ShortcutLens: A Visual Analytics Approach for Exploring Shortcuts in Natural Language Understanding Dataset

Zhihua Jin, Xingbo Wang, Furui Cheng, Chunhui Sun, Qun Liu, and Huamin Qu, Member, IEEE

Abstract—Benchmark datasets play an important role in evaluating Natural Language Understanding (NLU) models. However, shortcuts—unwanted biases in the benchmark datasets—can damage the effectiveness of benchmark datasets in revealing models' real capabilities. Since shortcuts vary in coverage, productivity, and semantic meaning, it is challenging for NLU experts to systematically understand and avoid them when creating benchmark datasets. In this paper, we develop a visual analytics system, *ShortcutLens*, to help NLU experts explore shortcuts in NLU benchmark datasets. The system allows users to conduct multi-level exploration of shortcuts. Specifically, Statistics View helps users grasp the statistics such as coverage and productivity of shortcuts in the benchmark dataset. Template View employs hierarchical and interpretable templates to summarize different types of shortcuts. Instance View allows users to check the corresponding instances covered by the shortcuts. We conduct case studies and expert interviews to evaluate the effectiveness and usability of the system. The results demonstrate that *ShortcutLens* supports users in gaining a better understanding of benchmark datasets.

Index Terms—Visual Analytics, Natural Language Understanding, Shortcut.

1 INTRODUCTION

BENCHMARK datasets play a fundamental role in Natural Language Understanding (NLU) by providing grounds for model evaluation and comparison [1], [2]. Successful examples, like GLUE [3] and SuperGLUE [4], have been widely accepted in assessing models' NLU capability, which can benefit downstream applications, such as sentiment analysis [5], [6] and fake news detection [7], [8].

Building high-quality benchmark datasets is challenging. The effort is far beyond randomly collecting and labeling samples. Instead, researchers have to assess the quality of the benchmark dataset from different perspectives carefully and fix the data problems accordingly [9]. What is a high-quality benchmark dataset? Except for the primarily used criteria, e.g., validity (whether the dataset is correctly labeled) and diversity, recent research argues that a high-quality benchmark dataset needs to be challenging and pertinent [1], [10]. Challenging means that the benchmark dataset should be able to reveal the gaps between different models [10]. For example, we do not want simple models to achieve human-level performance in the benchmark datasets easily. *Pertinent* indicates that the model's performance on the benchmark dataset should reflect its capability in the target task. Intuitively, an NLU model should leverage the task-related words and language structures to make predictions instead of spurious biases.

Benchmark datasets collected from the traditional data collection process [9] cannot guarantee such criteria. According to recent

 Qun Liu is with Huawei Noah's Ark Lab, Hong Kong, China. E-mail: qun.liu@huawei.com.

Manuscript received XX XX, 2022; revised XX XX, 2022.

literature, shortcuts–unwanted biases in the datasets–undermine the quality of NLU datasets. Shortcuts are easily captured by the machine learning (ML) models in making predictions [11], allowing even simple models to achieve good performance in challenging prediction tasks. In NLU benchmark datasets, shortcuts can be words, phrases, and language structures that occur with different frequencies in different categories of sentence instances. Such shortcuts should be identified and removed to ensure the quality of the NLU benchmark dataset.

Most existing approaches target building *challenging* datasets with fully-automated methods, such as adversarial filtering [12], [13] and counterfactual augmentation [10], [14]. However, these methods do not consider the pertinence of the dataset simultaneously [1], [15]. Simple but representative instances may be removed from the original datasets. And the newly constructed datasets may still contain task-irrelevant patterns that can be obtained by the models in making predictions. An alternative approach is to keep the experts in the dataset exploration and correction loop, where experts are informed of the potential shortcuts and decide whether to fix them according to their expertise. However, it is challenging for dataset creators to inspect and mitigate the shortcuts in the benchmark dataset by directly analyzing the instances one by one, which requires many human labor efforts. Interactive and efficient tools are desired, such as visualization tools, to explore shortcuts and gain insights into how to mitigate the shortcuts in the benchmark dataset.

In this paper, we collaborate with two NLU experts for four months and summarize a list of requirements for system design and develop the system called *ShortcutLens*, which can help dataset creators conduct a multi-level exploration of the shortcuts in NLU benchmark datasets, as shown in Fig. 1. *ShortcutLens* includes three views: Statistics View, Template View, and Instance View. The Statistics View displays necessary statistics about the benchmark dataset and the potential shortcuts within. Users can

Zhihua Jin, Xingbo Wang, Furui Cheng, and Huamin Qu are with Hong Kong University of Science and Technology, Hong Kong, China. E-mail: {zjinak, xwangeg, fchengaa, huamin}@cse.ust.hk.

Chunhui Sun is with Peking University, Beijing, China. E-mail: sch@pku.edu.cn.

IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS, VOL. XX, NO. XX, XX 2022



Fig. 1. ShortcutLens assists NLU experts in conducting the multi-level exploration of shortcuts in NLU benchmark datasets. (a) The Control Panel allows users to select the benchmark dataset. (b) The Statistics View helps users inspect the statistics about the benchmark dataset and shortcuts. It also allows users to conduct what-if analysis on shortcuts of interest. (c) The Template View enables users to check the relationship of shortcuts and inspect the statistics about individual shortcuts. (d) The Instance View displays the instances covered by selected shortcuts from the Template View.

also select a group of potential shortcuts and conduct the whatif analysis (Fig. 1(b)). The Template View employs hierarchical and interpretable templates to summarize the shortcuts. The glyph is proposed to visually represent the shortcuts (Fig. 1(c)). The Instance View allows users to check the instances covered by selected shortcuts from the Template View (Fig. 1(d)). Case studies and expert interviews are conducted to evaluate the effectiveness and usability of *ShortcutLens*. The results demonstrate that *ShortcutLens* can help dataset creators explore the shortcuts and gain insights into how to mitigate them in the benchmark dataset.

The contributions of this work can be summarized as follows:

- A visual analytics system to help dataset creators systematically explore the shortcuts in NLU benchmark datasets.
- Case studies on understanding and mitigating shortcuts in NLU benchmark datasets and expert interviews to demonstrate the effectiveness and usability of *ShortcutLens*.

2 RELATED WORK

In this section, we introduce related work on criteria and guidelines of natural language understanding datasets construction, visualization for natural language processing, and visualization for improving dataset quality.

2.1 Criteria and Guidelines of Natural Language Understanding Datasets Construction

In recent years, pre-trained language models like BERT [16] have achieved great success on NLU benchmarks, such as GLUE [3] and SuperGLUE [4]. Although models achieve comparable and even superior performance to human performance on standard benchmarks, they can easily fail in challenging or real-world cases [10], [15], [17]. It shows that these benchmarks are no longer adequate for evaluating human-like complicated and comprehensive language abilities. One reason for such a phenomenon is that many benchmarks suffer from spurious biases that incentivize the inflating model performance.

According to prior research [15], [18], [19], many spurious biases in NLU benchmarks relate to shortcuts. They can inflate the performance of models and mislead the researchers to improve it in the wrong direction. For example, Gururangan *et al.* [19] found that the annotator is inspired by the examples in the annotation guideline to introduce artifacts into the benchmark dataset by simply substituting words. Therefore, words such as "not" will be closely associated with specific labels, and such shortcuts can inflate the model performance. In addition, a line of research revealed that models tend to exploit simple functions or spurious statistical cues when performing NLU tasks [20], [21].

Facing these issues, many works have proposed guidelines on dataset construction to measure and improve the benchmark quality. For example, clear and detailed documentation of benchmark datasets should be carefully considered [22], [23]. They can help reveal the limitations of benchmark datasets and more precisely depict the properties of benchmark datasets. Kiela *et al.* [10] encouraged dynamically collecting data and evaluating models to help mitigate the biases in the benchmark datasets. Bowman *et al.* [1] proposed a more general guideline to help construct benchmark datasets. The benchmark should meet four criteria: (1) Good performance in the benchmark is an indicator of robust performance in the domain task. (2) The instances in the benchmark

datasets should be accurately and unambiguously annotated. (3) It should be significantly difficult or large enough. (4) Social biases can be revealed through the benchmark and they should not be incentivized by the benchmark.

However, these guidelines cannot be easily adopted in practice. Many criteria are coupled and impose great challenges for dataset creators to improve the dataset quality. For example, it is challenging to mitigate the bias in the benchmark dataset while preserving the capability of indicating the models' specific NLU ability. To facilitate this process, we follow the guidelines to design our system to meet its desired properties. We mainly consider two important properties: The first one is to make benchmark datasets more challenging. The second one is to reflect the models' NLU capability instead of shortcuts, which means that it can imply robust performance.

2.2 Visualization for Natural Language Processing

There is a line of related research focusing on visualization for NLP models [24], including model understanding [25], [26], [27], [28], [29], debugging [30], [31], [32], and refinement [33], [34]. For NLP model understanding, Strobelt et al. [25] proposed LSTMVis which can match hidden state changes with syntactic function changes in the text. Ming et al. [26] developed RNNVis that employed the coclustering technique to reveal the relationship between hidden states and words. Tenney et al. [28] further integrated different methods, such as projection and counterfactual generation, to assist model developers in investigating their models. For NLP model debugging, Laughlin et al. [30] utilized adversarial text generation to help researchers identify model deficits. However, their work cannot be utilized to summarize the potential shortcuts in the benchmark dataset. Wu et al. [31] proposed hierarchical structural templates to summarize the query dataset and enable model developers to find the error patterns of their models. One issue of their templates is that they are limited to short query datasets, which cannot be applied to long-sentence summarization. For NLP model refinement, NLIZE [33] allows users to interactively refine models via the attention matrix. Ming et al. [34] built ProtoSteer, which can help domain experts update a sequence model by interactively adding, deleting, and revising prototypes.

However, most of these works mainly focus on improving the model performance, neglecting to improve the dataset quality. Moreover, they do not provide sufficient help for dataset creators to find flaws in the benchmark dataset and gain insights into how to improve it.

2.3 Visualization for Improving Dataset Quality

A group of work focuses on improving image or tabular dataset quality although it neglects the textual data [35]. They can be categorized into two classes: data anomaly detection [36], [37], [38], [39], [40] and missing values detection [41]. In terms of data anomaly detection, Chen *et al.* [40] proposed OoDAnalyzer which can help improve OoD detection accuracy with a grid-based visualization. For missing value detection, Alemzadeh *et al.* [41] developed VIVID to diagnose the root causes of missing values through multiple coordinated visualizations. These works cannot be easily adapted to textual data because textual data is unstructured and not easy to explore, while tabular data is structured data and image data is easy to understand through a glimpse. A recent study called DQI [42] can help reveal shallow statistics of textual datasets. However, it is not effective for investigating benchmark datasets that are designed for evaluating the complicated reasoning ability of models, as models can exploit more complex shortcuts. Moreover, it does not integrate different types of statistics into compact visualizations, making the analysis inefficient.

Therefore, we aim to develop a tool to explore the shortcuts in NLU benchmark datasets to fill in the research gap. Current work fails to reveal the shortcut in a unified, hierarchical, and understandable way. Instead, our work adopts hierarchical and understandable templates to help dataset creators explore and identify potential shortcuts.

3 BACKGROUND

A wide range of tasks for evaluating the NLU capability is usually formulated as the single sentence classification task or multiple sentence classification task [3], [4]. In this paper, we mainly focus on the single sentence classification tasks. We will illustrate two examples, including spatial reasoning and grammatical acceptability classification.

Spatial reasoning. Spatial reasoning is a cognitive process to cope with environments, which requires constructing representations of spatial relationships and transformations between objects [43], [44]. To test the model's spatial reasoning capability, a team of researchers recently created a new benchmark dataset called SpaCE2021 to evaluate the textual spatial reasoning capability of the models. An assumption is that if the model has spatial reasoning capabilities, then it must be not only able to recognize the correct spatial information but also able to recognize abnormal and incorrect spatial information. For example, for the sentence "sign a name on all sides", people can realize that it is weird because "a name" is usually not signed on "all sides". For the sentence "walking under a train", people can clearly know that in most cases, no one walks under a train. Instead, the sentences "sign a name on a book" and "walking on the train" have correct spatial information.

Grammatical acceptability judgment. Warstadt *et al.* [45] built a benchmark dataset called CoLA to test models' ability on grammatical acceptability judgment. For example, the sentence "This building is than that one." is considered a grammatically unacceptable sentence, while a slightly modified sentence "This building is taller and wider than that one." will be regarded as a grammatically acceptable sentence.

4 ABSTRACTIONS

In this section, we introduce an abstraction of the domain problem, analysis workflows, and design requirements in NLU benchmark dataset construction.

We collaborated with a research team from a large company for four months in developing SpaCE2021, a Chinese spatial reasoning benchmark dataset. We closely worked with two team members (E1, E2), who are also our co-authors. E1 is a senior Ph.D. student with more than two years of experience in NLU dataset construction. E2 is the team leader with more than twenty years of experience in developing NLP models and benchmark datasets in commonsense reasoning and machine translation. Through our observation of their dataset construction and validation process and multiple rounds of interviews, we characterize the problems, their workflow, and requirements for designing an interactive tool as follows.

4.1 Domain Problem

The benchmark dataset development consists of three stages: data collection, augmentation or filtering, and validation. Taking SpaCE2021 as an example, experts (including both NLP experts and linguists) first collected sentences with spatial relationship descriptions from novels, essays, and fiction. Then they randomly replaced some locality words like "inside" with other locality words like "outside" in the sentence. They further distributed the modified sentences to a group of annotators to judge whether the spatial information in the modified sentences is correct or not. In the validation stage, they mainly focused on checking the annotator agreement on the labels and fixing data contamination issues. However, when applying this dataset for benchmarking, the experts found that simple models achieve human-level performance, which was unreasonable considering the difficulty of the prediction problems. Thus, they began to notice the shortcut problems in the benchmark dataset and expected to explore and summarize the shortcuts to gain actionable insights to improve the benchmark dataset quality.

4.2 Workflow

The experts used to find shortcuts based on their intuitions. They first inspect some challenging but correctly solved instances. From these instances, they summarize a set of words (usually nouns) that they think to be shortcuts. Then the experts test the frequency of these words in different categories of samples (i.e., true and false samples for a binary classification problem). When a word's distribution is biased in different categories, it is likely to be a shortcut. The experts will record these words and adjust the datasets (e.g., removing the related instances) to test whether the dataset's quality is improved. However, the existing workflow only considers narrow-scoped shortcuts and heavily relies on experts' intuitions. To address these limitations, we need an interactive visual analytics tool to extract and organize those shortcuts in an interpretable, hierarchical, and unified way.

4.3 Design Requirements

We summarize the requirements of designing an interactive visual analytics tool for exploring shortcuts in the NLU benchmark dataset.

R1: Provide an overview of shortcuts. An overview helps experts make sense of the potential shortcuts. The experts commented that certain statistical information is critical for them to understand the importance of each shortcut, including the number of instances covered by the potential shortcuts (i.e., coverage), the percentage of the correct predicted instances in covered instances (i.e., productivity) [14], [15], [31], [42]. Besides, the experts suggested that the shortcuts should be organized into a hierarchy. After checking the overview of potential shortcuts, they may select a group of potential shortcuts, and they further inspect them to understand the details behind them.

R2: Associate shortcuts with corresponding instances. Experts are willing to inspect individual shortcuts and quickly find their corresponding instances. They commented that it is essential for them to "*find concrete examples*" and "*understand the shortcuts based on the context*". Besides, they expect the system to highlight the words related to the shortcuts in the instances.

R3: Estimate the effects of removing shortcuts. After exploring and identifying shortcuts in datasets, the next step is to fix them, e.g., removing the corresponding instances [1], [10],

[17], [30]. However, removing these instances will sometimes bring other shortcuts, which cannot necessarily improve the dataset quality. Besides, conducting a formal test (including training large language models) can be time-consuming. Instead, experts expect first to estimate the influence of mitigating them. For example, suppose users want to remove specific instances covered by specific shortcuts. In that case, it is helpful to know the changes in the machine accuracy between the original set and the set after cleaning instances. Such functionality can help dataset creators to decide the priority of fixing those shortcuts.

5 SHORTCUTLENS

Based on the derived design requirements, we design *ShortcutLens* to support dataset creators in investigating potential shortcuts in NLU benchmark datasets. In this section, we introduce an overview of the system first. Then we introduce the scope and relationship of shortcuts and what-if analysis. Design choices of individual views are introduced next. Finally, we introduce the shortcut mining and aggregation algorithm.

5.1 System Overview

The system consists of three major modules: storage, shortcut mining module, and interactive visualization module. The storage module mainly manages the NLU benchmark datasets. The shortcut mining module is used to extract the potential shortcuts and parse the text for each instance. The interactive visualization module allows users to interact with the system to explore the potential shortcuts and conduct the what-if analysis. The storage and shortcut mining modules are built upon Python and then are integrated into a backend server built upon the Flask. Regarding the Chinese NLU benchmark dataset, we use the HanLP¹ to extract the Part-of-Speech (POS)² and embedding of each word in the instances. For the English NLU benchmark dataset, we use spaCy³ to extract the POS⁴ and word embedding. The interactive visualization module is implemented as the frontend supported with browser using React, Typescript, and D3.

5.2 Scope and Relationship of Shortcuts

To help extract the potential shortcuts in the dataset, it is desirable to have a suite of algorithms to mine the potential shortcuts from the benchmark dataset. We first assume that the potential shortcuts are matching-based shortcuts, that is, if the text contains the specific pattern, it will be considered covered by these potential shortcuts. Inspired by the defined shortcuts in DQI [42] and the structural templates which include POS and named entities from Tempura [31], we define the pattern as a set of spatially-related words. Specifically, we consider one word with its POS, or two words with their POS and their relative position (i.e., how many words are in between) to define the pattern or template. For example, we consider one pattern as follows: (1) The first word "The" is a determiner. (2) The second word "will" is a verb. (3) The first word is in front of the second word, and there is a word between them. Such a pattern can match sentences such as "The men will all leave." The components of the potential shortcuts can be abbreviated, that is, some of their components can appear optionally

2. https://hanlp.hankcs.com/docs/annotations/pos/pku.html

4. https://universaldependencies.org/u/pos/

^{1.} https://www.hanlp.com/

^{3.} https://spacy.io/

in the potential shortcuts. For example, in the template above, the scope of the second word can be extended by considering this word as a verb, and not just the specific word "will". The newly defined template is the parent of the original template. The mining algorithm selects shortcuts with high *productivity* and *coverage* [14], [15]. Coverage refers to the number of instances covered by the shortcut. The prediction label of the potential shortcut is defined as the dominant label in the label distribution. Productivity is calculated as the percentage of covered instances with the same label as the prediction label in covered instances. Details of the shortcut mining algorithm are illustrated in Sec. 5.5.1.

After running the shortcut mining algorithm, we find that there are a large number of potential shortcuts in the benchmark dataset. One of the reasons is that the number of different words is quite large. It imposes difficulty for users to explore such largescale potential shortcuts. Therefore, we aggregate similar potential shortcuts based on the semantic meaning of the words to reduce the difficulty for users to explore. Details of the shortcut aggregation algorithm are illustrated in Sec. 5.5.2.

5.3 What-If Analysis

In order to help dataset creators decide which potential shortcuts should be fixed and what actions should be taken to fix the potential shortcuts, a what-if analysis on shortcuts of interest is necessary for dataset creators to make decisions. Possible actions include constructing new instances, modifying the existing instances, or removing the existing instances to battle against a group of potential shortcuts. However, after taking action, it is hard for dataset creators to know whether it really resolves the issues. Therefore, a necessary what-if analysis should be considered. We will take an initial attempt to consider the influence of the removal of a group of instances.

After users have focused on a group of shortcuts of interest, the instances covered by these shortcuts will be found. We will consider the instances covered by them as the dirty set and the instances not covered by them as the clean set. When calculating the productivity for a group of potential shortcuts, it is possible that one instance is covered by multiple potential shortcuts with different prediction labels. Those instances will be considered as "disagreed" instances. The number of "disagreed" instances will be calculated. Then, the productivity of a group of shortcuts is defined as the percentage of covered "agreed" instances with the same prediction label as the corresponding potential shortcuts. Such metrics indicate the impact of those potential shortcuts.

If the machine prediction results are provided by the users, we will calculate the machine accuracy for the dirty set and clean set, respectively. We compute the machine accuracy on the dirty set and clean set by averaging the accuracy of all models on the corresponding sets. It is interesting to compare the machine accuracy between the whole set and the clean set to see whether the removal of potential shortcuts leads to the degrading of the performance metrics in the clean set. Such indicators can help dataset creators understand the potential benefits gained from the removal of the instances covered by those potential shortcuts.

5.4 Visualization

The *ShortcutLens* visualization module consists of Control Panel, Statistics View, Template View, and Instance View. The Control Panel allows users to select the benchmark dataset to explore the potential shortcuts (Fig. 1(a)). The Statistics View assists dataset creators in inspecting the statistics about the benchmark dataset, the potential shortcuts, and baseline models provided by the users (Fig. 1(b)). It also allows users to select a group of potential shortcuts to conduct what-if analysis and further inspect in the Template View. The Template View supports dataset creators in inspecting the individual potential shortcuts and the relationship between the potential shortcuts (Fig. 1(c)). After users select one potential shortcut in the Template View, the Instance View displays the corresponding instances for dataset creators to further inspect (Fig. 1(d)). Users can choose to check the neighborhood of the highlighted matched text, which can quickly determine the location and context of the potential shortcut.



Fig. 2. The glyphs are used to represent the shortcuts. (a) The Statistics View uses a circle-based glyph to encode the productivity, coverage, and prediction label of the shortcut. (b) The Template View uses a block-based glyph to encode the template of the shortcut. Coverage, productivity, and the number of aggregated words of the shortcuts will be displayed accordingly.

5.4.1 Statistics View

The Statistics View helps dataset creators gain an overview of the benchmark dataset, potential shortcuts, and baseline models, as shown in Fig. 1(b) (**R1**). It consists of three panels, Instances, Shortcuts, and Machine Accuracy. The Instances panel displays the statistical summary of the benchmark dataset. It shows the number of instances and the label distribution in the whole benchmark dataset and different sets of the benchmark datasets. The Machine Accuracy panel shows the accuracy of individual models on the specific set. The Shortcuts panel allows users to explore the shortcuts in the Projection plane. Users can filter the shortcuts by setting the minimum productivity and coverage in the Filters plane.

Visual design. We use a circle-based glyph to represent the potential shortcuts and encode the relevant metrics about the potential shortcuts, as shown in Fig. 2(a). The radius of the circle encodes the coverage of the shortcut. The outer ring of the circle encodes the productivity of the shortcut. The color of the circle encodes the prediction label of the shortcut. For the binary classification task, we use red color to encode the label "*false*" and dark blue color to encode the label "*true*". The experts said that shortcuts with similar properties should be displayed in a near region, and it will be useful for users to inspect them. Thus, we use the UMAP [46] projection algorithm with collision avoidance to calculate the layout of the glyph of potential shortcuts. The distance between different potential shortcuts a and b is calculated as follows:

$$Dist(a,b) = |Prod_a - Prod_b|^2 + |Norm(Cover_a) - Norm(Cover_b)|^2 + \mathbb{I}\{Pred_a \neq Pred_b\}$$
(1)

where *Prod* is the productivity of the potential shortcut, *Cover* is the coverage of the potential shortcut, *Norm*(*) is the normalization function, and *Pred* is the prediction label of the potential shortcut. $\mathbb{I}\{*\}$ is an indicator function. If the expression is true, the value is one. Otherwise, it is zero.

We considered one design alternative of circle-based glyphs. We can use two bars to encode the productivity and coverage of the shortcut and the color to encode the prediction label of the shortcut. However, the experts commented that it is more intuitive to use arc angles to present the percentage (i.e., productivity). Therefore, we decided to use a circle-based glyph for showing productivity and coverage, where coverage is encoded by circle radius, the angle of the outer arc represents productivity, and the circle color indicates the prediction label.

Interaction. After lasso-selecting a group of potential shortcuts, the Template View will display the corresponding potential shortcuts. Moreover, the what-if statistics will be displayed in the following spaces. It will demonstrate the coverage and productivity of the selected potential shortcuts. If the machine prediction results are provided by users, it will also display the difference in machine accuracy in the clean set or the dirty set. It can help dataset creators to estimate the influence of removing the instances covered by those shortcuts (**R3**).

5.4.2 Template View

To allow dataset creators to further inspect the selected potential shortcuts and the relationship between the potential shortcuts (**R1**), the Template View is designed to display more fine-grained information of potential shortcuts ordered in a hierarchical structure (Fig. 1(c)).

Visual design. The experts commented that a vivid representation of the shortcuts is desirable for an intuitive understanding of the semantic meaning of the shortcuts. Therefore, we use a block-based glyph to encode the matching content of the potential shortcuts, as shown in Fig. 2(b). Each block represents a word. If the background color of the block is white, it means it will match nothing but stands as the placeholder for the relative position. The other background color of the block encodes the POS. The words in the block are the specified matched words. For example, as shown in Fig. 2(b), it will match sentences with the pattern where the first word is the adposition "up", the second word in front of the first word is a pronoun which is similar to "he", and there is a word between those two words. The coverage and productivity of that potential shortcut are individually encoded by the bar width. If the shortcut is generated via the aggregation of similar words, it will have a triangle upon the representative of the word sets. Also, the number of aggregated words is displayed near the triangle. Such a glyph can help users understand the template behind the potential shortcuts.

Interaction. If users change the split, the coverage and productivity of the shortcut will only be calculated on the selected split. Users can click the shortcuts with thicker borders of the blocks to expand the children of that potential shortcuts to check more fine-grained information about that potential shortcut. Moreover, the relationship between the shortcuts is also displayed. Users can also click the radio button in one of the potential shortcuts to further inspect the instances covered by it in the Instance View (**R2**).

5.4.3 Instance View

To help dataset creators understand the template presented in the Template View, as shown in Fig. 1(d) ($\mathbf{R2}$), users can explore the

covered instances in the Instance View. The experts commented that it is better to check the neighborhood of the matched text. Thus, users can select the text style "Neighbor" to show only the neighborhood of the matched text if the text is too long to inspect. The three words to the left and three words to the right of the highlighted text are considered the neighborhood of the matched text. If the text is short, users can switch the text style to "Full" to check all the text in the instances. The matched text will be highlighted. The background color of matched text depends on the POS of individual words, which has the same encoding in the Template View. The table also displays the split, label, and machine accuracy for users to inspect. Users can search text, filter instances based on split or label, and sort the machine accuracy in the table to find the instances of interest.

5.5 Shortcut Mining and Aggregation

Shortcut mining and aggregation algorithm will be introduced in the following paragraphs. The overview of workflow of those algorithms is depicted in Fig. 3.



Fig. 3. Shortcut mining algorithm (a) is to extract, select, and organize shortcuts. Shortcut aggregation algorithm (b) is to cluster shortcuts and insert representative shortcuts of clusters into the hierarchy.

5.5.1 Shortcut Mining

The basic idea of the mining algorithm (Fig. 3(a)) is extracting potential shortcuts via an exhaustive method. Then it will filter potential shortcuts based on the productivity and coverage. Finally, it will organize the shortcuts into a hierarchical structure. The details of each step in the algorithm are illustrated as follows:

Step 1: Extract Shortcuts. Once we have the POS and words for a specific text, we will enumerate each word in the text to determine the POS of the first word and then its specific word. We will enumerate the possible relative position to get the position of the next word and its specific POS and word.

Step 2: Calculate Shortcuts Statistics. The label distribution for each potential shortcut will be updated accordingly. If the specific instance has been considered in the potential shortcut, we will only consider that the potential shortcut covers that instance once. Then, we will calculate the coverage and productivity for each potential shortcut.

Step 3: Select Shortcuts. We will filter the potential shortcuts by a preset minimum coverage and productivity in the whole set. Also, we will consider the minimum coverage and productivity in each set, like the training set, development set, and test set.

Step 4: Organize Hierarchy of Shortcuts. We will organize the filtered potential shortcuts into a hierarchical structure based on the definition of the parent relationship. To reveal the relationship between the filtered potential shortcuts, we will track the paths from the root of the hierarchy of the potential shortcuts to filtered potential shortcuts and the children of filtered potential shortcuts.

5.5.2 Shortcut Aggregation

The basic idea of the shortcut aggregation algorithm (Fig. 3(b)) is using hierarchical clustering to aggregate potential shortcuts with similar words based on word embedding into a new shortcut and insert this shortcut into the original hierarchy of shortcuts. The details are illustrated as follows:

Step 1: Group Mergeable Shortcuts and Calculate Distance Matrix. We first check whether the two potential shortcuts can be merged, and then we calculate the distances between every two potential shortcuts. If two potential shortcuts share the same parent, the same prediction label, and the final component of each potential shortcut is the word, we will calculate the distances for the two final words. The distance between two final words *ftok* is defined as $1 - Sim(ftok_a, ftok_b)$. The similarity between two words *Sim* is defined as the cosine similarity of word embeddings for those two words. Then the distances between two potential shortcuts will be set to the distances between two final words. If two potential shortcuts do not satisfy the conditions for merging, the distances between these two potential shortcuts will be set to infinity.

Step 2: Cluster Shortcuts. We adopt hierarchical clustering with complete linkage to cluster similar potential shortcuts based on the distances between two potential shortcuts [47]. We have empirically set the clustering methods as the complete linkage as we found that such a method can ensure that the distance between each shortcut will be under a certain threshold, which is empirically set as 0.75.

Step 3: Insert Representative Shortcuts into the Hierarchy. If a cluster of potential shortcuts has more than one potential shortcut, it will be converted into a new potential shortcut, and this new potential shortcut will be inserted into the hierarchy of the potential shortcuts. The final word of this new potential shortcut is defined as the final word of potential shortcuts with the smallest average distances to all other final words in the same cluster. This new potential shortcut covers the instances covered by all potential shortcuts in that cluster.

6 EVALUATION

In this section, we report the case studies with two experts and the results of expert interviews to demonstrate the effectiveness and usability of our system. Actionable insights in case studies are verified in Sec. 6.3.



Fig. 4. E1 set the filters in the Statistics View (a) and five shortcuts satisfy the conditions. E1 found that four shortcuts consist of locality words while one shortcut consists of other kinds of words in the Template View (b).

6.1 Case-I: SpaCE2021 Dataset

We conducted this case study with the dataset creator (E1) on spatial reasoning in Chinese. The background of E1 is depicted in Sec. 4. The team of E1 has made an initial attempt at constructing the benchmark dataset for SpaCE2021. The dataset description and collection process are depicted in Sec. 3 and Sec. 4.1 respectively. After training nine models on this benchmark dataset, which include three BERT-based models [16], [48], five RoBERTabased models [48], [49], and one ERNIE-based model [50], he found that the machine performance is higher than expected. He decided to further inspect whether the benchmark dataset is good enough to evaluate the models' capability in spatial reasoning. He was wondering whether the benchmark dataset suffers from the shortcuts.

E1 selected the SpaCE2021 benchmark dataset in the Control Panel (Fig.1(a)). In the Statistics View, E1 found that the number of all instances is 5837 and the number of instances in the training set, development set, and test set is respectively 4237, 806, and 794. The percentage of instances with the label "false" in the test set is 0.5416 while the average machine accuracy on the test set is 0.6924, which is significantly better than the majority baseline on the test set. Given the situation that the model performance is quite good, E1 decided to further explore the potential shortcuts in the benchmark dataset. E1 set the minimum productivity as 0.75, which surpasses the average model accuracy, and the minimum coverage as 100 to filter the potential shortcuts (Fig. 4(a1)). There are five potential shortcuts satisfying this condition as shown in Fig. 4(a2). E1 further checked the What-If Analysis plane. The productivity of those shortcuts is 0.8136. In terms of machine accuracy, the performance on the dirty set is 0.0807 higher than the performance on the whole test set. E1 thought that those potential shortcuts can have a high impact on the model performance. E1 found that four potential shortcuts consist of locality words but one of them consists of other kinds of words (Fig. 4(b)). Since spatial reasoning requires a deep understanding of locality words in the sentence, it is more interesting to see that there exist potential shortcuts with no locality words, as shown in Fig. 4(b2). To understand this

IEEE TRANSACTIONS ON VISUALIZATION AND COMPUTER GRAPHICS, VOL. XX, NO. XX, XX 2022



Fig. 5. E1 chose a shortcut in the Projection plane (a). The productivity of this shortcut is similar to machine accuracy in the dirty set in the What-If Analysis plane of the Statistics View (b). E1 found that only one child of that shortcut (c) covers the test instances (d). Most of the instances are with similar neighborhoods around the highlighted text (e).

phenomenon, E1 decided to further explore one of the potential shortcuts with locality words (Fig. 4(b1)) and the shortcut without locality words (Fig. 4(b2)) (**R1**, **R3**).

6.1.1 Analyzing Shortcuts with Locality Words

E1 selected the shortcut with the locality word "left" with noun word ahead in the Projection plane (Fig. 1(b1), Fig. 1(c1)) and checked the What-If Analysis plane (Fig. 1(b2)). E1 found that this potential shortcut covers 15 instances in the test set with productivity of 1. In the dirty set, which includes instances covered by this shortcut, the machine accuracy achieves 0.963. It indicates that the models may find that this kind of instance is easier to answer and tend to exploit this kind of shortcut. E1 further checked the Template View and expanded the children of that potential shortcut (Fig. 1(c1)). E1 found that there exist some children shortcuts with low productivity, like "temple left", which indicates that there exist some cases battling against this potential shortcut (Fig. 1(c2)). However, this potential shortcut covers 143 instances in total and has high productivity of 0.87, which means that other children of this shortcut tend to have high productivity. E1 further checked the Instance View and found that it indeed matches part of the original text's incorrect spatial information (Fig. 1(d)). Such biased distribution worried the E1 because he did not want such kinds of shortcuts to inflate the model performance, especially when checking the model performance on the dirty set.

Summary. Since this shortcut is not an intended solution and inflates model performance, E1 said that he would try to avoid this kind of shortcut in the new dataset construction process, for example, by reducing the number of locality words with "left" (**R1**, **R2**, **R3**).

6.1.2 Analyzing Shortcuts without Locality Words

Then, E1 further explored another shortcut without locality words (Fig. 4(b2)). E1 chose a shortcut that matches sentences where the preposition is "in", the second word in front of the word is a numeral, and there are no words between these two words (Fig. 5(a), Fig. 5(c)). E1 further checked the What-If Analysis plane and found that the productivity of this shortcut is similar to machine accuracy in the dirty set (Fig. 5(b)). It indicates that the model performance is somehow related to this shortcut. E1 further checked the relationship between shortcuts in the Template View. E1 found that only one child of potential shortcuts matched instances in the test set. This potential shortcut matches sentences where the preposition is "in" and the second word in front of the word is a numeral which is "two" (Fig. 5(d)). E1 selected it in the Template View and checked the instances in the Instance View. He found that all instances have the same neighborhoods around the highlighted text (Fig. 5(e)). It reminds him that many sentences are modified from the same original sentence. Therefore, when the label distribution is biased for the same original sentence, it allows potential shortcuts to achieve spurious performance on the test set.

Summary. Based on the observation that irrelevant text in instances modified from the same original sentence will be matched by shortcuts, it inspired him that when selecting instances in the test set, he should guarantee that the label distribution of the sentences modified from the same original sentence should be more balanced to avoid the potential shortcuts matching irrelevant places in the text (**R1**, **R2**, **R3**).

6.2 Case-II: CoLA Dataset

We conducted this case study with a senior NLP researcher (E2). The background of E2 is depicted in Sec. 4. E2 is working on one



Fig. 6. E2 selected one of the shortcuts with the prediction label "true" in the Projection plane and found that such kinds of irrelevant words can lead to high productivity (a). E2 found one shortcut with the prediction label "false" and the instances covered by this shortcut seem to be modified from similar sentences (b).

project which is related to the grammatical error detection dataset construction. He is very interested in a similar benchmark dataset called CoLA [45] in the GLUE [3] benchmark. One BERT-based model is finetuned on this benchmark dataset. He is willing to inspect whether there exist some shortcuts in CoLA and gain some insights into avoiding shortcuts in new dataset construction.

6.2.1 Analyzing Shortcuts with Prediction Label "true"

E2 selected the CoLA benchmark dataset in the Control Panel. He found that the benchmark dataset consists of 9594 instances, with 8551 instances in the training set and 1043 instances in the development set⁵. The percentage of instances with the Label "*true*" in the development set is 0.6913 while the machine accuracy on the development set. Since the models achieve high accuracy on the development set, E2 selected the minimum productivity as 0.90 and the minimum coverage as 50. There are only two remaining shortcuts, with the prediction label "*true*" (**R1**, **R3**).

E2 inspected one of them that matches sentences where a verb is "will", the second word ahead of it is a determiner and similar to "that", and there is one word between these two words (Fig. 6(a1), Fig. 6(a3)). This shortcut covers 10 instances with the same productivity as the machine accuracy (0.9) in the development set (Fig. 6(a2)). E2 is surprised that such kinds of irrelevant words can lead to high productivity. E2 further inspected the instances in the Instance View (Fig. 6(a4)). He believed that it is a potential shortcut unintended introduced when collecting data from the biased distribution. Such shortcuts cannot guarantee that the matched sentences are grammatically acceptable. For example, the grammatically acceptable sentence "The men will all leave" can

be modified as the grammatically unacceptable sentence "The men will all" while both sentences can be matched by this shortcut.

Summary. Since the shortcuts are not guaranteed that the sentences can be grammatically correct, E2 thought that negative examples can be constructed from grammatically acceptable sentences by simply removing some critical components of them $(\mathbf{R2}, \mathbf{R3})$.

6.2.2 Analyzing Shortcuts with Prediction Label "false"

E2 was curious about whether there exist shortcuts with the prediction label "false" because instances with the label "false" are minor samples compared to instances with the label "true". E2 set the minimum productivity as 0.9 and the minimum coverage as 15. He found that there are only two shortcuts with the same structure and prediction label "false" (Fig. 6(b1)). This shortcut matches sentences where the first word is punctuation ",", the second word after the first word is adposition "in", and there are two words between these two words (Fig. 6(b3)). E2 further inspected the instances covered by this potential shortcut (Fig. 6(b4)). He found that most sentences before the matched punctuation seem incomplete, then it will potentially lead to grammatically unacceptable sentences as a whole. The punctuation seems a very weak indicator of the incorrect sentence. The two words between the matched text mostly (13/15) are "this girl". If further considering these two words, all matched sentences are labeled as "false". In terms of this shortcut, E2 further carefully checked the matched sentences and found that they are mostly modified from similar sentences. But there are only two sentences in the development set (Fig. 6(b2)).

Summary. Based on the observation that there are a small number of instances matched by this shortcut. E2 thought that ignoring this shortcut seems not too problematic. E2 said that to solve this shortcut, he can take a simple fix like adding one more

^{5.} Since the labels of the test set of CoLA are not available, we do not consider the test set in this benchmark dataset and assume that the development set serves the role of evaluating the models.

positive sentence which can be matched by this shortcut in the development set $(\mathbf{R2})$.

6.3 Actionable Insights Verification

To verify actionable insights in case studies, we can use two ways to evaluate whether the dataset's quality has improved after taking action to fix shortcuts. The first way is to inspect the model performance before and after fixing shortcuts in the test dataset. If the model performance on the new test set is less than the model performance on the original test set, it indicates that the dataset has become challenging. The second way is to inspect the number of shortcuts before and after fixing shortcuts under certain filtering conditions. The shortcut mining algorithm will be run on the modified dataset to ensure that all potential shortcuts defined in the paper are taken into account. Suppose the number of shortcuts decreases and the target shortcut disappears under the original filtering conditions. In that case, it indicates that the shortcut has been fixed, and no more shortcuts are added to the dataset. That is, after taking action to fix the shortcut, many other shortcuts will also disappear. Another situation is that some shortcuts disappear, and new shortcuts appear. Since the number of shortcuts decreases. it indicates that the number of shortcuts disappearing is larger than the number of shortcuts appearing. When experts determine that the shortcuts should be fixed, it indicates that they are the source of task-irrelevant cues in the dataset. In other words, removing the non-pertinent instances or decreasing the number of shortcuts can bring a more pertinent dataset, improving the dataset quality.

We acknowledge the following limitations: 1) We consider only a limited number of models in the system. For models outside our scope, even if the model performance in the system decreases for a clean set, it is not guaranteed that the performance of the models outside our scope will decrease. 2) When adding or modifying instances in the dataset, it is necessary to know the model prediction results for the new instances. If the model inference is not available, the performance of models on the new dataset cannot be accurately calculated. 3) In terms of inspecting the number of shortcuts in the dataset, we only consider the potential shortcuts defined in the paper, which can be extracted using the shortcut mining algorithm. There are also other kinds of shortcuts that are outside of our scope, which also cannot be revealed by our shortcut mining algorithm. For those shortcuts, we cannot detect whether those shortcuts exist in the dataset via the current shortcut mining algorithm. Consequently, we cannot detect whether new shortcuts outside of our scope have been introduced into the dataset after we take action to fix one shortcut in the dataset.

In the following paragraphs, we report the verification results of actionable insights in case studies.

6.3.1 Case-I: Fixing Shortcuts in SpaCE2021 Dataset

When we set the minimum productivity as 0.75 and minimum coverage as 10, the number of shortcuts is 232. The coverage of those shortcuts is 308. If removing those instances and consequently removing the shortcuts, the machine performance on the clean set is 0.6417 (0.0507 less than the machine performance on the whole test set, which is 0.6924). It implies that if the experts investigate all those shortcuts and fix them, the machine performance can drop a lot and will make the dataset more challenging.

Fixing shortcuts with locality words. In terms of improving this dataset and fixing shortcuts with locality words, E1 suggested that we can remove those instances which are covered by the

target shortcut with the locality word "left" and noun word ahead in this dataset. Following E1's suggestions, we removed 15 instances which are covered by this target shortcut in the test set of SpaCE2021. After rerunning the shortcut mining algorithm in the new dataset and reloading the data into the system, we set the minimum productivity as 0.75 and the minimum coverage as 100. Compared to the original dataset, we found that two shortcuts disappear. One shortcut is the target shortcut. Another shortcut is matching the sentences with the locality word "right" and a verb which is two words behind the locality word "right". We further set the minimum productivity as 0.75 and minimum coverage as 10. Compared to the number of shortcuts in the original dataset, which is 232, the number of shortcuts in the new dataset is 211 (21 less than the number of shortcuts in the original dataset). The machine performance in the new dataset is 0.6872 (0.0052 less than the machine performance in the original dataset). Those results imply that this action can make the dataset more challenging and pertinent.

Fixing shortcuts without locality words. In terms of improving this dataset and fixing shortcuts without locality words, E1 suggested that we can remove some instances where the preposition is "in", the second word in front of the word is a numeral which is "two", and the label is "false" to balance the label distribution of the sentences modified from the same original sentence. Following E1's suggestions, we removed 15 instances with the label "false' to balance the label distribution of the sentences covered by the target shortcut. After rerunning the shortcut mining algorithm and reloading the data into the system, we set minimum productivity as 0.75 and minimum coverage as 100. The target shortcut has been removed. When we set the minimum productivity as 0.75 and the minimum coverage as 10, the number of shortcuts in the original dataset is 232 while the number of shortcuts in the new dataset is 224 (8 less than the number of shortcuts in the original datasets). The machine performance in the new dataset is 0.6882 (0.0042 less than the machine performance in the original dataset). Those results imply that this action can improve dataset quality.

6.3.2 Case-II: Fixing Shortcuts in CoLA Dataset

When we set the minimum productivity as 0.75 and minimum coverage as 10, the number of shortcuts is 754. The coverage of those shortcuts is 382. If removing those instances and consequently removing the shortcuts, the machine performance on the clean set is 0.7595 (0.0679 less than the machine performance on the whole test set, which is 0.8274). It implies that if the experts investigate all those shortcuts and fix them, the machine performance can drop a lot and will make the dataset more challenging.

Fixing shortcuts with prediction label "true". To solve this target shortcut which matches sentences where a verb is "will", the second word ahead of it is a determiner and similar to "that", and there is one word between these two words, E2 suggested that we can add negative sentences which can be matched by this shortcut in the development set via removing some critical components of grammatically acceptable sentences. Following E2's suggestions, we add two new sentences with the label as "false". One sentence is "The men will all." Another sentence is "This week will a difficult one for us." They can be matched by the target shortcut. After rerunning the shortcut mining algorithm and reloading the data into the system, we set the minimum productivity as 0.9 and the minimum coverage as 50. We found that the target shortcut has been removed. We further set the minimum productivity as 0.75 and the minimum coverage as 10. The number of shortcuts in the

original dataset is 754, and the number of shortcuts in the new dataset is 750 (4 less than the number of shortcuts in the original dataset). The machine performance in the new dataset is 0.8258 (0.0016 less than the machine performance in the original dataset). The results demonstrate that adding negative sentences can improve dataset quality.

Fixing shortcuts with prediction label "false". To solve this target shortcut which matches sentences where the first word is punctuation ",", the second word after the first word is adposition "in", and there are two words between these two words, E2 suggested that we can add one grammatically acceptable sentence which can be matched by this shortcut in the development set. Following E2's suggestions, we add one sentence which is "When we put a picture of Bill on your desk before tomorrow, this girl in the red coat will put a picture of Bill on your desk before tomorrow." The label of this sentence is "true". The target shortcut can match this sentence. After rerunning the shortcut mining algorithm and reloading the data into the system, we set the minimum productivity as 0.9 and the minimum coverage as 15. We found that two target shortcuts are removed. We further set the minimum productivity as 0.75 and the minimum coverage as 10. The number of shortcuts in the original dataset is 754, and the number of shortcuts in the new dataset is 749 (5 less than the number of shortcuts in the original dataset). The machine performance in the new dataset is 0.8266 (0.0008 less than the machine performance in the original dataset). The results demonstrate that constructing a positive sentence can improve dataset quality.

6.4 Expert Interviews

We recruited eight NLU experts (E3-E10, age: 23-50, 5 males, 3 females) by email and then conducted interviews with them to collect their feedback on the usefulness and usability of the system. They all have more than two years of experience in NLU dataset construction. E3 has participated in constructing various kinds of NLU datasets such as treebank, semantic roles labeling, spatial reasoning, and so on. E4 has proposed two NLP-related datasets. E6 mainly focuses on dataset quality evaluation during the NLU dataset construction process. E8 focuses on premise-based multimodal reasoning dataset construction. E3, E5, and E7 have worked on constructing the SpaCE2021 benchmark dataset and are going to build a better dataset to evaluate the spatial reasoning capability of NLU models. E9 has worked on dataset construction for error detection for the generated text from pre-trained language models. E10 has worked on constructing an adversarial dataset for the GLUE benchmark. None of them (E3-E10) are our co-authors nor tried the system (ShortcutLens) before the interviews.

For the procedures of the expert interviews, we first introduced *ShortcutLens* to them, including the shortcut mining algorithm and interface, and showed the experts the system usage using case one (in Sec. 6.1). Then we asked them to follow the workflow in the demonstrated case study to explore the system using the SpaCE2021 dataset. When exploring the system, they were asked to identify potential shortcuts in the benchmark dataset and figure out actionable insights to fix the shortcuts. They can freely ask questions and provide comments and suggestions on the system. We summarize the results of expert interviews in the following paragraphs.

Shortcut exploration. All experts confirmed that *ShortcutLens* could facilitate systematic and comprehensive shortcut exploration and inspire them to take action to fix shortcuts. Compared to single

gram analysis in their previously-used workflows, experts found that our system helps generalize their findings of shortcuts by considering different relationships (e.g., structural and semantic similarities) between data instances. For example, E3 said that previously he found that sentences with the words "palm left" are easy for models to solve. By checking the relationship between the shortcuts, it is interesting to find that the sentences where a locality word "left" is following a noun word are mostly labeled as "false". Moreover, E5 and E9 mentioned that associating the shortcuts with data instances assists them in identifying undesirable shortcuts. For example, E9 observed that there is one shortcut which matches sentences with the word "top". E9 checked the matched text for this shortcut. E9 found that most of the text is highly related to grammatical errors and should be fixed. In addition, what if analysis supported by the system were considered useful to estimate the severity of shortcuts and identify those that need immediate actions for fixing. For example, E4 noticed that there is one shortcut which matches sentences where the word "side". The machine performance of the dirty set of the shortcut is even higher than the productivity of the shortcut. It means that this shortcut is exploited by the models. E4 thought that it is urgent to fix this shortcut.

Visualization and interaction designs. Overall, experts agreed that visualizations and system interactions are useful and easy to learn and use. Particularly, experts favored the Template View for summarizing a group of potential shortcuts with relatively simpler representations. E10 said that "the potential shortcuts are somehow hard to understand in the beginning, but the glyph of them seems simple and intuitive to understand." However, E3 noticed that he sometimes needs to scroll long lists of templates in the Template View. Experts also appreciated the Instance View, which provides detailed context information about the data instances covered by potential shortcuts. The highlighting of components of potential shortcuts helps them quickly reason shortcut patterns with concrete example sentences. For the Statistics View, E7 and E10 mentioned sometimes he needs to adjust filtering to reduce visual clutter. Besides the general statistical results. E4 desired more descriptions of the glyphs in the system.

Suggestions for improvement. Experts were generally eager to use the system in their future work to inspect the newly constructed dataset and find and mitigate potential shortcuts in them. Nevertheless, they identified some system limitations and provided several valuable suggestions. E5, E8, E9, and E10 requested system support for editing and improving datasets (besides what-if analysis) after discovering undesirable shortcuts. For example, E9 commented that "for those texts with shortcuts including spatial words, we can use pre-trained language models to fix it, like picking a spatial word generated by pre-trained language models." E6 stated that it would be better to incorporate a higher-level shortcuts mining algorithm in the system. For example, incorporating more types of components, like named entities, dependency, and logical structure, in shortcuts will be interesting. In addition, other system functions were mentioned. For example, E7 suggested that the system should provide a function to switch between different languages to accommodate users from different countries. E3 and E4 asked us to provide more documentation and tutorials on the system to help them set up the system for their customized data.

7 DISCUSSION AND FUTURE WORK

In this section, we will summarize the lessons learned during the design and evaluation of *ShortcutLens*. We will also discuss the limitations and future work.

7.1 Lessons Learned

Contextualizing shortcuts with concrete instances. To present an overview of data instances covered by potential shortcuts, we have proposed glyphs to provide the abstraction of shortcuts using templates based on POS, relative positions between words, and semantic meaning. During expert interviews, the experts appreciated the compactness and intuitiveness of the visual designs, especially templates for summarizing shortcuts. Then, to decide whether the potential shortcuts are undesirable or not, they would like to see the context of those shortcuts in original sentences in the Instance View. Such information can help them understand the reason and logic behind the emergence of specific shortcuts. Therefore, we expect future visual analytics systems should tightly associate the abstract summary with detailed context to facilitate a deep and comprehensive understanding of complex concepts.

Focusing on data issues besides model issues. Recent machine learning models, especially deep learning models for NLP, achieve compelling performance on different benchmark datasets. However, they are far from being perfect since their performance could be inflated by dataset biases (e.g., shortcuts). In this paper, we design and build *ShortcutLens* to help dataset creators systematically explore shortcuts in NLU datasets and conduct what-if analyses. The case studies and expert interviews confirm that our system can help users discover shortcuts and gain valuable insights into mitigating them in the dataset construction. We hope that our work can draw attention to data issues in the AI community.

7.2 Limitations and Future Work

Investigate more complex shortcuts in NLU benchmark datasets. In this paper, we extract and visualize shortcuts considering POS, relative positions between words, and semantic meaning. In the future, we can further consider other linguistic properties of shortcuts, such as the named entities, dependency, and logical structure. Moreover, we adopt matching-based and similarity-based methods to decide the coverage of shortcuts. That is, if a data instance is covered by a shortcut, it must satisfy or resemble the condition of the shortcut. However, other types of decision conditions can be studied. For example, exclusion-based shortcuts cover instances that do not contain the components of the specific shortcuts. And count-based shortcuts cover the instances that contain specific components multiple times.

Support shortcut exploration for more diverse NLU tasks and benchmark datasets. In this paper, we demonstrate our system for shortcut exploration through two single sentence classification tasks, including grammatical acceptability judgment and spatial reasoning. We can further extend the scope of our system to support multiple sentence classification tasks (e.g., natural language inference). For example, we can investigate the co-occurrence of our defined shortcuts in multiple sentences. Moreover, it will be beneficial to improve the system to investigate shortcuts in text generation tasks, such as question answering [51] and story generation [52]. Last but not least, the current system supports analyzing benchmark datasets in Chinese (i.e., SpaCE2021) and English (i.e., CoLA). It is valuable for our system to support benchmark datasets in different languages.

Enable interactive shortcuts fixing. Currently, ShortcutLens facilitates multi-level exploration of shortcuts regarding productivity, coverage, and semantic and structural information. It helps experts discover shortcuts to be further fixed. Also, the system provides the what-if analysis to estimate the influences of removing potential shortcuts in the datasets. As suggested by the experts during the interview in the evaluation, we can further enhance the system usability with more choices to fix shortcuts. For example, we can enable users to add or modify instances to adjust the distribution of labels using adversarial text generation [53], [54], [55] and model-based text generation methods [56], [57], [58]. Besides, it is possible that when we solve one shortcut, another shortcut will appear. One possible solution is that we can visualize the label distribution shift of other related potential shortcuts in the system so that users understand the possible consequences and trade-offs of fixing shortcuts of interest. Furthermore, we can use statistics of shortcuts and machine performance to measure the quality of the dataset. Users can observe the benefits gained from their actions for improving the dataset.

Improve scalability of visual designs to handle a larger number of instances and classes. The current visual designs cannot scale well when the number of data instances and classes increases. For example, the Projection plane in the Statistics View only supports up to 300 shortcuts. To enable the functionality (e.g., data selection) of the Projection plane, users need to set the minimum coverage and productivity to filter the potential shortcuts. Moreover, in the Template View, if the number of selected shortcuts is large, users may need to scroll down for a while to find shortcuts of their interest. It would be useful to have a search function for users to look for shortcuts containing words and structures of their interest. Last but not least, since we use color to encode labels, if the number of categories is large (i.e., 100), different colors for different classes may be indistinguishable. The system should also support displaying the labels near the glyph of shortcuts if this situation occurs.

Conduct a long-term evaluation of the effectiveness of the system. In this paper, we have demonstrated case studies and expert interviews to evaluate the effectiveness of the system. Actionable insights verification is conducted to prove that the derived actionable insights could be utilized to improve the benchmark dataset quality. Currently, the verification only focuses on evaluating whether handling one shortcut can improve the benchmark dataset quality. The interventions currently are only simple demonstrations or executions, and the effects of the interventions are small. In the future, a long-term evaluation, which includes iteratively exploring and solving more types of shortcuts, can be conducted. We can invite experts to perform multiple meaningful and purposeful interventions for the benchmark dataset. The goal of these interventions would be to reduce the number of shortcuts in the dataset to satisfactory levels, that is, little shortcuts with high productivity and coverage in datasets under the scope of the system. Assessed quality results of final datasets, such as machine performance and the number of shortcuts, can be calculated and reported as the long-term evaluation results of the system.

8 CONCLUSION

NLU benchmarks are found to suffer the spurious bias and inflate the model performance in recent years. To further help NLU experts to build challenging and pertinent benchmark datasets, in this paper, we develop a system called *ShortcutLens* to support dataset creators in conducting multi-level exploration of shortcuts in the NLU benchmark dataset. Specifically, the Statistics View helps users grasp the overview of shortcuts in the benchmark dataset and conduct a what-if analysis of shortcuts of interest. The Template View employs hierarchical and interpretable templates to help summarize the benchmark dataset issues through shortcuts. The Instance View allows users to further explore the instances covered by shortcuts. We have conducted two case studies with the experts on understanding and mitigating shortcuts in NLU benchmark datasets, as well as expert interviews with eight experts to demonstrate the usefulness and usability of the *ShortcutLens*.

ACKNOWLEDGMENTS

We would like to thank external experts and anonymous reviewers for their valuable comments on this work. This research was supported in part by Hong Kong Theme-based Research Scheme grant T41-709/17N.

REFERENCES

- S. R. Bowman and G. E. Dahl, "What will it take to fix benchmarking in natural language understanding?" in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, 2021, pp. 4843–4855.
- [2] K. Ethayarajh and D. Jurafsky, "Utility is in the eye of the user: A critique of nlp leaderboards," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2020, pp. 4846–4853.
- [3] A. Wang, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman, "Glue: A multi-task benchmark and analysis platform for natural language understanding," in *International Conference on Learning Representations*, 2019.
- [4] A. Wang, Y. Pruksachatkun, N. Nangia, A. Singh, J. Michael, F. Hill, O. Levy, and S. R. Bowman, "Superglue: A stickier benchmark for general-purpose language understanding systems," in *Advances in Neural Information Processing Systems*, vol. 32, 2019, pp. 3261–3275.
- [5] C. Sun, L. Huang, and X. Qiu, "Utilizing bert for aspect-based sentiment analysis via constructing auxiliary sentence," in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, 2019, pp. 380–385.
- [6] H. Xu, B. Liu, L. Shu, and P. S. Yu, "Bert post-training for review reading comprehension and aspect-based sentiment analysis," in *Proceedings of* the Conference of the North American Chapter of the Association for Computational Linguistics, 2019, pp. 2324–2335.
- [7] R. K. Kaliyar, A. Goswami, and P. Narang, "Fakebert: Fake news detection in social media with a bert-based deep learning approach," *Multimedia Tools and Applications*, vol. 80, no. 8, pp. 11765–11788, 2021.
- [8] H. Jwa, D. Oh, K. Park, J. M. Kang, and H. Lim, "exbake: Automatic fake news detection model based on bidirectional encoder representations from transformers (bert)," *Applied Sciences*, vol. 9, no. 19, p. 4062, 2019.
- [9] A. Paullada, I. D. Raji, E. M. Bender, E. Denton, and A. Hanna, "Data and its (dis)contents: A survey of dataset development and use in machine learning research," *Patterns*, vol. 2, no. 11, p. 100336, 2021.
- [10] D. Kiela, M. Bartolo, Y. Nie, D. Kaushik, A. Geiger, Z. Wu, B. Vidgen, G. Prasad, A. Singh, P. Ringshia, Z. Ma, T. Thrush, S. Riedel, Z. Waseem, P. Stenetorp, R. Jia, M. Bansal, C. Potts, and A. Williams, "Dynabench: Rethinking benchmarking in nlp," in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, 2021, pp. 4110–4124.
- [11] R. Geirhos, J.-H. Jacobsen, C. Michaelis, R. Zemel, W. Brendel, M. Bethge, and F. A. Wichmann, "Shortcut learning in deep neural networks," *Nature Machine Intelligence*, vol. 2, no. 11, pp. 665–673, 2020.
- [12] K. Sakaguchi, R. Le Bras, C. Bhagavatula, and Y. Choi, "Winogrande: An adversarial winograd schema challenge at scale," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, 2020, pp. 8732–8740.
- [13] R. Le Bras, S. Swayamdipta, C. Bhagavatula, R. Zellers, M. Peters, A. Sabharwal, and Y. Choi, "Adversarial filters of dataset biases," in *Proceedings of the 37th International Conference on Machine Learning*, 2020, pp. 1078–1088.

- [14] T. Niven and H.-Y. Kao, "Probing neural network comprehension of natural language arguments," in *Proceedings of the 57th Conference of* the Association for Computational Linguistics, 2019, pp. 4658–4664.
- [15] R. Branco, A. Branco, J. Rodrigues, and J. Silva, "Shortcutted commonsense: Data spuriousness in deep learning of commonsense reasoning," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2021, pp. 1504–1521.
- [16] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, 2019, pp. 4171–4186.
- [17] B. Wang, C. Xu, S. Wang, Z. Gan, Y. Cheng, J. Gao, A. H. Awadallah, and B. Li, "Adversarial glue: A multi-task benchmark for robustness evaluation of language models," *arXiv preprint arXiv:2111.02840*, 2021.
- [18] A. Poliak, J. Naradowsky, A. Haldar, R. Rudinger, and B. Van Durme, "Hypothesis only baselines in natural language inference," in *Proceedings* of the Seventh Joint Conference on Lexical and Computational Semantics, 2018, pp. 180–191.
- [19] S. Gururangan, S. Swayamdipta, O. Levy, R. Schwartz, S. R. Bowman, and N. A. Smith, "Annotation artifacts in natural language inference data," in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, 2018, pp. 107–112.
- [20] Y. Lai, C. Zhang, Y. Feng, Q. Huang, and D. Zhao, "Why machine reading comprehension models learn shortcuts?" in *Findings of the Association for Computational Linguistics*, 2021, pp. 989–1002.
- [21] R. T. McCoy, E. Pavlick, and T. Linzen, "Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference," in *Proceedings of the 57th Conference of the Association for Computational Linguistics*, 2019, pp. 3428–3448.
- [22] E. M. Bender and B. Friedman, "Data statements for natural language processing: Toward mitigating system bias and enabling better science," *Transactions of the Association for Computational Linguistics*, vol. 6, pp. 587–604, 2018.
- [23] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. D. Iii, and K. Crawford, "Datasheets for datasets," *Communications of the ACM*, vol. 64, no. 12, pp. 86–92, 2021.
- [24] J. Choo and S. Liu, "Visual analytics for explainable deep learning," *IEEE Computer Graphics and Applications*, vol. 38, no. 4, pp. 84–92, 2018.
- [25] H. Strobelt, S. Gehrmann, H. Pfister, and A. M. Rush, "Lstmvis: A tool for visual analysis of hidden state dynamics in recurrent neural networks," *IEEE Transactions on Visualization and Computer Graphics*, vol. 24, no. 1, pp. 667–676, 2017.
- [26] Y. Ming, S. Cao, R. Zhang, Z. Li, Y. Chen, Y. Song, and H. Qu, "Understanding hidden memories of recurrent neural networks," in *IEEE Conference on Visual Analytics Science and Technology*, 2017, pp. 13–24.
- [27] J. F. DeRose, J. Wang, and M. Berger, "Attention flows: Analyzing and comparing attention mechanisms in language models," *IEEE Transactions* on Visualization and Computer Graphics, vol. 27, no. 2, pp. 1160–1170, 2020.
- [28] I. Tenney, J. Wexler, J. Bastings, T. Bolukbasi, A. Coenen, S. Gehrmann, E. Jiang, M. Pushkarna, C. Radebaugh, E. Reif *et al.*, "The language interpretability tool: Extensible, interactive visualizations and analysis for nlp models," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2020, pp. 107–118.
- [29] X. Wang, J. He, Z. Jin, M. Yang, Y. Wang, and H. Qu, "M2lens: Visualizing and explaining multimodal models for sentiment analysis," *IEEE Transactions on Visualization and Computer Graphics*, vol. 28, no. 1, pp. 802–812, 2022.
- [30] B. Laughlin, C. Collins, K. Sankaranarayanan, and K. El-Khatib, "A visual analytics framework for adversarial text generation," in *IEEE Symposium* on Visualization for Cyber Security, 2019, pp. 1–10.
- [31] T. Wu, K. Wongsuphasawat, D. Ren, K. Patel, and C. DuBois, "Tempura: Query analysis with structural templates," in *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 2020, pp. 1–12.
- [32] Z. Jin, Y. Wang, Q. Wang, Y. Ming, T. Ma, and H. Qu, "Gnnlens: A visual analytics approach for prediction error diagnosis of graph neural networks," *IEEE Transactions on Visualization and Computer Graphics*, 2022.
- [33] S. Liu, Z. Li, T. Li, V. Srikumar, V. Pascucci, and P.-T. Bremer, "Nlize: A perturbation-driven visual interrogation tool for analyzing and interpreting natural language inference models," *IEEE Transactions on Visualization* and Computer Graphics, vol. 25, no. 1, pp. 651–660, 2018.
- [34] Y. Ming, P. Xu, F. Cheng, H. Qu, and L. Ren, "Protosteer: Steering deep sequence model with prototypes," *IEEE Transactions on Visualization* and Computer Graphics, vol. 26, no. 1, pp. 238–248, 2019.

- [35] J. Yuan, C. Chen, W. Yang, M. Liu, J. Xia, and S. Liu, "A survey of visual analytics techniques for machine learning," *Computational Visual Media*, vol. 7, no. 1, pp. 3–36, 2021.
- [36] S. Kandel, R. Parikh, A. Paepcke, J. M. Hellerstein, and J. Heer, "Profiler: Integrated statistical analysis and visualization for data quality assessment," in *Proceedings of the International Working Conference on Advanced Visual Interfaces*, 2012, pp. 547–554.
- [37] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 6402–6413.
- [38] K. Lee, H. Lee, K. Lee, and J. Shin, "Training confidence-calibrated classifiers for detecting out-of-distribution samples," in *International Conference on Learning Representations*, 2018.
- [39] C. Bors, T. Gschwandtner, and S. Miksch, "Capturing and visualizing provenance from data wrangling," *IEEE Computer Graphics and Applications*, vol. 39, no. 6, pp. 61–75, 2019.
- [40] C. Chen, J. Yuan, Y. Lu, Y. Liu, H. Su, S. Yuan, and S. Liu, "Oodanalyzer: Interactive analysis of out-of-distribution samples," *IEEE Transactions* on Visualization and Computer Graphics, vol. 27, no. 7, pp. 3335–3349, 2020.
- [41] S. Alemzadeh, U. Niemann, T. Ittermann, H. Völzke, D. Schneider, M. Spiliopoulou, K. Bühler, and B. Preim, "Visual analysis of missing values in longitudinal cohort study data," in *Computer Graphics Forum*, vol. 39, no. 1, 2020, pp. 63–75.
- [42] S. Mishra, A. Arunkumar, B. Sachdeva, C. Bryan, and C. Baral, "Dqi: Measuring data quality in nlp," arXiv preprint arXiv:2005.00816, 2020.
- [43] D. H. Clements and M. T. Battista, "Geometry and spatial reasoning," *Handbook of Research on Mathematics Teaching and Learning*, pp. 420– 464, 1992.
- [44] R. Mirzaee, H. R. Faghihi, Q. Ning, and P. Kordjmashidi, "Spartqa: A textual question answering benchmark for spatial reasoning," in *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics*, 2021, pp. 4582–4598.
- [45] A. Warstadt, A. Singh, and S. R. Bowman, "Neural network acceptability judgments," *Transactions of the Association for Computational Linguistics*, vol. 7, pp. 625–641, 2019.
- [46] L. McInnes, J. Healy, and J. Melville, "Umap: Uniform manifold approximation and projection for dimension reduction," *arXiv preprint* arXiv:1802.03426, 2018.
- [47] B. Everitt, S. Landau, M. Leese, and D. Stahl, *Cluster Analysis*, 5th ed. Wiley, 2011.
- [48] Y. Cui, W. Che, T. Liu, B. Qin, and Z. Yang, "Pre-training with whole word masking for chinese bert," *IEEE ACM Transactions on Audio, Speech,* and Language Processing, vol. 29, pp. 3504–3514, 2021.
- [49] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," arXiv preprint arXiv:1907.11692, 2019.
- [50] Z. Zhang, X. Han, Z. Liu, X. Jiang, M. Sun, and Q. Liu, "Ernie: Enhanced language representation with informative entities," in *Proceedings of the* 57th Conference of the Association for Computational Linguistics, 2019, pp. 1441–1451.
- [51] A. Rogers, M. Gardner, and I. Augenstein, "Qa dataset explosion: A taxonomy of nlp resources for question answering and reading comprehension," *arXiv preprint arXiv:2107.12708*, 2021.
- [52] J. Guan, F. Huang, Z. Zhao, X. Zhu, and M. Huang, "A knowledgeenhanced pretraining model for commonsense story generation," *Transactions of the Association for Computational Linguistics*, vol. 8, pp. 93–108, 2020.
- [53] Y. Zang, F. Qi, C. Yang, Z. Liu, M. Zhang, Q. Liu, and M. Sun, "Wordlevel textual adversarial attacking as combinatorial optimization," in *Proceedings of the 58th Conference of the Association for Computational Linguistics*, 2020, pp. 6066–6080.
- [54] G. Zeng, F. Qi, Q. Zhou, T. Zhang, Z. Ma, B. Hou, Y. Zang, Z. Liu, and M. Sun, "Openattack: An open-source textual adversarial attack toolkit," arXiv preprint arXiv:2009.09191, 2020.
- [55] D. Jin, Z. Jin, J. T. Zhou, and P. Szolovits, "Is bert really robust? a strong baseline for natural language attack on text classification and entailment," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, 2020, pp. 8018–8025.
- [56] T. Schick and H. Schütze, "Generating datasets with pretrained language models," in *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, 2021, pp. 6943–6951.
- [57] M. Zhao, F. Mi, Y. Wang, M. Li, X. Jiang, Q. Liu, and H. Schütze, "Lmturk: Few-shot learners as crowdsourcing workers," *arXiv preprint* arXiv:2112.07522, 2021.

[58] A. Liu, S. Swayamdipta, N. A. Smith, and Y. Choi, "Wanli: Worker and ai collaboration for natural language inference dataset creation," *arXiv* preprint arXiv:2201.05955, 2022.



Zhihua Jin is currently a Ph.D. student at the Hong Kong University of Science and Technology (HKUST). He received his BEng degree in Computer Science and Technology from Zhejiang University in 2019. His research interests lie in the intersection of visualization and machine learning, especially explainable artificial intelligence (XAI).



Xingbo Wang is currently a postdoctoral fellow in the Department of Computer Science and Engineering at the Hong Kong University of Science and Technology (HKUST), where he also received his Ph.D. degree in 2022. He obtained a B.E. degree from Wuhan University, China in 2018. His research interests include humancomputer interaction (HCI), data visualization, natural language processing (NLP), and multimodal analysis. For more details, please refer to https://andy-xingbowang.com/.



Furui Cheng is a Ph.D. candidate in the Department of Computer Science and Engineering at the Hong Kong University of Science and Technology (HKUST). He obtained a B.E. degree from Beihang University, China, in 2018. His research interests include visual analytics, eXplainable AI (XAI), and biomedical AI.



Chunhui Sun is a Ph.D. candidate in the Department of Chinese Language and Literature at Peking University (PKU). He obtained a M.Ed. degree from Soochow University, China, in 2018. His research interests lie in the areas of Computational Linguistics, Cognitive Linguistics, Natural Language Understanding Benchmarks, and eXplainable AI (XAI), etc.



Qun Liu is the Chief Scientist of Speech and Natural Language Processing of Huawei Noah's Ark Lab. He was a Full Professor in Dublin City University and the Theme Leader of the ADAPT Centre, Ireland during July 2012 and June 2018. Before that, he was as a Professor in the Institute of Computing Technology (ICT), Chinese Academy of Sciences for 20 years, where he founded and led the ICT NLP Research Group. He obtained a B.Sc. in computer science in the University of Science and Technology of China, a

M.Sc. in Chinese Academy of Sciences, and a Ph.D. in Peking University. His research interests lie in the areas of Natural Language Processing, Machine Translation, Pre-trained Language Models, etc.



Huamin Qu is a professor in the Department of Computer Science and Engineering (CSE) at the Hong Kong University of Science and Technology (HKUST) and also the director of the interdisciplinary program office (IPO) of HKUST. He obtained a BS in Mathematics from Xi'an Jiaotong University, China, an MS and a PhD in Computer Science from the Stony Brook University. His main research interests are in visualization and human-computer interaction, with focuses on urban informatics, social network analysis, E-

learning, text visualization, and explainable artificial intelligence (XAI).