

An Interface for Assisted Curation of Knowledge Bases from Unstructured Text

Jeremy Wright
 Computer Information Systems
 University of the Fraser Valley
 Abbotsford, Canada
 Email: jeremy.wright@student.ufv.ca

Gabriel Murray
 Computer Information Systems
 University of the Fraser Valley
 Abbotsford, Canada
 Email: gabriel.murray@ufv.ca

Ben Hachey
 School of Information Technologies
 University of Sydney
 Sydney, Australia
 Email: ben.hachey@sydney.edu.au

Abstract—We present an interactive graphical tool for assisted curation of knowledge bases from unstructured text data. Given text input, the user can create a knowledge base from scratch, including sub-tasks of entity mention annotation, matching mentions that refer to the same entity, and extracting relations between entities. The interface is designed to enable organizations to extract valuable knowledge from text data that may otherwise remain unexploited. We evaluate the interface through a formative user study. The study results suggest several key directions for refinement. Results also highlight the efficacy of the interface: all participants were able to create a knowledge base from scratch.

Keywords—knowledge base population; entity matching; relation extraction; intelligent interfaces; natural language processing

I. INTRODUCTION

Many organizations possess vast amounts of textual data that is not exploited for decision-making and business intelligence purposes because of its unstructured nature. We create an interactive graphical tool that can facilitate the creation of a knowledge base (KB) from such unstructured textual input, a task known as knowledge base population (KBP). Consider the following example text:

Barack Hussein Obama II is the 44th and current President of the *United States*, and the first African American to hold the office. Born in *Honolulu, Hawaii*, *Obama* is a graduate of *Columbia University* and *Harvard Law School*, where he served as president of the *Harvard Law Review*. He was a community organizer in *Chicago* before earning his law degree. He worked as a civil rights attorney and taught constitutional law at *University of Chicago Law School* from 1992 to 2004. He served three terms representing the *13th District* in the *Illinois Senate* from 1997 to 2004, running unsuccessfully for the *United States House of Representatives* in 2000.

Obama was born on August 4, 1961, at *Kapi'olani Maternity & Gynecological Hospital* (now *Kapi'olani Medical Center for Women and Children*) in *Honolulu, Hawaii*, and would become the first President to have been born in *Hawaii*.

https://en.wikipedia.org/wiki/Barack_Obama

The tool allows the user to identify mentions of entities (e.g., *Barack Hussein Obama II*, *Honolulu*, *Hawaii*, *Illinois Senate*), label the type of each entity (e.g., person, location, organization), determine when multiple mentions are referring to the same entity (e.g., *Barack Hussein Obama II* and *Obama*), identify relations between entities (e.g., *Barack Obama* was born in *Honolulu, Hawaii*), record offsets for entities, and enter entity descriptions (e.g., *Barack Hussein Obama II* is the leader of the U.S.A.).

Assisted manual creation of a knowledge base is particularly valuable, as fully automated techniques for these sub-tasks still suffer from low accuracy [1]. Studies of assisted curation in the biomedical literature demonstrate that it can improve throughput [2], but suggest that user experience should drive design of both user interface and automated KBP techniques [3]. Previous work discusses design principles for annotation tools [4]. However, to our knowledge, no previous work explicitly evaluates interface design for such KBP annotation tools.

The main contributions of this work are as follows:

- We present an interactive graphical interface for creating a knowledge base from unstructured text data.
- We validate the KBP curation tool through a formative user study.
- We are making the software and source code publicly available upon publication.

The formative user study consists of a self-guided cognitive walkthrough in which participants were asked to complete eight tasks. The results are encouraging in that all users were able to successfully create and save a knowledge base from text input. Key areas for improvement include providing better guidance and more explicit feedback. Participants familiar with NLP found the tool particularly easy to use.

The structure of the paper is as follows. Section II describes related work on knowledge base population and document-level annotation. Section III describes our interactive tool and provides several sample screenshots. Section IV describes the structure of the user study, and the results are presented in Section V. In Section VI we describe our planned future work and extension of the interactive tool. Conclusions are presented in Section VII.

II. RELATED WORK

Knowledge base population: KBP builds a repository of formal knowledge by extracting information from text collections [5]. The resulting KB can be used for semantic queries, e.g., retrieving Honolulu, Hawaii for the query “Where was Barack Obama born?”. The US National Institute for Standards and Technology has run shared tasks on KBP as part of the Text Analysis Conference (TAC) since 2009. The primary tasks have been entity discovery and linking (EDL) [6], and slot filling (SF) [1]. EDL has evaluated query-driven and document-level named entity handling – identifying mentions of entities in text, and clustering the mentions that refer to the same entity (e.g. “Barack Obama” and “President Obama”). SF has evaluated query-driven and collection-level relation extraction – identifying all possible facts for a given entity (e.g. the fact that Barack Obama is president of the United States). EDL and SF can be combined for end-to-end KBP.

KBP is part of the larger field of natural language processing (NLP), which aims to enable computers to analyze, represent and generate human languages (as opposed to programming languages or formal mathematical languages). In turn, NLP is part of the larger field of Human-Computer Interaction (HCI), where the medium of interaction is natural human language.

KBP for retrieval: Guo et al. [7] calculate that 71% of web search queries include entities. In addition to building a formal knowledge repository, KBP is also a means for semantic indexing of documents. It allows users to retrieve documents that mention specific entities instead of relying on keyword searches based on ambiguous name strings. Users can review evidence supporting facts in the knowledge base. And, with the right tools, users can correct facts.

Interactive KBP: Research in KBP to date focuses on fully automated techniques [1]. However, accuracies are not high enough to deploy in most settings. In the TAC slot-filling task, for example, inter-annotator agreement is 70% F-score. The best system uses active learning to source additional human annotation, but still scores 37% [8]. This motivates interactive approaches that incorporate intelligent algorithms to streamline knowledge curation with human accuracy. Text mining is used in document curation for real-world biomedical databases [9] and elsewhere [10]. Previous work demonstrates that this can improve throughput for annotation of genes, alleles, and relations [2], selection of documents for curation [11]; and gene indexing [12]. But [3], [13] caution that user experience should drive design of automated KBP techniques as well as the user interface.

Large-scale KBP: In recent years, systems have exploited massive amounts of data to extract more facts with confidence [14]–[16]. Google reports very high coverage and precision when using very large web data [17]. However, these approaches do not generalise well to smaller document

collections that cannot exploit redundancy at the same scale. In interactive KBP, an organization can control the tradeoff between cost and coverage. They can leverage algorithms to inform annotation of rare facts, rather than ignoring them.

Annotation tool design: There are a number of tools for document-level annotation in support of KBP [18]–[20]. Alex et al. [3] briefly summarise interface design issues for assisted curation of protein interactions in biomedical research papers: “what information is displayed to the curator, in what form, and what kind of manipulations can the curator carry out?” Pontus et al. [4] state design principles for the Brat annotation tool:

We believe that intuitive and user-friendly interfaces as well as the judicious application of NLP technology to support, not supplant, human judgements can help maintain the quality of annotations, make annotation more accessible to non-technical users such as subject domain experts, and improve annotation productivity, thus reducing both the human and financial cost of annotation.

However, to our knowledge, no previous work evaluates interaction design or user experience for KBP annotation.

III. DESCRIPTION OF THE INTERACTIVE TOOL

The knowledge management tool presented here has a number of features that make the extraction and management of unstructured text documents easier. The main goal of this interactive tool is to make knowledge base curation as quick and easy as possible by providing intuitive and simple features that will reduce the learning curve significantly. Another objective is to limit human error as much as possible.

There are four important areas of this interface as seen in Figure 1: The “Mentions” tree, the “Entities” tree, the text viewing area and the button bar. The “Mentions” tree is where all annotations from the text are stored for easy access. The tree has four columns which represent the important parts of a mention: ID, Name, Class, and Description. The ID is a unique key for each mention. The numbering starts at 1001 and increments every time a new mention is created. The Name is the non-unique name of the mention; in some cases, there could be multiple mentions of the same entity, and they can all be uniquely stored. The class column contains the class information of the mention. Class can be either PER, for a person, GPE for a geo-political entity, or an ORG for an organization. The last column holds the description, which is unrestricted text provided by a user to give a meaningful description of the entity. Each mention also gets a character offset within the text. This offset information is stored in a dictionary related to each unique key, and when the user clicks on a mention, the text area shows where the mention is from, by highlighting the entity within the document. At the same time, the “Entities” tree will highlight the entity mention matches.

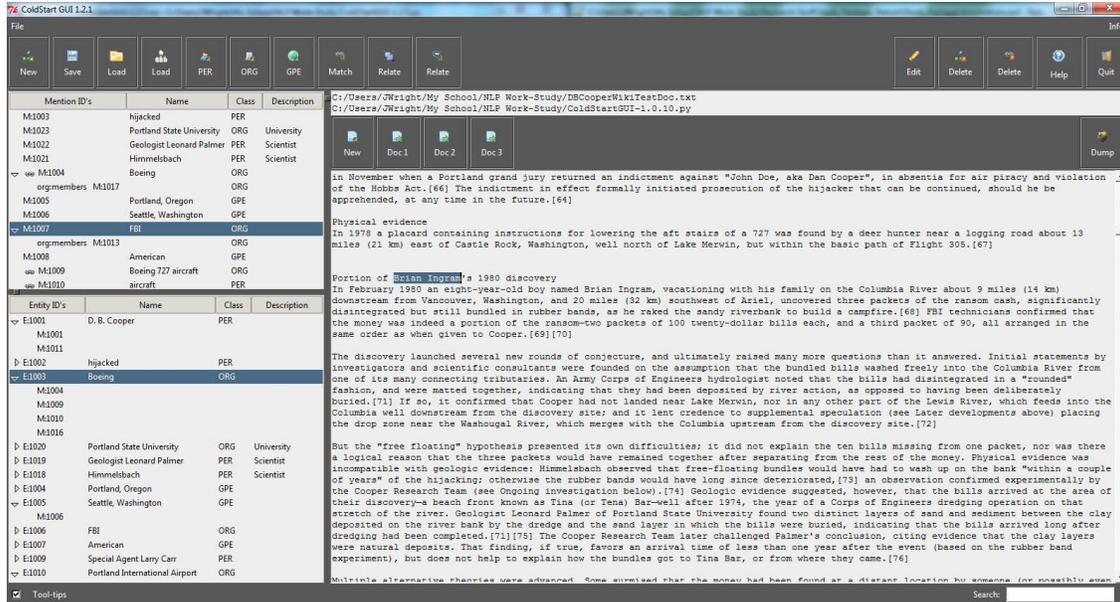


Figure 1. This image shows the four main sections, the “Mentions” tree, the “Entities” tree, the text viewing area and the button bar.

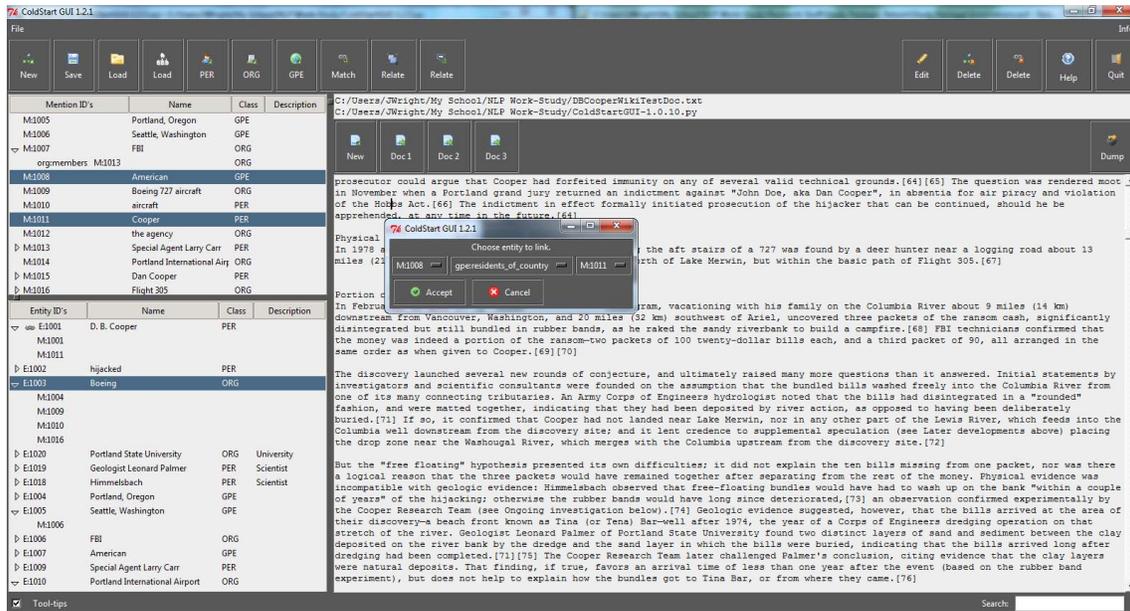


Figure 2. This image shows a *resident* relation being added between mentions American (M:1008) and Cooper (M:1011)

Right-clicking on any mention brings up a tool bar that includes edit and delete options. The edit option allows users to change a value in the mention, so it does not need to be deleted if there is a typographical error. The most important information that is stored in this area is the relationships between mentions. Each mention acts as a parent node for the relationships that are associated with it. Each relationship acts as a child node, and holds the unique ID of the mention that shares the relationship.

The “Entities” tree is very similar in design to the “Mentions” tree, with a few key differences. The “Entities” tree stores unique entities from the text. Each entity has the same columns and options as mentions, but instead of holding relationships as child nodes, it holds unique mention identifiers for each occurrence of the mention. If a user clicks an entity, all of its mentions are highlighted in the “Mentions” tree, and if a user right-clicks, the edit and delete options appear just like the “Mentions” tree.

The text area displays the textual content of the loaded document. This is where the user selects text and either drag and drops it into the “Mentions” or “Entities” trees, or drags and drops it into a class button. This is how the user can create a relation between an entity and a non-entity string. A string can be selected in the text and added to a mention as a title or an age, to give two examples. The current document names are stored above the text. A button appears above the text area for each document that is loaded. These buttons can be used to quickly jump to different areas of the text.

The button bar provides easy access for most of the tasks this interface supports. There are some intuitive functions, such as New, Save, Load, Edit, Delete and Quit, as well as unique features like PER, ORG, GPE, Match and Relate. The PER, ORG, and GPE buttons are for creating new mentions with the chosen class type. If a user selects a name in text, then clicks one of these buttons, they will get a new mention with a unique ID, the name from text, the class from the button pressed, and a description that is manually entered. This also creates a new entity, unless the user chooses to add to existing entity. The match button will take all selected mentions and resolve them to the user-defined entity. Multiple mention selections can be made by holding down CTRL while clicking the mouse button. The relate button will make a relationship between two mentions. The user selects two mentions and presses the relate button. A prompt, as seen in Figure 2, will display asking which relationship to assign the mentions. There is also a help button that displays a list of frequently asked questions that, when clicked, display an answer. This help section can be seen in Figure 3.

The interface was created using Python 2.7.7. The modules used include Tkinter and ttk. Tkinter is a graphical user interface tool kit that includes some widgets such as buttons, list menus, and text areas. The module ttk, (Tkinter tool-kit) is an add on to the base Tkinter that provides some more

advanced features and widgets. Python includes packages such as the Natural Language Tool-Kit (NLTK) that will be helpful in the automatic entity extraction process. Windows 7 was the operating system used for development, but testing was done in both Windows 7 and Ubuntu Linux 12.04.

The eventual goal for this system is to allow for automatic entity and relation extraction. It will also be a useful tool for semi-manual curation. This means users will be assisted by the natural language processing techniques while still extracting entities manually.

IV. USER STUDY

To evaluate the interactive tool, we carried out a formative user study that has aspects of a heuristic evaluation and aspects of a cognitive walkthrough [21]. In a heuristic evaluation [22], participants are typically experts in HCI and they rate how well the interface satisfies particular design principles, known as design *heuristics*. In a cognitive walkthrough [23], each participant is presented with a series of sub-tasks and they determine how easy or difficult it is to accomplish the task with the given interface. This is often done in the presence of the user study coordinator.

Because this was a formative user study rather than a large-scale task-based evaluation, participants were not given extensive tool training or demonstrations of its features. They were given a brief introductory text about the task, and presented with the interface, which includes help functionality. Most participants were experts in NLP and/or HCI. A defining characteristic of heuristic evaluations and cognitive walkthroughs is that they involve consultation with experts who may not be representative of an intended end-user.

Our user study was self-directed, with participants given several sub-tasks and asked questions about the ease or difficulty of accomplishing the task given the provided interface. All of the tasks concerned a short sample text excerpted from President Barack Obama’s Wikipedia biography.

The tasks they were given were:

- **Task 1.** Load the provided text file.
- **Task 2.** Select all mentions of President Barack Obama (e.g. “President Obama,” “Barack Obama”, “Obama”, etc.) and label each of them as PER (a person).
- **Task 3.** Ensure that all mentions of Barack Obama are matched to a single entity.
- **Task 4.** Select all mentions of Honolulu (e.g. “Honolulu” “Honolulu, HI”, “Honolulu, Hawaii”, etc.) and label each of them as GPE (geopolitical).
- **Task 5.** Ensure that all mentions of Honolulu are matched to a single entity.
- **Task 6.** For any portion of the text that describes Barack Obama as being born in Honolulu, add a *city-of-birth* relation between the Obama mention and the Honolulu mention in that portion of text.

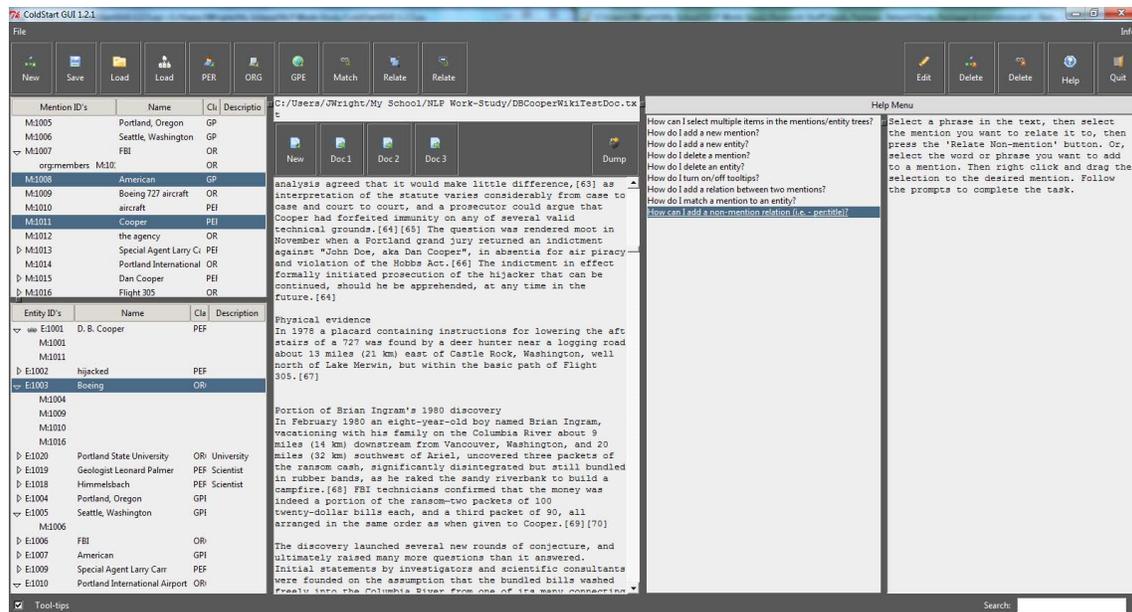


Figure 3. This image shows the help area. Several questions are displayed with their answers.

- **Task 7.** Locate a portion of the text that describes Barack Obama as having the title President of the United States. Add a *title* relation between the mention of Obama and the string “President of the United States.”

- **Task 8.** Save the annotations.

For each of the above tasks, the participants were asked to answer the following questions:

- Is the task clear?
- Does the interface provide appropriate tools for accomplishing the task easily?
- Is the interface missing any tools for accomplishing this task?
 - If so, what?
- Is the system feedback appropriate? For example, were you confident that you had successfully finished the task?
- Do you have any suggestions for making this task easier?

Additionally, they were asked to provide several Likert-scale ratings:

- How easy to use is the user interface, on a scale of 1-7?
- How attractive is the user interface, on a scale of 1-7?
- How efficient are the tasks overall using this interface, on a scale of 1-7?

On the 1-7 scale for each of the criteria above, 1 corresponds to *not at all* while 7 is *very*.

We also asked our participants how much experience they have in NLP and in HCI. Finally, we asked if they had any further open-ended comments or suggestions.

The user study was carried out in two locations: Sydney, Australia and Abbotsford, BC, Canada. Educational levels ranged from first-year undergraduate to doctoral students. Not all participants were native English speakers, but all were very proficient in English.

V. RESULTS

Figure 4 summarizes results. The tasks varied on how easy they were to complete. All participants completed trivial tasks like loading and saving easily, but other tasks were varied in their ease. For example, task 7, described in Section IV and requiring the participant to relate the Barack Obama entity to a string representing his title, was not completed easily by most participants. The challenge is that this is a relationship between an entity and a string rather than a relationship between two entities, and most participants asked for “a more explicit procedure”, or that the interface prompts be provided “more intuitively.” This shows that better feedback is needed and more steering as far as leading users in the correct direction. One participant gave a suggestion about the provided help area: “improve the help menu with a step-by-step description of how to do the tasks.” This shows that users require a more detailed explanation of the task, not just which button they need to press.

Other tasks were completed easily by some and not by others, including task 3, which asked users to “Ensure that

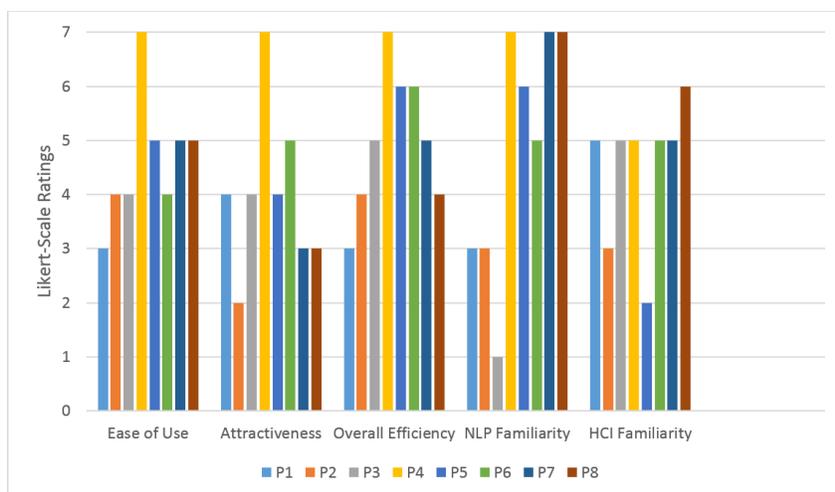


Figure 4. Likert-Scale Ratings by Participant

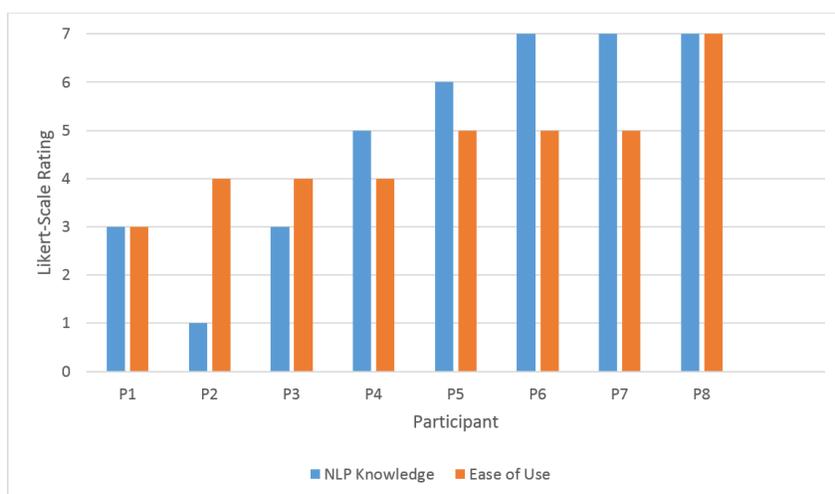


Figure 5. Ease of Use vs knowledge of natural language processing

all mentions of Barack Obama are matched to a single entity.” Some participants could quickly see that the mentions were properly matched to their corresponding entity, and that when they clicked on the entity, a link icon is displayed beside all matched mentions. Other participants could not complete this task, saying simply “I got lost. If there was an easier pathway to do it, that would be great.” This is proof that even with the help button present on the main surface of the interface, not all users will search the help area for information about their task. In the case of task 3, there was a help section labeled “How do I match a mention to an entity?”

Another finding from this study is that a good grasp of natural language processing and the KBP curation task is very helpful when using the system. This can be seen in Figure 5. Many of the participant comments dealt with

the terminology and structure of the tasks and what the task requires. One participant asked the question, “should the President of the United States string be added as an Organization entity before adding the title relation?” If the user had known that an attribute like *age* or *title* is not a relationship between two named entities, then they would have found the “relate string” button. The purpose of this interactive tool is not to teach users about the terminology or the curation of KBP, but instead, it is designed to aid users who are already doing this task.

Most participants wanted more specific system feedback and success prompts, such as “Mention matched successfully to entity” or “New relation created.” This would give users a better sense that they are carrying out the tasks correctly. More specifically, some users were not happy with the number of pop-ups they were interacting with. The challenge

here is to create better feedback, while reducing the number of prompts. One solution that was suggested by participants in the study was the use of drag-and-drop to match mentions and entities and also to relate mentions. Drag and drop reduces the number of times the interface needs to ask for mention or entity choices, and also speeds up the process of matching or relating mentions.

A strong point of the interactive tool is that regardless of the issues that the participants were having, most of them were still able to produce a correct mention and entity tree. This is evident from the saved files that some of the participants made, and the comments from the others. This suggests that the system is robust enough to handle user experimentation and still prevent enough errors that the user eventually will find the correct method without losing work they have already done. The only task that some participants eventually gave up on was task 6, where it asks users to add a *city-of-birth* relation between the Obama mention and the Honolulu mention. Some users never found that to select more than one mention, the user has to hold the CTRL button. This led to comments like, “I tried selecting 2 entities to relate them but it only let me select 1.”

As is evident in Figure 5, the higher the participant’s knowledge of natural language processing, the easier the tasks were to complete.

VI. FUTURE WORK

The results of the user study suggest a number of refinements for the document annotation interface.

Assisted curation: We are currently adapting automatic methods for entity and relationship annotation. We will incorporate these into the user interface as suggestions for human correction, turning our interactive tool into an intelligent user interface (IUI). Rather than annotating everything manually, users will correct output from the machine. This raises new questions about presentation and prioritisation. How should we compute (lack-of) confidence from the machine and convey that to users in the interface? Can machine confidence be used to confidently add some annotations automatically with acceptable precision? Can machine confidence be used to confidently ignore some annotations with acceptable recall? What is the best way to prioritise the remaining machine annotations for human review?

KB-centric curation for cold-start KBP: We will extend the user interface from a document-centric tool for entity linking to a KB-centric tool for building a KB from scratch. This requires new functionality, e.g., processing and browsing multiple documents, prioritising curation tasks across the document collection, presentation and navigation of curation tasks, functionality to apply curation actions across mentions and documents, and KB-update and tool update mechanisms. We will evaluate effectiveness using the cold-start entity discovery and slot filling tasks from

this year’s Text Analysis Conference.¹ Specifically, we will explore the tradeoff between cost, coverage, and precision. We will also benchmark interactive KBP against state-of-the-art methods for automatic KBP.

Task-Based Evaluation: The formative user study has provided guidance for subsequent large-scale task-based evaluations we will carry out. In those evaluations, we will ensure participants are given training time and demonstration of the tool features. They will also be thoroughly versed in the knowledge base population task. In those studies, we will aim to capture more detailed information such as the time taken to complete a task, and the number of errors made while completing a task. We also plan to instrument the interface in order to capture and log information about where the participants are clicking.

Querying: In addition to facilitating the creation of a knowledge base from unstructured text, we want to allow the user to query the knowledge base in order to satisfy an information need. To continue our running example, a sample query may be “Who has the title President of the United States?” Some queries may require the IUI to reason about several relations in combination with one another.

Finally, we will make the interface software and source code available upon publication of this paper.

VII. CONCLUSION

We have described the development of an interactive tool to assist manual creation of a knowledge base from an input consisting of unstructured text. The tool can be used for identifying mentions of entities, labeling their entity type, matching mentions that refer to the same entity, and adding relations or facts about entities. This is a valuable and needed resource for a domain where fully-automated systems still struggle with low accuracy.

The formative user study results were encouraging in that all participants were able to create and save a knowledge base from a sample text input. The participants who were most familiar with NLP found the tool to be appropriate and efficient for the task, indicating that this will be an effective annotation tool for NLP researchers.

Outside of the NLP research community, we foresee this software being useful in business intelligence and enterprise data management. Many organizations possess unstructured data that is not fully exploited. Being able to build a knowledge base from that data, and having the functionality to query that knowledge base, would allow organizations to incorporate the data into their decision-making processes.

ACKNOWLEDGMENT

Ben Hachey is the recipient of an Australian Research Council Discovery Early Career Researcher Award (DE120102900).

¹<http://www.nist.gov/tac/2015/KBP/ColdStart/>

REFERENCES

- [1] M. Surdeanu and H. Ji, "Overview of the English slot filling track at the TAC2014 knowledge base population evaluation," in *Proceedings of the Text Analysis Conference*, 2014.
- [2] N. Karamanis, I. Lewin, R. Seal, R. Drysdale, and E. Briscoe, "Integrating natural language processing with Flybase curation," in *Proceedings of the Pacific Symposium on Biocomputing 12*, 2007, pp. 245–256.
- [3] B. Alex, C. Grover, B. Haddow, M. Kabadjov, E. Klein, M. Matthews, S. Roebuck, R. Tobin, and X. Wang, "Assisted curation: does text mining really help?" in *Proceedings of the Pacific Symposium on Biocomputing 13*, 2008, pp. 556–567.
- [4] P. Stenetorp, S. Pyysalo, G. Topić, T. Ohta, S. Ananiadou, and J. Tsujii, "BRAT: a web-based tool for NLP-assisted text annotation," in *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics*, 2012, pp. 102–107.
- [5] H. Ji and R. Grishman, "Knowledge base population: Successful approaches and challenges," in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, 2011, pp. 1148–1158.
- [6] H. Ji, J. Nothman, and B. Hachey, "Overview of TAC-KBP2014 entity discovery and linking," in *Proceedings of the Text Analysis Conference*, 2014.
- [7] J. Guo, G. Xu, X. Cheng, and H. Li, "Named entity recognition in query," in *Proceedings of the 32nd International Conference on Research and Development in Information Retrieval*, 2009, pp. 267–274.
- [8] G. Angeli, J. Tibshirani, and J. Y. Wu, "Combining distant and partial supervision for relation extraction," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing*, 2014, pp. 1556–1567.
- [9] Z. Lu and L. Hirschmann, "Biocuration workflows and text mining: overview of the BioCreative 2012 Workshop Track II," *Database*, 2012, accessed 15 June 2015 from <http://database.oxfordjournals.org/content/2012/bas043.full>.
- [10] D. McDonald and U. Kelly, "Value and benefits of text mining," Jisc, Tech. Rep., 2012, accessed 15 June 2015 from <http://www.jisc.ac.uk/reports/value-and-benefits-of-text-mining>.
- [11] L. Hirschman, G. A. P. C. Burns, M. Krallinger, C. Arighi, K. B. Cohen, A. Valencia, C. H. Wu, A. Chatr-Aryamontri, K. G. Dowell, E. Huala, A. Lourencço, R. Nash, A.-L. Veuthey, T. Wieggers, and A. G. Winter, "Text mining for the biocuration workflow," *Database*, 2012, accessed 16 June 2015 from <http://database.oxfordjournals.org/content/2012/bas020.full>.
- [12] C.-H. Wei, B. R. Harris, D. Li, T. Z. Berardini, E. Huala, H.-Y. Kao, and Z. Lu, "Accelerating literature curation with text-mining tools: a case study of using PubTator to curate genes in PubMed abstracts," *Database*, 2012, accessed 16 June 2015 from <http://database.oxfordjournals.org/content/2012/bas041.full>.
- [13] C. N. Arighi, B. Carterette, K. B. Cohen, M. Krallinger, W. J. Wilbur, P. Fey, R. Dodson, L. Cooper, C. E. V. Slyke, W. Dahdul, P. Mabee, D. Li, B. Harris, M. Gillespie, S. Jimenez, P. Roberts, L. Matthews, K. Becker, H. Drabkin, S. Bello, L. Licata, A. Chatr-aryamontri, M. L. Schaeffer, J. Park, M. Haendel, K. V. Auken, Y. Li, J. Chan, H.-M. Muller, H. Cui, J. P. Balhoff, J. C.-Y. Wu, Z. Lu, C.-H. Wei, C. O. Tudor, K. Raja, S. Subramani, J. Natarajan, J. M. Cejuela, P. Dubey, and C. Wu, "An overview of the BioCreative 2012 Workshop Track III: interactive text mining task," *Database*, 2012, accessed 16 June 2015 from <http://database.oxfordjournals.org/content/2013/bas056.long>.
- [14] F. Niu, C. Zhang, C. Ré, and J. Shavlik, "Elementary: large-scale knowledge-base construction via machine learning and statistical inference," *International Journal on Semantic Web & Information Systems*, vol. 8, no. 3, pp. 42–73, 2012.
- [15] A. Carlson, J. Betteridge, B. Kisiel, B. Settles, E. R. Hruschka, and T. M. Mitchell, "Toward an architecture for never-ending language learning," in *Proceedings of the 24th AAAI Conference on Artificial Intelligence*, 2010, pp. 1306–1313.
- [16] N. Nakashole, M. Theobald, and G. Weikum, "Scalable knowledge harvesting with high precision and high recall," in *Proceedings of the 4th ACM International Conference on Web Search and Data Mining*, 2011, pp. 227–236.
- [17] X. Dong, E. Gabrilovich, G. Heitz, W. Horn, N. Lao, K. Murphy, T. Strohmann, S. Sun, and W. Zhang, "Knowledge vault: a web-scale approach to probabilistic knowledge fusion," in *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2014, pp. 601–610.
- [18] K. Bontcheva, I. Roberts, L. Dercynski, and D. Rout, "The GATE crowdsourcing plugin: crowdsourcing annotated corpora made easy," in *Proceedings of the Demonstrations at the 14th Conference of the European Chapter of the Association for Computational Linguistics*, 2014, pp. 97–100.
- [19] S. M. Yimam, I. Gurevych, R. E. de Castilho, and C. Biemann, "WebAnno: A flexible, web-based and visually supported system for distributed annotations," in *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 2013, pp. 1–6.
- [20] W. Radford, D. Tse, J. Nothman, B. Hachey, G. Wright, W. C. James R. Curran, T. O'Keefe, M. Honnibal, D. Vadas, and C. Loxley, "The Computable News project: research in the newsroom," in *Proceedings of the WWW Workshop on Web and Data Science for News Publishing*, 2015, pp. 903–908.
- [21] Y. Rogers, H. Sharp, and J. Preece, *Interaction Design*, 3rd ed. West Sussex, UK: Wiley, 2011.
- [22] J. Nielsen, "Heuristic evaluation," *Usability inspection methods*, vol. 17, no. 1, pp. 25–62, 1994.
- [23] ———, "Usability inspection methods," in *Conference companion on Human factors in computing systems*. ACM, 1994, pp. 413–414.