Facilitating Knowledge Sharing from Domain Experts to Data Scientists for Building NLP Models

Soya Park
Massachusetts Institute of Technology
USA
soya@mit.edu

April Wang
University of Michigan
USA
aprilww@umich.edu

Ban Kawas
Q. Vera Liao
David Piorkowski
Marina Danilevsky
IBM Research
USA
bkawas@us.ibm.com
[vera.liao,david.piorkowski}@ibm.com
mdanile@us.ibm.com

ABSTRACT

Data scientists face a steep learning curve in understanding a new domain for which they want to build machine learning (ML) models. While input from domain experts could offer valuable help, such input is often limited, expensive, and generally not in a form readily consumable by a model development pipeline. In this paper, we propose Ziva, a framework to guide domain experts in sharing essential domain knowledge to data scientists for building NLP models. With Ziva, experts are able to distill and share their domain knowledge using domain concept extractors and five types of label justification over a representative data sample. The design of Ziva is informed by preliminary interviews with data scientists, in order to understand current practices of domain knowledge acquisition process for ML development projects. To assess our design, we run a mix-method case-study to evaluate how Ziva can facilitate interaction between domain experts and data scientists. Our results highlight that (1) domain experts are able to use Ziva to provide rich domain knowledge, while maintaining low mental load and stress levels; and (2) data scientists find Ziva's output helpful for learning essential information about the domain, offering scalability of information, and lowering the burden on domain experts to share knowledge. We conclude this work by experimenting with building NLP models using the Ziva output for our case study.

CCS CONCEPTS

• **Human-centered computing** → *Empirical studies in HCI; Interactive systems and tools.*

KEYWORDS

Human-in-the-loop machine learning, CSCW, Multi-disciplinary collaboration

ACM Reference Format:

Soya Park, April Wang, Ban Kawas, Q. Vera Liao, David Piorkowski, and Marina Danilevsky. 2021. Facilitating Knowledge Sharing from Domain Experts to Data Scientists for Building NLP Models. In *26th International Conference on Intelligent User Interfaces (IUI '21)*, April 14–17, 2021, College Station, TX, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3397481.3450637>

1 INTRODUCTION

In recent decades, machine learning (ML) technologies have been sought out by an increasing number of professionals to automate their work tasks or augment their decision-making [83]. Broad areas of applications are benefiting from integration of ML, such as healthcare [15, 17], finance [22], employment [49], and so on. However, building an ML model in a specialized domain is still expensive and time-consuming for at least two reasons. First, a common bottleneck in developing modern ML technologies is the requirement of a large quantity of labeled data. Second, many steps in an ML development pipeline, from problem definition to feature engineering to model debugging, necessitate an understanding of domain-specific knowledge and requirements. Data scientists therefore often require input from domain experts to obtain labeled data, to understand model requirements, to inspire feature engineering, and to get feedback on model behavior. In practice, such knowledge transfer between domain experts and data scientists is very much ad-hoc, with few standardized practices or proven effective approaches, and requires significant direct interaction between data scientists and domain experts. Building a high-quality legal, medical, or financial model will inevitably require a data scientist to consult with professionals in such domains. In practice, these are often costly and frustrating iterative conversations and labeling exercises that can go on for weeks and months, which usually still do not yield output in a form readily consumable by a model development pipeline.

In this work, we set out to develop methods and interfaces that facilitate knowledge sharing from domain experts to data scientists for model development. We chose to focus on natural language processing (NLP) modeling tasks, and we are especially motivated by real-world cold-start scenarios where labeled data is small or nonexistent. Informed by a formative interview with data scientists regarding current practices and challenges of learning from domain

This work is licensed under a Creative Commons Attribution International 4.0 License.

IUI '21, April 14–17, 2021, College Station, TX, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-8017-1/21/04.
<https://doi.org/10.1145/3397481.3450637>

experts, we developed a domain-knowledge acquisition interface **Ziva** (With Zero knowledge, How do I deVelop A machine learning model?). Instead of a data-labeling tool, Ziva intends to provide a diverse set of elicitation methods to gather knowledge from domain experts, then present the results as a repository to data scientists to serve their domain understanding needs and to build ML models for specialized domains. Ziva scaffolds the knowledge sharing in desired formats and allows asynchronous exchange between domain experts and data scientists. It also allows flexible re-use of the knowledge repository for different modeling tasks in the domain.

Specifically, informed by findings from the formative interview and requirements of NLP modeling tasks, Ziva focuses on eliciting key concepts in the text data of a domain (`concept creation`), and rationale justifying a label that a domain expert gives to a representative data instance (`justification elicitation`). In the current version of Ziva, we provide five different justification elicitation methods – `bag of words`, `simplification`, `perturbation`, `concept bag of words`, and `concept annotation`.

To evaluate and inform future development of Ziva, we conducted a case study in assessment of its coupled design goals: 1) to provide an efficient and user-friendly experience for domain experts to supply domain knowledge; 2) to support data scientists building NLP models, especially in cold-start scenarios.

We performed a lab study (N=12) and a crowd-deployment study (N=88) for participants to act as domain experts of a restaurant reviewing domain, and use Ziva to provide concepts and justification-based knowledge. We found the completion time and subjective workload using different elicitation methods varied. Interestingly, the popular keywords based justification (`bag of words`) approach led to higher self-reported task success but was considered more stressful.

We conducted an interview study with 7 data scientists to investigate whether and how Ziva could help them build NLP models. Through the study, we identified design requirements for domain knowledge-sharing tools in ML development workflow – scalability of information and lowering workload for domain experts. Participants also reflected on how the shared domain knowledge facilitated by Ziva may be utilized, including bootstrapping labels, supporting feature engineering, improving explainability, and training few-shot learning models. Based on these suggestions, we experimented with building a rule-based model using the data from our user study, and report the outcomes using knowledge elicited with different methods. In summary, the contributions of the paper are as follows:

- Through a formative interview with data scientists who built models in a specialized domain, we identified their under-supported needs to learn about a domain from domain experts.
- We developed Ziva, a tool providing `concept creation` and five kinds of `justification elicitation` to gather domain knowledge from domain experts in formats that could help data scientists build NLP models.
- We conducted a case study using Ziva to elicit domain knowledge then presented the output to data scientists in an interview study. Their feedback validated the utility of Ziva and provided design insights for tools that support knowledge sharing and collaboration between domain experts and data scientists.

- We also investigated the experience of domain experts using Ziva. We believe that our analysis could inform the design of knowledge elicitation methods for domain experts.

2 RELATED WORK

We are informed by recent studies of data science practices, as well as ML and HCI work that leverages domain experts' input to train or improve models, and research to facilitate knowledge sharing in teams and organizations.

2.1 Data Science practices and collaboration

Recently the data science domain has spurred great research interest in the HCI community. Besides developing numerous tools to support specific data science tasks (e.g. [4, 34, 35, 86]), an emerging area of research has focused on studying the practices of data scientists in model development work. Many have recognized the collaborative nature of data science projects, with both intra- (among data scientists) [43] and multi-disciplinary (with domain experts) collaboration [59, 84]. In particular, data scientists rely heavily on domain experts during core modeling building stages, such as data access and feature extraction. Domain experts also feature prominently in latter stages of data science projects such as model evaluation and communication of results [59]. However, data scientists' work faces significant challenges as such collaborative activities are currently not well supported [50, 57], and they are often left with no choice but to rely on "an intuitive sense of their data and processes" [53].

Computational notebooks are positioned as a potential solution to both support collaborative coding and communicating results to stakeholders [78]. However, a recent study reported reluctance for data scientists to directly communicate the in-progress model work in notebooks [65]. While there are tools emerging to address the technology gaps to support collaborative data science practices, to our knowledge they tend to focus on supporting teams of data scientists and place domain experts with limited elicitation. In this work, we explore the approach of providing interfaces in which domain experts can create a knowledge repository for a sophisticated domain, so that it could be consumed by data scientists asynchronously and flexibly when the availability of domain experts is limited.

2.2 ML with domain experts

There has been a long-standing desire to increase the involvement of domain experts in model building in both the ML and HCI communities. For example, tasks like text annotation, image annotation involve massive input from domain experts to provide domain-related feedback. Ziva interface is inspired by NLP text annotation tools [54, 60]. We take this design further to acquire domain knowledge for model development and data scientists. Recognizing the challenge of having domain experts label a large quantity of data, many ML works have explored more efficient learning algorithms to reduce the workload [68, 73], or utilize domain experts input as rules [45], constraints [19, 56], prior information [11, 70], or feedback to re-weight features [24, 61]. Given the prominence of label-hungry ML algorithms, weak supervision has become a popular approach to bootstrap labels based on feedback from domain experts [23, 63, 64].

The HCI community is further concerned with the isolation of domain experts from the model development process, requiring thus data scientists to go through lengthy and asynchronous iterations to get their input [9, 58]. To tackle the problem, the sub-field of interactive machine learning (iML) is motivated to enable domain experts or end-users to directly drive model behaviors [9, 36, 80]. Since domain experts might not have training in ML or programming, iML systems elicit their input through intuitive and interactive interfaces (e.g., visualization [39], graphic user interfaces [9], conversational interfaces [18]), and a tight feedback loop for them to adjust their input. A variety of user input have been explored in prior work for different tasks in model development, including unseen training data to help correct the model's mistakes [18, 25], provide new feature-level input [44] or adjustment to feature weights [68], assessment of model performance [10, 26], error preferences [41], parameter choices [52], model ensemble [75], etc.

Research on iML has been especially fruitful for NLP modeling tasks, partly because text data and features (e.g., bag of words) are often comprehensible to people, increasing the likelihood of obtaining effective feedback from domain experts or end-users [46]. For example, the tools solicit feedback for learned features [69] or support features ideation by the users [14]. Interactive topic modeling is another well-explored area to incorporate domain experts' input [21, 37, 38, 71], for example, by moving documents around or adding words, to refine clusters of topics.

Our work is informed by prior work on iML but takes a complementary approach by facilitating knowledge sharing from domain experts to data scientists. iML is not a panacea to effectively leverage domain experts' input. There are known issues with letting ML novices directly adjust models [72], such as lacking generalization or over-fitting [81]. In practice it is not always feasible to set up an iML system for domain experts to work with. Currently most ML projects still rely on data scientists to write code and set up the pipelines [59]. Moreover, having data scientists mediate the knowledge input offers the flexibility to apply it to different kinds of ML algorithms, and allow domain experts to provide reusable knowledge not constrained by a particular modeling task.

In general, it is possible to elicit diverse kinds of knowledge from people, not all of which could be consumed directly by a given ML model. For example, Stumpp et al. [74] and Ghai et al. [28] explored what kind of feedback people naturally want to give seeing model explanations. Only a small subset of the various forms of feedback is readily consumable by existing ML algorithms. However, as the ML field rapidly advances, many novel usages of domain knowledge are being explored. For example, since ML models might use low-level features that are not human understandable (e.g., pixels of an image), interpretable ML works explored eliciting human-interpretable *concepts* in the domain (e.g., an object in the image) and use the concepts to explain the model decisions [29, 42]. Elicited domain concepts have also been used to create sub-groups for labeled data to enable "structure labeling", which could lower the re-labeling burden when a target class changes [47]. We further envision elicited domain concepts could help data scientists head start their model building, as revealed in our preliminary interview. By facilitating knowledge sharing from domain experts, we also hope to inspire novel algorithmic work that could leverage such a knowledge repository.

2.3 Technologies for knowledge sharing

Ziva is also motivated by prior work on technologies that facilitate knowledge sharing in enterprise and organizations. Knowledge sharing has been long studied in the Computer Supported Collaborative Work (CSCW) community focusing on building collective knowledge repositories and locating related experts [5, 67]. Knowledge repository tools elicit various formal and informal information including manuals, best practices, common questions, and so on. For example, Goldberg et al. studied collaborative tagging and filtering mechanisms for workers to construct a knowledge repository [30]; Answer Garden is a system to build a repository through people asking and answering questions [4]; Terveen et al. designed a memory framework for large-scale software engineering where groups collectively build a shared memory [76]; Nam and Ackerman studied methods for elicitation of informal information into more organized forms [55].

Knowledge sharing in ML projects poses unique challenges [8, 16] to make the knowledge transferrable into ML specifications. The challenges are amplified in sophisticated domains. For example, for a medical ML model, a clinician may have to help data scientists understand complex drug information. We inform the design of Ziva both by prior work on involving domain experts in data science projects and model development, and a preliminary interview study to understand how data scientists learn from domain experts. Meanwhile, studies have warned that knowledge sharing and repository tools often fail in practice [33, 82, 88], if the design fails to take into account the social dynamics, including what benefits and demands these technologies bring for both the knowledge providers and the knowledge consumers [5, 31]. Thereby we evaluate Ziva by involving both the knowledge **consumers**—data scientists, and the knowledge **providers**—domain experts.

3 PRELIMINARY INTERVIEW

To understand how data scientists grasp a domain, we conducted semi-structured interviews with four data scientists working on NLP models (2 females, 2 males). Each interview was 45 minutes long and guided by a script that asks participants about their recent projects with domain experts and their typical interactions with those domain experts. We recruited participants via posting on slack channels of an international technology company. Each interviewee was compensated \$15 for their time. We summarized our interviewees' projects and challenges in Table 1. As a result, we identified the current practices of learning from domain experts and design requirements for our tool.

3.1 Limited time and limited best practices

All of our data scientist interviewees indicated they often need domain experts' help and feedback. However, domain experts are busy and have little time to spare. One said: *"The first issue is getting hold of their time... I think hardly I was getting one day a week, you can say one hour a week, not even an entire day."* So data scientists try to extract as much knowledge as they can in the limited time they have. They have to spend significant time preparing for these discussions. For example, they often manually curated examples such as mis-classified instances and instances that contain the unfamiliar keywords to ground the discussion during the meetings with

Table 1: Interviewees information

Pn (domain)	Model (reasons of choosing the model)	Methods of knowledge exchange
P1 (Legal, law)	Rule-based (transparency, few labels)	Instance perturbation
P2 (Disaster recovery)	Supervised neural net (sufficient labelers)	Education session of domain overview domain experts labeling
P2 (")	Rule-based (transparency, few labels)	Domain experts think aloud labeling data
P3 (Customer categorization)	Random forest (transparency)	Pair-authoring (Go over analysis together with domain experts [78])
P4 (Sports)	AutoML (time)	Brute-force model building

domain experts. Even though there is no standard way to extract domain knowledge across different domains, but through *mutual effort* they find what works best for a project. We identified the following approaches to learn domain knowledge from domain experts:

Example-driven conversation: The first field of approaches is domain knowledge sharing based on examples. By inquiring about how and why domain experts would label or make decisions for these examples, data scientists learn *rationales* of how the model should behave for the instances. There are three tactics mentioned by our interviewees. P2 observed domain experts during labeling to learn the domain experts' thought process: *"They would go line by line in front of me so that I can also see what their brain is looking at classifying them."* P1 initially took P2's approach, but due to the complexity of the law domain, explaining rationale required extensive background knowledge. Oftentimes, it is unclear to data scientists how to connect the explanations provided by the domain experts to model specifications. P1 used a strategy called *instance perturbation* – for a given instance, the domain experts were asked to minimally change the instance until the model changes the label and discuss the reasons. With this, data scientists were able to narrow in on the parts of the instance that should be the most important to the model's decision. Instead of aiming to build a perfect model right off the bat, P4 deployed their model first and incrementally improved the model upon domain experts' request. Whenever domain experts encountered mis-classified results, they shared the instances with data scientists and discussed why they were mis-classified.

General background knowledge acquisition: Concepts are key units of information for a given domain, such as notions, entities, components or properties. A set of domain concepts can be seen as a *taxonomy*. Understanding them could help data scientists make sense of the domain. Participants reported approaches to learn concepts in an unfamiliar domain. P2 and P3 said domain experts in one of their projects offered a *lecture* to explain key concepts their domain. For P2, domain experts gave an overview and touched on the basic concepts of each class. P3 *pair-authored* [78] with domain experts to bridge concepts and a mathematical formula that encapsulates the information. With this iterative learning process, data scientists were able to kick start model building. P2 said *"I think that was very helpful because after that, my dependency reduced a bit. I could myself assess that what category they belong to."*

Summary: From our interview, we derived several design requirements to design Ziva. We found that the usage of domain knowledge is not only limited to labeling but also other parts in ML development, sometimes open-ended learning. Thus, interfaces of Ziva ought to facilitate domain-knowledge learning of data scientists in general throughout the development workflow. More specifically, we found that the tool should scaffold domain experts to efficiently elicit domain knowledge within short amount of time (R1). Next, a tool should help data scientists to extract basic domain concepts (R2). Lastly, data scientists indicated that they often learn from domain experts' rationale, especially how they justify a decision or label. Hence, the tool needs to facilitate label justification sharing (R3).

4 ZIVA: INTERFACE FOR ELICITING DOMAIN KNOWLEDGE

This section introduces the interface of Ziva. Ziva provides features for domain experts to create domain concepts and elicit justification from representative instances that are automatically curated by Ziva. We discuss Ziva's different components and how they meet the design requirements in detail.

4.1 Representative sampling for instances creation

As highlighted in the our formative interview, domain experts have limited time for labeling or sharing domain knowledge (R1). Hence, it is important to ask them to review only a few of instances and the sample can cover most concepts in the domain. Ziva extracts such a representative sample of m instances from a large training set of N text instances by the simple method of transforming the original text into 'tf-idf' space, clustering the result using an algorithm such as k-means (setting $k = m$), and, for each cluster, returning the text instance closest to the cluster center. This method is not deterministic, but provides a reasonable set of representative instances, for cases where $m \ll N$.

4.2 Concept creation

Creating a taxonomy is an effective way of organizing information [20, 48]. Ziva provides an interface where SMEs can extract domain concepts (R2). Users are asked to categorize each example instance, presented as a card, via a card-sorting activity. Users first group cards by topic (general concepts of the domain such as atmosphere, food, service, price). Cards in each topic are then further

Figure 1: To facilitate domain knowledge sharing, Ziva presents representative instances and to interfaces to review the instances to domain experts, then which will be used by data scientists.

divided cards into descriptions referencing specific attributes for a topic (e.g., cool, tasty, kind, high). The interface (Figure 2) was implemented as a drag-and-drop UI using LMDD [2].

4.3 Justification-elicitation interface

Once a domain expert finishes the concept extraction, they review each instance using one of elicitation interfaces, which ask the domain expert to justify an instance's label (this information is then intended for consumption by data scientists (R3)). We used Materialize to implement the justification elicitation conditions.

The justification elicitation interfaces were designed through an iterative process of paper prototyping, starting with initial designs inspired by our preliminary interviews. As we conducted paper prototyping, we examined if (1) the answers from different participants were consistent and (2) the information from participants' answers were useful to data scientists. We now describe the five different justification elicitation methods that we created and evaluated, and highlight the design rationale where appropriate.

Bag of words. This base condition reflects the most common current approach. Given an instance and a label (e.g., positive, negative), the domain experts are asked to highlight the text snippets that justify the label assignment.

Instance perturbation. Inspired by one of our data scientists in the formative study, this condition asks a domain expert to *perturb* (edit) a part of the instance such that the assigned label is no longer justifiable by the resulting text. For example, in the restaurant domain, “our server was kind”, can be modified to no longer convey a positive sentiment by either negating an aspect (e.g., “our server was not kind”) or altering it (e.g., “our server was rude”).

This strategy is also inspired by the research area of generating natural language adversarial examples [7]. Such approaches algorithmically alter training examples to create similar adversarial examples that fool well-trained NLP models. In our scenario, the domain expert is seeking to alter training examples in order to point out the most salient characteristics to the data scientist; the latter learns from this information, combining it with syntactic and semantic analysis of the original and perturbed instances.

Instance simplification. This condition asks domain experts to shorten an instance as much as possible, leaving only text that justifies the assigned label of the original instance. For example, “That’s right. The red velvet cake... ohhhh.. it was rich and moist”, can be simplified to “The cake was rich and moist”, as the rest of

the content does not convey any sentiment, and can therefore be judged irrelevant to the sentiment analysis task.

This condition is inspired by the plethora of methods for sentence simplification used in extractive text summarization [77]. In particular, the domain expert is performing *sentence reduction* as in [40]. The output can be considered to be a concise summary of the original instance, keeping only that content which is directly relevant to the sentiment analysis task. The result for the data scientist is clean, compact, and fully relevant high quality training examples.

Concept bag of words. This condition incorporates the concept extracted in the prior step. Similar to the **Bag of words** condition, domain experts are asked to highlight relevant text within each instance to justify the assigned label; however, each highlight must be grouped into one of the concepts. If, during Concept creation, the domain expert copied a card to assign multiple topics and descriptions, then the interface prompts multiple times to highlight relevant text for each one. For example, if they classified the instance, “That’s right. The red velvet cake... ohhhh.. it was rich and moist”, into the concept “food is tasty”, they can select *rich*, *moist* and *cake* as being indicative words for that concept.

Concept annotation. This condition is similar to the above **Concept bag of words** condition. However, when annotating the instance text, domain experts are directed to distinguish between words relevant to the topic and words relevant to the description. Given the above sample instance, the domain expert would need to indicate which part of the sentence applies to *food* (e.g., *cake*) and which to *tasty* (e.g., *rich and moist*). Both this and the previous concept condition are motivated by the well-established knowledge that a variety of NLP tasks, such as relation extraction, question answering, clustering and text generation can benefit from tapping into the the conceptual relationship present in the hierarchies of human knowledge [85]. Learning taxonomies from text corpora is a significant NLP research direction, especially for long-tailed and domain-specific knowledge acquisition [79].

In the rest of the paper, we present a case study to evaluate the utility of the Ziva interface in two parts. In Section 5, we conduct a lab experiment and a crowd experiment in which participants acted as domain experts using Ziva. We choose the domain of restaurant reviews (Yelp Open Dataset[3]) and the NLP task of sentiment analysis, as being extremely familiar and easy enough to understand for most people to qualify as domain experts. In Section 6, we conduct an interview study with data scientists to

Figure 2: Ziva interface: domain experts first extract domain knowledge with curated instances. Then they review each instance one by one using one of justification-elicitation interfaces.

evaluate the utility of domain knowledge collected in the above experiments. We instruct the data scientists to assume no previous knowledge of the domain, so we could use the elicited knowledge about restaurant reviewing as proxy to understand how Ziva could help them build NLP models.

5 EVALUATION ON DOMAIN EXPERTS' EXPERIENCE

We recruited participants to act as domain experts of restaurant reviewing to use Ziva. In a lab study (N=12), we compared participants' task completion and experience with all concept and justification elicitation methods, and gathered their qualitative feedback. To allow quantitatively compare the results of different justification elicitation methods, we conducted a follow-up crowd experiment (N=88).

5.1 Lab study

Study protocol: To avoid noisiness in labeling, we pre-labeled the set of yelp reviews instance so we could focus on comparing the elicitation methods. We created binary labels based on ground-truth ratings: if the number of stars is 1 or 2 for a review, we labeled it as negative, 4 or 5 as positive [87]. We then took a random balanced sample of 10,000 instances. 8,000 were used as a (balanced) 'training set', from which we extracted ten representative instances to use in the study (see Section 4.1.) We set the other 2,000 (balanced) instances to use as a test set for analyzing the performance of models built from the study output (see Section 6).

We recruited participants (5 female, 7 male), who self-report little or no knowledge in ML via posting on slack channels of an international technology company. Participants are designers, graduate students, researchers, trained professionals, skilled laborers, software engineers and project managers. To compensate their time, we ran a \$30 raffle.

Participants were given introduction to the project and a tutorial of the Ziva interface. They were also given a practice task in a different domain, i.e., clothing. For the concept extraction task, all participants used the same interface. For the justification interface,

we randomly assigned each participant one treatment from the elicitation methods without concepts (bag of words, label perturbation and simplification), and one from those with concepts (concept bag of words and concept annotation). Thus, each participant experienced two elicitation interfaces and reviewed 5 instances each. After each interface, participants were asked to fill out the NASA TLX form [32] to evaluate their subjective workload and share their feedback. The entire session lasted about for up to one hour.

Task Results: One participant could not complete the second justification interface. We reported the summary of concepts generated by participants, as well as quantitative and qualitative experience using justification methods.

Concept creation. Participants took 879.7 seconds on average ($\sigma=385.4$). They created 3.92 topics on average ($\sigma=1.11$). Everyone included *Food quality* and *Customer service* in their topics. To assess taxonomy from each domain expert, we examined the consistency between domain experts and coverage of the restaurant domain.

- **Consistency between domain experts:** The union of all topics across all participants includes following 10 topics: *ambiance*, *cuisine*, *food quality*, *customer service*, *additional service*, *complaint*, *speciality*, *reservation*, *location*, and *price*. For each topic, we rated whether each taxonomy intersects with the topic or not. Thus, the inter-rater reliability (IRR) across all domain experts was 58% using Fleiss' κ .
- **Coverage of the domain:** We selected 3 additional instances which were not shown to the participants. We used our curation method to pick another set of representative instances. We then inspected how many instances can be categorized using each taxonomy resulting in a coverage of 69% (25 out of 36 instances).

Justification elicitation. The average task completion time in each condition is summarized in Table 2. Since each participant was assigned two out of five justification elicitation, there were only a few data points per elicitation technique (3 to 5 per technique). To further investigate in a larger population, we deployed Ziva on a crowd platform described in the following section.

Table 2: Average task completion time (standard deviation) of lab study participants

Bag of words	Simplification	Perturbation	Concept bag of words	Concept annotation
39.2 s (20.7)	106.6 s (86.8)	107.2 s (48.0)	36.8 s (14.0)	81.9 s (40.5)

Most participants found the **bag of words** condition easy to complete. One participant said: *“This was easy because a lot of words were clearly positive or negative, such as “terrible” or “delicious”*”. However, some considered it tricky to identify words that are indicative of the overall sentiment. For example, one participant said, *“this can be just descriptive without any positive or negative feelings without the context. So it’s difficult to isolate the context out of the words.”*

For the **simplification** task, participants indicated the task was straightforward. Participants said *“easy as it had eliminated redundant and unnecessary words”* and *“quite easy and intuitive, paraphrasing keeping the original intent is what I usually do as part of minutes of meetings”*. One participant said sometimes the task became hard because some instances could not be obviously shortened and instead need to be entirely rewritten.

Participants said **perturbation** is also straightforward but it required them to understand the entire instance thoroughly. One participant commented, *“It was kind of hard because I don’t know some of the words”*. Another participant suggested that if the interface suggested antonyms, it would be easier to finish the task.

With **concept bag of words**, participants said it allowed subjective and nuanced elicitation, as they could pick words associated with a concept without judging their sentiment. However, it led to more varied results among participants. For example, for the concept *Food is tasty*, and the instance *Ohhhh... The red velvet cake is rich and moist*, most participants selected rich and moist. One participant said *“Even red velvet cake could be the indicative words if you personally like the cake”*. Others said *“maybe ohhhh part can be included”* and *“moist doesn’t necessarily mean delicious”*.

For the **concept annotation** task, participants said it is straightforward to choose words directly mapped to each token. On the other hand, it complicated the articulation to have to label in such fine granularity. One participant commented, *“slightly tedious as it required me to comprehend on how best to label the words accordingly”*.

5.2 Crowd Experiments

To assess different justification methods on larger population, we deployed the Ziva interface on a crowd platform.

Study Protocol: Using our representative sampling method, we extracted 10 reviews from the datasets used in the lab study. We pre-populated a taxonomy. In order to provide a representative sampled concept, we recruited 5 volunteers and asked them to extract concepts of the restaurant domain using the concept extraction interface and two of the authors aggregated the taxonomy.

We installed 5 test questions for each condition with ground-truth created by the authors. If a crowd worker did not pass more than half of test questions, they can not continue to the Human Intelligence Task (HIT). Each worker was given one of the five justification elicitation interfaces, and reviewed 10 instances.

We recruited our study participants from Appen [1]. We compensated them with \$0.5 per HIT, they are rewarded \$2.5 in addition

for test questions. From the lab study, we observed each HIT took less than 2 minutes on average, which makes hourly wage of \$15. After the tasks, we asked them to fill out the same NASA TLX form to report on their subjective workload. Participants were rewarded additional \$3 for the survey. A total of 88 crowd workers completed our study resulting in 857 instances with elicited data.

Result: We analyzed participants’ survey responses using an one-way Kruskal–Wallis ANOVA as summarized in Table 3. There was marginal difference in self-reported success of task accomplishment and significant difference in stress level across justification elicitation methods.

As a post-hoc analysis, we ran a one-tailed Mann-Whitney U Test. The result revealed that participants completed the tasks using **bag of words** perceived higher success in accomplishing the tasks than participants with **simplification** ($U=76.5$, $z=2.31$; $p=.01$) and **concept annotation** ($U=97.5$, $z=2.02$; $p=.02$). **Concept bag of words** users also perceived higher success than **simplification** ($U=75.5$, $z=2.34$; $p=.009$) and **concept annotation** users ($U=103$, $z=1.85$; $p=.03$).

As for the stress level, **bag of words** users reported significantly higher stress than **perturbation** ($U=55$, $z=2.90$; $p=.002$), **concept bag of words** ($U=84.5$, $z=-2.24$; $p=.01$), and **concept annotation** users ($U=81.5$, $z=-2.34$; $p=.01$).

6 DATA SCIENTISTS INTERVIEW STUDY

To investigate what and how domain knowledge extracted from Ziva helps data scientists, we conducted an interview study with data scientists. We showed them concept and different justification results extracted by domain experts and asked them how they could use them in their ML development workflow.

Study Protocol: Participants were given introduction to the project, prompts shown to domain experts and corresponding outputs of each part of the interface. Each interview was 1 hour long and driven by a questionnaire that posed questions related to compare domain knowledge extracted by domain experts using Ziva to their current practice. Finally, they were asked to rank usefulness of justification interface to their workflow.

We recruited 7 data scientists who have between 4 and 20 years of experience building models with domain experts in sophisticated domains, using the slack channels of an international technology company and word-of-mouth.

Results: We re-ranked the scores on a linear scale, with a data scientist’s favorite at 5 points, the second-most favorite at 4 and so on. If two techniques were tied for N-th rank, we averaged the scores for the both of techniques (e.g., if two techniques are 4th, they are given 1.5). As a result, the **concept annotation** technique scored the most (30), then **concept bag of words** and **perturbation** (22.5), **simplification** (17.5), and **bag of words** (12.5). Data scientists had several reasons why they prefer one justification technique to another and applications for different techniques. Through the metrics, we were able to identify the design requirements and important factors of domain-knowledge learning.

Table 3: Crowd experiment Likert result. H statistics (p-value) in significance level 0.05

Mentally demanding	Successfully accomplishing	Hard to accomplish	Insecure, Stressed
2.0825 (.72059)	8.7959 (.06641)	8.0609 (.08937)	9.9411 (.04143)

Figure 3: Post-question responses in NASA TLX (1- Very low, 7- Very high) (Crowd experiment participants, n=88)

Standardized protocols. As revealed in our preliminary interview and previous work [50], there is no set protocol of communications or *common ground* between two parties, and expressed a need for a *protocol of communication with domain experts*. The steep learning curve of a specialized domain and lack of guidance for how to extract domain knowledge exacerbates the collaboration with domain experts. Three of interviewees said that having such a concept and examples upfront provided by domain experts has helped them to build a model in prior projects. One said, “They describe what are the component information and examples. It was not very difficult to understand after reading the documents.” In light of this, interviewees preferred justification techniques to inform them about the domain. For example, interviewees found *concept annotation* helpful because it is tightly connected with the concept, hence they can learn from examples how different components of the domain is expressed in the instances. *Simplification* is also helpful, as it is a simpler version of instances without rhetoric.

One interviewee suggested to use justification techniques to *explain model decisions*. They said, “I work on active learning ML a lot. So I work with users. And so far all the interactions I expect for the user, fairly simple, either binary feedback – correct or incorrect. I have any incorporated explanation of when the user provide feedback. What’s the explanation behind this feedback? I think that that would be very useful to generate some explanation or learn how to generate explanation.” While model explanation is not the intended usage of the Ziva’s justification techniques, the data scientist found the techniques helpful for debugging a model.

Scalability of domain knowledge. Interviewees were also interested in how they would *scale* the Ziva output. Since they only received only 10 labeled instances and justification, it was too small for data scientists to train a model.

One application of Ziva output mentioned by data scientists is to label more instances by generalizing concept and justification, so called *weak-supervised learning* [62]. One interviewee said, “They are trying to give me guidance on how to propagate the labels. So one, the concept is going to be able to give me some notion on how to bucket my data, right, like, just in an unsupervised fashion.”

Interviewees stressed the importance of domain knowledge in *feature engineering*. During meetings with domain experts, they focus on identifying features for their model: “I immediately start

looking at what are the different features or abstractions of features that seem to be important to the domain expert.” However, data scientists expressed the difficulty of feature ideation in building models in a specialized domain. Repeat meetings were required to go over many instances together in hope of covering the complete set of features. 3 of interviewees said they would use Ziva output to facilitate feature engineering by using the concepts created by domain experts as features. A participant explained: “Vector that we can convert each restaurant record into a some feature vector.” When it comes to the best justification techniques for feature engineering, one said “The one with the highest resolution would be more beneficial for feature learning potentially because it allow me to generalize better”. One data scientist suggested that they can propagate the feature across different components (e.g., food/food quality, service) of the domain expert’s concepts using *distributional signature* [13]. For instance, in a restaurant domain, once they identified positive-sentiment words related to food, they can find similar sentiment of words related to service using the distribution of words.

Reduced burden on domain experts. We also found data scientists were being mindful of domain experts’ cognitive load when they generated Ziva output, because domain experts were often busy. Another reason is if the eliciting justification is difficult, data scientists would not get a reliable result. One interviewee said: “I would say there’s also the question of what I think would be more easier for people, if it’s difficult, then they’re probably not going to do it very well. I wouldn’t give it to them because I would think it’s going to be more noisy.”

Elicitation and learning outcomes. To demonstrate the feasibility of translating the Ziva output into useful features for model building, we constructed a real implementation. Inspired by a use case suggested by a data scientist in our study, we built 5 models for weak-supervised learning, mimicking a real-world cold start scenario with extremely limited labeling resources and no pre-trained model available. With such constraints, no one can expect state-of-the-art performance after a few training examples. Instead, a valuable characteristic at this early stage is intra-class consistency, demonstrating parallel improvement in precision and recall performance on various classes (here, positive and negative sentiment). This would suggest that the model is indeed learning something relevant to the entire task rather than guessing wildly, and hints at

a good robustness that can be reliably improved upon with additional examples. We therefore worked with rule models, which are both consistent and explainable, relatively simple for a human to construct, and need very little labeled data to generate candidate rules. However, the interviewed data scientists agreed that at scale, the extracted elicitations could serve as features for a variety of learning models.

Excepting the **bag of words** condition, the models primarily focused on recognizing the semantic pattern of ‘Noun is [not] Adjective’. Of course, this can take several forms (‘food is good’, ‘good food’, ‘food is not bad’, etc.) We built rule-based models that extend a generic semantic role labeling model [6] which can easily handle such variations. The generic model identifies all existing semantic roles, and the ten instances, annotated in each condition, are used to populate the dictionaries that those roles should match on (e.g., ‘food’ and ‘good’). In general, we were careful during model construction to not make any use of additional external knowledge (e.g., we do not know that ‘hot wings’ and ‘burgers’ are both a type of food, if this information was not in the output of Ziva.) Below we describe the details of each elicitation method and discuss the results, which are summarized in Table 4:

Bag of words. This was simple keyword matching on the terms identified in this condition. The positive terms output from this condition were mostly generic (‘amazing’, ‘delicious’) whereas many negative terms were very specific (‘over-hyped’, ‘small quantities’). This is an artifact of both the domain (restaurant reviews) and the labels. The performance on the two classes reflects this: the positive class has pretty bad precision but great recall, as it severely over-generalizes, whereas the negative class has amazing precision but barely finds any examples, because it is so specific.

Perturbation. The perturbed parts of the instances were treated as local training instances. All possible ‘Noun is Adjective’ signals were extracted from those instances to populate the relevant dictionaries. If a verb was negated, or an adjective transformed into an antonym (e.g., changing ‘delicious’ to ‘disgusting’ in ‘There were delicious burgers’, assigned a positive label), this meant that the topic (‘burgers’) is highly relevant, the original text (‘delicious burgers’) was a training example for the given label, and the perturbed result (‘disgusting burgers’) was for the opposite label.

Simplification. The simplified instances were treated as high quality training instances. All possible ‘Noun is Adjective’ signals were extracted from those instances to populate the relevant dictionaries. These signals did not overlap much in content, so the model could do little generalizing. Much like the **perturbation** condition, the recall for both classes is therefore extremely low, and the precision is respectable for only 10 training examples. **Perturbation** recall results are slightly better because each perturbed instance yields both a positive and a negative signal.

Concept bag of words and Concept annotation The concept taxonomy described in Section 5.2 follows the ‘Noun is Adjective’ format by definition, so it was encoded accordingly for both of these conditions. The outputs of each condition were then used to extend the possible dictionaries. For *concept bag of words*, each annotation was added to both the ‘Noun’ and ‘Adjective’ dictionaries (whenever grammatically possible). For *concept annotation*, the ‘Noun’ and ‘Adjective’ elements were elicited separately, and thus were added

to their respective dictionaries. It is unclear that either of these conditions is more successful than the other, at this stage. The recall is markedly better than for *simplification* and *perturbation* owing to the well-structured concept taxonomy, that lends itself well to generalization. But this comes at a price, as the delta in performance between the classes is similarly worse.

7 DISCUSSION

Capturing nuanced domain knowledge. While Ziva provides some basic components in a domain, data scientists pointed out there is information that the current design of Ziva does not reflect. For instance, domain experts provide insight about data, such as sparsity of a certain column. Data scientists find such information helpful, but it can not be captured in the Ziva output. More investigation is required on how to extract such nuanced knowledge. One possible direction is to leverage proposed documentation for data [27] or for models [12, 51]. Another tactic is to take a set of guided questions similar to the ones proposed in [66] in discussions between domain experts and data scientists. Structure provided by such artifacts can facilitate domain knowledge transfer and get teams on the same page quickly.

Re-evaluating the old normal: Bag of words. Bag of words is one of the dominant ways in the NLP domain to elicit signals. It appears to be most simple and straightforward task for domain experts. However, to our surprise, our work suggests otherwise. Participants in our user study indicated that **bag of words** is in fact more mentally demanding, harder to accomplish and more stressful for them than other justification techniques. Furthermore, in our exercise of building a rule-based model with different justification methods, the other methods outperformed the **bag of words**. This suggests that both domain experts and data scientists can benefit from our justification techniques during collaboration. We believe our justification method could be used throughout the ML development workflow and provide an outlet for stakeholders to efficiently communicate during model building.

Limitations. Various use cases of the Ziva output validated the efficacy of our interfaces drawn from our preliminary study and literature review, demonstrating that domain experts’ elicited knowledge can facilitate model building. This paper only directly considers the concrete setting of a sentiment classifier for Yelp restaurant reviews. Nevertheless, the overall approach described in this paper is domain-agnostic, and extremely relevant to real-life scenarios with complex tasks, specialized domains, and significant constraints on the resources to generate large amounts of labeled data. Further, although only rule-based models were built, the semantic role-based features constructed from the output are quite appropriate as input to other learning approaches. Future work should examine the generalizability of the approach for other tasks (e.g., document classification, clustering, machine translation, and question answering), other domains (e.g., education, health science), and other learning models. We therefore believe we have identified a number of interesting design requirements of domain-knowledge sharing in the ML development workflow that are not currently addressed, and are applicable across tasks and domains.

Table 4: Performance of Rule-Based Models on 2,000 test instances, for different justification conditions on 10 training instances. Because the test dataset is balanced, the Recall (R) value is equivalent to Accuracy. The last three columns are the really meaningful ones, as they highlight the absolute differences in Precision/Recall/F1 between the two classes (lower is better; values below 0.10 are highlighted). The Trivial model, which always assigns a positive label to each instance, is shown for reference.

	Positive Class			Negative Class			Delta Between Classes		
	P	R	F	P	R	F	P	R	F
<i>Trivial (Always Pos)</i>	0.5	1.0	0.667	0.0	0.0	0.0	0.5	1.0	0.667
Bag of Words	0.641	0.886	0.744	0.968	0.03	0.058	0.327	0.856	0.686
Perturbation	0.768	0.076	0.138	0.891	0.041	0.078	0.123	0.035	0.060
Simplification	0.775	0.069	0.127	0.857	0.030	0.058	0.082	0.039	0.069
Concept Bag of Words	0.735	0.219	0.337	0.836	0.102	0.182	0.101	0.117	0.155
Concept Annotation	0.723	0.245	0.366	0.806	0.112	0.197	0.083	0.133	0.169

8 CONCLUSION

In this paper, we presented a system and a case study on how data scientists can get help from domain experts in the ML development lifecycle. Along the way, we identified the current practice of how data scientists acquire domain knowledge. Inspired by the existing workarounds, we designed an interface that facilitates the sharing of expert domain knowledge. We presented the interface output to ML practitioners to reflect their experience building an ML model in a specialized domain, from which we learned that scalability of a piece of domain knowledge and low cognitive load of domain experts are important factors for any domain knowledge-bootstrapping tool. We continued the work by investigating the cognitive load of different methods in our interface. We found that the traditional and most-used elicitation method “bag of words” is actually the least preferred by domain experts in terms of mental load and stress level, and provides the least knowledge scalability compared to other elicitation methods.

ACKNOWLEDGMENTS

We thank Dakuo Wang, David Karger and Ranit Aharonov for their feedback.

REFERENCES

- [1] 2021. Appen. <https://appen.com>.
- [2] 2021. Lean-Mean-Drag-and-Drop. <https://supraniti.github.io/Lean-Mean-Drag-and-Drop/>.
- [3] 2021. Yelp Open Dataset. <https://www.yelp.com/dataset>.
- [4] Mark S Ackerman. 1998. Augmenting organizational memory: a field study of answer garden. *ACM Transactions on Information Systems (TOIS)* 16, 3 (1998), 203–224.
- [5] Mark S Ackerman, Juri Dachtera, Volkmar Pipek, and Volker Wulf. 2013. Sharing knowledge and expertise: The CSCW view of knowledge management. *Computer Supported Cooperative Work (CSCW)* 22, 4-6 (2013), 531–573.
- [6] A. Akbik and Yunyao Li. 2016. K-SRL: Instance-based Learning for Semantic Role Labeling. In *COLING*.
- [7] Moustafa Alzantot et al. 2018. Generating Natural Language Adversarial Examples. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2890–2896. <https://doi.org/10.18653/v1/D18-1316>
- [8] Saleema Amershi et al. 2019. Software engineering for machine learning: A case study. In *2019 IEEE/ACM 41st International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)*. IEEE, 291–300.
- [9] Saleema Amershi, Maya Cakmak, William Bradley Knox, and Todd Kulesza. 2014. Power to the people: The role of humans in interactive machine learning. *Ai Magazine* 35, 4 (2014), 105–120.
- [10] Saleema Amershi, Max Chickering, Steven M Drucker, Bongshin Lee, Patrice Simard, and Jina Suh. 2015. Modeltracker: Redesigning performance analysis tools for machine learning. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. 337–346.
- [11] David Andrzejewski, Xiaojin Zhu, and Mark Craven. 2009. Incorporating domain knowledge into topic modeling via Dirichlet forest priors. In *Proceedings of the 26th annual international conference on machine learning*. 25–32.
- [12] Matthew Arnold, Rachel KE Bellamy, Michael Hind, Stephanie Houde, Sameep Mehta, A Mojsilović, Ravi Nair, K Natesan Ramamurthy, Alexandra Olteanu, David Piorkowski, et al. 2019. FactSheets: Increasing trust in AI services through supplier’s declarations of conformity. *IBM Journal of Research and Development* 63, 4/5 (2019), 6–1.
- [13] Yujia Bao, Menghua Wu, Shiyu Chang, and Regina Barzilay. 2019. Few-shot text classification with distributional signatures. *arXiv preprint arXiv:1908.06039* (2019).
- [14] Michael Brooks, Saleema Amershi, Bongshin Lee, Steven M Drucker, Ashish Kapoor, and Patrice Simard. 2015. FeatureInsight: Visual support for error-driven feature ideation in text classification. In *2015 IEEE Conference on Visual Analytics Science and Technology (VAST)*. IEEE, 105–112.
- [15] Carrie Jun Cai et al. 2019. Human-Centered Tools for Coping with Imperfect Algorithms during Medical Decision-Making. <https://arxiv.org/abs/1902.02960>
- [16] Carrie J Cai and Philip J Guo. 2019. Software Developers Learning Machine Learning: Motivations, Hurdles, and Desires. In *2019 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*. IEEE, 25–34.
- [17] Carrie Jun Cai, Samantha Winter, David Steiner, Lauren Wilcox, and Michael Terry. 2019. “Hello AI”: Uncovering the Onboarding Needs of Medical Practitioners for Human-AI Collaborative Decision-Making.
- [18] Maya Cakmak and Andrea L Thomaz. 2012. Designing robot learners that ask good questions. In *2012 7th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 17–24.
- [19] Ming-Wei Chang, Lev Ratinov, and Dan Roth. 2007. Guiding semi-supervision with constraint-driven learning. In *Proceedings of the 45th annual meeting of the association of computational linguistics*. 280–287.
- [20] Lydia B Chilton, Greg Little, Darren Edge, Daniel S Weld, and James A Landay. 2013. Cascade: Crowdsourcing taxonomy creation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1999–2008.
- [21] Jaegul Choo, Changhyun Lee, Chandan K Reddy, and Haesun Park. 2013. Utopian: User-driven topic modeling based on interactive nonnegative matrix factorization. *IEEE transactions on visualization and computer graphics* 19, 12 (2013), 1992–2001.
- [22] Robert Culkin and Sanjiv R Das. 2017. Machine learning in finance: The case of deep learning for option pricing. *Journal of Investment Management* 15, 4 (2017), 92–100.
- [23] Mostafa Dehghani, Hamed Zamani, Aliaksei Severyn, Jaap Kamps, and W Bruce Croft. 2017. Neural ranking models with weak supervision. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 65–74.
- [24] Gregory Druck, Burr Settles, and Andrew McCallum. 2009. Active learning by labeling features. In *Proceedings of the 2009 conference on Empirical methods in natural language processing*. 81–90.
- [25] Jerry Alan Fails and Dan R Olsen Jr. 2003. Interactive machine learning. In *Proceedings of the 8th international conference on Intelligent user interfaces*. 39–45.
- [26] James Fogarty, Desney Tan, Ashish Kapoor, and Simon Winder. 2008. CueFlik: interactive concept learning in image search. In *Proceedings of the sigchi conference on human factors in computing systems*. 29–38.
- [27] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. 2018. Datasheets for datasets. *arXiv preprint arXiv:1803.09010* (2018).

- [28] Bhavya Ghai, Q Vera Liao, Yunfeng Zhang, Rachel Bellamy, and Klaus Mueller. 2020. Explainable Active Learning (XAL): An Empirical Study of How Local Explanations Impact Annotator Experience. *arXiv preprint arXiv:2001.09219* (2020).
- [29] Amirata Ghorbani, James Wexler, James Y Zou, and Been Kim. 2019. Towards automatic concept-based explanations. In *Advances in Neural Information Processing Systems*. 9273–9282.
- [30] Yaron Goldberg, Marilyn Safra, and Ehud Shapiro. 1992. Active mail—a framework for implementing groupware. In *Proceedings of the 1992 ACM conference on Computer-supported cooperative work*. 75–83.
- [31] Jonathan Grudin. 1988. Why CSCW applications fail: problems in the design and evaluation of organizational interfaces. In *Proceedings of the 1988 ACM conference on Computer-supported cooperative work*. 85–93.
- [32] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in psychology*. Vol. 52. Elsevier, 139–183.
- [33] Sven Hoffmann et al. 2019. Cyber-Physical Systems for Knowledge and Expertise Sharing in Manufacturing Contexts: Towards a Model Enabling Design. *Computer Supported Cooperative Work (CSCW)* 28, 3-4 (2019), 469–509.
- [34] Fred Hohman, Andrew Head, Rich Caruana, Robert DeLine, and Steven M Drucker. 2019. Gamut: A design probe to understand how data scientists understand machine learning models. In *Proceedings of the 2019 CHI conference on human factors in computing systems*. 1–13.
- [35] Fred Hohman, Minsuk Kahng, Robert Pienta, and Duen Horng Chau. 2018. Visual analytics in deep learning: An interrogative survey for the next frontiers. *IEEE transactions on visualization and computer graphics* 25, 8 (2018), 2674–2693.
- [36] Andreas Holzinger. 2016. Interactive machine learning for health informatics: when do we need the human-in-the-loop? *Brain Informatics* 3, 2 (2016), 119–131.
- [37] Enamul Hoque and Giuseppe Carenini. 2015. Convisit: Interactive topic modeling for exploring asynchronous online conversations. In *Proceedings of the 20th International Conference on Intelligent User Interfaces*. 169–180.
- [38] Yuening Hu, Jordan Boyd-Graber, Brianna Satinoff, and Alison Smith. 2014. Interactive topic modeling. *Machine learning* 95, 3 (2014), 423–469.
- [39] Liu Jiang, Shixia Liu, and Changjian Chen. 2019. Recent research advances on interactive machine learning. *Journal of Visualization* 22, 2 (2019), 401–417.
- [40] Hongyan Jing. 2000. Sentence reduction for automatic text summarization. In *Proceedings of the sixth conference on Applied natural language processing (ANLP '00)*. Association for Computational Linguistics, 310–315. <https://doi.org/10.3115/974147.974190>
- [41] Ashish Kapoor, Bongshin Lee, Desney Tan, and Eric Horvitz. 2010. Interactive optimization for steering machine classification. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1343–1352.
- [42] Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viegas, et al. 2018. Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In *International conference on machine learning*. PMLR, 2668–2677.
- [43] Miryung Kim, Thomas Zimmermann, Robert DeLine, and Andrew Begel. 2016. The emerging role of data scientists on software development teams. In *Proceedings of the 38th International Conference on Software Engineering*. ACM, 96–107.
- [44] Josua Krause, Adam Perer, and Enrico Bertini. 2014. INFUSE: interactive feature selection for predictive modeling of high dimensional data. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 1614–1623.
- [45] Rajasekar Krishnamurthy, Yunyao Li, Sriram Raghavan, Frederick Reiss, Shivakumar Vaithyanathan, and Huaiyu Zhu. 2009. SystemT: a system for declarative information extraction. *ACM SIGMOD Record* 37, 4 (2009), 7–13.
- [46] Todd Kulesza, Margaret Burnett, Weng-Keen Wong, and Simone Stumpf. 2015. Principles of explanatory debugging to personalize interactive machine learning. In *Proceedings of the 20th international conference on intelligent user interfaces*. 126–137.
- [47] Todd Kulesza, Denis Charles, Rich Caruana, Saleema Amin Amershi, and Danyel Aharon Fisher. 2019. Structured labeling to facilitate concept evolution in machine learning. US Patent 10,318,572.
- [48] David Laniado, Davide Eynard, Marco Colombetti, et al. 2007. Using WordNet to turn a folksonomy into a hierarchy of concepts. In *Semantic Web Application and Perspectives-Fourth Italian Semantic Web Workshop*. 192–201.
- [49] James Manyika, Michael Chui, Mehdi Miremadi, et al. 2017. A future that works: AI, automation, employment, and productivity. *McKinsey Global Institute Research, Tech. Rep* 60 (2017).
- [50] Yaoli Mao et al. 2019. How Data Scientists Work Together With Domain Experts in Scientific Collaborations: To Find The Right Answer Or To Ask The Right Question? *Proceedings of the ACM on Human-Computer Interaction* 3, GROUP (2019), 1–23.
- [51] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. 2019. Model cards for model reporting. In *Proceedings of the conference on fairness, accountability, and transparency*. 220–229.
- [52] Thomas Mühlbacher, Lorenz Linhardt, Torsten Möller, and Harald Piringer. 2017. Treepod: Sensitivity-aware selection of pareto-optimal decision trees. *IEEE transactions on visualization and computer graphics* 24, 1 (2017), 174–183.
- [53] Michael Muller, Ingrid Lange, Dakuo Wang, David Piorkowski, Jason Tsay, Q. Vera Liao, Casey Dugan, and Thomas Erickson. 2019. How Data Science Workers Work with Data: Discovery, Capture, Curation, Design, Creation. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems* (Glasgow, UK) (CHI '19). ACM, New York, NY, USA, Forthcoming.
- [54] Hiroki Nakayama et al. 2018. doccano: Text Annotation Tool for Human. <https://github.com/doccano/doccano>
- [55] Kevin K Nam and Mark S Ackerman. 2007. Arkose: reusing informal information from online discussions. In *Proceedings of the 2007 international ACM conference on Supporting group work*. 137–146.
- [56] Radu Stefan Niculescu, Tom M Mitchell, and R Bharat Rao. 2006. Bayesian network learning with parameter constraints. *Journal of machine learning research* 7, Jul (2006), 1357–1383.
- [57] Samir Passi and Steven J Jackson. 2018. Trust in data science: collaboration, translation, and accountability in corporate data science projects. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–28.
- [58] Claudio Pinhanes. 2019. Machine Teaching by Domain Experts: Towards More Humane, Inclusive, and Intelligent Machine Learning Systems. *arXiv preprint arXiv:1908.08931* (2019).
- [59] David Piorkowski, Soya Park, April Yi Wang, Dakuo Wang, Michael Muller, and Felix Portnoy. 2021. How AI Developers Overcome Communication Challenges in a Multidisciplinary Team: A Case Study. *arXiv:2101.06098 [cs.CY]*
- [60] prodigy. 2018. prodigy. <https://prodigy.ai>
- [61] Hema Raghavan, Omid Madani, and Rosie Jones. 2006. Active learning with feedback on features and instances. *Journal of Machine Learning Research* 7, Aug (2006), 1655–1686.
- [62] Alex Ratner, Stephen Bach, Paroma Varma, and Chris Ré. 2019. Weak supervision: the new programming paradigm for machine learning. *Hazy Research. Available via https://dawn.cs.stanford.edu/2017/07/16/weak-supervision/*. Accessed (2019), 05–09.
- [63] Alexander Ratner, Stephen H Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. 2017. Snorkel: Rapid training data creation with weak supervision. In *Proceedings of the VLDB Endowment. International Conference on Very Large Data Bases*, Vol. 11. NIH Public Access, 269.
- [64] Alexander J Ratner, Christopher M De Sa, Sen Wu, Daniel Selsam, and Christopher Ré. 2016. Data programming: Creating large training sets, quickly. In *Advances in neural information processing systems*. 3567–3575.
- [65] Adam Rule, Ian Drosos, Aurélien Tabard, and James D Hollan. 2018. Aiding collaborative reuse of computational notebooks with annotated cell folding. *Proceedings of the ACM on Human-Computer Interaction* 2, CSCW (2018), 1–12.
- [66] Shems Saleh, William Boag, Lauren Erdman, and Tristan Naumann. 2020. Clinical Collabsheets: 53 Questions to Guide a Clinical Collaboration. In *Machine Learning for Healthcare Conference*. PMLR, 783–812.
- [67] A Th Schreiber, Guus Schreiber, Hans Akkermans, Anjo Anjewierden, Nigel Shadbolt, Robert de Hoog, Walter Van de Velde, Bob Wielinga, R Nigel, et al. 2000. *Knowledge engineering and management: the CommonKADS methodology*. MIT press.
- [68] Burr Settles. 2009. *Active learning literature survey*. Technical Report. University of Wisconsin-Madison Department of Computer Sciences.
- [69] Burr Settles. 2011. Closing the loop: Fast, interactive semi-supervised annotation with queries on features and instances. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*. 1467–1478.
- [70] Patrice Y Simard, Saleema Amershi, David M Chickering, Alicia Edelman Pelton, Soroush Ghorashi, Christopher Meek, Gonzalo Ramos, Jina Suh, Johan Verwey, et al. 2017. Machine teaching: A new paradigm for building machine learning systems. *arXiv preprint arXiv:1707.06742* (2017).
- [71] Alison Smith, Varun Kumar, Jordan Boyd-Graber, Kevin Seppi, and Leah Findlater. 2018. Closing the loop: User-centered design and evaluation of a human-in-the-loop topic modeling system. In *23rd International Conference on Intelligent User Interfaces*. 293–304.
- [72] Alison Smith-Renner, Varun Kumar, Jordan Boyd-Graber, Kevin Seppi, and Leah Findlater. 2020. Digging into user control: perceptions of adherence and instability in transparent models. In *Proceedings of the 25th International Conference on Intelligent User Interfaces*. 519–530.
- [73] Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*. 4077–4087.
- [74] Simone Stumpf, Vidya Rajaram, Lida Li, Margaret Burnett, Thomas Dietterich, Erin Sullivan, Russell Drummond, and Jonathan Herlocker. 2007. Toward harnessing user feedback for machine learning. In *Proceedings of the 12th international conference on Intelligent user interfaces*. 82–91.
- [75] Justin Talbot, Bongshin Lee, Ashish Kapoor, and Desney S Tan. 2009. EnsembleMatrix: interactive visualization to support machine learning with multiple classifiers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 1283–1292.
- [76] Loren G Terveen, Peter G Selfridge, and M David Long. 1995. Living design memory: framework, implementation, lessons learned. *Human-Computer Interaction*

- 10, 1 (1995), 1–37.
- [77] Rafaella Vale et al. 2020. An Assessment of Sentence Simplification Methods in Extractive Text Summarization. In *Proceedings of the ACM Symposium on Document Engineering 2020 (DocEng '20)*. Association for Computing Machinery, Article 9, 9 pages. <https://doi.org/10.1145/3395027.3419588>
- [78] April Yi Wang, Anant Mittal, Christopher Brooks, and Steve Oney. 2019. How Data Scientists Use Computational Notebooks for Real-Time Collaboration. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–30.
- [79] Chengyu Wang, Xiaofeng He, and Aoying Zhou. 2017. A Short Survey on Taxonomy Learning from Text Corpora: Issues, Resources and Recent Advances. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 1190–1203. <https://doi.org/10.18653/v1/D17-1123>
- [80] Malcolm Ware, Eibe Frank, Geoffrey Holmes, Mark Hall, and Ian H Witten. 2001. Interactive machine learning: letting users build classifiers. *International Journal of Human-Computer Studies* 55, 3 (2001), 281–292.
- [81] Tongshuang Wu, Daniel S Weld, and Jeffrey Heer. 2019. Local Decision Pitfalls in Interactive Machine Learning: An Investigation into Feature Selection in Sentiment Analysis. *ACM Transactions on Computer-Human Interaction (TOCHI)* 26, 4 (2019), 1–27.
- [82] Chi-Lan Yang, Chien Wen Yuan, and Hao-Chuan Wang. 2019. When Knowledge Network is Social Network: Understanding Collaborative Knowledge Transfer in Workplace. *Proceedings of the ACM on Human-Computer Interaction* 3, CSCW (2019), 1–23.
- [83] Qian Yang, Aaron Steinfeld, and John Zimmerman. 2019. Unremarkable ai: Fitting intelligent decision support into critical, clinical decision-making processes. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. 1–11.
- [84] Amy X Zhang, Michael Muller, and Dakuo Wang. 2020. How do Data Science Workers Collaborate? Roles, Workflows, and Tools. *arXiv preprint arXiv:2001.06684* (2020).
- [85] Hao Zhang et al. 2016. Learning Concept Taxonomies from Multi-modal Data. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 1791–1801. <https://doi.org/10.18653/v1/P16-1169>
- [86] Jiawei Zhang, Yang Wang, Piero Molino, Lezhi Li, and David S Ebert. 2018. Manifold: A model-agnostic framework for interpretation and diagnosis of machine learning models. *IEEE transactions on visualization and computer graphics* 25, 1 (2018), 364–373.
- [87] Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-Level Convolutional Networks for Text Classification. *arXiv:1509.01626 [cs]* (Sept. 2015).
- [88] Xiaomu Zhou, Mark Ackerman, and Kai Zheng. 2011. CPOE workarounds, boundary objects, and assemblages. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. 3353–3362.