Supporting Annotators with Affordances for Efficiently Labeling Conversational Data

AUSTIN Z. HENLEY, University of Tennessee, USA

DAVID PIORKOWSKI, IBM Research, USA

Without well-labeled ground truth data, machine learning-based systems would not be as ubiquitous as they are today, but these systems rely on substantial amounts of correctly labeled data. Unfortunately, crowdsourced labeling is time consuming and expensive. To address the concerns of effort and tedium, we designed CAL, a novel interface to aid in data labeling. We made several key design decisions for CAL, which include preventing inapt labels from being selected, guiding users in selecting an appropriate label when they need assistance, incorporating labeling documentation into the interface, and providing an efficient means to view previous labels. We implemented a production-quality implementation of CAL and report a user-study evaluation that compares CAL to a standard spreadsheet. Key findings of our study include users using CAL reported lower cognitive load, did not increase task time, users rated CAL to be easier to use, and users preferred CAL over the spreadsheet.

Additional Key Words and Phrases: labeling data, cognitive load, user interface, machine learning, conversational data

1 INTRODUCTION

Without well-labeled ground truth data, today's machine-learning-based systems would not nearly be as ubiquitous as they are today. These systems rely to substantial amounts of correctly labeled data. In some cases, this data is relatively straightforward to label correctly, as is the case of using data from clinical trials' or log data from users' interactions. However, for some labeling tasks, such as those that are qualitative instead of quantitative, human judgment is required. Due to the relatively large amounts of data required for ground truth, the most common approach for labeling is to crowdsource such tasks.

Unfortunately, crowdsourced labeling is fraught with difficulties. The fundamental problem with labeling is that it often produces noisy data, as labeling is colored by the perspectives of the individual annotator who may label differently than his or her peers. It is also time-consuming, expensive, and error-prone, especially in cases where annotators label without specialized tools and instead use commonly-available software. In particular, it has been reported that annotators often use spreadsheets [14], which we also found from our informal surveys at a large technology company. Consequently, reducing the amount of labeled data necessary via new machine learning techniques has been a popular research topic for several years [1, 2, 5, 9, 16, 20], but even these approaches require some amount of consistently labeled data in order to be effective.

Researchers have already begun developing formative methods and tools to improve labeling efficiency and quality. One notable tool is Aeonium, a visual-analytics interface designed to support qualitative labeling tasks through affordances such as visual overviews and keyword distributions to support multiple coders working on the same task [8]. Another tool is Revolt, which leverages disagreements between annotators to improved label quality [4].

Despite these advancements in tools for supporting labeling tasks, many barriers regarding the usability of such tools still remain.

To address the concern of effort and tedium, we designed CAL, a novel interface to aid annotators in efficiently labeling conversational data. We made several key design decisions for CAL focusing on reducing errors that annotators make and providing easy access to information to help improve annotator recall. Affordances that we developed in CAL include preventing inappropriate labels from being selected, in-situ guidance for selecting a label, incorporation of

Fig. 1. An example spreadsheet used for labeling conversational data. It includes the (a) transcript with the human utterance in the first column and the chatbot's response in the second column. The annotator enters *y* or *n* into the (b) three columns representing different categories in this example code set.

labeling guidelines into the interface, and access to prior labels for similar examples. These key design decisions were selected to address the established challenges annotators face when labeling large amounts of data.

In this paper, we present a production quality implementation of CAL and report a user-study evaluation that compares annotators labeling using CAL to those using a commonly used alternative, a spreadsheet. In particular, our user study addressed the following research questions:

- RQ1: Did CAL users have lower cognitive load during labeling tasks than spreadsheet users?
- RQ2: Do CAL users complete labeling tasks in less time than spreadsheet users?
- RQ3: Did CAL users rate CAL to be easier to use than the spreadsheet?
- RQ4: Did users prefer using CAL over the spreadsheet?

2 BACKGROUND & RELATED WORK

2.1 Data Labeling

Machine learning-based systems require large sets of data to have *labels* that form the ground truth of supervised machine learning models. (Other similar terms include annotations, tags, codes, and classifications.) To produce these labels, human *annotators* manually look at the data, consisting of many *examples* and select an appropriate label or set of labels for each example. The annotators follow established rules and criteria for their labeling task, to decide which label is the most appropriate from the available *label options*. The types of label options may vary. For example, annotators could be asked to make a binary choice like true or false or they may be asked for more open-ended labels such as entering any noun that most accurately describes the data. Additionally, each example may require multiple *categories* that an annotator must select a label for. For example, in the context of a chatbot, an annotator might label whether the chatbot correctly answered a user's question and whether the response was easy for the user to understand. In this case, there are two categories, correct response and understandable response, and each have two label options, *yes* and *no*, that the annotator must select from.

While performing labeling tasks, annotators often record the labels in a spreadsheet. Figure 1 shows an example based on an actual labeling task at a large technology company. The spreadsheet contains the data to be labeled in the left two columns (Figure 1a), which is the transcript between a human and a chatbot. The labels are found on the right (Figure 1b), which consist of three different categories: relevance, quantity, and manner. Each category has the option of being labeled *y*, *n*, or *s* (yes, no, skip). During the task, the annotator reads the transcript, enters the appropriate labels, and continues scrolling down the spreadsheet to label the remainder of the conversation. Moreover, the annotator may refer to the code set, often in the form of text documents, whenever they are unsure of what label option to choose.

2.2 Studies on Data Labeling

Prior research has looked at data labeling and annotation tasks and focused on methods to reduce labeling effort in various ways. One approach is to estimate the difficulty per example to label and using that to reduce the number of examples an annotator would have to label [1, 5]. Another popular approach is by developing better algorithms for selecting which examples to present to the annotator in an active learning scenario. Several prior works have surveyed such algorithms and suggested improvements [2, 9, 16, 20]. Our work differs since we focus on providing relevant information to help annotators making a labeling decision in the environment instead of quantifying labeling difficulty or choosing high-information examples to present to the labeler.

In contrast with approaches that have looked at reducing the number of examples annotators have to label, others have focused on soliciting extra information from labelers to increase the amount of information provided with the label. Zaiden et al. developed a framework to extract rationales in addition to labels from annotators [23, 24]. Others extended that approach to work in other domains [6, 11]. Building on [23], Yessenalina et al. developed a method to automatically generate rationales, thereby bypassing the need to solicit the rationales from the annotators [22].

An alternative approach to reducing labeling effort and soliciting extra labeling info is to determine and leverage the most consistently accurate annotators. Donmez et al. developed a technique to select the most accurate annotators among a crowdsourced group in an effort to reduce noise in the labeled data [7]. In a similar vein, Sheng et al. developed a method for determining which crowdsourced-labeled examples to send back to the crowd to relabel [17]. Snow et al. showed that crowdsourced workers were as good as experts for natural language labeling tasks when bias correction approaches were applied [18].

2.3 Tools to Support Labeling

As supervised machine-learning techniques continue to grow in popularity, a number of tools have emerged to connect people with data to label with annotators willing to label that data. Tools such as Figure Eight[1] (previously CrowdFlower) and Hive[2] provide data labeling services to researchers and businesses. Similarly, Mechanical Turk[3] enables researchers to hire people to act as annotators for small tasks. These platforms are efficient at connecting annotators to a labeling task, but they have little to no explicit support built in to support the annotator. Our work differs because it addresses the question of what sorts of affordances do annotators need to label more efficiently and more reliably.

Researchers have also studied challenges regarding data labeling and have developed tools to supporting labeling tasks. In the domain of chat logs, TextPrizm [15] and Aeonium [8] both provide affordances to the annotator to ease the task of labeling data. Aeonium provides several types of feedback to the annotator to help the annotator better assess

---

[1]https://www.figure-eight.com
[2]https://thehive.ai/
[3]https://www.mturk.com/

how well their labeling conforms to the rest of the crowd. In particular, Aeonium supports groups of annotators in identifying data that is unclear to overcome inconsistencies in labeling across the group. Brat, a tool to support natural language processing labeling tasks, allows annotators to select and label sections of text [19]. To help annotators make consistent decisions, it uses a machine learning-based system to recommend semantic classes to the annotator, along with a probability estimate.
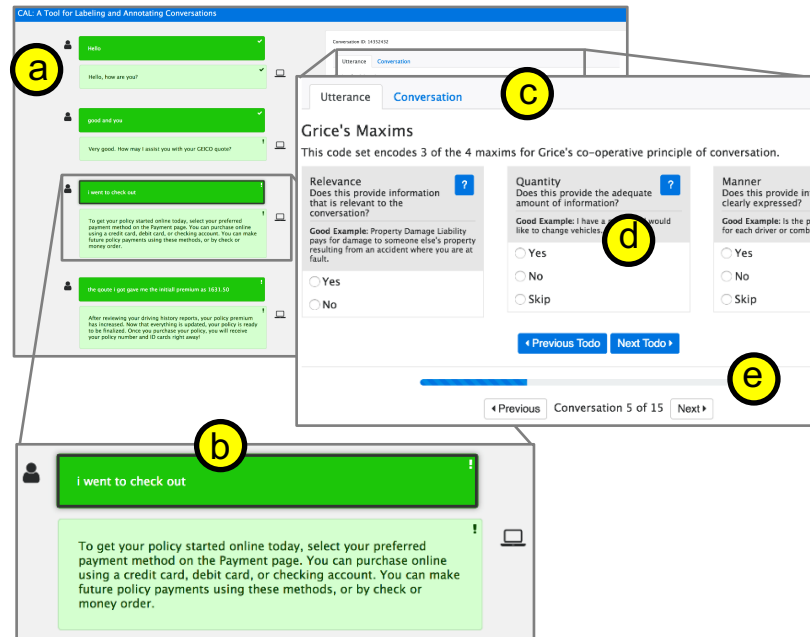

Fig. 2. (a) The CAL web application's data labeling view, including the (b) conversation transcript. To label the utterances, a user selects an individual utterance and uses the (c) labeling interface to (d) select the applicable labels. The user can track their progress with the (e) progress bar and navigate to the next/previous conversation.

3 TOOL DESIGN: CAL

To support annotators in efficiently labeling conversational data, we designed CAL, a web application designed to enable groups of data labelers to efficiently label conversational data. This section first details the basic features of CAL. Next we describe the design rationale for the novel features of CAL. Lastly, we describe the implementation details for CAL.

3.1 Basic Features

The essential feature of CAL is to enable an annotator to label an example from a data set. In this paper, CAL had been configured to label conversations, so the examples to be labeled were utterances from a chat bot's logs. First, the user clicks on the utterance that they are labeling from the conversation transcript (Figure 2b), which displays the labeling interface on the right (Figure 2c). Next, the user can select the appropriate labels via radio buttons, check boxes, and text boxes. Which label sets and label options appear for each example are defined in a configuration file supplied with the project when it is created. The task may require them to select choices for multiple categories (e.g., Figure 2d shows

three categories with radio buttons: Relevance, Quantity, and Manner. After completing their selections, they can move on to labeling the next utterance, either by clicking the Next button, using a keyboard shortcut, or by clicking the next utterance in the conversation transcript. Utterances that have not been completely coded have a small exclamation mark, while completed ones have checkmarks. The user continues doing this for every utterance that requires labels in the conversation. Additionally, CAL supports annotation for the entire conversation. If defined in the configuration file, after completing the utterance labels, the annotator will be presented with the appropriate label options for the entire conversation. The interface for annotating conversation-level labels is similar to the utterance-level one. The annotators continues applying labels to utterances and conversations until the progress bar (Figure 2e) shows 100%, then they can click or use a keyboard shortcut to proceed to the next conversation.

At any time, the annotator can view the status of their labeling tasks via a project status page. A progress bar depicts each annotator's progress towards completing the labeling tasks. Optionally, measurements of agreement (e.g., Cohen's Kappa and Jaccard Index) can be presented, either to all annotators or to the creator of the project as shown in Figure 5. CAL stores where each annotator left off, so the status page provides a quick link to easily get the annotators back on task.

There are two main levels of configuration in CAL: projects and annotations. A project configuration contains a set of data to be labeled, the annotators to do the labeling, and the assignment of code sets to utterances or conversations. The annotation configuration defines the categories and corresponding label options, the type of label (radio buttons, check boxes, and text boxes), definitions, examples and the rules for disabling or auto selecting related labels. If there is a flow-chart-like rule set, that is also encoded in the rule set. Configuration is stored in a JSON format and dynamically read by the CAL application and appropriately displayed to the annotator during the labeling task.

3.2 Design Rationale and Novel Features

During our formative investigations in how annotators perform data labeling tasks, we identified a number of issues that annotators faced. First, and especially at the beginning of the annotation task, annotators often referred back to documentation on their labeling rules to better inform their decisions or the recall nuances of a particular label. Doing so required them to repeatedly alternate between applications (e.g., a spreadsheet for their labels and a text document for the documentation), or refer to a printed rule set, which took considerable time. Second, annotators made mistakes entering labels. For example, we observed annotators entering an invalid label for a category (e.g., entering "Y" or "N" when "T" or "F" is expected). Another example, was in instances where a code set's rules included dependencies between labels (e.g., skip categories 2 and 3 if category 1's label is "F"). These mistakes could be caused by the annotator not being able to remember all of the labeling or simply unintentionally entering the wrong label and not realizing it. Third, even when annotators did not make a mistake, it was time consuming and tedious for them to follow the labeling documentation in order to provide a well-informed label. Fourth, annotators often referred back to their previous labels, requiring them to scroll back-and-forth between their labels, and not only increasing cognitive load as they had to remember why a label was given to a particular item, but potentially also causing additional issues (e.g., inadvertently changing a label while scrolling or confusing a previous label for another).

As a result of these findings, we identified four key principles that CAL was based upon:

- Provide the label ruleset and documentation in the interface.
- Prevent the user from selecting inapt labels.
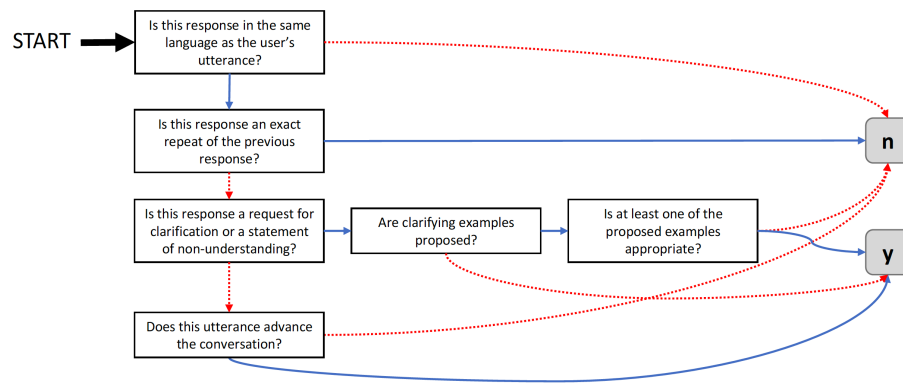- Guide the annotator in choosing the best label based on documentation.

Fig. 3. An example of labeling documentation that helps labelers in choosing the correct label by answering a series of yes/no questions. The wizard feature presents these questions to the user one question at a time. After answering all of the questions, CAL will select the label for the user and notify the user that the label was selected.

- Provide an affordance to efficiently view previous labels chosen by the annotator.

3.2.1 Integrate labeling code set. To eliminate the need for annotators to alternate between applications or between the computer screen and paper, CAL integrates the code set's documentation into the user interface. CAL displays definitions along with examples for each of the categories (see Figure 2d for Quantity's definition). This eliminates the need for annotators to refer to documentation outside the tooling. A button to toggle this information is also provided to reduce any unnecessary clutter.

3.2.2 Prevent inapt labels. For some code sets, additional rules restrict which code sets, categories or even labels should be applied to a given utterance. For example, for code sets that define label dependencies (i.e., if certain categories should be excluded due to a label option being selected, or if label options in other categories should be automatically be selected due to label option being selected), CAL provides mechanisms for enforcing those sorts of rules. First, when selecting an utterance, CAL will only show relevant categories and label options for that utterance. Second, after selecting a label that depends on other categories' labels, CAL will disable any inapt label options from being selected in combination with the current selection and automatically trigger other label options to automatically be selected if appropriate. For example, selecting the "Not Applicable" label for one category could automatically label other categories as "Skip." By enforcing these rules in the application, CAL prevents annotators from making common errors, such as selecting invalid labels.

3.2.3 Assist in choosing the best label. As mentioned above, to aid annotators in selecting the most appropriate label, CAL provides a wizard to guide the annotator in their selection. The wizard consists of a sequence of yes/no questions regarding the utterance, which after answering, will automatically select the appropriate label and display a notification that the label was automatically selected. This feature was inspired by documentation already used by labelers that acts as a set of heuristics by breaking down the labeling criteria into easier to understand questions. For example, Figure 3 shows a visualization of such a wizard from an existing project's code set's documentation (CAL does not display the visualization to users). The questions and the outcomes must be defined in CAL when configuring the project, which
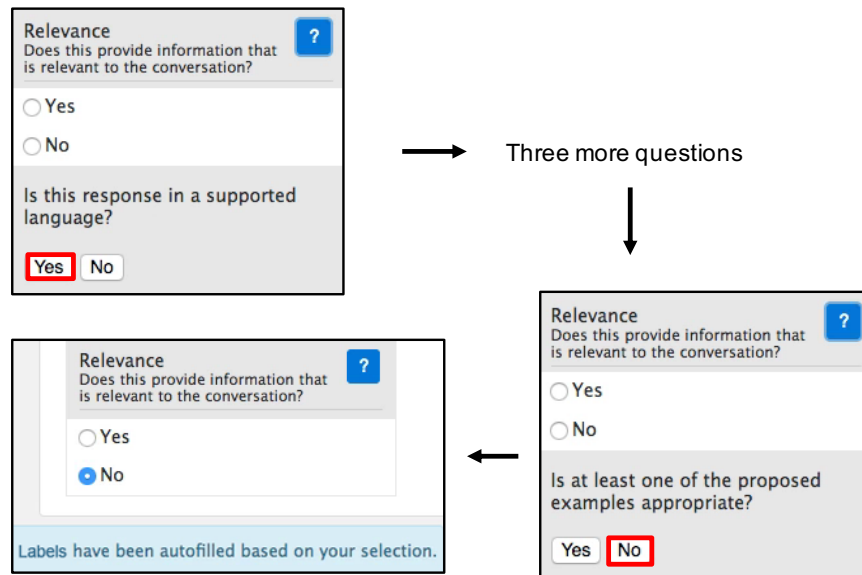
Fig. 4. An example of using the Wizard feature to aid an annotator in selecting the most appropriate label. After clicking the "?" button to initiate the wizard, the annotator is asked a series of binary questions. Depending on the answers provided to these questions, a label is automatically selected on behalf of the annotator. The annotator can revisit the wizard or modify the selected label if needed.

allows CAL to generate the wizard. The questions are presented one at a time in CAL's labeling interface, as depicted in Figure 4.

3.2.4 Efficient viewing of previous labels. From our formative investigations, we often observed labelers referring back to previous labels. The labelers did this to remind themselves of what they labeled in a particular case and also in an attempt to be consistent in their labeling. CAL provides an affordance to compare the current utterance label to a previously labeled utterance. To see this comparison, the labeler hovers over a potential label, which reveals a "View Previous" button. Clicking the button will show a recent previous utterance that this labeler had labeled with this specific label. For example, hovering over "Yes" for the Relevance and clicking the "View Previous" button will show a comparison view of previous utterance that the labeler marked "Yes" in the Relevance category to the currently selected utterance. This view is displayed below CAL's labeling interface (the otherwise unused area of Figure 2a in the bottom right).

3.3 Implementation Details

CAL is a written as a Node.js[4] application. The front-end is built using HTML 5, Embedded JavaScript[5], and Bootstrap[6]. Data is stored on a MongoDB[7] backend and accessed via the Flask[8] microservice framework. For the purposes of the

---

[4]http://nodejs.org/
[5]http://ejs.co/
[6]https://getbootstrap.com/
[7]https://www.mongodb.com/
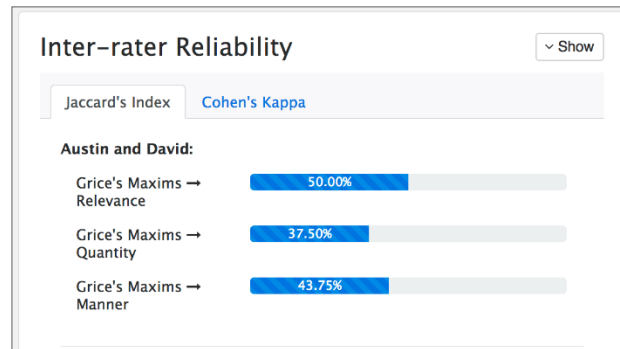[8]http://flask.pocoo.org/

Fig. 5. The inter-rater reliability view that displays the amount of agreement between data labelers. In this case, the two labelers have 37.5% to 50% agreement on three different categories using Jaccard's Index.

Table 1. Background and demographic information of study participants.

Participant	Gender	Age	Education	Job role	Native language	Labeling experience
P1	M	20-29	MS	Soft. Engr.	Chinese	No
P2	M	30-39	BS	Researcher	English	No
P3	F	20-29	MS	Soft. Engr.	English	No
P4	M	20-29	BS	Soft. Engr.	English	No
P5	F	30-39	PhD	Researcher	English	Yes
P6	F	20-29	BS	Soft. Engr.	English	No
P7	F	20-29	PhD	Researcher	Chinese	Yes
P8	F	50-59	MS	Soft. Tester	English	No
P9	F	20-29	BS	UX Designer	Japanese	No
P10	F	20-29	MS	Soft. Engr.	Gujarah	No

study, the application was running a VirtualBox[9] virtual machine configured via Vagrant[10]. To guarantee that each participant received the same treatment, we spun up a clean pre-configured Vagrant instance of the virtual machine for each participant.

4 EVALUATION METHOD

To address our research questions, we conducted a laboratory study of people performing data labeling tasks. The study had a within-subjects design in which each participant experienced two treatments. The control treatment was a standard spreadsheet using Microsoft Excel based on our prior observations in the field. The experimental treatment was the CAL web application. To account for order effects, we blocked and balanced based on the treatment orderings.

---

[9]https://www.virtualbox.org/
[10]https://www.vagrantup.com/

4.1 Participants

Our participants were recruited by advertising on Slack groups (a corporate chat application) at a large technology company. We did not require any particular experience or skills, but the groups we advertised to consisted mostly of researchers and software engineers. They were provided a free lunch voucher for their participation. After recruiting for one week, 10 participants (7 female, 3 male) agreed to take part in the study. All participants had experience using Excel, but only two participants had experience with data labeling: one used Excel to label images and one used Excel to label chatbot data. See Table 1 for background and demographic information regarding the participants.

4.2 Labeling Tasks

For our study, the participants performed labeling tasks on conversations between a human and a chatbot. These conversations consisted of real customers seeking help while attempting to purchase a product on a publicly available website. The nature of the conversations is proprietary and cannot be disclosed. We filtered this dataset to find recent conversations within the last 3 months and to restrict the minimum and maximum conversation length (10-20 utterances). After filtering, we randomly selected 30 conversations.

The participants were tasked with labeling the utterances in the conversations. In this context, conversations consisted of alternating utterances between the chatbot and the human. The labels were based on a subset of Grice's maxims [10], a set of established properties that are needed in an effective conversation. For our participants, this meant labeling "Yes" or "No" for three different properties (Quality, Quantity, and Manner) regarding each utterance. An additional label, Topic Change, was used for human utterances that indicated whether or not the human changed conversation topics. After labeling all of the utterances in a conversation, the participants moved on to the next conversation.

4.3 Procedure

Each participant took part in an individual session that lasted approximately 75 minutes. All participants began the session by signing a consent form, listening to a summary of the study session, and filling out a background questionnaire. They were also instructed on how to think aloud during the labeling tasks. Next, each participant performed labeling for 25 minutes with one treatment and then performed labeling for 25 minutes with the other treatment. Five random participants used the control spreadsheet first and five participants used the CAL application first. The participants were given a training session on the labeling task with explanations on the labeling criteria. They were also provided documentation on the labels (Grice's maxims [10] and topic change) consisting of definitions and examples in the form of PDFs that they could refer back to at any time. They were given a demonstration on each tool before using it. After completing each 25 minute task, they were a given a cognitive load questionnaire regarding the task and a usability questionnaire regarding the tool they used. At the end of the session, all participants took part in a semi-structured interview to discuss benefits, challenges, and possible improvements for both CAL and the spreadsheet.

4.4 Data Collection

During the study, the computer screen and audio of the participants was recorded. The CAL application also logged the participant's interactions (e.g., label selections and button clicks). For the cognitive load questionnaire, we used a well-validated short shelf-report instrument based on Cognitive Load Theory [12]. The instrument measures a person's cognitive load during a task with a 7-point Likert scale and has been shown to correlate with complex measurements, such as NASA TLX [21], and with physiological measurements, such as heart rate [13]. For the usability questionnaire,
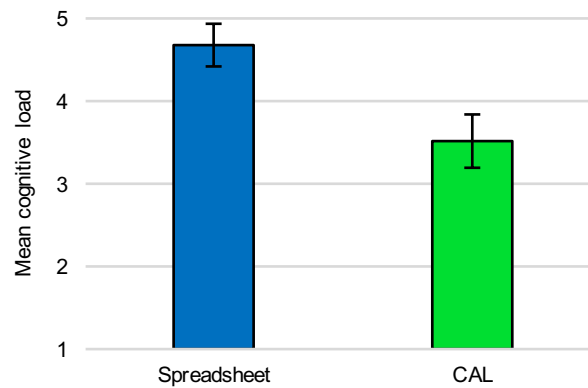
Fig. 6. CAL users reported significantly lower cognitive load than spreadsheet users (lower is better). Whiskers denote standard error.
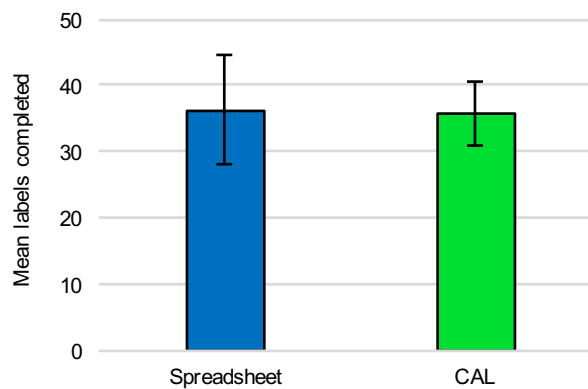

Fig. 7. Overall, spreadsheet and CAL users did not exhibit a significant different in the number of labels completed. Whiskers denote standard error,

we used the widely used system usability scale (SUS) [3]. This questionnaire consists of 10 Likert-scale questions that produce a single score that can be used to compare systems in terms of usability.

5 RESULTS

The 10 participants completed two sequences of labeling, each using CAL and the spreadsheet. Collectively, the participants made 839 labels for utterances in 30 different conversations. The remainder of this section describes the results for our 4 research questions.

5.1 RQ1 Results: Cognitive Load

As Figure 6 shows, participants using CAL reported considerably lower cognitive load than spreadsheet users. Indeed, every individual participant reported a lower or equal cognitive load using CAL than when they used the spreadsheet. The results of a Mann-Whitney U test showed that CAL users reported a significantly lower cognitive load than spreadsheet users ($U = 19$, $Z = 2.3$, $p < 0.05$).
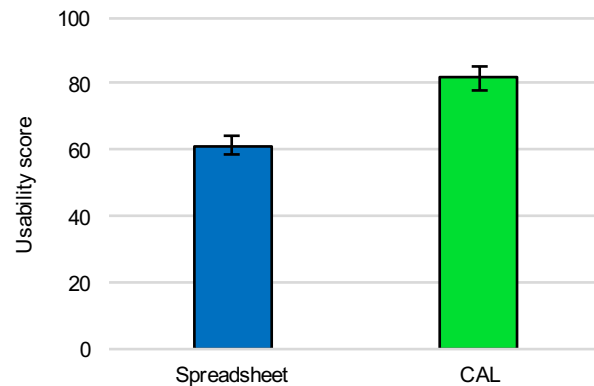
Fig. 8. Participants rated CAL to be significantly more usable than the spreadsheet on the SUS questionnaire [3] (higher is better). Whiskers denote standard error.

5.2 RQ2 Results: Task Progress

As Figure 7 shows, CAL users labeled a similar number of utterances as spreadsheet users. In fact, there is less than a 2% difference in the average number of labels completed between the treatments, despite all participants having experience with Excel and no experience with CAL. The number of labels completed did not differ significantly between CAL and spreadsheet users. Each individual participant did complete more labels during their second task than their first task, regardless of the treatments, which suggests a learning effect (see Section 6.5 for a discussion of the study limitations).

5.3 RQ3 Results: Usability

As Figure 8 shows, CAL users rated the usability of CAL considerably higher than spreadsheet users on the SUS questionnaire [3], which consists of 10 Likert-scale questions regarding the tool's effectiveness, efficiency, and satisfaction. In fact, all 10 participants rated CAL higher than they rated the spreadsheet in each of the questions on the SUS questionnaire. The results of a Mann-Whitney U test shows that CAL users reported a significantly higher usability score than spreadsheet users ($U = 5.5, Z = 3.3, p < 0.001$).

5.4 RQ4 Results: Preference

As Figure 9 shows, the majority of the participants stated they would prefer using CAL over the spreadsheet in the future. Of the 10 participants, 7 participants said they would prefer to use CAL and 3 participants said "it depends" on the context. No participant preferred using the spreadsheet over CAL. The details of their preferences are explained in Section 6.

6 DISCUSSION

Overall, the quantitative results from our user evaluation of CAL were highly favorable. CAL significantly reduced users cognitive load during the labeling tasks. Furthermore, CAL users took roughly the same amount of time on tasks as spreadsheet users. After the labeling tasks, users rated CAL to be significantly easier to use based on the SUS questionnaire and when asked which tool they preferred, the majority of users chose CAL.
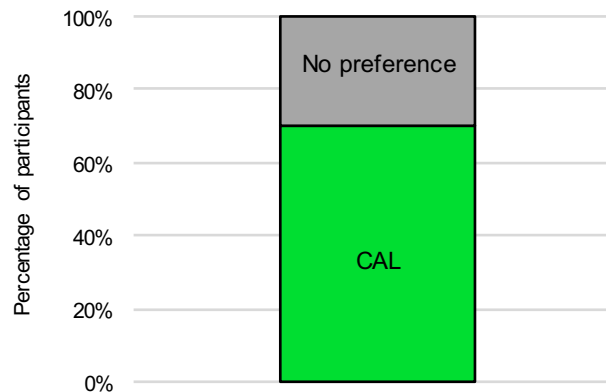
Fig. 9. When asked which tool did the participant prefer using, the majority chose CAL. No participant chose the spreadsheet.

The remainder of this section will report our analysis of the qualitative evidence from the think aloud data and interview responses in an effort to better understand our results.

6.1 CAL eliminates typing

The CAL interface eliminated the need for typing the specified label. P2 and P6 both expressed their satisfaction with how CAL eliminated the need to type on the keyboard. P2 continued by indicating it was beneficial to not have to continuously switch between keyboard and mouse, like he had to do with the spreadsheet. P3 had a similar thought, saying that he could no longer make a mistake while typing because CAL didn't require any typing.

However, the reliance on using the mouse did have performance implications. P10 said she believed that using CAL took more clicks than the spreadsheet did. P5 and P7 responded to interview questions by indicating that the spreadsheet is faster to input the labels into. P5 continued by saying that she felt faster using the spreadsheet but would prefer CAL if there were keyboard shortcuts (although shortcuts were implemented, our tutorial on CAL did not focus on them).

6.2 CAL provides assistance for labeling

CAL users overwhelmingly found the wizard feature to be beneficial. In fact, participants 6 of the participants made verbal remarks about the usefulness of the feature, including:

P2: "I like the thing with the question mark where you can decide whether it is relevant or not. With the spreadsheet, you have to go look at the documentation."

P3: "I liked that I got the help when I needed to by clicking on those question marks if I was ever stuck. It was a lot easier than going back to the documentation."

P4: "Having the integrated help so that you can step through the questions if you're unsure about something, it helps you get there."

P8: "I liked that if I didn't know what to pick, it would take me through a series of questions and pick the answer for me."

The other labeling features provided by CAL were also found to be useful for participants. The autofilling of labels was deemed helpful by both P7 and P8. Although the View Previous feature was used sparingly by our participants, P5 expressed multiple positive sentiments:

P5: " I think this View Previous is really good... I think it is a really good idea. Even when I filled out the first one, I wanted to view my previous labels... I think that is super valuable."

On the other hand, P7 was observed manually going back to view previous labels without using this feature, which indicates future work to make this feature easier to use.

6.3 Spreadsheet was not user friendly

Participants complained about various ways in which the spreadsheet was difficult to use. Most commonly, they complained about how difficult the spreadsheet was to read. In fact, 5 participants verbally expressed this sentiment throughout their session, such as:

P1: "It is difficult for me to read."

P6: "The cells do bombard me, then I have to remember which column is going for which one."

P10: "I think the spreadsheet hurts your eyes after a while. I think it is because it is not user friendly."

P4 and P5 drew comparisons to their experience with using CAL:

P4: "Being able to see it laid out in text conversation rather than having to read cells, that made it easier."

P5: "This one is like ahhh because I don't have to strain my eyes [like in the spreadsheet]."

6.4 Spreadsheet showed more information

The spreadsheet shows more information on the screen at any given time than CAL. Based on the participants' feedback, one of the biggest advantages that the spreadsheet had over CAL is that it allows the user to quickly scan the labels for multiple utterances without any actions. This was summarized well by P8:

P6: "For the spreadsheet, it was easier to compare all my responses overall."

During the interviews, 5 of the participants suggested that CAL could be improved by showing the label for the utterance and the response simultaneously. P3 also requested that items in CAL be closer together to reduce the amount of scrolling needed.

6.5 Limitations

Our study had a number of limitations to generalizability that are common to laboratory studies. First, the tasks may not generalize to other situations of labeling data, but we did base the tasks on actual labeling tasks being conducted for industrial use and we used data from an actual product. Second, it is not well understood how our sample of participants represents our target population. To partially mitigate this, we recruited participants from a large technology company. Third, using spreadsheets may not be representative of how researchers do conversation labeling tasks, but our formative investigations found it to be the most used tool. Fourth, the participants may have concluded that CAL was our design and thus biased their behavior and responses. Fifth, since this was a within-subjects study, there may be order effects of the treatments and tasks; however, we counterbalanced the treatment order to control for this effect.

7 CONCLUSION

In this paper, we introduced the novel CAL tool to support annotators in labeling data more efficiently. We designed the tool based on 4 key design goals and implemented a production-quality web application. An evaluation study comparing our CAL tool to a standard spreadsheet made the following key findings:

- RQ1: CAL users reported significantly lower cognitive load during labeling tasks.
- RQ2: CAL had no noticeable effect on the time taken to label data.

- RQ3: Users rated CAL to be significantly more usable than the spreadsheet.
- RQ4: Most users preferred using CAL, while two reported "it depends" (none preferred the spreadsheet).

CAL aims to make considerable advances in supporting human annotators in the tedious task of manually labeling data to be used in machine learning-based systems. In the future, CAL could utilize machine learning to provide suggestions of which label to select based on the annotator's recent labels. Such a technique could also be applied to CAL's *View Previous* feature, by using an intelligent recommendation system to display relevant data that the annotator recently labeled. Additionally, our study revealed that annotators begin showing signs of fatigue after an hour of annotating, which suggests that CAL should monitor for behavior that is indicative of fatigue, such that it can ensure the labels are accurate. These future works have tremendous potential to transform the way in which human annotators label data.

REFERENCES

[1] Shilpa Arora, Eric Nyberg, and Carolyn P Rosé. 2009. Estimating annotation cost for active learning in a multi-annotator environment. In *Proceedings of the NAACL HLT 2009 Workshop on Active Learning for Natural Language Processing*. Association for Computational Linguistics, 18–26.

[2] Avrim L Blum and Pat Langley. 1997. Selection of relevant features and examples in machine learning. *Artificial intelligence* 97, 1-2 (1997), 245–271.

[3] John Brooke et al. 1996. SUS-A quick and dirty usability scale. *Usability evaluation in industry* 189, 194 (1996), 4–7.

[4] Joseph Chee Chang, Saleema Amershi, and Ece Kamar. 2017. Revolt: Collaborative crowdsourcing for labeling machine learning datasets. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. ACM, 2334–2346.

[5] Aron Culotta and Andrew McCallum. 2005. Reducing labeling effort for structured prediction tasks. In *AAAI*, Vol. 5. 746–751.

[6] Jeff Donahue and Kristen Grauman. 2011. Annotator rationales for visual recognition. In *Computer Vision (ICCV), 2011 IEEE International Conference on*. IEEE, 1395–1402.

[7] Pinar Donmez, Jaime G Carbonell, and Jeff Schneider. 2009. Efficiently learning the accuracy of labeling sources for selective sampling. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 259–268.

[8] M. Drouhard, N. C. Chen, J. Suh, R. Kocielnik, V. Peña-Araya, K. Cen, Xiangyi Zheng, and C. R. Aragon. 2017. Aeonium: Visual analytics to support collaborative qualitative coding. In *2017 IEEE Pacific Visualization Symposium (PacificVis)*. 220–229. https://doi.org/10.1109/PACIFICVIS.2017.8031598

[9] Yifan Fu, Xingquan Zhu, and Bin Li. 2013. A survey on instance selection for active learning. *Knowledge and information systems* 35, 2 (2013), 249–283.

[10] H Paul Grice. 1975. Logic and conversation. *1975* (1975), 41–58.

[11] Tyler McDonnell, Matthew Lease, Mucahid Kutlu, and Tamer Elsayed. 2016. Why is that relevant? Collecting annotator rationales for relevance judgments. In *Fourth AAAI Conference on Human Computation and Crowdsourcing*.

[12] Fred Paas, Juhani E. Tuovinen, Huib Tabbers, and Pascal W. M. Van Gerven. 2003. Cognitive Load Measurement as a Means to Advance Cognitive Load Theory. *Educational Psychologist* 38, 1 (2003), 63–71.

[13] Fred GWC Paas, Jeroen JG Van Merriënboer, and Jos J Adam. 1994. Measurement of cognitive load in instructional research. *Perceptual and Motor Skills* 79, 1 (1994), 419–430.

[14] Marcus Renner and Ellen Taylor-Powell. 2003. Analyzing qualitative data. *Programme Development & Evaluation, University of Wisconsin-Extension Cooperative Extension* (2003), 1–10.

[15] Taylor Jackson Scott. 2017. *A Framework of Distributed Affect in Text-Based Communication*. Ph. D. Dissertation. University of Washington.

[16] Burr Settles and Mark Craven. 2008. An analysis of active learning strategies for sequence labeling tasks. In *Proceedings of the conference on empirical methods in natural language processing*. Association for Computational Linguistics, 1070–1079.

[17] Victor S Sheng, Foster Provost, and Panagiotis G Ipeirotis. 2008. Get another label? improving data quality and data mining using multiple, noisy labelers. In *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 614–622.

[18] Rion Snow, Brendan O'Connor, Daniel Jurafsky, and Andrew Y Ng. 2008. Cheap and fast—but is it good?: evaluating non-expert annotations for natural language tasks. In *Proceedings of the conference on empirical methods in natural language processing*. Association for Computational Linguistics, 254–263.

[19] Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun'ichi Tsujii. 2012. BRAT: a web-based tool for NLP-assisted text annotation. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 102–107.

[20] Meng Wang and Xian-Sheng Hua. 2011. Active learning in multimedia annotation and retrieval: A survey. *ACM Transactions on Intelligent Systems and Technology (TIST)* 2, 2 (2011), 10.

[21] D Windell and EN Wiebe. 2007. Measuring cognitive load in multimedia instruction: A comparison of two instruments. *Annual meeting of the American Educational Research Association* (2007).

[22] Ainur Yessenalina, Yejin Choi, and Claire Cardie. 2010. Automatically generating annotator rationales to improve sentiment classification. In *Proceedings of the ACL 2010 Conference Short Papers*. Association for Computational Linguistics, 336–341.

[23] Omar Zaidan, Jason Eisner, and Christine Piatko. 2007. Using "annotator rationales" to improve machine learning for text categorization. In *Human Language Technologies 2007: The Conference of the North American Chapter of the Association for Computational Linguistics; Proceedings of the Main Conference*. 260–267.

[24] Omar F Zaidan, Jason Eisner, and Christine Piatko. 2008. Machine learning with annotator rationales to reduce annotation cost. In *Proceedings of the NIPS* 2008 workshop on cost sensitive learning*. 260–267.