

XNLI: Explaining and Diagnosing NLI-based Visual Data Analysis

Yingchaojie Feng, Xingbo Wang, Bo Pan, Kam Kwai Wong, Yi Ren, Shi Liu, Zihan Yan,
Yuxin Ma, Huamin Qu, and Wei Chen.

Abstract—Natural language interfaces (NLIs) enable users to flexibly specify analytical intentions in data visualization. However, diagnosing the visualization results without understanding the underlying generation process is challenging. Our research explores how to provide explanations for NLIs to help users locate the problems and further revise the queries. We present XNLI, an explainable NLI system for visual data analysis. The system introduces a Provenance Generator to reveal the detailed process of visual transformations, a suite of interactive widgets to support error adjustments, and a Hint Generator to provide query revision hints based on the analysis of user queries and interactions. Two usage scenarios of XNLI and a user study verify the effectiveness and usability of the system. Results suggest that XNLI can significantly enhance task accuracy without interrupting the NLI-based analysis process.

Index Terms—Natural language interface, visual data analysis, explainability.

1 INTRODUCTION

NATURAL Language Interfaces (NLIs) have received widespread attention in academia [1], [2], [3], [4], [5] and industry [6], [7]. In visual data analysis, typical NLIs [1], [4], [5] allow users to pose natural language (NL) queries to specify visualizations from data attributes, analytical tasks, and visual encodings. To better meet user expectations, previous studies have enhanced the query interpretation capability of NLIs, such as parsing long compound statements [3], resolving ambiguities [4], inferring underspecification [8], and understanding contextual references [3].

However, as the visualization generation process contains multi-step data transformations and visual encodings [9], users have difficulty diagnosing the final visualization results without understanding the underlying generation process. Considering a movie exploration scenario, a user poses a query to select the movie records within a rating range and then visualize their average budgets across genres in a bar chart. As the rating values are not encoded in the resulting chart, the user cannot assess the correctness of the filtering task. In addition, the lack of knowledge about the NLI impedes users from correcting unexpected results. Thus, users may fail to derive desired insights from data and lose trust in NLI tools. Both of these issues hinder the process of NLI-based data analysis and decision-making.

Prior studies [2], [3], [4], [10] have introduced interactive

widgets to help users identify ambiguities in the queries and support adjustments. Nevertheless, diagnosing the unexpected results caused by failures of query interpretation in the NLI pipeline or impractical visualizations specified in the queries remains a challenge. In this study, we investigate the major causes of the unexpected results and summarize them with the key aspects of the NLI process (*i.e.*, attributes, tasks, and visual encodings). To help users understand the NLI process, identify problems, and possibly fix them during data analysis, we formulate a set of design requirements and develop XNLI, an explainable NLI that elucidates the visualization generation process, supports interactive adjustments, and provides hints for query revision.

In addition to interpreting user queries and generating visualization results, the system reveals the process from the aspects of the attributes, tasks, and visual encodings and leverages data provenance to facilitate user understanding. The relevant trigger words in the original queries and the inference uncertainty during the query interpretation are visualized for examining the potential errors in the process and locating their origins. Based on this information, the system provides interactive widgets, allowing users to intervene and adjust the visualization process. To facilitate the use of NLI for the desired data analysis, the system helps users refine their queries by providing two types of hints. Rule-based hints point out the potential problems in the queries regarding attributes, tasks, and visual encodings. Interaction-based hints are generated based on the analysis of users' interactive adjustments, and contain query revision suggestions as well as valid query examples that can be applied in the subsequent query formulation.

Finally, we introduce two usage scenarios to illustrate how our system helps users identify and correct different kinds of problems in the process, such as unexpected attributes and tasks, and gain knowledge about the NLI capability (e.g., how to phrase system-supported queries). We also conduct a controlled user study with a baseline to evaluate the effectiveness and usability of our system

- Y. Feng, B. Pan, Y. Ren, S. Liu, W. Chen are with The State Key Lab of CAD & CG, Zhejiang University, Hangzhou, Zhejiang, China. W. Chen is also with the Laboratory of Art and Archaeology Image (Zhejiang University), Ministry of Education, China. W. Chen is the corresponding author. E-mail: {fycj, bopan, rymaster, zju_ls, chenvis}@zju.edu.cn.
- X. Wang, KK Wong, H. Qu are with the Hong Kong University of Science and Technology, Hong Kong, China. Email: {xwang, kkwong, huamin}@cse.ust.hk.
- Z. Yan is with the MIT Media Lab, Cambridge, MA, USA. Email: yzihan@media.mit.edu.
- Y. Ma is with the Department of Computer Science and Engineering, Southern University of Science and Technology, Guangdong, China. Email: mayx@sustech.edu.cn.

Manuscript received April 19, 2005; revised August 26, 2015.

and the impact of explanations on the analysis process. The study results suggest that our system is easy to use, and the explanations effectively improve the analytical efficiency without interrupting the analysis process. In summary, our contributions include:

- A pilot study for NLI-based visual analysis that identifies challenges in process understanding and summarizes the major problems from the NLI aspects.
- An NLI-based visual analysis system that explains the NLI process, supports interactive error correction, and provides hints to help query formulation and revision.
- Two usage scenarios and a controlled user study to evaluate the effectiveness and usability of our system.

2 RELATED WORK

2.1 NLI for Data Visualization

NLIs are widely studied as a promising means of interaction for visual analytics [1], [11]. NLIs for data visualization generate visual representations in response to NL queries, which help reveal data insights [12], [13], [14], [15], [16], [17], [18], [19]. Cox et al. [20] aimed to integrate NLI into existing visualization systems. Subsequently, extensive research has focused on understanding complex and free-form NL queries, such as incomplete or referential utterances [2], [3] and conversational queries [2], [3], [21], [22]. NL4DV [5] is a general NLI toolkit that integrates a set of NLP techniques (e.g., dependency parsing) to interpret NL data queries and recommends visualizations to generate results. Snowy [23] and QRec-NLI [24] recommend utterances in NL-based visual analysis, leveraging NL to provide analytical guidance. In recent years, NLIs have been applied to various scenarios and tasks, such as flow data exploration [25], visualization authoring [26], and comparative analysis [27]. Other studies [28], [29], [30], [31] build query corpora for linguistic property analysis and construct labeled benchmark datasets for NLI evaluation and deep learning model training [32], [33].

In addition to enhancing the capability of query interpretation, revealing the NLI process and enabling feedback for error correction are also necessary. DataTone [4] introduces ambiguity widgets highlighting ambiguities in query interpretation and supporting interactive specification. Such interactive widgets have been widely integrated into later works [2], [3], [10], [13]. Commercial NLI systems, such as Power BI [7] and Tableau [6], present the intermediate parsing results as structured text and enable interactive adjustments of the chart results. Our study provides systematic process explanations from the key aspects of the general NLI pipeline [1], [5] and demonstrates data provenance to help users understand the process. Besides supporting widget-based adjustments, our system provides suggestions for query rephrasing and revision, facilitating the use of NLI in the subsequent analysis process.

2.2 Provenance

Provenance records the evolution of the analysis process and has been widely studied in the visualization field. Ragan et al. [34] summarize an organizational framework for the provenance of information (data, visualization, and interaction) and purposes (collaborative communication,

presentation, and meta-analysis). Heer et al. [35] abstract the provenance as a node-link diagram and propose a graphical history tool with provenance to support analysis, communication, and evaluation. Provenance can be used in data wrangling [36], [37], [38] and analytical reasoning [39]. Datamations [40] maps the data operations in the provenance to various types of animations and demonstrates the benefits of data animation to explain the Simpson's paradox.

Some works employ provenance to explain complex codes. QueryVis [41] generates logic-based diagrams to represent the meaning of deeply nested SQL statements. Berant et al. [42] design a cell-based provenance with NL explanations to accelerate user understanding of complex queries. Data Tweening [43] introduces a grammar to generate incremental visualizations for queries of data transformation. A recent work, DIY [44], leverages the representative sample data to demonstrate the transformation process during the execution of SQL statements.

We base our work on provenance to facilitate user understanding of the NLI process and support interactive adjustments. Moreover, we highlight the inference uncertainties and ambiguities during the query interpretation to guide users in investigating the potential problems in the process, further improving the efficiency of problem identification.

3 DESIGN PROCESS

The target users of our system are data analysts who leverage NLIs to perform visual analysis. The design process lasted 9 months and started with a pilot study to investigate the challenges of understanding and diagnosing visualization results in NLI-based analysis scenarios. We summarized the unexpected system results in the open-ended task exploration and categorized them into the major cause types from the perspective of the NLI process. We iteratively refined the design requirements based on user feedback and built corresponding system prototypes for verification.

3.1 Experiment Design

Participants. We recruited 9 data analysts (two females and seven males) from a local university. Seven participants have experience in visual analysis, with one having prior experience in using NLI. The remaining two are beginners in visual analysis.

Settings. We developed an NLI prototype system for visual analysis. It is powered by NL4DV [5], a typical NLI toolkit for tabular data analysis. The system enables users to preview the dataset and pose NL queries. We used the IMDb movie dataset, which contains 710 movie items with nine attributes. Participants were requested to perform open-ended task exploration [45] over the dataset to provide investment advice.

Procedure. We first introduced the study settings and workflow of the system prototype. After the participants familiarized themselves with the system, they performed the open-ended task exploration. We provided minimal support during the sessions and let participants obtain their desired analytical results through trial and error. Then, we conducted a post-study interview with participants to reflect on the exploration experience.

3.2 Results Analysis

During the post-study interview, we explained to the participants the provenance of the visualization results obtained during the exploration process, from which we identified some system results that did not match the user’s expectations. For these cases, we asked the participants to revise their queries to get the desired visualizations with a time constraint. Finally, following the bottom-up approach in the prior study, [6], [46], [47], [48], we summarized the major causes of the unexpected system results from the aspects of attributes, tasks, and visual encodings. Furthermore, we analyzed the participants’ ability to locate and correct unexpected system results. The results are shown in Table 1.

TABLE 1

Major types of unexpected results that we summarized from three key aspects (*i.e.*, attributes, tasks, visual encodings) of the NLI workflow.

Aspects	Unexpected Behaviors	Total	Locate	Correct
Attributes	Fail to infer	17	12	9
	Unexpected types	14	5	5
Tasks	Fail to infer	9	6	1
	Unexpected types	15	9	4
	Impractical parameters	12	5	4
Visual Encodings	Chart types	8	6	4
	Encoding channels	5	5	3

For **attributes**, NLI usually *failed to infer* the attributes. Participants inputted typos or used synonyms, and NLI failed to match them with the corresponding attributes. Fortunately, these issues could usually be identified due to the absence of attributes in the resulting charts and were easy to correct. *Unexpected types* of attributes were typically caused by incorrect implicit inference. For example, the vague modifier “high” in the query was treated by the NLI as the specific item of the “Title” attribute, which also led to an unexpected filtering (*i.e.*, only choosing the movies with “high” in the title). Participants felt something wrong but always could not locate the problems or correct them.

For **tasks**, NLI might *fail to infer* their types, including unsupported tasks (*e.g.*, finding the extremes) and task-related keywords (*e.g.*, failing to understand “at least”). *Unexpected types* of tasks occurred when multiple task schemas were available. A participant posed, “show the correlation of gross and IMDB rating”, expecting the average gross of the binned IMDB rating in a bar chart instead of all data items in a scatterplot. Participants sometimes specified tasks with *impractical parameters* (*e.g.*, out-of-range filter conditions) due to a lack of knowledge of the data. As a result, no data records matched the criteria, and empty charts were returned, which always confused the users when trying to locate the problems. Another reason for *impractical parameters* was the incorrect mappings of tasks and parameters constructed based on the parsing dependency of NL queries. Participants had difficulty noticing the issues if the task-related attributes were not encoded in the visualization results.

For **visual encodings**, unexpected *chart types* made up the majority. Participants might specify an unsupported chart type and get nothing back. They could locate the problem and expected the system to suggest some available chart types. Unexpected *chart types* also occurred when they were

underspecified (*e.g.*, participants expected a pie chart to visualize the data proportion instead of the provided bar chart.) A few cases of unexpected *encoding channels* were observed (*e.g.*, participants preferred encoding a quantitative attribute by color instead of size in the scatter plot).

In summary, the user challenge of process understanding is twofold. First, it can be difficult to locate the problems of the unexpected results, especially for the underlying attributes and tasks which are not encoded in the final chart results. Second, it would be time-consuming for inexperienced users to correct the errors and obtain the desired results through query revision. Due to their lack of knowledge about the NLI’s mechanism and capability, users might fail to revise the queries with the correct trigger words and phrases that are well supported by the NLI system.

3.3 Design Requirements

To help users understand the visualization process and address the aforementioned issues, we conducted regular bi-weekly meetings with all participants (*P1-P9*) from the pilot study and refined our prototype system iteratively. We also conducted a literature survey of the prior work on NLI. Finally, we summarized a set of design requirements.

R1: Provide an overview of the visualization process from the key aspects of the NLI pipeline. The typical NLIs interpret the user queries and formulate the visualization responses based on three key aspects (*i.e.*, attributes, tasks, and visual encodings). As observed in the pilot study, it is challenging for users to identify unexpected results without knowing the NLI pipeline. Most participants stated that the system should present these key aspects as an overview of the visualization process so that they could go through the process and dive into some aspects of interest.

R2: Leverage the representative data items to demonstrate the visualization process. *P2* pointed out that explaining the process with data could promote understanding and increase trust. However, presenting all the contents of a large dataset causes serious visual clutter beyond human recognition affordance [38], [44]. The system should choose the representative data items from relevant data columns and use them to demonstrate the visualization process. Moreover, visual cues (*e.g.*, highlighting) are also necessary to guide the user’s attention to the differences between steps.

R3: Reveal the connections between the visualization process and the original query. Each aspect of the visualization process is specified according to the original query. These potential connections constructed by NLI [16] are usually inaccessible to the users. *P1* and *P5* commented that they wondered how their queries were parsed by the system, especially when obtaining unexpected results. The system should link the visualization process to the related parts of the query to help users to investigate the causes of unexpected results and gain inspiration for query revision.

R4: Show the inference types of NLI to highlight the uncertainty of query interpretation. NLI employed multiple strategies for query interpretation [4], [10], [12], [49] because of the diversity of NL expression. For example, data attributes may be explicitly or implicitly mentioned in the query, which require various approaches for inference [1] and thus introduces different levels of uncertainty [5]. Presenting

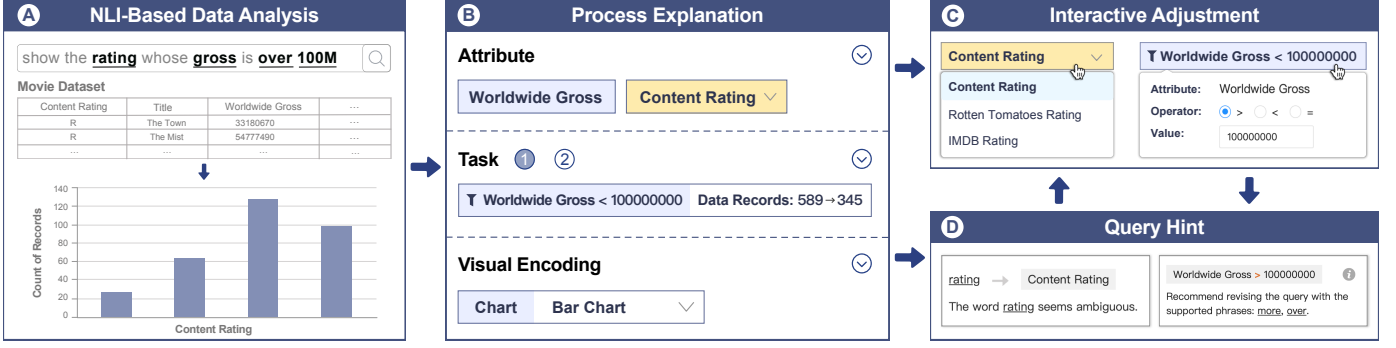


Fig. 1. The workflow of our system consists of four components that (A) support NLI-based data analysis, (B) provide process explanations for visualization results from the aspects of attributes, tasks, and visual encodings, (C) provide interactive widgets to adjust the process from each aspect, and (D) provide hints to help problem localization and query revision.

the inference types and ambiguities makes users aware of the uncertainty in the process and prioritizes attention for further inspection.

R5: Support adjustment of the visualization process through directed manipulation. A considerable proportion of cases from the pilot study show that even when users locate the problems of unexpected results, they still face challenges in query revision for desired results. Directed manipulation can serve as a complementary interaction modality [13], [50] to support the further expression of user expectations. Prior works [2], [3], [4], [13] have shown the great potential of interactive widgets in resolving ambiguity. The system should provide interactive widgets to support user adjustments to the visualization process.

R6: Provide hints and query examples for query phrasing and revision. Compared to interactive adjustments of current query results, providing hints for query revision has a more positive effect on the subsequent iteration of NL interaction. This requirement leverages unexpected cases as teaching opportunities and spares users from the tedious task of guessing supported utterances. *P6* suggested that the system should inform the users about potential errors in the query proactively. *P1* stated that providing smart and effective suggestions and valid query examples would be better to help him revise the original queries.

4 SYSTEM DESIGN

To meet the derived requirements, we develop XNLI, a prototype system that provides visual explanations for NLI results to facilitate understanding of the visualization generation process and gain knowledge of query phrasing. In this section, we first describe an overview of the system workflow. Then, we introduce the visual and functional designs of the user interface to demonstrate the usage of explanations for NLI-based visual analysis. Finally, we proceed to the detailed implementation of major system components.

4.1 System Workflow

XNLI is implemented as a web-based system with a graphical user interface to enable easy access. The system workflow is summarized in Figure 1. As the basic feature of NLI, our system supports **NL-based data analysis** (Figure 1A). Users can pose data-related queries through NL and conduct visual analysis based on the generated visualization

results. The query interpretation capability of our system is powered by NL4DV [5], a general and representative toolkit for NLI construction. To facilitate user understanding of the visualization results, our system provides **process explanations** (Figure 1B) regarding three key aspects of the NLI pipeline, *i.e.*, attributes, tasks, and visual encodings (R1-R2). The process explanations build a basis for users to identify unexpected results (R3-R4). The system further provides **interactive widgets** (Figure 1C) for each visualization aspect, leveraging directed manipulation as a complementary interaction modality to support error correction (R5). To help users benefit more from NLI, the system provides **query hints** (Figure 1D) as feedback to indicate potential errors in user queries and suggest valid query phrasing (R6). The hints are generated based on the analysis of visualization aspects and logs of interactive adjustments.

4.2 User Interface

The user interface (Figure 2) consists of three modules. The NLI module allows users to overview the explored dataset (Figure 2A₁, A₂) and pose NL queries for visual data analysis (Figure 2B₁, B₂). The Explanation module reveals the visual transformation process (Figure 2B₃) from aspects of the attributes, tasks, and visual encodings. Each aspect supports interactive adjustments for error correction. To further help the user derive the desired queries for data analysis, the Hint module provides suggestions for query revision and pops up (Figure 2B₄) when the user clicks the icon (Ω).

4.2.1 NLI Module

After users choose or upload a dataset for exploration (Figure 2A₁), the Data Overview panel presents metadata for each attribute in the dataset (Figure 2A₂), including the attribute name, type, and typical values. Attribute types include Quantitative, Nominal, Ordinal, and Temporal [51]. The typical values of the attributes are chosen based on the occurrence rate. To the right of each attribute name is a filter icon that users can click to add a filtering task for the attribute to the visualization process (R5).

The input box (Figure 2B₁) at the top of Query View allows users to type NL queries. After the query interpretation, the keywords in the sentence will be encoded in different styles according to their usage. The keywords related to attributes, tasks, and visual encoding are encoded in black and underlined, while other words are encoded in grey.

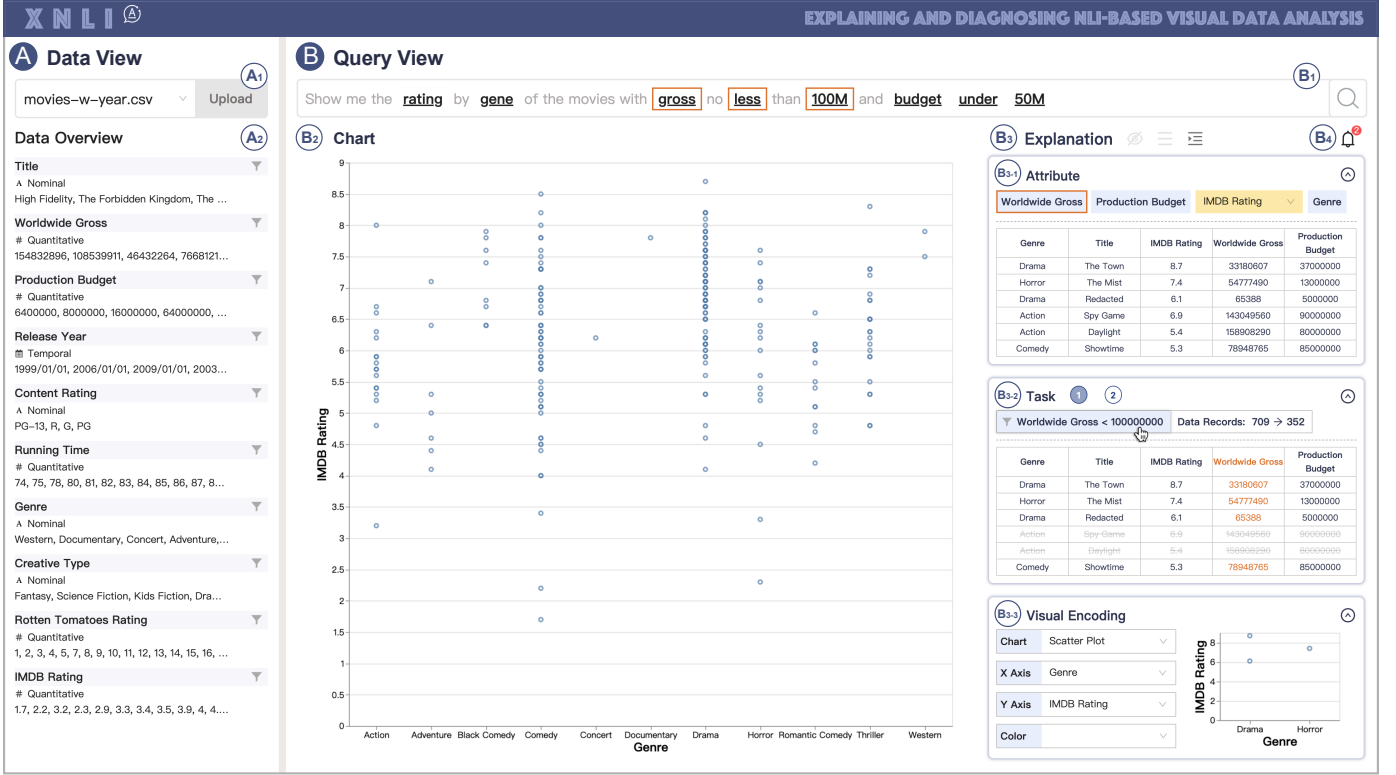


Fig. 2. The user interface of XNLI consists of two views, including (A) the Data View and (B) the Query View. Users can (A1) select or upload the dataset, (A2) view the data attributes, (B1) input natural language queries, (B2) use charts for data analysis, (B3) check the explanations of the NLI process and correct the unexpected results through interactive widgets, and (B4) gain hint feedback for query revision.

The Chart panel (Figure 2B₂) presents the visualization results in response to user queries. It is the main window for data analysis. The charts are specified by Vega-Lite and support user interactions, such as hovering and selection.

4.2.2 Explanation Module

The Explanation panel (Figure 2B₃) presents the visualization process from aspects of the attributes, tasks, and visual encodings (R1). Each aspect supports overview and details-on-demand. The detailed level can be collapsed (⊖) and opened (⊕). Users can choose the default display modes of the explanations in the title bar, including *No Explanation* (⊖), *Overview Explanation* (⊕), and *Detailed Explanation* (⊕).

The Attribute aspect (Figure 2B₃₋₁) lists all the attribute names involved and visualizes the inference types in different background colors (R4). The attributes that are explicitly inferred (e.g., Genre in Figure 2B₃₋₁) are encoded in light blue, while the attributes that are implicitly inferred (e.g., Title in Figure 5A₃) or ambiguous (e.g., IMDB Rating in Figure 2B₃₋₁) are encoded in yellow. The ambiguous attributes are augmented with interactive selectors (Figure 1C) to support adjustments (R5). User preference for ambiguities indicated by explicit adjustments or implicit agreement will be used for subsequent query disambiguation. The detailed level of the Attribute aspect provides a sample dataset to demonstrate the visualization provenance (R2).

The Task aspect (Figure 2B₃₋₂) shows the data transformation tasks and the corresponding changes in the actual amount of data. For example, the amount of data is reduced from 709 to 352 during the filtering task. The background color of the task “Worldwide Gross < 100,000,000” visualizes

the inference types from queries (R4), just like the attributes. Hovering on the task triggers the highlighting of the related keywords “gross”, “less”, and “100M” in the query (R3), helping users discover that the unexpected task is caused by a misinterpretation of “no less”. Users can click the task (Figure 1C) and fix the error in the pop-up box (R5). The detailed level of the Task aspect leverages visual cues to emphasize the changes of sample data provenance (R2). For example, the task is performed on the highlighted attribute Worldwide Gross, and some data items are filtered out.

The Visual Encoding aspect (Figure 2B₃₋₃) presents chart type and specific encoding channels. Users can change any mappings through directed manipulation (R5). The encoding results of sample data (after data transformation tasks) are shown on the right to aid in understanding and validation.

4.2.3 Hint Module

The Hint panel (Figure 3) provides two types of hints for query revision, i.e., rule-based hints and interaction-based hints. Rule-based hints (Figure 3A) indicate the potential errors in the queries (R6), such as input typos and ambiguities. The system detects these errors during the query interpretation and provides the hints alongside the generated visualization results. Other errors, such as unexpected types of attributes and tasks, are due to unexpected query interpretations that are unknown to the system. Therefore, the system provides interaction-based hints (Figure 3B) after user adjustments through interactive widgets. The interaction-based hints contain some targeted suggestions for local query modifications, as well as valid query examples

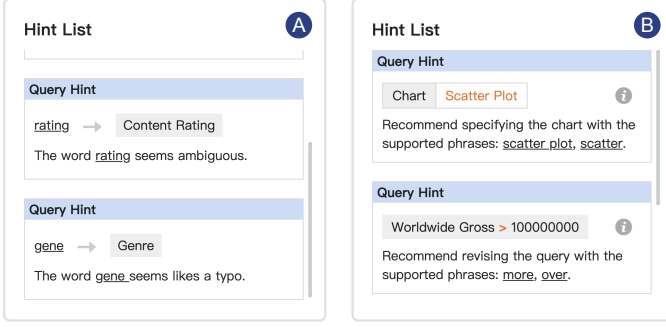


Fig. 3. Hint module presents (A) rule-based hints for issues in queries (e.g., input typos and ambiguities) and (B) interaction-based hints after user adjustments through interactive widgets.

that correspond to the adjusted visualization results and can be used in subsequent data analysis (R6).

4.3 Provenance Generator

The visualization results of NL4DV are specified with Vega-Lite grammar (Figure 4A), which is popular in the field of NLI [5], [23], [24], [28], [29], [30], [32], [52]. We leverage provenance to depict the three aspects of the visualization process and introduce the Provenance Generator, which consists of three major components to construct the visualization provenance from the Vega-Lite grammar, select representative sample data for demonstration, and apply visual cues to highlight the changes in provenance.

4.3.1 Vega-Lite Parsing

Vega-Lite is a high-level grammar for interactive data visualization [51]. This grammar supports declarative specification of unit visualization through a JSON object, which can be formally represented as the following four-tuple:

$$unit := (data, transforms, mark-type, encodings)$$

In NLI systems, the four-tuple reflects the query interpretation results of three aspects. Therefore, we parse the JSON object to reveal the visualization process.

According to the execution logic of the grammar, the initial dataset with attributes is first loaded from the location declared in the *data* field. Then, transformation tasks (if specified) are performed in the order specified in the *transformation* field. Aggregation tasks (e.g., sum, mean) are sometimes specified in the *encoding* fields, so we extract them from the visual encoding. To do so, we first identify the encoded attributes without an *aggregation* field or with a *bin* field, and use these attributes as dimensions to group data records, then aggregate the remaining attributes to obtain the derived values. To visualize the transformed data, we extract the chart type in the *mark* field and the encoding scheme in the *encoding* field to generate a new version of the Vega-Lite specification. This new version of visualization keeps the connection with the transformed data and supports visual cues (e.g., highlighting) and user interactions.

4.3.2 Data Selection

Inspired by the prior work [44], we add smart constraints to choose data subsets with relevant columns to achieve an

efficient presentation. For the filtering task in Figure 4B, smart constraints require the data subset to contain both data records that satisfy and do not satisfy the filtering condition. As a result, some data records in the data subset are filtered out while others are retained, making the task visible by comparing the data tables before and after the task.

The attributes specified in the Vega-Lite specification are chosen as relevant columns, allowing users to focus on useful information. However, presenting only the relevant columns is not intuitive for users to trace the data items between tasks. The system chooses one attribute as a key attribute and presents it throughout the sequence of data provenance. The key attribute needs to meet two requirements:

- 1) Uniquely identifies the data item, distinguishing itself from other data items.
- 2) Provides semantic information that is easy to understand (e.g., each person's name).

Therefore, we select the nominal or ordinal attribute with the lowest repetition rate as the key attribute. In Figure 4B, the Title field is used as the key attribute of the IMDb movie dataset to identify each data item.

4.3.3 Visual Cues

The system leverages visual cues to emphasize the changes in sample data provenance. For filtering tasks, the table columns of relevant attributes are highlighted in orange, and the filtered data items are displayed in translucent gray with a strike-through line. For aggregation and distribution tasks that group the data records by an attribute to calculate the new derived values, the same values of the grouped attribute are merged into the same cell. Visual cues are also used to connect the data tables and the visualization. The corresponding record in the data table will be highlighted when hovering on an element in the visualization charts.

4.4 Hint Generator

Hints are designed to help users in query phrasing and revision. Based on the results analysis of the pilot study (Section 3.2) and survey of prior work, we design the Hint Generator to produce hints during the data analysis. The system integrates a set of well-known rules to detect errors regarding attributes, tasks and visual encodings, and provides hints alongside the returned charts. For the problems from unexpected query interpretation, the system infers the potential errors according to visualization adjustments through widgets and provides hints after user interactions.

4.4.1 Ruled-based Hints

Rule-based hints are geared towards the problems in users' queries. The system first examines the query sentences and alerts users to the unused but potentially important keywords (e.g., *at least*). Then, the system checks the rest of the query sentence from the following aspects.

For **attributes**, misspellings and ambiguities are common during the input. The word in the query is identified as a misspelling if it is "close" but not equal to an attribute name (measured by Levenshtein distance [53]). We set the threshold of "close" to $0.2 * len(attributeName)$ based on prior work [12]. The ambiguities can be detected when multiple attributes are inferred from the same word.

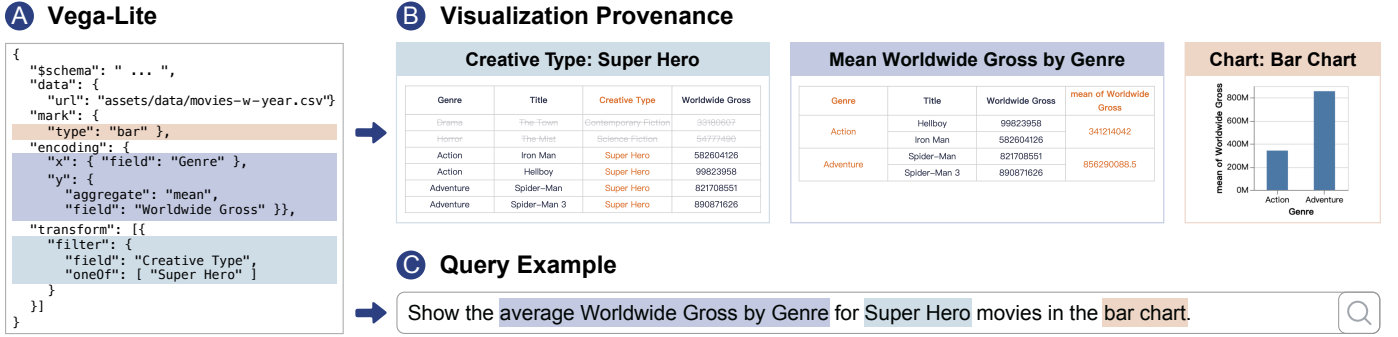


Fig. 4. XNLI leverages the (A) Vega-Lite specification to generate (B) visualization provenance to explain the process of visualization generation and (C) query examples as hints to help users revise the queries.

For **tasks**, sometimes users specify *impractical parameters* because of a lack of knowledge of the dataset. For example, an out-of-range filter criteria results in an empty chart with no data item. Our system records the changes in the data column during the provenance and informs users when the dataset becomes empty in the process.

For **visual encodings**, the system integrates a set of well-known rules to check the encoding scheme, including the chart type and encoding channel. When users pose unsuitable schemes for visual encoding (e.g., specifying a bar chart for two quantitative attributes), the system generates the requested chart type (e.g., bar chart) and provides a hint for the recommended chart type (e.g., scatterplot).

4.4.2 Interaction-based Hints

Since the interactive widgets enable users to correct errors, the system provides hints as feedback from user interactions to help them gain knowledge of query phrasing. The hints are generated based on the analysis of interaction logs and the original queries and consist of two major components, including target suggestions for local revision to the original query and valid query examples that can be used for subsequent analysis. The generation of query examples is introduced in Section 4.4.3.

For **attributes**, when a new attribute is added to the process and is not extracted before, the system considers this a *failure of inference* and suggests users explicitly mention the attribute in the query. If the attribute is extracted but used for other tasks, which means the NLI constructs an unexpected parsing dependency of sentence structure, the system suggests users rephrase the sentence and provide query examples as alternatives. When unexpected attributes are implicitly extracted and removed by users, the system infers the error of *unexpected types* and gives hints to users to avoid the corresponding words in the original query.

For **tasks**, multiple alternatives may be available when it is underspecified in the query. For instance, the NLI chooses *Mean* as the default aggregation type. If users switch to other types (e.g., *Sum*), the system considers this *unexpected type* of task and suggests adding related words (e.g., *total*) to the query for explicit specification. Filtering tasks are usually partially correct in terms of operators (e.g., *>*) and parameters (e.g., *100,000*), and users can easily modify the unexpected parts. After user correction, the system suggests supported phrases (e.g., *more* and *over*) to revise the related parts of the queries.

For **visual encodings**, users often modify the unexpected *chart types* using interactive widgets, which is intuitive and efficient for obtaining the desired results. The system suggests an explicit specification of the desired chart type in the query when it is not the default option for the analytical intents. Since *encoding channels* cannot be specified through NL queries in NL4DV, the system does not provide hints for the exchange of encoding channels.

4.4.3 Query Example Generation

After user adjustments through widgets, the system considers the final visualizations as the desired results and generates query examples to help the users in query phrasing. The query examples are generated by a template-based method, which has been shown to be effective and controllable in specifying known analysis intents [23].

Our method takes the Vega-Lite object of the final result as input, extracts the visualization process from *transform*, *mark*, and *encodings* fields, and translates them into linguistic expression. We prepare a set of query templates with placeholders based on the utterances collected in the prior NLI study [28], [30]. For example, multiple templates are available for a scatterplot with two quantitative attributes.

- Show *attribute1* and *attribute2* in scatter plot.
- What is the relationship between *attribute1* and *attribute2*?
- How are *attribute1* and *attribute2* correlated?

The query templates can be further extended by phrases (e.g., *for Action movies*) and clauses (e.g., *whose release year is after 2002*) to specify additional filtering tasks. A generated query example is shown in Figure 4C. Such translation ensures a complete expression of the visualization results and can be easily understood by the users.

To assess the feasibility of the generated examples, we employ NL4DV as a discriminator to parse the example queries and compare their results with the current result adjusted by the users. Finally, we randomly recommend one of the valid examples to the users.

5 USAGE SCENARIO

This section presents two usage scenarios. The first scenario (Figure 5) presents how the system helps users increase their knowledge about the NLI capability. The user inputs queries that NLI cannot support well, and the system helps him revise the query formulation. The second scenario (Figure 6) shows how the system helps users increase the knowledge

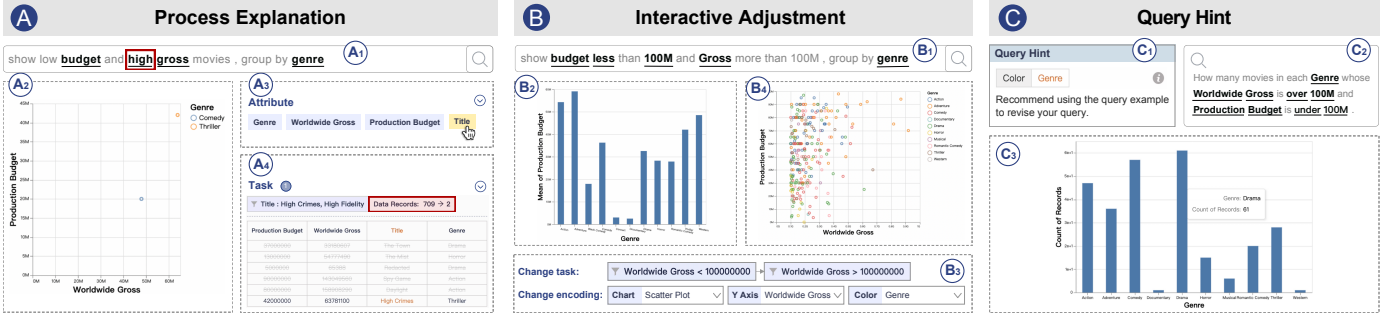


Fig. 5. The first scenario illustrates how the system helps the user increase his knowledge of NLI capability. (A) From the process explanations, the user identifies an unexpected filtering task and learns that the NLI can not understand vague modifiers. (B) He tries a new query and corrects errors through interactive adjustments. (C) He leverages the query example in the hint to formulate a new query and verify the data insight. Through the query adjustments, he learns how to better phrase the queries.

about the data to be explored. In this scenario, the user may pose queries that contain inaccurate expressions of attributes and wrong parameters of tasks. The system helps him locate and correct the unexpected results.

5.1 Scenario 1: Increase the Knowledge of NLI

In this scenario, the analyst is a novice user of NLI tool and wants to analyze the IMDb movies and gain insights to inform movie investment. He first loads the movie dataset and spends some time familiarizing himself with this dataset. Being interested in the cost-effective movie genres, he inputs the query “show low budget and high gross movies, group by genre” (Figure 5A₁). In response, the system generates a scatter plot to visualize the correlation between Production Budget, Worldwide Gross, and Genre (Figure 5A₂). However, the chart has only two points. The analyst feels confused about the result, so he turns to the Explanation panel to understand the visualization process. He finds an attribute Title with uncertainty in the process (Figure 5A₃), and it is inferred from the keyword “high” in the query (Figure 5A₁), which means that “high” is treated as a specific value of movie Title instead of the data range of Worldwide Gross. As a result, only two movies containing “high” in the Title are left (Figure 5A₄) and visualized in the resulting chart.

With the help of the explanations, the analyst realizes that the NLI can not understand vague modifiers of the data range, so he clarifies the concepts of “high” and “low” with specific data ranges: “show budget less than 100M and Gross more than 100M, group by genre” (Figure 5B₁). The system generates a bar chart to show the average Production Budget across different Genres (Figure 5B₂). Besides the unexpected chart type, the analyst also notices an incorrect filtering task “Worldwide Gross < 100,000,000” from the Explanation panel. However, the analyst fails to get the desired result through query revision. So he turns to using interactive widgets to modify the process (Figure 5B₃), including modifying the filtering tasks and switching the chart type to scatterplot. From the adjusted chart (Figure 5B₄), the analyst observes that the Drama genre seems to have the most movie items with high gross and low budget.

To verify this finding, the analyst leverages the query example on the Hint panel (Figure 5C₁) and makes simple edits to “How many movies in each Genre whose Worldwide Gross is over 100M and Production Budget is under 100M”

(Figure 5C₂). The system interprets the query and returns a bar chart (Figure 5C₃) showing that the Drama genre has the largest number of movies that meet the filter criteria. From this, the analyst also learns that such sentence structures are more likely to be correctly parsed by the NLI.

5.2 Scenario 2: Increase the Knowledge of Data

In this scenario, the analyst is not entirely familiar with the dataset. Being intrigued by the correlation between movie rating and revenue, he inputs “show the rating and box office” (Figure 6A₁) and expects a scatterplot. However, the system renders a bar chart that only presents the data distribution of Content Rating (Figure 6A₂). The analyst considers it an unexpected result and tries to fix the problems.

From the query input box (Figure 6A₁), he finds that the keyword “box office” is not identified. The analyst realizes that NLI fails to infer the attribute Worldwide Gross, so he revises the query with the correct attribute name (Figure 6B₁). Meanwhile, a hint for “Content Rating” reminds him to resolve the ambiguity of “rating” in the NL query. He chooses “IMDB Rating” through an interactive widget and receives a new scatterplot (Figure 6B₂). From the chart, the analyst does not find the apparent correlation between gross and ratings, so he wants to dive into some movies from creative types of interests and released in later years. He modifies the query to “show the rating and worldwide gross of super hero movies released after 2009” (Figure 6C₁).

The system has learned user preference for “rating” and takes IMDB Rating by default (Figure 6C₂). However, the chart is empty with no data point. Fortunately, the Hint panel helps the analyst to address the problem: “No records satisfy the filter criteria.” The Task aspect of the Explanation panel indicates that only 20 Super Hero movies are in the dataset, and none of them are released after 2009. Therefore, the analyst deletes the filtering task for “Release Year” and gets the non-empty scatterplot in return (Figure 6C₃). Based on the chart analysis, the analyst finds that the movie with the highest Worldwide Gross (i.e., The Dark Knight) gets the highest IMDB rating. By contrast, in the Spider-Man series, the movies with higher gross have lower ratings.

6 EVALUATION

We conduct a user study to evaluate the system’s effectiveness and usability in explaining and diagnosing NLI-based




Fig. 6. The second scenario presents how XNLI helps the user increase the knowledge about the dataset to be explored.

data analysis. We also examine the impact of the explanations on the analysis process, including whether the explanations increase the knowledge about the NLI, disrupt the analysis process, and increase the trust in NLI-based data analysis.

6.1 Participants and Setup

We recruited 12 participants (P1-P12, three females, nine males, aged 24-32) from a technology company’s business intelligence department and the national laboratory of a local university. Four participants (P1-P4) were senior data analysts with more than five years of working experience and proficiency in programming languages (e.g., Python, SQL) and data analysis software (e.g., Tableau [6], Microsoft Power BI [7]). Four participants (P5-P8) were data analysis researchers with over four years of visualization experience and have published related research papers. The remaining four participants (P9-P12) were graduate students at a local university. They rated themselves as intermediate users or beginners when we asked them about their expertise in data analysis and NLI tools. Each participant was compensated with \$15 in cash after completing our studies.

To conduct a comparative study, we used the system in the pilot study as the baseline which is the ablated version of XNLI without the explanation components. Both systems share the same functions to support dataset selection (Figure 2A₁), data overview (Figure 2A₂), and NL-based data analysis (Figure 2B₁, B₂). The baseline does not provide users with the Explanation panel about the NLI process and the query input box (Figure 2B₁) does not provide visual cues to highlight keywords in the query.

6.2 Procedure and Tasks

The whole pilot study consisted of four stages as follows and lasted for about 2 hours. We introduce two types of tasks (i.e., target replication and open-ended exploration) for different analysis scenarios, which have been widely used in prior works [2], [3], [4], [45], [54]. Target replication leverages the visualization targets as the ground truth of user expectations and helps us to quantify the helpfulness of our system in process understanding and problem solving. Open-ended exploration mimics the real-world analysis process in which users perform multi-step exploration to find data insights. It helps evaluate the usability of our system and the impact of explanations on the analysis process.

Introduction and Training (30 min). We first briefed participants about our study. We collected their demographic

information (e.g., name, sex, and age) and sought their consent to record the system interactions and the conversation for further analysis. Then, we provided a tutorial for each system with several analysis cases on an example dataset. After the demonstration, participants could familiarize themselves with the systems using the example dataset.

Target Replication (30 min). Participants were required to use our system and baseline to reproduce 12 visualization targets for the 2 given datasets (3 tasks for each system and dataset). In addition to the *IMDb movie* dataset in the pilot study, we also included a *car* dataset that contained 304 data items with eight attributes (e.g., Cylinders, Displacement, and Origin). We used different tasks for two systems to mitigate the learning effect on the tasks. The difficulty of the tasks in each system was balanced according to the number of attributes involved. The order of the systems and datasets was counterbalanced.

Open-Ended Exploration (30 min). Participants were required to use our system and the baseline to explore the *housing* dataset to provide advice for the clients to purchase homes. The *housing* dataset contains 1,460 data items with 15 attributes (e.g., Lot area, Home type, Year, and Price). The advice should be supported by the data insights found in the exploration. Participants were encouraged to explore the data following a realistic and natural workflow. We did not restrict the number of attempts or the details of exploration.

Semi-structured Interview (30 min). After the exploration tasks, the participants needed to complete a questionnaire with 10 questions. We provided five-point Likert-scale questions that required participants to express their opinions ranging from “strongly disagree” (1) to “strongly agree” (5). The questions are listed in Figure 7A. Then, we conducted a semi-structured interview to collect user feedback regarding their exploration experience using NLIs.

6.3 Results and Analysis

Considering the quantitative analysis of the target replication tasks, the user behaviors in response to system results, the ratings of the questionnaire, and the user feedback from the semi-structured interviews, we discuss the effectiveness and usability of our system, as well as the impact of the explanations on the NLI-based analysis process.

6.3.1 Task Completion Quality

For the target replication tasks, we conducted a quantitative comparison between our system and the baseline in terms of task accuracy and completion time. In Figure 7B₁, we

A Questionnaire and Score Distribution

Question	Average	Score Distribution
Q1 The explanations help me assess the correctness of visualization results.	4.58	1 3 8
Q2 The interactive widgets help me correct the unexpected results.	4.50	6 6
Q3 The query hints helps me phrase and revise the query statements.	4.08	3 5 4
Q4 The overall workflow of the system is reasonable.	4.83	2 10
Q5 It is easy to learn the system.	4.67	1 2 9
Q6 It is easy to use the system.	4.75	3 9
Q7 I will use this system again.	4.92	1 11
Q8 The system's explanations help me understand the NLI process.	4.50	6 6
Q9 The system's explanations do not disrupt the NLI-based data analysis flow.	4.75	3 9
Q10 The system's explanations increase my trust in NLI-based data analysis.	4.58	1 2 9

■ 1-Strongly Disagree
 ■ 2-Disagree
 ■ 3-Neutral
 ■ 4-Agree
 ■ 5-Strongly Agree

B Task Completion Quality

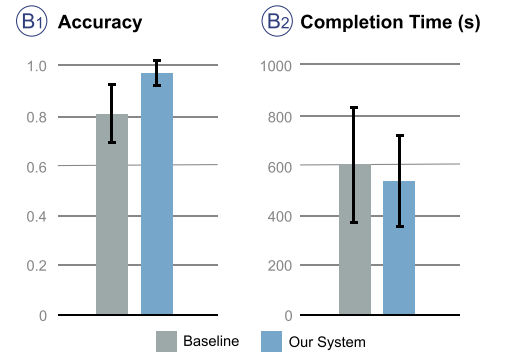


Fig. 7. The results of user interviews and task completion analysis. (A) The questions in the questionnaire, together with the average values and the distribution of the ratings. (B) The task completion quality in terms of accuracy and completion time.

calculated the average task **accuracy** (the proportion of the completed tasks) at the individual level with 95% confidence interval. On average, our system ($\mu = 0.96$, $\sigma = 0.07$) helped users to achieve a higher accuracy compared with the baseline system ($\mu = 0.81$, $\sigma = 0.11$). The Wilcoxon signed-rank test shows a significant difference ($p = 0.0066$) in average accuracy between the two systems. For task **completion time** in Figure 7B₂, it generally took less time for participants to complete the tasks using our system ($\mu = 540.08$, $\sigma = 174.42$) than using the baseline ($\mu = 599.75$, $\sigma = 222.21$). However, no significant difference ($p = 0.2036$) was observed between the two systems. This result indicates that the explanations of our system did not hinder task efficiency.

6.3.2 User Behaviors

Together with the participants, we reviewed the recordings of system operations and oral comments during the tasks to analyze the use of explanations. We identified three kinds of user behaviors in response to system results.

- **Use (view only):** Participants viewed the process explanations only and revised the queries on their own if they identified problems.
- **Use (view and adjust):** Participants viewed the process explanations and adjusted the process with the help of interactive widgets or query hints.
- **Not use:** Participants neither viewed the process explanations nor adjusted them.

As shown in Table 2, the participants preferred using the explanation components to diagnose the process and fix the problems. The process explanations were frequently used ($80+45=125$ out of 149) to confirm the expected system results, and sometimes (45 out of 149) even supported interactive adjustments for follow-up exploration. When encountering unexpected system results, participants relied on the process explanations to locate the problems ($37+42=79$ out of 95) and adopted different correction strategies depending on the complexity of the problems. Specifically, they preferred using interactive widgets if the problem required a minor change (e.g., only need to modify the parameter of a filtering task), while they tended to revise the query on their own when the problem required multi-step modification operations. However, they might switch to interactive widgets after several failed attempts at query revision. Query hints were very helpful in case of empty results (caused by empty data

TABLE 2
Distribution of participants' use of explanation components in response to system's parsing results.

Parsing results	Use (only view)	Use (view and adjust)	Not use	Total
Expected	80	45	24	149
Unexpected	37	42	16	95
Total	117	87	40	244

records or unsupported chart types). Sometimes participants were confident in the system results or tried to locate and correct the problems on their own, so they did not use the explanation components (16 out of 95).

6.3.3 Effectiveness

According to the questionnaire and interviews, most participants (8/12 strongly agree, 3/12 agree) agreed that the **process explanations** helped them assess the correctness of the visualization results (Q1). P5 said "There is no doubt that the explanations can help me to identify problems in the process, especially for some tasks 'behind' the visualizations." P10 mentioned that the data provenance is intuitive and helps her understand how the tasks operate on the data, such as the distribution and trend tasks. Some participants appreciated the visualization of the query-visualization connections and the inference uncertainty for problem location. P7 commented "Once I get the explanations of the process, I always check the uncertainty in the process first and compare it with the highlighted content in the query. It really helps me improve the efficiency of problem checking."

Participants also provided positive feedback (6/12 strongly agree, 6/12 agree) to the **interactive widgets**, praising its ability to support flexible adjustments to get the desired results (Q2), even some that could not be obtained through natural language queries. P4 mentioned "I particularly like using interactive widgets to change the visual encoding of results, as I find it impossible to specify a specific visual encoding channel for an attribute through natural language." P3 suggested extending interactive widgets to support more complex scenarios, "It would be better to support the set operations for multiple tasks, such as union ('or' operation)."

Most participants thought that the **query hints** helped them phrase and revise the queries (Q3), especially helping users "avoid unreasonable errors" (P6) and "improve the efficiency of query modification" (P5). Participants were eager to use

the query examples for the following exploration. P3 and P5 mentioned that they preferred leveraging the suggested query examples for new analysis rather than entering the queries from scratch, which improved the input efficiency and accuracy. P9 said *"I can just modify the content of the query example while keeping the syntactic structure. Such queries are more likely to be parsed correctly."* We also discussed with the participants the quality of the query examples. Participants thought that the provided query examples were *"easy to understand"* (P11) and *"easy to follow"* (P7). Regarding the naturalness of the query examples, P8 commented *"The query examples are clear in expression, but sometimes a little wordy."*

The **overall workflow** of the system is considered reasonable (10/12 strongly agree, 2/12 agree) according to user feedback (Q4). P8 specified *"The process explanations help me to identify unexpected results, so I can fix simple problems by myself and solve more complex ones with the help of widgets and hints. The system workflow is very fluent."* The correctness of data analysis is highly valued in the industry. P1 pointed out *"The workflow improves the rigor of NLI analysis and will facilitate the application of NLI analysis in industrial analysis scenarios."*

6.3.4 Usability

The system usability centered on user experience regarding learning cost and workload. Most participants (except P9) agreed that the system is **easy to learn** (Q5) and *"intuitive"* (P1, P5, P10). P1 added that *"I feel familiar with the system because it reminds me of Jupyter Notebook, where I can debug the codes block by block based on the intermediate results and system logs, such as errors and warnings."* However, P9 (with a rating of 3) posed his concerns about the learning cost because *"the system introduces a series of features that took me some time to fully understand."* Nevertheless, all participants (9/12 strongly agree, 3/12 agree) agreed that the system is **easy to use** (Q6) and most of them mentioned that the system design is not complex. P9 reported that *"the system is very user-friendly, logical, and beautifully designed. I feel very comfortable using this system."* Finally, all participants (11/12 strongly agree, 1/12 agree) are willing to **use our system again** in the future.

6.3.5 Impact on Visual Analysis

We further investigate the impact of explanations on the NLI-based visual analysis in terms of understanding, fluidity, and trust. According to participants' rating (Q8) and feedback, the explanations help **understand the NLI process** (6/12 strongly agree, 6/12 agree) in two main ways. First, our system improves the users' knowledge about the NLI pipeline. P10 stated *"I realized that (inference of) attributes, tasks, and visual encodings are important steps for NLI to process queries, so I organize my query statements according to these perspectives."* Second, our system reveals the interpretation capability of NLI, including the supported tasks (e.g., correlation) and the corresponding triggered keywords (e.g., *correlate, relationship*). P4 commented *"I think the NLI is not perfect, and the explanations helped me to identify its strengths and weaknesses, so I can use it better for my analysis."*

As an introduced feature, the explanations (Q9) **do not disrupt NLI-based data analysis** (9/12 strongly agree, 3/12 agree). The results analysis of the target replication task shows that our system improves the efficiency of task completion compared to the baseline system. In the interview,

P7 mentioned that *"the functions of explanations are integrated into the same area and can be turned off, so I can use only the NLI part of the system when I am confident in the analysis."* In contrast, the explanations improved the efficiency of the analysis. P3 reported that *"it would have taken me a lot of time to understand the query output without explanations. The system used a storytelling-like approach to explain the visualization output, which significantly accelerates my analysis progress."* In addition, he mentioned that the explanations also give users heuristic guidance for subsequent exploration.

Finally, most participants (except P4) agreed that the explanations **increased the trust** in NLI-based data analysis (Q10). The explanations *"make the visualization process transparent"* (P6) and make the participants *"feel confident after the confirmation of the intermediate process"* (P5). However, P4 raises the concern of trusting NLI (without explanations) for data analysis. He pointed out that he remains suspicious of the current NLI-based analysis tools because they still require human understanding and intervention.

7 DISCUSSION

In this section, we first summarize the lessons that we learned from the user study, including providing a preview of query results for NLI debugging and using explanations for NLI to promote system discoverability. Then, we discuss the generalizability as well as the limitations and future work.

7.1 Lessons Learned

Provide a preview of query results for NLI debugging. We observed some interesting patterns during the evaluation. Some users could not determine the correctness of the results in the baseline system, so they tried to revise the query to test the system response. For example, one user changed the range of filter criteria in the query to test whether the filtering task worked in the process. However, he was still not sure whether the filtering task was correct. When receiving unexpected results, some users had to simplify the query sentences by removing some words to locate the problems in the queries. These debugging strategies are similar to "black-box testing" and cost a considerable amount of time during the analysis process. This finding also motivates us to further study how to help users quickly preview and compare the results of different revision strategies.

Use explanations for NLI to promote system discoverability. The explanations increase the NLI discoverability (i.e., the users' awareness and understanding of system-supported commands). In contrast to much prior research [23], [24], [55] that recommends supported utterances for next-step data analysis, the explanations reveal the NLI process, enabling users to explore more alternative expressions that adapt to users' preferences and can be understood by the system. For example, one user in our evaluation was accustomed to using parentheses or commas in specifying multi-step data operations (e.g., filters). From the system explanations, he verified that such syntax can be well parsed by NLI and preferred to use this syntax in the subsequent queries.

7.2 Generalizability

Our work is based on NL4DV, a typical toolkit following the general pipeline of NLI [1], [4], and the major components of

our system (i.e., Provenance Generator, interactive widgets, and Hint Generator) are independent of the specific details of NLI algorithms and can be directly applied to the existing NLIs. The Provenance Generator reveals the visualization process based on the Vega-Lite specification, which is widely used by many NLIs [5], [23], [24], [28], [29], [30], [32], [52]. The Hint Generator integrates a set of well-known rules [4], [5], [12] to detect problems in the query and can be extended to support more scenarios in different application domains. Other features, such as the highlighting of the algorithm uncertainty and the connection between visualizations and NL queries are based on the query interpretation in the NLI process. They can also be used by other NLIs that follow the typical NLI pipeline. Therefore, our work is general and can be easily used for other NLIs.

7.3 Limitations and Future Work

Improve the query examples. The query examples in query hints are generated based on the templates collected in the prior NLI study [28], [30]. Despite its effectiveness and controllability, the approach may generate long structured sentences for complex cases with multiple tasks which do not seem as natural as those produced by humans. One of our future works is improving the quality of the generated query examples. Recently, generative language models [56] trained on a large corpus have been proposed to facilitate various downstream tasks [57], [58]. They open up the possibility of enhancing the diversity and naturalness of linguistic expression while maintaining the accurate representation of the users' intentions.

Improve the system evaluation. We show the effectiveness of our system through a comparative study with a baseline NLI that only maintains the basic functions without explanation components. However, it would be better to design and implement a stronger baseline (e.g., NLI with simple ambiguity widgets [2], [3], [4], [6]) for comparative study. And the comparison can further help reflect the values of the systematic explanations for the NLI process.

Provide explanations for follow-up queries. XNLI provides explanations for one-off queries that specify the entire visualization process in a single sentence. Several studies [2], [3], [21], [22], [59] have supported conversational data analysis, enabling users to pose follow-up queries (e.g., "Show only movies released after 2008") in the context of prior queries (e.g., "show the distribution of production budget in a bar chart"). It is a flexible means of expressing analytical intents but also introduces a larger risk of incorrect query interpretation [3]. The system should provide explanations for follow-up queries to present how the NLIs leverage the context information to interpret the follow-up queries.

Learn from users. The explanations accelerate the user's learning of NLI by revealing the visualization process and increasing the discoverability of the system-supported queries. The system provides interactive widgets to support further specification of user expectations, which also provide great opportunities for NLI to learn from users. We have enabled the system to learn the user preference to resolve the ambiguities in the following queries. A beneficial task is exploring how to support user explanations for their queries and then translate them into personalized features to better interpret the queries and generate desired responses.

8 CONCLUSION

This study focuses on explaining and diagnosing NLI-based visual analysis. We first identify and summarize the problems of unexpected visualization results during the NLI-based visual analysis that require process explanations. Then we present XNLI, an explainable NLI system to illustrate the generation process of visualization results. Specifically, it reveals the process from the aspects of the attributes, tasks, and visual encodings and generates sample data provenance based on the Vega-Lite specification. To help users correct errors, the system provides interactive widgets to adjust the process and generates hints to refine the queries. We introduce two usage scenarios and a user study to demonstrate the effectiveness and usability of our system in improving the task accuracy and efficiency of NLI-based data analysis.

In the future, we plan to enhance the explainability of XNLI, including revealing more details of query interpretation, supporting a higher level of control and personalization of the NLI process, and providing explanations for follow-up queries in conversational visual analysis. We also aim to improve the quality of the generated query examples to better assist users in query revision. Moreover, we plan to conduct longitudinal studies to further investigate the impact of systematic explanations in building trust in NLI results and facilitating human-machine collaboration.

ACKNOWLEDGMENT

We would like to thank Wei Zhang and Jiacheng Pan for their kind help. We also thank the anonymous reviewers for their insightful comments. This paper is supported by National Natural Science Foundation of China (62132017) and Fundamental Research Funds for the Central Universities (226-2022-00235).

REFERENCES

- [1] L. Shen, E. Shen, Y. Luo, X. Yang, X. Hu, X. Zhang, Z. Tai, and J. Wang, "Towards natural language interfaces for data visualization: A survey," *arXiv preprint arXiv:2109.03506*, 2021.
- [2] V. Setlur, S. E. Battersby, M. Tory, R. Gossweiler, and A. X. Chang, "Eviza: A natural language interface for visual analysis," in *Proceedings of ACM UIST*, 2016, pp. 365–377.
- [3] E. Hoque, V. Setlur, M. Tory, and I. Dykeman, "Applying pragmatics principles for interaction with visual analytics," *IEEE TVCG*, vol. 24, no. 1, pp. 309–318, 2017.
- [4] T. Gao, M. Dontcheva, E. Adar, Z. Liu, and K. G. Karahalios, "Datatone: Managing ambiguity in natural language interfaces for data visualization," in *Proceedings of ACM UIST*, 2015, pp. 489–500.
- [5] A. Narechania, A. Srinivasan, and J. Stasko, "NL4DV: A toolkit for generating analytic specifications for data visualization from natural language queries," *IEEE TVCG*, vol. 27, no. 2, pp. 369–379, 2020.
- [6] M. Tory and V. Setlur, "Do what I mean, not what I say! design considerations for supporting intent and context in analytical conversation," in *Proceedings of IEEE VAST*. IEEE, 2019, pp. 93–103.
- [7] "Microsoft power bi," <https://powerbi.microsoft.com/en-us/>.
- [8] V. Setlur, M. Tory, and A. Djalali, "Inferencing underspecified natural language utterances in visual analysis," in *Proceedings of ACM IUI*, 2019, pp. 40–51.
- [9] M. Card, *Readings in information visualization: using vision to think*. Morgan Kaufmann, 1999.
- [10] A. Srinivasan, B. Lee, and J. Stasko, "Interweaving multimodal interaction with flexible unit visualizations for data exploration," *IEEE TVCG*, vol. 27, no. 8, pp. 3519–3533, 2020.
- [11] A. Srinivasan and J. Stasko, "Natural language interfaces for data analysis with visualization: Considering what has and could be asked," in *Proceedings of Eurographics: Short papers*, 2017, pp. 55–59.

- [12] B. Yu and C. T. Silva, "Flowsense: A natural language interface for visual data exploration within a dataflow system," *IEEE TVCG*, vol. 26, no. 1, pp. 1–11, 2019.
- [13] A. Srinivasan and J. Stasko, "Orko: Facilitating multimodal interaction for visual exploration and analysis of networks," *IEEE TVCG*, vol. 24, no. 1, pp. 511–521, 2017.
- [14] V. Setlur, E. Hoque, D. H. Kim, and A. X. Chang, "Sneak Pique: Exploring autocompletion as a data discovery scaffold for supporting visual analysis," in *Proceedings of ACM UIST*, 2020, pp. 966–978.
- [15] V. Setlur, S. Battersby, and T. Wong, "GeoSneakPique: Visual autocompletion for geospatial queries," in *Proceedings of IEEE VIS*, 2021, pp. 166–170.
- [16] Z. Chen and H. Xia, "CrossData: Leveraging text-data connections for authoring data documents," in *Proceedings of ACM CHI*, 2022, pp. 1–15.
- [17] Z. Chen, Q. Yang, X. Xie, J. Beyer, H. Xia, Y. Wu, and H. Pfister, "Sporthesia: Augmenting sports videos using natural language," *IEEE TVCG*, 2022.
- [18] J. Wu, D. Liu, Z. Guo, and Y. Wu, "Rasipam: Interactive pattern mining of multivariate event sequences in racket sports," *IEEE TVCG*, pp. 1–11, 2022.
- [19] R. Chen, X. Shu, J. Chen, D. Weng, J. Tang, S. Fu, and Y. Wu, "Nebula: A coordinating grammar of graphics," *IEEE TVCG*, vol. 28, no. 12, pp. 4127–4140, 2022.
- [20] K. Cox, R. E. Grinter, S. L. Hibino, L. J. Jagadeesan, and D. Mantilla, "A multi-modal natural language interface to an information visualization environment," *International Journal of Speech Technology*, vol. 4, no. 3, pp. 297–314, 2001.
- [21] A. Kumar, J. Aurisano, B. Di Eugenio, A. Johnson, A. Gonzalez, and J. Leigh, "Towards a dialogue system that supports rich visualizations of data," in *Proceedings of the Special Interest Group on Discourse and Dialogue*, 2016, pp. 304–309.
- [22] E. Fast, B. Chen, J. Mendelsohn, J. Bassen, and M. S. Bernstein, "Iris: A conversational agent for complex tasks," in *Proceedings of ACM CHI*, 2018, pp. 1–12.
- [23] A. Srinivasan and V. Setlur, "Snowy: Recommending utterances for conversational visual analysis," in *Proceedings of ACM UIST*, 2021, pp. 864–880.
- [24] X. Wang, F. Cheng, Y. Wang, K. Xu, J. Long, H. Lu, and H. Qu, "Interactive data analysis with next-step natural language query recommendation," *arXiv preprint arXiv:2201.04868*, 2022.
- [25] J. Huang, Y. Xi, J. Hu, and J. Tao, "Flownl: Asking the flow data in natural languages," *IEEE TVCG*, 2022.
- [26] Y. Wang, Z. Hou, L. Shen, T. Wu, J. Wang, H. Huang, H. Zhang, and D. Zhang, "Towards natural language-based visualization authoring," *IEEE TVCG*, 2022.
- [27] A. Gaba, V. Setlur, A. Srinivasan, J. Hoffswell, and C. Xiong, "Comparison conundrum and the chamber of visualizations: An exploration of how language influences visual design," *IEEE TVCG*, 2022.
- [28] S. Fu, K. Xiong, X. Ge, S. Tang, W. Chen, and Y. Wu, "Quda: Natural language queries for visual data analytics," *arXiv preprint arXiv:2005.03257*, 2020.
- [29] Y. Luo, N. Tang, G. Li, C. Chai, W. Li, and X. Qin, "Synthesizing natural language to visualization (nl2vis) benchmarks from nl2sql benchmarks," in *Proceedings of International Conference on Management of Data*, 2021, pp. 1235–1247.
- [30] A. Srinivasan, N. Nyapathy, B. Lee, S. M. Drucker, and J. Stasko, "Collecting and characterizing natural language utterances for specifying data visualizations," in *Proceedings of ACM CHI*, 2021, pp. 1–10.
- [31] Y. Luo, J. Tang, and G. Li, "nvBench: A large-scale synthesized dataset for cross-domain natural language to visualization task," *arXiv preprint arXiv:2112.12926*, 2021.
- [32] Y. Luo, N. Tang, G. Li, J. Tang, C. Chai, and X. Qin, "Natural language to visualization by neural machine translation," *IEEE TVCG*, vol. 28, no. 1, pp. 217–226, 2021.
- [33] C. Liu, Y. Han, R. Jiang, and X. Yuan, "Advisor: Automatic visualization answer for natural-language question on tabular data," in *Proceedings of Pacific Visualization Symposium*, 2021, pp. 11–20.
- [34] E. D. Ragan, A. Endert, J. Sanyal, and J. Chen, "Characterizing provenance in visualization and data analysis: an organizational framework of provenance types and purposes," *IEEE TVCG*, vol. 22, no. 1, pp. 31–40, 2015.
- [35] J. Heer, J. Mackinlay, C. Stolte, and M. Agrawala, "Graphical histories for visualization: Supporting analysis, communication, and evaluation," *IEEE TVCG*, vol. 14, no. 6, pp. 1189–1196, 2008.
- [36] C. Bors, T. Gschwandtner, and S. Miksch, "Capturing and visualizing provenance from data wrangling," *IEEE CG&A*, vol. 39, no. 6, pp. 61–75, 2019.
- [37] S. Kasica, C. Berret, and T. Munzner, "Table scraps: an actionable framework for multi-table data wrangling from an artifact study of computational journalism," *IEEE TVCG*, vol. 27, no. 2, pp. 957–966, 2020.
- [38] K. Xiong, S. Fu, G. Ding, Z. Luo, R. Yu, W. Chen, H. Bao, and Y. Wu, "Visualizing the scripts of data wrangling with somnus," *IEEE TVCG*, no. 01, pp. 1–1, 2022.
- [39] Y. B. Shrinivasan and J. J. Van Wijk, "Supporting the analytical reasoning process in information visualization," in *Proceedings of ACM CHI*, 2008, pp. 1237–1246.
- [40] X. Pu, S. Kross, J. M. Hofman, and D. G. Goldstein, "Datamations: Animated explanations of data analysis pipelines," in *Proceedings of ACM CHI*, 2021, pp. 1–14.
- [41] A. Leventidis, J. Zhang, C. Dunne, W. Gatterbauer, H. Jagadish, and M. Riedewald, "QueryVis: Logic-based diagrams help users understand complicated sql queries faster," in *Proceedings of ACM SIGMOD*, 2020, pp. 2303–2318.
- [42] J. Berant, D. Deutch, A. Globerson, T. Milo, and T. Wolfson, "Explaining queries over web tables to non-experts," in *Proceedings of IEEE ICDE*. IEEE, 2019, pp. 1570–1573.
- [43] M. Khan, L. Xu, A. Nandi, and J. M. Hellerstein, "Data tweening: incremental visualization of data transforms," *Proceedings of the VLDB Endowment*, vol. 10, no. 6, pp. 661–672, 2017.
- [44] A. Narechania, A. Fourney, B. Lee, and G. Ramos, "DIY: Assessing the correctness of natural language to sql systems," in *Proceedings of ACM IUI*, 2021, pp. 597–607.
- [45] A. Srinivasan and J. Stasko, "How to ask what to say?: Strategies for evaluating natural language interfaces for data visualization," *IEEE CG&A*, vol. 40, no. 4, pp. 96–103, 2020.
- [46] T. Wambsganss, C. Niklaus, M. Cetto, M. Söllner, S. Handschuh, and J. M. Leimeister, "AL: an adaptive learning support system for argumentation skills," in *Proceedings of ACM CHI*, 2020, pp. 1–14.
- [47] M. Xia, Q. Zhu, X. Wang, F. Nei, H. Qu, and X. Ma, "Persua: A visual interactive system to enhance the persuasiveness of arguments in online discussion," *arXiv preprint arXiv:2204.07741*, 2022.
- [48] X. Wang, Y. Ming, T. Wu, H. Zeng, Y. Wang, and H. Qu, "DeHumor: Visual analytics for decomposing humor," *IEEE TVCG*, 2021.
- [49] K. Dhamdhere, K. S. McCurley, R. Nahmias, M. Sundararajan, and Q. Yan, "Analyza: Exploring data with conversation," in *Proceedings of the ACM IUI*, 2017, pp. 493–504.
- [50] A. Saktheeswaran, A. Srinivasan, and J. Stasko, "Touch? speech? or touch and speech? investigating multimodal interaction for visual network exploration and analysis," *IEEE TVCG*, vol. 26, no. 6, pp. 2168–2179, 2020.
- [51] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer, "Vega-lite: A grammar of interactive graphics," *IEEE TVCG*, vol. 23, no. 1, pp. 341–350, 2016.
- [52] M. R. Islam, I. Razzak, X. Wang, P. Tilocca, and G. Xu, "Natural language interactions enhanced by data visualization to explore insurance claims and manage risk," *Annals of Operations Research*, pp. 1–19, 2022.
- [53] V. I. Levenshtein et al., "Binary codes capable of correcting deletions, insertions, and reversals," in *Proceedings of Soviet physics doklady*, vol. 10, no. 8. Soviet Union, 1966, pp. 707–710.
- [54] A. Srinivasan, B. Lee, N. Henry Riche, S. M. Drucker, and K. Hinckley, "InChorus: Designing consistent multimodal interactions for data visualization on tablet devices," in *Proceedings of ACM CHI*, 2020, pp. 1–13.
- [55] A. Srinivasan, M. Dontcheva, E. Adar, and S. Walker, "Discovering natural language commands in multimodal interfaces," in *Proceedings of ACM IUI*, 2019, pp. 661–672.
- [56] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell et al., "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [57] T. Wu, M. Terry, and C. J. Cai, "Ai chains: Transparent and controllable human-ai interaction by chaining large language model prompts," in *ACM CHI*, 2022, pp. 1–22.
- [58] E. Jiang, E. Toh, A. Molina, K. Olson, C. Kayacik, A. Donsbach, C. J. Cai, and M. Terry, "Discovering the syntax and strategies of natural language programming with generative language models," in *ACM CHI*, 2022, pp. 1–19.

- [59] R. Mitra, A. Narechania, A. Endert, and J. Stasko, "Facilitating conversational interaction in natural language interfaces for visualization," *arXiv preprint arXiv:2207.00189*, 2022.



Yingchaojie Feng is currently a Ph.D. candidate in the College of Computer Science and Technology, the Zhejiang University, China. He received the B.E. degree in software engineering from the Zhejiang University of Technology, China in 2020. His research interests include natural language processing and visual analysis.



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 human-computer interaction (HCI), data visualization, natural language processing (NLP), and multimodal analysis. For more details, please refer to <https://andy-xingbowang.com/>.



Bo Pan received the BS degree in Electrical and Computer Engineering from the University of Illinois Urbana-Champaign and Zhejiang University in 2022. He is currently working towards the Ph.D. degree at the State Key Lab of CAD&CG, Zhejiang University. His research interests include visualization and deep learning.



Kam Kwai Wong is a Ph.D. candidate in the Department of Computer Science and Engineering at the Hong Kong University of Science and Technology (HKUST). He received his B.E. in HKUST. His main research interests are in data visualization, visual analytics and data mining. For more details, please refer to <https://jasonwong.vision>.



Yi Ren received a BS degree in geographic information science from Southeast University in China in 2016. He is currently working towards the master's degree at the College of Software Engineering, the Zhejiang University. His research interests include information visualization and visual analytics



Shi Liu received his B.S. in software engineering from the Central South University, China in 2021. He is currently a Master student in the School of Software Technology at the Zhejiang University. His research interests include information visualization, visual analysis and human-computer interaction.



Zihan Yan received the BS degree in computer science department from the Zhejiang University, China, in 2015. She is currently working toward the master's degree in MIT Media Lab, United States. Her research interests include human-computer interaction (HCI), multimodal analysis and deep learning. For more details, please refer to <https://yzihan.github.io/>.



Yuxin Ma is a tenure-track assistant professor in the Department of Computer Science and Engineering, Southern University of Science and Technology (SUSTech), China. He received his B.Eng. and Ph.D. degrees from Zhejiang University. Before joining SUSTech, he worked as a Postdoctoral Research Associate in VADER Lab, CIDSE, Arizona State University. His primary research interests are in the areas of visualization and visual analytics, focusing on explainable AI, high-dimensional data, and spatiotemporal data.



Huamin Qu is a chair 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.



Wei Chen is a professor in the State Key Lab of CAD&CG at Zhejiang University. His current research interests include visualization and visual analytics. He has published more than 80 IEEE/ACM Transactions and IEEE VIS papers. He actively served in many leading conferences and journals, like IEEE PacificVIS steering committee, ChinaVIS steering committee, paper co-chairs of IEEE VIS, IEEE PacificVIS, IEEE Lдав and ACM SIGGRAPH Asia VisSym. He is an associate editor of IEEE TVCG, IEEE TBG, ACM TIST, IEEE T-SMC-S, IEEE TIV, IEEE CG&A, FCS, and JOV. More information can be found at: <http://www.cad.zju.edu.cn/home/chenwei>.