore, six different participants in our experiment evaluated each document.

At the end of the study, participants were asked to fill a questionnaire for the qualitative analysis of the proposed system. The questionnaire responses were followed up with a small interview with each participant.

4.1 Results and Discussion

We report our findings in terms of the following:

1) Accuracy (by total number of relations identified correctly)
2) User satisfaction (by users feedback)

4.1.1 Accuracy

For our uPick scheme we measure the accuracy in terms of the total number of relations correctly identified by the participants for the selected set of four documents. Table 4 shows the performance of the participants and the accuracy of the system achieved by filtering the relations with majority votes.

Table 4: Accuracy of uPick scheme considering majority votes of the participants

	D_1	D_2	D_3	D_4
Total number of presented relations	37	39	40	33
Correctly identified valid relations	19	18	19	15
Correctly identified invalid relations	12	12	16	15
Incorrectly identified valid relations as invalid	5	6	4	1
Incorrectly identified invalid relations as valid	1	3	1	2
Accuracy (Correctly identified relations / total relations)	84%	77%	87%	91%

In the table, the total number of relations generated from the automated system contains both valid and invalid relations and this number varies for all the documents. As discussed, for a given challenge users give their judgments by marking the correct relations and leaving the incorrect relations as unmarked. However, it is possible that a user marks an invalid relation as valid or leaves a valid relation unmarked considering it invalid. In both possibilities a user is not able to identify the relation correctly. Therefore, we analyze the accuracy of the uPick system by the total number of correctly filtered facts, both valid and invalid, considering the majority vote of the participants.

In Table 4, we can observe that the responses of the participants were fairly accurate for all the four documents and give an insight that the users are able to perform the task of filtering in a significantly efficient manner. Therefore, we can conclude that the uPick scheme, a system with human intervention, can improve the accuracy of an automated system that is without human support.

4.1.2 User satisfaction

At the end of the user study, we collected oral feedback from every participant about the presented scheme. Seventy-five percent of users (9 out of 12) found the system simple and easy to use and they were willing to contribute their time and energy for such tasks, provided the presented documents are of their personal interest.

In our scheme, the user has full freedom to perform the task in a manner, which please them. We observed that, forty-one percent of the users (5 out of 12) performed the challenge after reading the complete document; thirty-three percent of the users (4 out of 12) preferred to read each paragraph and then performed the challenge; and twenty-five percent of the users (3 out of 12) located the sentences of the given document based on the presented relations and then found out the validity of the relations. When asked how such a system can be helpful to them, ten participants replied that it would help them in remembering facts related to the document concerned. Four users suggested that such a scheme would be helpful to avid users to verify and extend their knowledge in an entertaining way.

Two users didn't appreciate the presentation of the challenge at the end of the reading; instead they wanted a flexible scheme wherein the challenge related to a paragraph of the document is available with any random combination of paragraph. According to them, such randomness will stimulate the task even more, in terms of finding challenge related to each paragraph and will yield more learning environment. However, three participants didn't find the present game design particularly engaging and suggested a few alternative designs as puzzles and object finding games.

4.2 Benefits

Below we mention the two essential benefits of the scheme.

4.2.1 Effectiveness

The game is designed in such a way that it requires minimum human cognitive effort and time. The users only need to give their responses in form of clicks to the facts they find related to the document. Simultaneously, clicking the options rather than writing makes the task easy and interesting. Also, the approach does not depend on any external resource, and can therefore easily scale for any domain and corpus.

4.2.2 Generalization

The proposed approach is independent of the language of the document and to any data corpus. This generalization will only require changes in the rules (dependent on structure of a particular language) to extract relations from the document. Also, the proposed approach can be implemented to validate information of different types, e.g., validate detected anaphora's from a document, sentiments related to a documents and many more by providing a similar crowdsourcing environment.

5. CONCLUSIONS AND FUTURE WORK

This paper proposes a crowdsourcing-based scheme to improve the accuracy of the existing extraction techniques. Our idea is to gather the judgments on the extracted relations of an article (system generated) from the interested readers, and thereby filter out the valid relations from them. At present we were able to test our approach on only a limited number of users and therefore, could not provide proper justification to our proposed hypotheses. We are planning a more extensive study for the next set of evaluations, which we plan to conduct on social networks such as Facebook, to connect many users and to garner their judgments. Furthermore, the idea to provide a competitive environment using leaderboard is not implemented completely. So far, users are provided a comparison of the scores they achieve after playing the challenge, but we will facilitate our scheme with leaderboards in the near future to invite more contribution.

Though our approach provides a challenging environment to the users, the user opinion we gathered pointed out that the fun element related to the associated task was not sufficient. Amongst the factors to persuade human participation, we observed that interestingness of the task is the bigger factor. Therefore, as future work, we plan to extend our system to an interesting game for performing the task of filtering Named Entity relations in a Crowdsourcing environment.

The proposed approach can be extended in various other ways. One possible extension is to provide a Question-Answering system based on individual documents. The answers collected by this approach could be used to answer user queries during a search, for example.

By considering uPick as the learning phase of the system, we can build an interactive environment for learning, where the performance of a user is scored and presented as her report card. This application could be useful in a remote classroom environment.

This approach may also be useful to deal with both temporal (actions and events) and spatial data (location and orientation information). Often, spatial information associated with the named entities derived from the document are more robust than temporal information, which is usually found embedded in new documents. The uPick approach could also be used to correct such temporal anomalies.

6. REFERENCES

[1] Amazon Mechanical Turk. http://amt.com. Last accessed November 2011.

[2] Brin, S. Extracting patterns and relations from World Wide Web. In Proc. of WebDB Workshop at 6th International Conference on Extending Database Technology (WebDB'98),172–183.

[3] Brun, C. and Hagège, C. Semantically-Driven Extraction of Relations between Named Entities. US Patent, Xerox Research Centre Europe, 2008.

[4] Chang, C. H., Kayed, M., Girgis, M. R., and Shaalan. K. 2006. A Survey of Web Information Extraction Systems. IEEE Trans. on Knowl. and Data Eng. 18, 10 (2006), 1411-1428.

[5] Chen, P. English sentence structure and entity relationship diagrams. In Proc. *International Conference on the Entity-Relationship Approach to Systems Analysis and Design, Information Science*, 127-149.

[6] Cusack, C., Largent, J., Alfuth, R. and Klask, K. 2010 Online Games as Social-Computational Systems for Solving NP-complete Problems. In Proc. *Meaningful Play*, 2010.

[7] Freitag, D., and McCallum A. Information extraction with HMM structures learned by stochastic optimization. In Proc. *AAAI 2000 and IAAI-2000*, 584-589.

[8] Games with a purpose http://gwap.com. Last accessed November 2011.

[9] Hasegawa, T., Sekine, S. and Grishman, R. Discovering relations among named entities from large corpora. In Proc. *ACL 2004*, Article 415.

[10] Heer, J. and Bostock, M. Crowdsourcing Graphical Perception: Using Mechanical Turk to Assess Visualization Design. In Proc. *CHI 2010*.

[11] Howe, J. Crowdsourcing: Why the Power of the Crowd Is Driving the Future of Business. *Crown Business, 2008.*

[12] Kambhatla, N. Combining lexical, syntactic and semantic features with Maximum Entropy models for extracting relations. In Proc. *42th Annual Meeting of the Association for Computational Linguistics, 2004.*

[13] Khanna, S., and Davis, J. Evaluating and Improving the Usability of Mechanical Turk for Low-Income Workers in India. In Proc. *Dev 2010.*

[14] Kittur, A., Chi, E. H., and Suh, B. Crowdsourcing User Studies With Mechanical Turk. In Proc. *CHI 2008.*

[15] Lafferty, J., McCallum, A., and Pereira. F. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. In Proc. *18th International Conf. on Machine Learning, 2001*. 282-289.

[16] Law, E., and von Ahn, L. 2011. Human Computation. Synthesis Lectures on Artificial Intelligence and Machine Learning. *Morgan Clay 2011.*

[17] Mason, W., and Watts, D. J. 2009. Financial incentives and the "performance of crowds". In Proc. *HCOMP 2009*. ACM Press, 77-85.

[18] Mcnamee, P., Dang H. T., Simpson, H., Schone, P., and Strassel, S. M. 2010. An Evaluation of Technologies for Knowledge Base Population. In Proc. *LREC 2010.*

[19] MindNet. http://research.microsoft.com/apps/pubs/default.aspx?id=69647. Last accessed March 2012.

[20] Molla, D., Zaanen M. V., and Smith, D. Named Entity Recognition for Question Answering. In Proc. *ALTW, 2006*. 51–58.

[21] Named entity recognition http://en.wikipedia.org/wiki/Named-entity_recognition. Last accessed November 2011.

[22] Part-of-Speech Tagging. http://en.wikipedia.org/wiki/Part-of-speech_tagging. Last accessed March 2012.

[23] Quinn, A. J., and Bederson. B. B. Human Computation: A Survey and Taxonomy of a Growing Field. In Proc. *CHI 2011.*

[24] Sekine S. Named Entity: History and Future. *New York University, 2004.*

[25] Semantic Web. http://en.wikipedia.org/wiki/Semantic_Web. Last accessed March 2012.

[26] Siorpaes, K., and Hepp, M. 2008. Games with a Purpose for the Semantic Web. *IEEE Intelligent Systems 23, 2008.* 50-60.

[27] Stanford POS tagger http://nlp.stanford.edu/software/tagger.shtml. Last accessed March 2012.

[28] Syed, Z., Viegas, E., and Parastatidis, S. Automatic Discovery of Semantic Relations using MindNet. In Proc. *LREC- 2010*.

[29] User Generated Content. http://en.wikipedia.org/wiki/User-generated_content. Last accessed November 2011.

[30] von Ahn, L., Liu, R., and Blum, M. Peekaboom: a game for locating objects in images. In Proc. *CHI 2006,* ACM Press, 55-64.

[31] von Ahn, L., and Dabbish, L. Labeling images with a computer game. In Proc *CHI 2004,* ACM Press, 319-326.

[32] WordNet. http://en.wikipedia.org/wiki/WordNet. Last accessed March 2012.

[33] Yuen, M. C., King, I., and Leung K. S. A Survey of Crowdsourcing Systems. In Proc *CSE 2009*, Vol. 4, IEEE Computer Society, 723-728.

[34] Zhou G., Su, J., Zhang, J., and Zhang, M. Combining Various Knowledge in Relation Extraction. In Proc. *43th Annual Meeting of the Association for Computational Linguistics 2005*.

[35] Zhu, J., Nie, Z., Liu, X., Zhang, B., and Wen, J. StatSnowball: a statistical approach to extracting entity relationships. In Proc. *WWW 2009,* ACM Press, 101-110.