

An Interactive Tool for Natural Language Processing on Clinical Text

Gaurav Trivedi

Intelligent Systems Program
University of Pittsburgh
trivedigaurav@pitt.edu

Rebecca Hwa

Dept. of Computer Science
& Intelligent Systems Program
University of Pittsburgh
hwa@cs.pitt.edu

Phuong Pham

Dept. of Computer Science
University of Pittsburgh
phuongpham@cs.pitt.edu

Janyce Wiebe

Dept. of Computer Science
& Intelligent Systems Program
University of Pittsburgh
wiebe@cs.pitt.edu

Wendy Chapman

Dept. of Biomed. Informatics
University of Utah
wendy.chapman@utah.edu

Harry Hochheiser

Dept. of Biomed. Informatics
& Intelligent Systems Program
University of Pittsburgh
harryh@pitt.edu

ABSTRACT

Natural Language Processing (NLP) systems often make use of machine learning techniques that are unfamiliar to end-users who are interested in analyzing clinical records. Although NLP has been widely used in extracting information from clinical text, current systems generally do not support model revision based on feedback from domain experts.

We present a prototype tool that allows end users to visualize and review the outputs of an NLP system that extracts binary variables from clinical text. Our tool combines multiple visualizations to help the users understand these results and make any necessary corrections, thus forming a feedback loop and helping improve the accuracy of the NLP models. We have tested our prototype in a formative think-aloud user study with clinicians and researchers involved in colonoscopy research. Results from semi-structured interviews and a System Usability Scale (SUS) analysis show that the users are able to quickly start refining NLP models, despite having very little or no experience with machine learning. Observations from these sessions suggest revisions to the interface to better support review workflow and interpretation of results.

Author Keywords

Clinical text analysis; electronic medical records; visualization; interactive machine learning.

ACM Classification Keywords

H.5.2. Graphical User Interface: User Interfaces; H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous; I.2.7 Natural Language Processing: Text analysis.

INTRODUCTION AND BACKGROUND

Electronic Health Records (EHRs) are organized collections of information about individual patients. They are designed

such that they can be shared across different settings for providing health care services. The Institute of Medicine committee on improving the patient record has recognized the importance of using EHRs to inform *decision support systems* and support data-driven quality measures [7]. One of the biggest challenges in achieving this goal is the difficulty of extracting information from large quantities of EHR data stored as unstructured free text. Clinicians often make use of narratives and first-person stories to document interactions, findings and analyses in patient cases [15]. As a result, finding information from these volumes of health care records typically requires the use of NLP techniques to automate the extraction process.

There has been a long history of research in the application of NLP methods in the clinical domain [4]. Researchers have developed models for automatically detecting outbreak of diseases such as influenza [23], identifying adverse drug reactions [12, 17, 13], and measuring the quality of colonoscopy procedures [10], among others. Due to the complexity of clinical text, the accuracy of these techniques may vary [14]. Current tools also lack provision for end-users to inspect NLP outcomes and make corrections that might improve these results. Due to these factors, Chapman et. al. [4] have identified “lack of user-centered development” as one of the barriers in NLP adoption in the clinical domain. There is a need to focus on development of NLP systems that are not only generalizable for use in different tasks but are also usable without excessive dependence on NLP developers. In this paper, we have explored the design of user-interfaces for use by end users (clinicians and clinical researchers) to support the review and annotation of clinical text using natural language processing.

We have developed an interactive web-based tool that facilitates both the review of binary variables extracted from clinical records, and the provision of feedback that can be used to improve the accuracy of NLP models. Our goal is to close the natural language processing gap by providing clinical researchers with highly-usable tools that will facilitate the process of reviewing NLP output, identifying errors in model prediction, and providing feedback that can be used to retrain or extend models to make them more effective.

Paste the appropriate copyright statement here. ACM now supports three different copyright statements:

- ACM copyright: ACM holds the copyright on the work. This is the historical approach.
- License: The author(s) retain copyright, but ACM receives an exclusive publication license.
- Open Access: The author(s) wish to pay for the work to be open access. The additional fee must be paid to ACM.

This text field is large enough to hold the appropriate release statement assuming it is single spaced.

Every submission will be assigned their own unique DOI string to be included here.

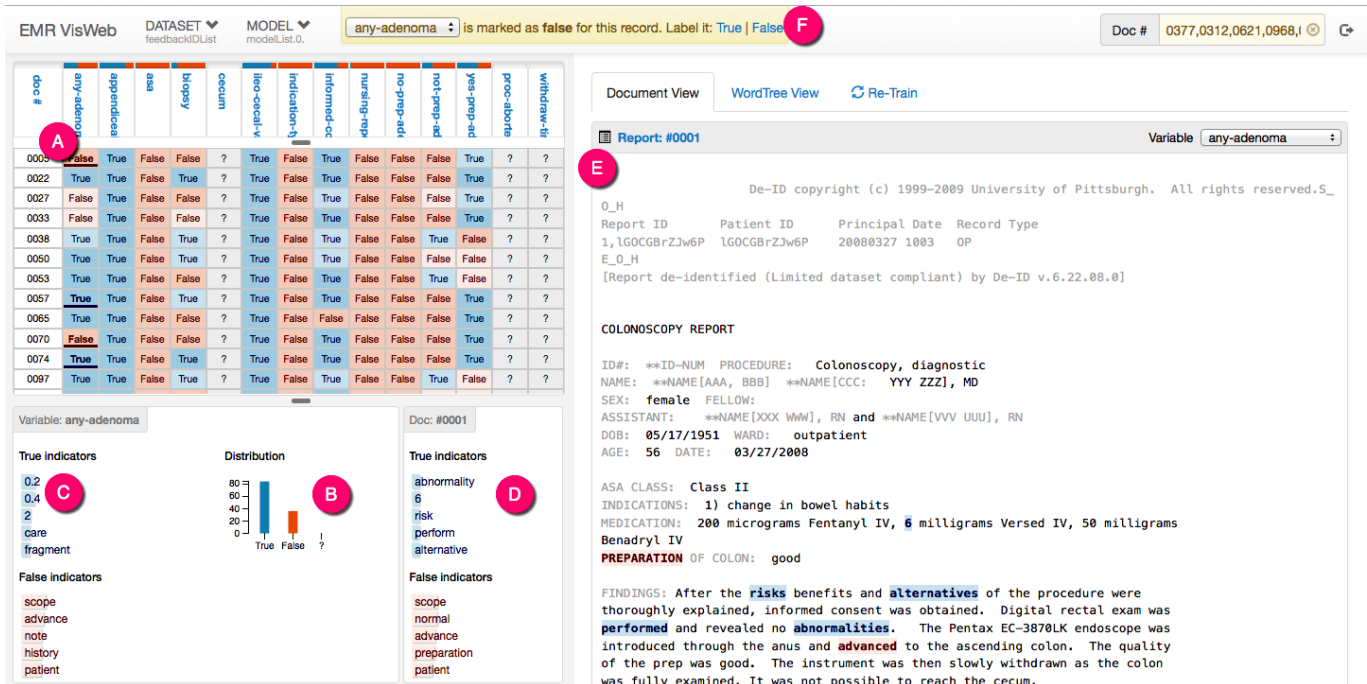


Figure 1. (a) The grid view shows the extracted variables in columns and individual documents in rows, providing an overview of NLP results. Below the grid, we have statistics about the active variable with (b) the distribution of the classifications and (c) the list of top indicators aggregated across all the documents in the dataset. (d) Indicators from the active report are shown on the right. (e) The document view shows the full-text of the patient reports with the indicator terms highlighted. (f) Feedback can be sent using the yellow control bar on the top, or by using the right-click context menu.

RELATED WORK

During the process of developing our interactive text analysis tool for clinical domain, we studied relevant work from multiple research areas spanning across different domains, such as *Visualization*, *Interactive Machine Learning* and *Interface Design*. We have built upon the following work in these areas for the design of our tool.

Visualization and Sensemaking

Visualization tools such as WordTree [21] and Tiara [22] help in providing a visual summary of large amount of text data. While Tiara focuses on content evolution of each topic over time, WordTree provides a keyword in context method of exploring the text. Other tools such as Jigsaw [20] help users interpret document collections by visualizing documents in multiple graph, cluster and list views. Our task in reviewing clinical documents is somewhat different, in that our goals are to understand common textual patterns and to use those patterns to improve NLP models. We have adapted elements of these views - in particular, WordTree's phrase view and Jigsaw's document view to support our goals. The purpose of these visualizations would be to provide both detailed document-level views and also dataset-level overviews.

Interactive Machine Learning

There have been many efforts to develop user-centric tools for machine learning and NLP making it easier for the end users to build models. D'Avolio et. al. [6] have described a prototype that combines several existing tools such as Knowtator [16] for creating text annotations, and cTAKES [18] for deriving NLP features, within a common user interface that can

be used to configure the machine learning algorithms and export their results. Our present work complements this effort, focusing instead on facilitating expert review of NLP results and provision of feedback regarding the accuracy and completeness of details extracted from NLP data.

Other efforts have taken this idea even further to build interactive machine learning systems that learn iteratively from their end-users. Sometimes referred to as "human-in-the-loop" methods, these techniques involve a learning system whose output is used by the end-user to further inform the system about the learning task. This forms a closed loop that can be used to build continuously improving models of prediction. Some examples include applications in interactive document clustering [5], document retrieval [11], image segmentation [8], bug triaging [1] and even music composition [9]. These successes suggest that it may be promising to use feedback to improve machine learning models in the clinical domain.

DESIGN REQUIREMENTS

We assume that the users of our tool are domain experts who are familiar with the contents of the documents being reviewed, but not with machine learning. Our approach focuses on designing interaction methods and novel data visualizations for the user to interact with and correct the learning models. Further, while most of the focus in previous work has been towards developing usable interaction methods with the learning algorithms, more often in real world applications, we find that obtaining reliable labels for the training examples is very difficult, costly or time-consuming. In domains such as medicine, we require the help of skilled domain experts.

Labeled data are important to support training automated systems; yet, large amounts of training data do not exist for new use cases or for applications that may arise in the future. It is therefore of great practical interest to develop methods for obtaining good quality labels efficiently. Such methods are even more in need for NLP applications because it is time consuming for annotators to obtain the contextual information from the text before labeling. Lastly, we need to design techniques for the users to review the output of the NLP models. They should allow the users to find errors in predictions and make changes to build revised models. This would form a closed loop that would allow the users to iteratively create more accurate models that can be useful in their analysis. These requirements are summarized as follows:

- R1: The tool should facilitate interactive review of clinical text and make it easier for machine learning non-experts to work with NLP models.
- R2: Visual presentation and interactive feedback components should facilitate the building of accurate NLP models. This includes providing efficient techniques for annotation and labeling the training examples and also for providing feedback that is consistent and informative.
- R3: The interactive components should support the entire interactive machine learning loop - i.e. a *review, feedback* and *retrain* cycle that can be used to revise NLP models iteratively.

INTERFACE DESIGN

To demonstrate our tool, we have used an example dataset of colonoscopy reports by building on work done by Harkema et. al. [10]. They have described an NLP system to extract values against a set of boolean variables from these reports. We have included a subset of 14 of these variables for the demo. Each patient record in the example dataset can include multiple linked reports from endoscopy and pathology. We have considered such reports together as a single document for learning and making predictions.

Figure 1 shows a screenshot of our web-based tool. We have also uploaded a demo video of the tool at <http://vimeo.com/trivedigaurav/emr-demo>. In the following sections, we describe the individual components of the tool's user-interface, relating to the three requirements discussed above.

Review

An interactive machine learning cycle begins with the review step where the output from the learning model is shown to the user. Initial models can be trained on a few hand-annotated training examples. We have designed the following views in our tool to help the user inspect the prediction results.

Grid View

The grid-view is a table with columns showing the 14 variables and rows representing the individual documents. Each cell in the table shows the predicted value – true or false – corresponding to the particular document-variable pair. This table is scrollable to accommodate all the documents in the dataset and extends beyond what is visible in Figure 1(a).

We also have some cells with a question mark (?), where the model is unsure about the classification. This might happen for one of two reasons: either the classification algorithm does not identify a clear answer, or the learning system does not have sufficient examples in the training data to make any predictions as yet. Subsequent feedback may tilt the classification in either direction.

If the user hovers the mouse over a particular cell, a pop-up appears below it that shows the prediction probability, or how confident the learning system is in making that prediction. For example, the probability of a particular cell being true may be 75 percent. The grid also doubles up as a way to navigate through the documents. When a user clicks on a particular cell, the corresponding document-variable pair is activated in the all other views. The document view opens up the active document on the right-half of the screen. The highlighted cell in the grid indicates the currently active document-variable pair. Whenever the user clicks on a cell, we also mark it as visited to keep track of them. Visited cells in the grid are denoted by an asterisk symbol (*).

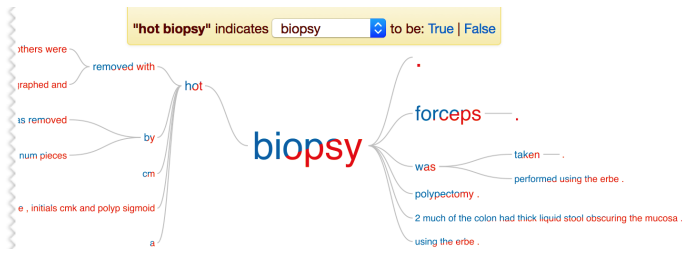
An overview bar at the top of each column displays the true-false distribution (skew) of each variable. Exact distribution percentages are shown when the user mouses over the variable name.

We have followed a uniform color scheme throughout the tool. Everything shaded blue represents a true value while the orange shades stand for false values. The colors were selected from a colorblind-safe palette. In the grid view, the cells with higher probability have a darker background color. For example, a light blue cell indicates a low probability about a true classification, and a darker blue for a higher probability.


Views for displaying Keywords and Statistics

Below the grid, we have views showing statistics about the currently active document and the variable. We show a histogram with a distribution of the true, false and unknown values over all the documents in the grid for the activated variable. Again, to reveal the exact counts under each prediction class, a user may hover the mouse over chart. This display is similar to the overview bar above the grid but is more detailed and changes dynamically when the user uses the search box or the WordTree view to filter the document collection.

Our NLP pipeline uses a bag of words feature-set and a support vector machine (SVM) learning model for every variable, but it can be extended for use with different kinds of models and complement other existing tools as well. It works by identifying more informative features from a document (top terms) to make predictions. Informative terms are highlighted when present in the current document in the right half, with overall distributions presented on the left-hand side of the screen. Terms are color-coded to indicate their contribution towards assigning the value of true or false against a variable, using colors from the document-variable grid. A mouse over each top term will reveal the feature weights from the learning system. Note that the current implementation consists of only unigram features but the same idea can be extended to n -grams as well.



(a) The *WordTree* view provides the ability to search for and explore word sequence patterns found across the documents in the corpus, and to provide feedback that will be used to retrain NLP models. In this example, we built the tree by searching for the word ‘*biopsy*’ and then drilled down upon the node ‘*hot*’. The word tree now contains all the sentences in the dataset with the phrase ‘*hot biopsy*’ in them. It allows the user to get an idea of all the scenarios in which *hot biopsy* has been used. Hovering over different nodes in the tree will highlight specific paths in the tree the selected term.



(b) The *Re-Train* view lists user-provided feedback, including any potential inconsistencies, and specifies changes in variable assignments due to retraining.

Figure 2. Screenshots of the *WordTree* and *Retrain* views.

Document View

On the right hand side of the tool, we show the full-text of the reports. The linked documents from a patient record, such as endoscopy and pathology reports are listed on the top of this view with shortcuts to jump to any of them. The top terms, both true and false (as seen in the grid view), are highlighted in the document. The keyword lists document the last view can be used to navigate through the report as well. Clicking on a keyword from the list causes the document to view to scroll to and highlight the first appearance of the term in the current document. This can also be done from the top terms list for the variables. However, since the term list contains the aggregate of the terms from all the documents together, there is a chance that a particular term doesn’t appear in the open document. When such a term is clicked the keyword will be animated with a brief jitter to indicate that it cannot be found. The highlighted top-terms in the document view follow the same color scheme for true and false indicators.

Clinical reports also contain boilerplate text, or portions that can be considered to be having no effect on the predictions. These include de-identification headers, footers, and report’s template text. These portions are dimmed in gray to improve the readability of the reports.

WordTree View

We have discussed views that provide detailed document-level visualization of the health records. But, we still need a visualization that could give a quick overview of the complete dataset. The *WordTree* [21] visualization offers a visual search tool for unstructured text that makes it easy to explore repetitive word phrases. The main advantage of using the wordtree is that it offers a complete data-set level visualization while retaining sentence level contextual information. A wordtree for a particular keyword consists of all the sentences in the dataset having that word or phrase. If one thinks of this keyword as the root of the tree, the branches represent the phrases that precede or follow that word. All the nodes of the tree are built recursively in this fashion. The font size of a particular node is decided according to the total proportion of sentences that have it as a common starting phrase.

We have made several improvements to the original wordtree design by Wattenberg and Viegas [21]. Their design restricts the root phrases to be present at the beginning or the end of a sentence. This allows the tree to grow only in one direction. We have used a modified design of the wordtree (as proposed in [19]) to construct a bi-directional tree that can grow in both directions. A sentence reads from left to right with the root phrase in the middle of the tree. Ends of sentences are denoted by a period (.) node. This information is also conveyed in plain-text and numbers as a tooltip on one corner of the *WordTree* view on mouse over. These gradients are dynamically updated according the variable selected by the user and also as the models change prediction upon retraining. The gradients provide an insight into how the machine learning model’s prediction changes as different words or phrases are present or absent from the documents.

To start using the *WordTree* view, the user must enter a keyword or a phrase in the search bar. The tool creates a wordtree with the search query as the root after scanning all the sentences in the dataset. One can interact with the wordtree and navigate through its branches by clicking on individual nodes as shown in Figure 2(a). Doing this prunes the tree and drill down into the details by adding the clicked node to be a part of the root node along with the search term. Clicking on the same node again reverts the view to the previous state of the tree. The gray bar below the wordtree, shows the number documents and the percentage of the dataset represented in the tree. The *WordTree* view also has a full-screen mode which hides the other views in the tool when required.

We have extended the wordtree to use color-coded gradients to encode class distribution information. Each word is painted in a gradient, with the extent of the blue/orange color indicating the percentage of active documents in the grid classified as true/false, using the previously-described color palette. The grid view is also linked to the wordtree: pruning the branches in the tree filters the set of documents to display in the grid to contain only those that are represented in the tree. The document ID list on the top right of the screen and the statistics views are similarly coordinated. If a user wishes to read more than what is available in the tree, they could just click on the corresponding cell in the grid to switch back to

the document view and review the complete document. In the following section we will describe how the wordtree can be used for providing annotations as well.

Feedback

Feedback from the user is used by the system to improve upon the machine learning models by providing labels for documents that were previously not part of the training set, or by correcting any misclassified documents. Useful feedback to the machine learning helps improve the prediction accuracy of the NLP models. A user can provide feedback by simply changing the prediction class of a document for a given variable. The learning system would then be able to use this as a training example to learn from its features. The marginal benefits may be greater if a user classifies a group of documents together instead of annotating them one by one. To classify a group of documents, we could search for a selected text span and label all the matching documents as belonging to a certain class. For this approach to work correctly, selected text spans must convey consistent meanings across different usage scenarios in the dataset, and feedback based on these selections should not imply any contradictory classifications.

Our prototype supports both multiple mechanisms for providing feedback and a review display that will alert the user to any potential inconsistencies associated with them. To provide feedback for a particular document-variable pair, the user can select either true or false on the yellow control bar above the document view (Figure 1(f)). The currently active variable from the grid is pre-selected but the control also allows the user to quickly change the variable of interest by choosing from the drop-down options or by activating a different cell in the grid.

Users can also provide more specific feedback by manually highlighting relevant text spans which could support document's classification. Like most other text annotators, the selection span automatically moves to ends of a word boundary if it is left hanging in the middle of a word. Multiple words forming a phrase can also be highlighted to be sent as a feedback.

Since the users are free to select their own text spans, there is a scope for the feedbacks to be inconsistent or prove to be less useful to the learning system. As a result, we have designed a feedback mechanism using the wordtree to provide them with some guidance in selecting these spans. Here the root phrase takes the role of the highlight span. Phrases are added before and after the root word as the user drills down the tree and prunes it. We believe that the wordtree may be useful while providing feedback step for it allows the users to give feedback on several documents together. The users are able to explore the different use cases of a phrase in all the documents in the dataset with a single glance. It also helps in identifying tighter and more generic feedback phrases with its click to drill down design. If the user is able to make a choice without having to view the complete sentence, we can identify phrases that may be more important for the machine learning system. The varying font sizes of the phrases provide strong visual cues about their frequency of use in the

dataset and thus encourages the users to attend to more useful phrases for training first. Further feedback for multiple documents based on a single phrase may help avoid potential conflict scenarios where the user may have highlighted similar keywords but selected different classes for feedback. To summarize, the wordtree not only provides an overview of the entire dataset but the provision of interacting with it allows the users to work directly with phrases and sentences in the dataset. It helps them to browse the data easily, send feedback actions to the learning systems, and see results with the help of color gradients.

All of three kinds of feedback can be submitted from the yellow bar present at the top of the screen that shows available options depending on the context. For example, an option for providing a feedback like this appears as soon as a text span is selected in the document. The document view also provides a right click menu as an additional affordance for the users to send feedback.

Re-Train

Re-training is the final step of the interactive machine learning cycle. The Re-Train view tab keeps a count of the number of feedbacks sent by the user and can be selected to view the list of proposed revisions to the model (Figure 2(b)). The list includes all three kinds of feedbacks. Clicking on the re-train button launches the re-training process. Once the retraining is complete, a new model is created and the system updates the predictions in the grid and the linked views. The grid view indicates all the differences between the old and the new model predictions. One can spot these changes in bold. These cells will also have a bold underline in the bottom. This allows the user to identify changes made in the model as a result of their feedback.

The Re-Train view also provides guidance for resolving potentially contradictory feedback items. For example, a user may provide a particular text span indicating that a given document-variable pair should be set to be true, even as they label it false in another feedback setting. In these cases, the system will return an error message specifying the problem, and highlight conflicting feedback items in red. These items can be revised or deleted from the Re-Train view, with red highlights disappearing when conflicts are resolved. Another conflict scenario involves the submission of suggested changes that undo the effects of earlier revisions to the model. These items are highlighted in yellow, and accompanied by an override option that will allow the newer input to take precedence over the earlier feedback. The repeated re-training steps allow the users to build the learning models over several iterations.

IMPLEMENTATION AND DEPLOYMENT

The system is implemented as a client-server architecture with communication over HTTP(S) using JSON. The user-interface has been built using the Angular (angularjs.org), D3 [2] and jQuery (jquery.com) javascript frameworks and libraries. The NLP learning system manages the model building and is deployed as a Tomcat Server Application. The tool

incorporates several other open-source libraries and packages, a list of which is available along with the source code at request.

USER STUDY

We conducted a formative user study to gain insight into usability factors of the tool that may be associated with errors or confusion, and to identify opportunities for improvement via re-design or implementation of new functionality.

Participants

We adopted a snowball sampling technique starting with clinicians identified by our colleagues to recruit participants for the user study. We conducted a total of 4 (+1 pilot) studies lasting between 60 to 90 minutes. Our participants worked as both clinicians and clinical researchers and had at least an MD degree. All participants were experienced with both clinical text and the colonoscopy procedures. Their positions varied from research faculty members to physician scientists. Three out of the four participants had between 5-10 years of experience in that position. They had limited experience with machine learning algorithms with average self-reported proficiency being 5.0 on a scale of 1 to 10 (Individual ratings: 2, 5, 6, 7), where 1 is for “No knowledge at all”, 5 – “Some idea about the algorithms”, and 10 being “Can read and understand current research”.

The pilot study helped us with some initial comments about the tool. This was done to identify any unnoticed bugs in our software prototype or any problems with our study protocol. Since we followed the same protocol in the pilot study as well and fixed only a couple of minor problems with the tool after it, we have also included its results with the rest of the studies.

Study Protocol

We began with a pre-study survey to gauge background information about the participants and their expectations from the tool. We gave a short 15-minute walkthrough of the interface before handing over the control to them. During the study, the participants were asked review documents using the tool, and revise NLP predictions by providing feedback wherever required. We asked the participants to work and build models for only one of the variables – *biopsy*, indicating whether or not the report discussed a sample biopsy. Actual interactions with the tool lasted between 20-30 minutes for the 4 studies but was longer for the pilot. The participants worked with 280 documents for providing feedback for a model built against a set of 30 hand-annotated documents. We followed the “think aloud” method to record their comments and reactions while using the tool. Sessions were conducted over web-conferencing software, which was also used to capture audio, screen content, and mouse interactions. At the end of the study, we asked users to complete the System Usability Scale [3] and to answer some questions regarding their understanding of the tool.

Results

System Usability scale

We used the System Usability Scale consisting of 10-questions on a 5-point Likert scale to help get a global view of

Category	Recommendation
Workflow	1. Allow sorting (or filtering) of the documents in the grid based on the prediction probabilities. This would make it easier for the users to prioritize documents to review.
	2. Add a button to open the next-in-line document for review. The order may be decided either trivially based on ID number or by using an active learning approach. This would save the users to navigate through the grid when they don't have their own strategy for selecting documents for review.
WordTree	1. Change the layout of the tool to show the WordTree view along with the document view. This would allow the user to quickly go through the full report text when the wordtree tree is unable to provide sufficient contextual information.
	2. Allow selection of multiple branches in the tree to give feedback on multiple paths in the tree at once.
Feedback	1. Provide a feedback mechanism to specify that a text span does not indicate either of the classes. This would allow the user to remove non-informative but possibly misleading features in re-training.
Re-Training	1. Perform auto-retraining in the background when a sufficient number of feedback items have been provided by the user.
	2. Provide a built-in mechanism to validate and generate a performance report for the current model against a held-out test set.

Table 1. Summary of recommendations inferred from the user study for improving the interface design.

subjective assessments of usability. The average SUS score was 70.5 out of 100. Individual scores are provided in Table 2.

Semi-structured interviews after the study

We classified the collected observations from the think-aloud sessions and comments from the semi-structured interviews into four categories: 1) *Workflow*: Comments and observations as the participants navigated through the documents for review, 2) *WordTree*: While selecting search queries and browsing the wordtree, 3) *Feedback*: While providing feedback to the learning system, and 4) *Re-training*: Upon seeing changes after re-training. Some of these comments also include requests for new features by the participants which are also summarized in Table 1.

- **Workflow**: Participants used the grid view to select documents of interest and the document view to navigate through the text. They found this part of the workflow to be tedious. Some participants requested a “Next” button that could be used to quickly move to a new document, instead of clicking on the cells in the grid. However, one of the participants also provided a contrasting view, expressing appreciation for the flexibility offered by the tool in selecting the documents for labeling. They also made use of the color shades representing the prediction probability numbers to prioritize documents for inspection. They requested a sort feature in the grid view that could arrange the documents according to these probability scores as well.

#	Question	Strongly Disagree	Disagree	Neither	Agree	Strongly Agree
1.	I think that I'd like to use this system frequently	0	1 [†]	0	3	1
2.	I found this system unnecessarily complex	0	3	2	0	0
3.	I thought the system was easy to use	0	0	0	5	0
4.	I think that I'd need the support of a technical person to be able to use this system	1	2	0	2	0
5.	I found the various functions in this system were well integrated	0	1	0	4	0
6.	I thought there was too much inconsistency in this system	3	1	1	0	0
7.	I'd imagine that most people would learn to use this system very quickly	0	0	1	3	1
8.	I found the system very cumbersome to use	1	4	0	0	0
9.	I felt very confident in using the system	0	0	1	3	1
10.	I needed to learn a lot of things before I could get going with this system	0	3	0	2	0

Table 2. Frequency of responses used for SUS score [3] calculation with 5 participants. Average SUS score was 70.5. a) We received poorer scores for questions about the learnability of the tool (Q4, Q10) as we were only able to provide a limited walkthrough of the interface to our participants due to time constraints. The participants acknowledged we could have done better if we were to spend more time on it. b) [†] When asked whether they would use the system frequently, the participant remarked that he gave a low score because he did not see the methods being directly useful in his current work.

- WordTree:** Perceptions about the wordtree were mixed. Some concerns regarding the wordtree appear to stem from the tabbed display that makes the user choose between the document view and the wordtree. While all participants found the wordtree to be a faster way to provide feedback, they felt that providing the feedback without being able to see the full document text at the same time might be error-prone. Although we were able to provide sentence long phrases in the tree and to show links to the full text of the relevant documents in the grid view, the participants were in favor of having a quicker way to access the complete reports. We have proposed a re-design to address this concern for future work. Our proposed redesign includes a provision for the user to make the WordTree view pop-out from its tab so that it can be used with the document and the grid views simultaneously.

Participants discovered an unexpected use case for this view. In addition to giving feedback, the wordtree allowed users to verify the quality of their models. This was a consequence of the gradient colors in it, which showed how the presence of individual keywords affect the classification of documents. By looking at how the gradient colors changed for the different keywords, the user could understand how well the model performed in predicting the values depending on the phrases contained in the document.

A common problem was that the users left the wordtree's search filter on even after they were done using it. The tool would filter the documents in the grid as the users navigated through the wordtree but would forget to clear it for the next round of analysis.

- Feedback:** Physicians indicated that they were accustomed to thinking in terms of rules suggesting a direct link between the feature and classification rather than the probabilistic associations used in our tool. As a result they were unsure at times about assigning a classification for a text span that serves as indicator in most but still not all of the cases. We address this problem in tool's design by encouraging models to be built iteratively. The user need not focus on building a completely accurate model at once but has an option to refine it for more specific cases in the future iterations. From the user study, however, we could not

recommend any further design improvements that could make the users more comfortable with this workflow.

One missing feature pointed out by the users was the ability to select a phrase and say that it didn't contribute towards the classification of the documents, when it was being picked either as a true or a false feature by the learning system. Otherwise the participants found the tool's features very usable for sending feedback to the machine learning system.

- Re-Training:** We had suggested that participants could build as many models as they like, which led them to have doubts about the optimal frequency of retraining. Future work may use NLP metrics to automatically determine when to retrain. The participants indicated that they were pleased to see the grid show changes in predictions after their feedback. Another suggestion was to provide a built-in option to test their model against a held-out hand annotated testing set.

Overall we received very encouraging responses from the participants. Four out of five (including the pilot) expressed interest in having the tool made available for their own work right away. The remaining participant was not involved in any research involving study clinical text. During the pre-study interview, we asked participants about their ideas on such a tool before showing our prototype. One of the participants who is actively working on related colonoscopy research requested features like a web-based interface for collaborating with people who are at geographically separated locations, flexibility in selecting documents to annotate, and a feedback mechanism for NLP. Our prototype tool was able to satisfy his needs in all of these aspects.

DISCUSSION AND FUTURE WORK

The initial feedback from the usability study provided both preliminary validation of the usability of the tool and guidance for improving the design of the tool. While we have not identified any major hurdle that would require a comprehensive re-design of any interface component, there are several extensions to the current set of features we believe might improve usability and will be promising in future work (Table 1).

One of the aims of this project was to explore the feasibility of using interactive review as a means of lowering the training requirements. We hypothesize that the manual review supported by this tool will enable rapid convergence on highly accurate models even by starting with smaller training sets. Testing this out in a statistically compelling manner has been left for a future empirical evaluation study. This would involve observing efficiency measures such as overall time spent and accuracy measures like F-Measure etc. under different variations of the tool. We may control the tool's presentation capabilities, types of feedback allowed and the number of training documents as the independent variables during this study. Another promising future direction would be to evaluate the use of the tool by several users in a collaborative work setting.

CONCLUSION

Despite the promising results shown by repeated studies involving NLP on clinical records, the benefits of NLP are all too often inaccessible to the clinicians and practitioners. Moreover, we have seen from previous studies that extracting structured insights from clinical text is hard. Although NLP techniques work well they have been put into limited use by the researchers in the field. In particular, without access to usable tools for clinicians that can make it easier for them to review and revise NLP findings, it is difficult apply these techniques.

We have built a candidate tool to help address these problems. The interactive components of the tool along with novel visualization techniques support the entire interactive machine learning cycle with review, feedback and retraining steps. We conducted a user-study with prospective users as study participants to validate our design rationales. We also identified opportunities for improvement that will be addressed before we move forward with an empirical evaluation of the system.

ACKNOWLEDGMENTS

We thank our user study participants. We would also like to thank Dr. Ateev Mehrotra for providing the colonoscopy reports dataset. This research was supported by NIH grant 5R01LM010964.

REFERENCES

- Amershi, S., Lee, B., Kapoor, A., Mahajan, R., and Christian, B. Cuet: Human-guided fast and accurate network alarm triage. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, ACM (New York, NY, USA, 2011), 157–166.
- Bostock, M., Ogievetsky, V., and Heer, J. D3: Data-driven documents. *IEEE Trans. Visualization & Comp. Graphics (Proc. InfoVis)* (2011).
- Brooke, J. SUS: a quick and dirty usability scale. In *Usability evaluation in industry*, P. W. Jordan, B. Weerdmeester, A. Thomas, and I. L. Mclelland, Eds. Taylor and Francis, London, 1996.
- Chapman, W. W., Nadkarni, P. M., Hirschman, L., D'Avolio, L. W., Savova, G. K., and Uzuner, O. Overcoming barriers to NLP for clinical text: the role of shared tasks and the need for additional creative solutions. *Journal of the American Medical Informatics Association* 18, 5 (2011), 540–543.
- Chau, D. H., Kittur, A., Hong, J. I., and Faloutsos, C. Apollo: Making sense of large network data by combining rich user interaction and machine learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, ACM (New York, NY, USA, 2011), 167–176.
- D'Avolio, L. W., Nguyen, T. M., Goryachev, S., and Fiore, L. D. Automated concept-level information extraction to reduce the need for custom software and rules development. *Journal of the American Medical Informatics Association* 18, 5 (2011), 607–613.
- Dick, R. S., Steen, E. B., Detmer, D. E., and Institute of Medicine (U.S.). Committee on Improving the Patient Record. *The computer-based patient record: an essential technology for health care*. National Academy Press, Washington, D.C, 1997.
- Fails, J., and Olsen Jr, D. Interactive machine learning. In *Proceedings of the 8th international conference on Intelligent user interfaces*, ACM (2003), 39–45.
- Fiebrink, R., Cook, P. R., and Trueman, D. Human model evaluation in interactive supervised learning. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, ACM (New York, NY, USA, 2011), 147–156.
- Harkema, H., Chapman, W. W., Saul, M., Dellon, E. S., Schoen, R. E., and Mehrotra, A. Developing a natural language processing application for measuring the quality of colonoscopy procedures. *JAMIA* 18, Supplement (2011), 150–156.
- Heimerl, F., Koch, S., Bosch, H., and Ertl, T. Visual classifier training for text document retrieval. *Visualization and Computer Graphics, IEEE Transactions on* 18, 12 (2012), 2839–2848.
- Iyer, S. V., Lependu, P., Harpaz, R., Bauer-Mehren, A., and Shah, N. H. Learning signals of adverse drug-drug interactions from the unstructured text of electronic health records. *AMIA Summits Transl Sci Proc* 2013 (2013), 83–87.
- Melton, G. B., and Hripcsak, G. Automated Detection of Adverse Events Using Natural Language Processing of Discharge Summaries. *Journal of the American Medical Informatics Association* 12, 4 (2005), 448–457.
- Nadkarni, P. M., Ohno-Machado, L., and Chapman, W. W. Natural language processing: an introduction. *Journal of the American Medical Informatics Association* 18, 5 (2011), 544–551.
- Nair, V., Kaduskar, M., Bhaskaran, P., Bhaumik, S., and Lee, H. Preserving narratives in electronic health records. In *Bioinformatics and Biomedicine (BIBM), 2011 IEEE International Conference on* (Nov 2011), 418–421.
- Ogren, P. V. Knowtator: a protégé plug-in for annotated corpus construction. In *Proceedings of the 2006 Conference of the North American Chapter of the Association for Computational Linguistics on Human Language Technology*, Association for Computational Linguistics (Morristown, NJ, USA, 2006), 273–275.
- Ramirez, E., Carcas, A. J., Borobia, A. M., Lei, S. H., Pinana, E., Fudio, S., and Frias, J. A pharmacovigilance program from laboratory signals for the detection and reporting of serious adverse drug reactions in hospitalized patients. *Clin. Pharmacol. Ther.* 87, 1 (Jan 2010), 74–86.
- Savova, G. K., Masanz, J. J., Ogren, P. V., Zheng, J., Sohn, S., Kipper-Schuler, K. C., and Chute, C. G. Mayo clinical Text Analysis and Knowledge Extraction System (cTAKES): architecture, component evaluation and applications. *Journal of the American Medical Informatics Association* 17, 5 (2010), 507–513.
- Shrikumar, A. A quick-and-dirty wordtree visualization in JS, HTML and SVG. <https://github.com/silverasm/wordtree>, 2011.
- Stasko, J., Görg, C., and Liu, Z. Jigsaw: Supporting investigative analysis through interactive visualization. *Information Visualization* 7, 2 (Apr. 2008), 118–132.
- Wattenberg, M., and Viegas, F. B. The word tree, an interactive visual concordance. *IEEE Transactions on Visualization and Computer Graphics* 14, 6 (2008), 1221–1228.
- Wei, F., Liu, S., Song, Y., Pan, S., Zhou, M. X., Qian, W., Shi, L., Tan, L., and Zhang, Q. Tiara: A visual exploratory text analytic system. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '10, ACM (New York, NY, USA, 2010), 153–162.
- Ye, Y., Tsui, F. R., Wagner, M., Espino, J. U., and Li, Q. Influenza detection from emergency department reports using natural language processing and bayesian network classifiers. *Journal of the American Medical Informatics Association* (2014).