

Towards Natural Language-Based Visualization Authoring

Yun Wang, Zhitao Hou, Leixian Shen, Tongshuang Wu, Jiaqi Wang,
He Huang, Haidong Zhang, and Dongmei Zhang

Abstract—A key challenge to visualization authoring is the process of getting familiar with the complex user interfaces of authoring tools. Natural Language Interface (NLI) presents promising benefits due to its learnability and usability. However, supporting NLIs for authoring tools requires expertise in natural language processing, while existing NLIs are mostly designed for visual analytic workflow. In this paper, we propose an authoring-oriented NLI pipeline by introducing a structured representation of users' visualization editing intents, called *editing actions*, based on a formative study and an extensive survey on visualization construction tools. The editing actions are executable, and thus decouple natural language interpretation and visualization applications as an intermediate layer. We implement a deep learning-based NL interpreter to translate NL utterances into editing actions. The interpreter is reusable and extensible across authoring tools. The authoring tools only need to map the editing actions into tool-specific operations. To illustrate the usages of the NL interpreter, we implement an Excel chart editor and a proof-of-concept authoring tool, VisTalk. We conduct a user study with VisTalk to understand the usage patterns of NL-based authoring systems. Finally, we discuss observations on how users author charts with natural language, as well as implications for future research.

Index Terms—Visualization authoring, Natural language interface, Natural language understanding

1 INTRODUCTION

Modern visualization authoring systems have emerged to enable creation of expressive visualizations. Nevertheless, they involve complicated GUIs and incur a steep learning curve. In recent years, as a complementary input modality to traditional WIMP interaction, Natural Language Interfaces (NLI) are adopted to lower the barrier of using advanced visualization tools [53]. In contrast to WIMP interfaces, which require complex menu items and mouse interactions, natural language-based systems require less prior knowledge of user interfaces, and users are not restricted to the locations of menus and buttons to author visualizations [20, 51].

While there has been active research into natural language interfaces for visualization systems [14, 17, 51, 61, 72], these systems are primarily designed for analyzing and exploring data. As shown in Figure 1(a), an analysis-oriented NLI parses NL queries (e.g., “find the relationship between player goals and salaries across player foot.”) into analytic tasks and data attributes, which are then translated into visualization specifications according to visual design constraints [39, 53]. These specifications may meet the intended analysis [9, 54], but they may not necessarily be the most *preferred* ones. In fact, users typically need to change the underlying data, specify visual encodings, and adjust visual presentations like axes, legends, marks, and layouts (e.g., “move the legend to the right of the chart”, “set mark to woman icon”, and “change color to pink”). Various modern visualization tools [12, 49] and authoring systems [24, 32, 44, 47, 66, 70] have recognized the importance of rich and flexible data binding and visual configurations; however, analysis-oriented NLIs do not fully support these diverse editing intents.

In this paper, we aim to lower the barrier to supporting NLI in visualization authoring tools. We design a pipeline (Figure 1(b)) that decouples the natural language understanding (with a *natural language interpreter*) and visualization editing command execution (by a *visualization application*). At the core of the pipeline is a set of *editing actions*. These actions are machine-executable commands for modeling

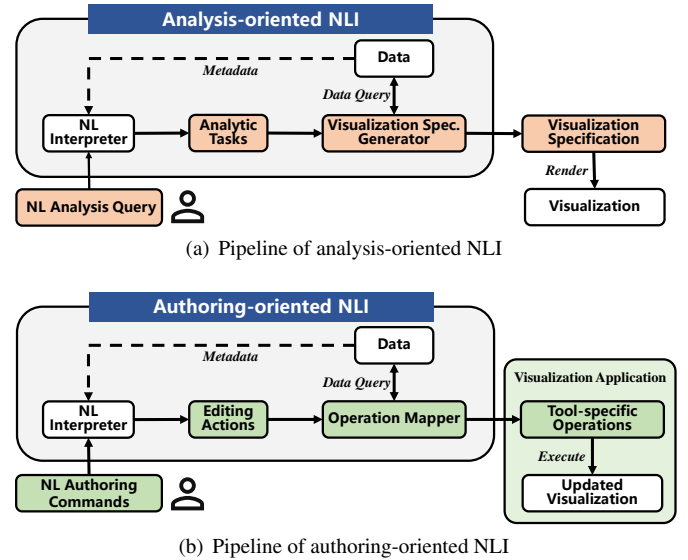


Figure 1. Analysis-oriented NLIs generate one-shot visualizations to satisfy users' analytic tasks, covering subsets of editing operations for visualization authoring. In our authoring-oriented NLI pipeline, we model users' editing intents as *editing actions*, which decouple natural language interpreter and visualization applications as an intermediate layer.

the aforementioned visualization editing intents. They bridge users and visualization applications: the natural language interpreter parses users' utterances into a sequence of such editing actions, and the actions are mapped into tool-specific operations. The visualization applications can then adapt and execute the operations to update the visualization.

We make our pipeline realistic by designing two primary building blocks. First, the formalization of *editing actions*. The role of editing actions in our authoring-oriented NLI pipeline is obvious; however, there is still a lack of comprehensive guidance on how to model visualization editing intents. To fill in the gap, we conduct a formative study and literature surveys on visualization construction tools and explicate editing actions as mappings between a series of well-defined *editing operations*, their *target objects*, and the corresponding *parameters*. Second, a *multi-stage NL interpreter*, for parsing users' NL queries into editing actions. It first recognizes data entities and replace them with abstract arguments; then, it uses a deep sequence-labeling model to extract intents and entities from the abstracted utterances; finally, it synthesizes the extracted information into a sequence of editing actions. With the deep-learning model, we envision that the interpreter can be

- Y. Wang, Z. Hou, H. Huang, H. Zhang, D. Zhang are with Microsoft Research Asia (MSRA). E-mail: {wangyun, zhith}@microsoft.com.
- L. Shen is with Tsinghua University. E-mail: slx20@mails.tsinghua.edu.cn.
- T. Wu is with Carnegie Mellon University. E-mail: sherryw@cs.cmu.edu.
- J. Wang is with Oxford University. E-mail: jiaqi.wang@cx.ox.ac.uk.
- Work done during L. Shen and J. Wang's internship at MSRA.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org.
Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxx

easily extended to cover wider range of editing actions as we collect a larger corpus of editing utterances.

Both the editing actions and the NL interpreter can be reused across multiple visualization applications, eliminating the needs to re-design NL interfaces for each one. Authoring tool developers only need to build an operation mapper to map the editing actions sequence into tool-specific operations. The only step required for authoring tool developers is to build an operation mapper for mapping editing actions sequences to tool-specific operations. We enumerate key desiderata in building these mappers, and provide empirical guidance on how to best reuse the provided two building blocks.

We demonstrate the utility of our design with two example applications, *i.e.*, an Excel chart editor and a proof-of-concept authoring tool, VisTalk. We further conduct an exploratory user study, where participants complete chart reconstruction tasks with natural utterances. Our observations shed light on future work: we should consider how the interpreters can cope with users with different backgrounds, and how the applications can resolve ambiguities and make recommendations when users' initial command is not supported.

The contributions of this work are summarized as follows:

- We propose an authoring-oriented NLI pipeline by formulating users' visualization editing intents as *editing actions*, which are atomic units executable by applications.
- We develop a multi-stage natural language interpreter to parse NL utterances into a sequence of editing actions. The interpreter can be reused across visualization applications.
- We design an Excel chart editor and a proof-of-concept authoring tool, VisTalk, to demonstrate the utility of the NL interpreter. We further conduct a user study with VisTalk to understand how users can construct and edit data charts through natural language.

2 RELATED WORK

Our work draws upon prior efforts in visualization creation tools and natural language interfaces for data visualization.

2.1 Visualization Creation Tools

A variety of visual authoring approaches have been proposed to facilitate visualization creation. For example, Polaris [63], Tableau [7], and Many Eyes [4] help users conduct data encoding with visual channels in a visualization. These tools enable users to create visualizations in a short time, but they are less flexible. The resulting visualizations are usually standard visual charts. More advanced techniques, such as Lyra [47] and iVisDesigner [43], have been proposed to enable more expressive visualization designs. With these systems, users can easily specify properties of graphical marks and bind them with data. However, these systems only support the modifications of properties by adjusting a set of style parameters. Later, Data Illustrator [32] uses repetition and partition operators for multiplying marks and generating data-driven expressive visualizations. Charticulator [44] also adopts a bottom-up approach to bind data fields to vector graphics and build visualizations with an emphasis on visualization layouts. Recently, Data-Driven Guides [24] has allowed users to draw and design graphical shapes to achieve creative designs. Similarly, DataInk [70] and InfoNice [66] also bind graphical elements with data fields to facilitate the generation of creative visualizations.

Researchers have explored the design of visualization authoring systems and supported creation from different perspectives. To support visualization creation, the systems have different internal logics. With these systems, we summarize and categorize most common functions used. We further formulate the functions as editing actions, which is an abstract layer that bridges human intents and editing functions. The editing actions can be explained and executed by the functions supported by different tools.

2.2 Natural Language Interfaces for Visualization

Natural language interfaces have been adopted extensively to improve the usability of visualization systems [53]. Commercial tools like IBM Watson Analytics [4], Microsoft Power BI [5], Tableau [7], ThoughtSpot [64], and Google Spreadsheet [14] automatically translate

the natural language questions to data queries and present query results with visualizations. However, these systems limit natural language interactions to data queries and corresponding standard charts.

Extensive research has been devoted to a better experience of using natural language for visual analytics [53]. To complete the analytical tasks, they treat user input utterances as natural language queries, and translate them to logic query languages [10, 28, 74]. Researchers further study the interactive visual analytic systems that support natural language queries. Wen *et al.* [67] design a system to update the visualizations based on user queries dynamically. Articulate [11] generates a graph to select proper visualizations to answer users' natural language questions and complete analytic tasks. ADVISor [31] and ncNet [34] leverage deep learning models to translate an NL query to a visualization. Going beyond data exploration, FlowSense [72] applies natural language techniques to dataflow visualization systems. Gridbook [55] eases the difficulty of writing formulas on the spreadsheet grid by supporting formulas expressed in natural language. Vis-Annotator [26] accepts textual descriptions and outputs annotated visualizations while Kim *et al.* [23] focus on answering questions about the given visualization. Srinivasan *et al.* integrate natural language interfaces to facilitate multimodal interaction for visual exploration [46, 56–58, 62].

Researchers also try to address ambiguity problems through different means. On the one hand, researchers attempt to address ambiguities through the design of interactive systems to invite users to clarify their requests. For example, DataTone [17] introduces interactive widgets to address ambiguity problems. Eviza [51] further enhances interactions by providing graphical answers that can be directly manipulated. Iris [16] is a conversational interface that enables users to combine commands through many rounds of question-answering. Sneak pique [3] explores autocompletion to help users formulate analytical questions while Snowy [60] recommends utterances for conversational visual analysis. On the other hand, design principles are also summarized to improve the understanding of users' utterances. To enhance natural language interactions, Hoque *et al.* [20] apply pragmatics principles and propose a theoretical framework for interaction with visual analytics. Hearst *et al.* [19] conduct an empirical study to explore default visualizations for vague expressions in natural language queries. Similarly, Setlur *et al.* [52] explore inferring underspecified natural language queries and propose a systematic approach to resolve partial utterances.

Existing work enables the natural language interfaces in a case-by-case manner, and mainly focuses on visual analysis. In comparison, we formalize users' diverse visualization editing intents into *editing actions*, covering a richer set of flexible visual configurations. The editing actions bridge the gap between the NL interpreter and visualization applications. The design and implementation of the natural language interpreter can be reused across applications, alleviating the burden of re-designing NLI per application.

3 FORMATIVE STUDY

User utterances are apt to be free and variant among usage scenarios. To better understand how visualizations are created with natural language, we conduct a formative study in the form of one-hour meetings with six participants on how they express their intentions for visualization authoring. Understanding usage patterns is conducive to the design of our natural language interpreter. All the participants had previously created visualizations for presentation purposes (e.g., data reports, presentation slides) in their daily work. They have used GUI-based (e.g., Excel and Power BI) or code-based (e.g., Python) tools to create visualizations. All of them have conducted configurations to standard visual charts or created expressive visualizations such as pictographs. Two create charts for data analysis as their daily work. Three other participants have used design tools (e.g., Adobe PhotoShop and Illustrator).

3.1 Procedure

The meeting was started with a semi-structured interview. During the interview, the participants were asked to describe their process when developing visualizations, as well as the major challenges they face when they want to ask a system to do the job for them. Then, they were

asked to list example natural language sentences to describe the operations and commands they will give if an intelligent assistant can do the job for them. Then, we collected six visualizations by sampling charts from systems like Excel, Tableau and more expressive chart designs from InfoNice [66] and Charticulator [44]. We told the participants to imagine using an authoring system that would correctly understand all the commands and update the visualization step by step. We asked them to articulate a list of natural language commands that they would use to clearly describe the steps to create these visualizations. These exercises helped us understand our users’ natural way of thinking and workflow through concrete examples.

3.2 Results

All of our participants expressed interests in having a natural language-based tools. They felt a natural language-based authoring tool is useful especially when they are not familiar with the user interfaces, which are more likely to happen when they need to create bespoke visualizations and involves a large number of steps in design tools. Here we summarize the common design practices.

General Workflow. All of the participants showed similar preferences of the general workflow of creating visualizations with natural language. In general, they preferred a top-down method for basic charts such as bar chart, column chart, pie chart, and so on. More specifically, they usually started with the chart type and encoding, then followed up with more configurations. The names of the charts give them the convenience of defining data encoding and basic properties of a visualization with several simple words, such as “Show me the sales by year with a line chart.” Based on the chart generated by the system, users could do more customization. At the same time, two of the participants also tried to describe and control the data encoding by themselves. For example, “Bind sales to y axis and year to x axis.” They felt it especially useful when they were not sure of the name of the charts to create, or when they thought they were creating a bespoke chart that the system may not support by default. As the final step, participants preferred to further fine tune the visualizations by adding or removing chart components, or modifying the default formatting of charts.

Intention Description. When conducting authoring, users expected the systems to react directly to each of their natural language commands. Sometimes, they used simple sentences, for example, “change the color of the bars”, or “please make the title bold”. They might also use complex utterances with more than one intent is included. For example, the sentence “I want a chart with rounded rectangles, showing me the distribution of the data” implies users wanted a column or bar chart with rectangles replaced with rounded shapes. Besides, the complexity of users’ input may relate to users’ experience of using natural language agents and visual chart creation. In our study, the more experienced the user is, the simpler the commands are. This might be because experienced users have lower expectations for parsing capabilities of natural language agents.

Visual Property Setting. When referring to the objects, they also used data labels. For example, “make the China bar red.” Some users prefer to use some properties to refer to an object. For example, they may use “red line” or “dashed line” to refer to a red reference line on a bar chart. Participants also tried to use pronoun (e.g., it, them) to refer to the objects, especially when they wanted to specify several properties of one visual element successively. When describing visual properties, our participants felt it easy to describe general properties like “blue” for colors, “> 200” for filters, “dashed line” for formatting, etc. They might directly use common adjectives, or the category name for these properties. By contrast, it is harder to describe degree, value, extents or positions precisely. As a result, our participants preferred to use relation to describe degree, level, extent, value, position, etc, such as “on the top of A”, “larger than B”, “lower than C”, “darker than current one”, “wider than D”.

The results of our formative study show the diversity of users’ expressions when they are not given guidance on their NL input. A structured representation is required to execute users’ intents. Further, we need to formulate them into intent units that can be combined so that users’ utterances can be represented with flexibility.

4 EDITING ACTIONS

As illustrated in Figure 1(b), the core of our pipeline is an intermediate layer called *editing actions*, which are a list of structured descriptions of user intents that can be executed by visualization applications. We design editing actions to bridge users and applications: the natural language interpreter parses users’ utterances into a sequence of these actions, and the actions can then be adapted and executed by the applications. On the one hand, these editing actions are somewhat akin to the declarative languages of visualizations (e.g., Vega-Lite), in the sense that they offer an implementation-independent bridge between the NL and the visualization. On the other hand, the editing actions differ from existing declarative grammars of visualizations, because they model the *delta* of visualization (e.g., “Turn the dashed lines black and thicker”), and is a potential first step towards authoring-oriented declarative language. Below, we formalize the model of editing action representation based on the formative study and surveys of prior work, and discuss the necessary designs to address vague, under-specified, or context-dependent descriptions in user utterances.

4.1 Editing Action Representation

Various systems have covered a wide range of functions to enable the composition of both standard charts and highly customized interactive visualizations. However, there was no common design model that can let us consistently describe visualization editing intents [48]. To address the lack of modeling, we survey and unify intent descriptions from our formative study and modern visualization construction systems, which all involve gradual editing and refinement. In particular, following the visualization guidelines [15, 22, 50], we examine editing actions involved in creating and editing visualization designs in commercial products or online demos¹, including Microsoft Excel, Tableau, Voyager [68], Data Illustrator [32], ChartAccent [42], Charticulator [44], and InfoNice [66].

These systems widely cover the three common categories of visualization construction tools [48], namely, template editor, shelf-construction, and visual builders. Each of the categories construct charts with different editing workflows: (1) In *template editor* tools like Microsoft Excel and RAWGraphs [37], users frequently start with high-level command and queries to generate template charts before making further extensive customization. These commands omit details on visualization channels, and instead stress chart types and targeted data entries. For example, to query a bar chart with y-axis encoded with Sales and x-axis encoded with different countries from these tools, the most natural driving utterance might be “Give me a bar chart for *Sales* by *Country*”. (2) Tools supporting *shelf-construction* (e.g., Tableau, Voyager [68]) builds charts by mapping data fields to encoding channels (e.g., x, y, color, shape). To construct the same example above, the utterance that best aligns with these tools would be “Bind *Sales* to y-axis” and “Bind *Country* to x-axis”. (3) *Visual builders* tools like Data Illustrator [32], Charticulator [44], ChartAccent [42], and InfoNice [66] would gradually tune marks, glyphs, coordinate systems, and layouts. Users might start with “draw a rectangle”, and then, “repeat it horizontally on *Country*”, and finally “bind height to *Sales*” to arrive at a similar bar chart.

Comparing across the aforementioned chart editing actions, we make two observations: (1) The levels of editing intent vary greatly from single element editing to integrated construction; and (2) despite the varying intents, the operation and its targeted data entries remain present in all utterances. Therefore, we consider *operations*, *objects*, and *parameters* as the core entries, and formally propose the notion of an editing action as a combination of the three:

editing action := {operation, objects, parameters}

The *operation* identifies the type of editing intent (e.g., “change color”, “add annotation”). With a single notion of “operation”, we maintain flexibility across varying levels of editing intent (e.g., low-level operations like “Bind *Sales* to y-axis” versus high-level ones like

¹The editing actions are collected from the tools’ official introduction web-sites and third-party online tutorials.

Operations	Objects	Parameters	Example utterances
Data op. <code>filter</code> Encoding op. <code>setChartType</code> <code>bindX</code> <code>bindY</code>	<code>data field</code> <code>chart</code> <code>x axis</code> <code>y axis</code>	<code>"2012"</code> <code>bar chart</code> <code>"Brand"</code> <code>"Sales"</code>	"Give me a bar chart of Sales by Brand in 2012"
Mark op. <code>setShape</code>	<code>mark</code>	<code>"woman icon"</code>	"Set the marks to woman icons"
Layout op. <code>moveDirection</code>	<code>*</code>	<code>direction=down</code>	"Make it lower"
	<code>relativePlace</code>	<code>position=chart[right]</code>	"Move the legend to the right of the chart"
Styling op. <code>setColor</code>	<code>[mark=bar, data="Ford"]</code>	<code>color=red</code>	"Set color of the Ford bar to red"
	<code>[shape=line, type=dashed]</code> <code>[shape=line, type=dashed]</code>	<code>color=black</code> <code>width=self[+]</code>	"Turn the dashed lines black and thicker"
Annotation op. <code>addTrend</code>	<code>trend line</code>	<code>*</code>	"Add a trendline to the chart"

Figure 2. An utterance can be mapped into a series of editing actions. An editing action consists of three parts: operation, objects, and parameters. The examples show the usages of object selector, vague parameters, and relational parameters. * denotes unknown properties.

"Give me a bar chart for *Sales by Country*"), and therefore can fit into different styles of visualization tools. Meanwhile, the *objects* are the targets that the operation applies to, such as the canvas area, the bars in a bar chart, the title text, etc, whereas the *parameters* indicate the degree or configurations of the operations. Figure 2 is an example that illustrates the process to author an annotated bar chart through natural language. We will describe them in detail below.

4.1.1 Operations

We first go through the category of Vega-Lite grammar [49], and a large number of editing operations related to *data*, *encoding*, and *mark*. To support more expressive changes for communication and presentation, we further survey ChartAccent [42], Charticulator [44], Data Illustrator [32], and InfoNice [66] to develop three more categories related to *layout*, *styling*, and *annotation*, which are not covered by analysis-oriented NLI. They form the eventual six main categories, which we use as guidance for designing the concrete operations (examples are listed in Figure 2):

- **Data operations** are those that retrieve and calculate data from datasets. Typical operations include filter, aggregate, bin, set time unit, and sort.
- **Encoding operations** are the ones that bind data fields to different encoding channels. Users may specify encodings by explicitly binding elements, or by assigning chart types.
- **Mark operations** refer to the configurations related to the symbols that encode data in a visualization. By changing the attributes of marks like shape, color, and style, users may implicitly customize the encoding styles applied to all related data points (e.g., use circles rather than rectangles in shape encodings).
- **Styling operations** include graphical and textual edits. Graphical operations can change the color, size, shape, icon, and stroke of graphical elements (e.g., axis lines), whereas textual operations are acted on text-related properties such as font and content. Compared to mark operations, styling specifically refine on elements without data encoding.
- **Layout operations** concern the positions and offsets of chart elements (e.g., annotations and legends). Users can place an object at a designated position, or move it along a certain direction.
- **Annotate operations** manipulate annotations that enhance the charts. Common types of annotations include labels, annotation text, reference line/band, trend line, average line, etc. They can be bounded with data; for example, trend lines can show the trend between two time points.

4.1.2 Objects

Objects are the components that operations target. For basic visual charts, the objects include marks, axes, titles, legends, gridlines, etc. More advanced objects include annotations such as trend lines, reference lines, reference bands, text annotations, and embellishments. These objects correspond to different operations. For example, the objects of *data operations* include data fields, data points, data type, data

range, etc. The objects of *encoding operations* include the data fields, chart components (e.g., axes, legends, and title), and mark channels (e.g., mark size, mark shape, and mark position). The objects of *styling operations* include visual properties such as opacity, stroke, text font, text color, etc.

Identifying objects from user utterances can be nontrivial for two reasons. First, the objects in an editing action can be underspecified. For example, a user may say "give me a chart" after uploading a dataset. We directly model users' intent and use a special notation "*" to represent the under-specified objects, such that the explicit reference can be deferred to the visualization system implementation.

Further, rather than using standard terms like "mark", users may refer to objects with descriptive languages on *object properties* ("red line", where it is not clear whether the line is a mark or an additional visual shape), or with names created on-the-fly (e.g., users may simply name the chart as "US2020" after applying a filter "US" and "2020" to the chart) [20, 27]. We help capture such reference patterns with object selectors and naming:

Object Selectors. We use object selector to process implicit filters. Consider the intention "turn the red line blue". While it does not directly specify an object, we use the selector to present it as `[shape=line,color=red]`, i.e., to select objects with the properties "line" and "red."

Object Naming. We allow dynamically assigning names to chart to reflect the aforementioned use case of "US2020". Besides easy reference, object naming further enables nested designs of charts. For example, users may put two charts together by opening up a new canvas and saying "put *US2019* and *US2020* side by side." This also helps enable creative visualization design. Figure 7(c) shows an example where the user uses two icons to create a pictograph and reuses it in another chart to create novel visualization designs.

4.1.3 Parameters

Parameters are specific configurations of the operations. They correspond to the operations and objects, describing to what extent the operations are performed, or how operations are applied to objects. As a result, the parameter types vary with object properties. They can be numerical values for chart width and height, enumerations for font style and encoding channels, and boolean values for adding or removing components. Figure 2 shows example parameters corresponding to different chart objects. While the parameters can be definite in some cases (e.g., the *absolute pixel quantity* in "set the chart 10px wide and 13px height."), they sometimes take more *qualitative* or *relative* forms:

Vague Parameters. Utterances usually contain free-form parameters that sound natural to humans, yet are hard-to-decode for the machines [19]. We collect these keywords as part of the parameter library, and expect the application tool to further interpret and map them into machine-understandable values based on design environments. For example, for color values, we support a list of color names (e.g., `color=red`, `color=navy blue`). We also define a number of extent keywords, ranging from "extremely", "very", "moderate", "little", to

Utterances	Editing actions
(a) Add a trendline to the chart	{AddTrend, trend line, *}
(b) Move legend to the right of the chart	{RelativePlace, legend, position = chart[right]}
(c) Make it lower	{MoveDirection, *, moving direction = down}
(d) Turn the dashed lines thicker	{SetWidth, [shape = line, type = dashed], line width = self[+]}
(e) Set color of the Ford bar to red	{SetColor, [data = "Ford", mark shape = "bar"], color = red}
(f) Set mark to pink woman icon	{BindChannel, mark shape, shape = "woman"} {SetColor, mark, color = pink}
(g) Give me a bar chart of Sales by Brand in 2012	{SetChartType, chart, chart type = bar chart} {BindX, x axis, data field = "Brand"} {BindY, y axis, data field = "Sales"} {Filter, *, data value = 2012}

Table 1. Example utterances (left) and corresponding editing actions (right). One utterance may correspond to multiple editing actions.

“very little”. Users can therefore describe the adjustments of size as “make it really large”, which will be mapped to “very large” and can be parsed as {setSize, *, size=very.large}.

Relational Parameters. Users also tend to make relative statements anchoring on some existing visual elements, e.g., “make it larger”, or “set the color of the title darker than the bars”. For these relational parameters, we expect users to specify the objects it compared to, and the direction of change. For example, “make the title darker” is explained as make *itself* darker, where the object is the title; “place the legend on the right of the plot area” is explained as *right* to the *plot*. Formally, the utterances are interpreted to {setColor, title, color=self[darker]} and {place, legend, position=plot[right]}. self[darker] means the color darker than it self while plot[right] means on the right of the plot. Applications should further determine the exact values of these parameters after receiving these vague parameters.

4.2 Example Utterances

Combining operations, objects, and parameters, authoring actions can represent various authoring utterances. An utterance may corresponds to multiple editing actions. Here we show example utterances and the corresponding editing actions to further illustrate the use of editing actions. In Table 1, the utterances are represent with one to four editing actions. These examples include object selectors (d, e), vague parameters (b, c, d, e, f), and relational parameters (b, d).

5 NL INTERPRETER

To demonstrate the efficacy of the editing actions and enable visualization applications to assess our modeling, we provide proof-of-concept design and implementation of an NL interpreter. For an input utterance, NL interpreter aims to translate it to a sequence of editing actions that will then be passed to visualization applications. As shown in Figure 3, we design a multi-stage interpreter to parse a natural language utterances into editing actions. First, it identifies user intents in the utterance as operations; then, it extracts useful parts as objects and parameters for the targeting operations; finally, it recognizes the relations between the identified operations, objects and parameters, and organizes them into editing action tuples.

5.1 Stage 1: In-context Data Entry Abstraction

To reduce the sparsity and uncertainty of intent recognition, we use a data entry recognizer to identify and abstract dataset-related data entries in the utterance. The recognizer enumerates through the N-grams in the utterances, checking for the similar entities (words or short phrases) that can be matched with the data attributes, including the *table names*, *column names*, and *cell values* of a dataset table. To resolve ambiguities, we compute the similarity between the tokenized entry and the data attributes in a similar way to NL4DV [39]. Once we completed data entry recognition, we replace the entities with placeholders like <column>, <value> to obtain the *abstracted utterances*. In the case of

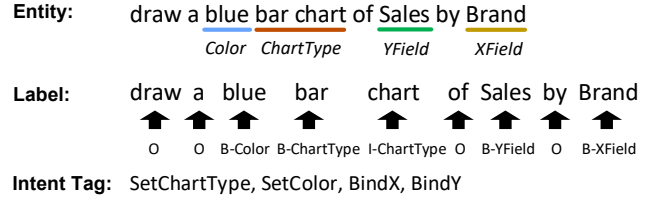


Figure 3. A multi-stage natural language interpreter: A natural language utterance is first abstracted by replacing data-related entries. Then, we extract operation categories. Following the BIO format [41], we can map the words in the sentence into a sequence of labels, where B-, I-, and O represent begin, inside, and outside. The extracted parameters, objects, and operations are further synthesized into editing actions.

special numbers and dates, we further specify value into <integer>, <float>, <date>, and <year>.

5.2 Stage 2: Information Extraction for the Editing Action

At the core of the interpreter is the ability to parse the utterances into operations, objects, and parameters. Operations are usually summarized or abstracted from utterances, whereas objects and parameters are directly extracted. As such, while there might be other alternatives, we treat this step as two NLP sub-tasks: (1) a multi-label classification task, to detect potentially multiple operations from an utterance; and (2) a sequence labeling task using BIO sequence tagging [41], to recognize token chunks that represent objects and parameters from one sentence. Specifically, we chose to frame extraction step as sequence tagging (as oppose to e.g., end-to-end NL to code translation) for its richer interpretability. With fine-grained semantic annotation on each entity, it is easier for people to inspect and potentially amend the parsing result of a given utterance.

As shown in Stage 2 of Figure 3, we design a deep-learning model to simultaneously performs the aforementioned two tasks, outputting separated lists for operations, objects, and parameters. While various prior work used rule-based methods (e.g., both FlowSense [72] and NL4DV [39] use lexical and dependency parsing structures) for precise recognition, we argue that heuristic rules usually only handle limited forms of utterance. In comparison, deep-learning models are more capable of flexibly interpreting diverse utterances. One can easily extend the capability of the NL interpreter by adding training examples that express intended operations, objects, and parameters [18,30,71,73].

In specific, as shown in Figure 4, the neural model is based on the encoder-decoder framework [30], attention mechanism [65], and Conditional Random Fields (CRF) algorithm [25]. We chose the model for its balanced quality and efficiency. On the one hand, Bi-LSTM-CRF is commonly used for similar sequence tagging tasks like Named Entity Recognition, and empirically we found its accuracy to be sufficiently high; On the other hand, the model is lightweight enough that it can be deployed and integrated into various different platforms without causing much latency (unlike some pre-trained models, e.g., BERT).

We train the model on a dataset containing a mix of real and synthetic utterances expressing certain visualization editing intents. We first crowdsourced real utterances on Amazon Mechanical Turk (AMT). To do so, we created 75 pairs of charts such that in each pair, the second chart can be made from the first one through up to three editing operations. We present these pairs to crowdworkers, and ask them to describe in natural language how they would make the edits (e.g., “change the color in the bar chart, and then rescale”). We collected 100 descriptions per chart pair, resulting in ~5.4k utterances, and kept 4.7k after manual validation. We configured the AMT HIT such that each crowdworker would describe ten chart pairs. We collected the dataset from ~400 unique workers. We further augment the dataset with synthetic examples by (1) paraphrasing these utterances through back translation and crowdsourcing [21], and (2) creating syntactic templates inspired by Malandrakis et al. [36]. The augmentation is a common technique for accelerating early collections of user intents [36].

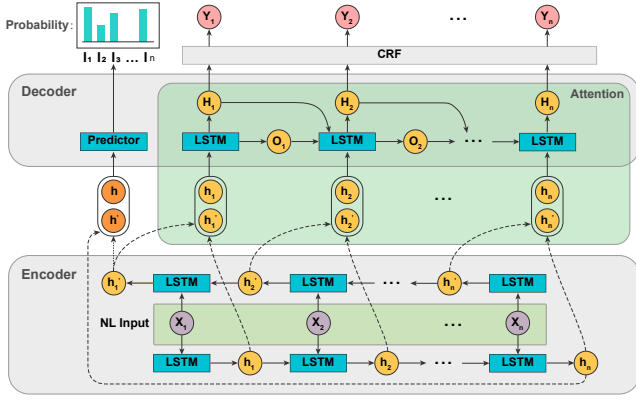


Figure 4. The architecture of the deep learning model for intent and entity extraction, which is based on the encoder-decoder framework, attention mechanism, and CRF algorithm.

In total, our dataset contains 10.7k utterances². We use 80% of the data for model training and the rest for testing. In the training process, we use the Adam optimizer. The batch size is 32, and the epoch is 150. Our model worked well on the test set: the operation classification and the tagging F1 scores for objects and parameters were 94.75%, and 97.34%, respectively. The sequence labeling F1 score is evaluated at entity-level using seqeval python library [6].

The performance is promising, and partially benefits from the fact that the data binding and data abstraction in Stage 1 (Section 5.1) eliminates noise in entity recognition. However, another primary factor driving the high performance would be our modest coverage of the utterances. The chart pairs we created to collect data involves 1-3 operations, which were easy to be expressed by the crowdworkers and learned by the model. It should be noted that the accuracy of the model may decrease when the user utterances become more complex. In fact, being “data hungry” is a significant bottleneck for deep learning models: to reach the best performance, a large amount of training data is required to provide enough training signal on more rare and complex patterns. Limited by the training data, we only see our implementation of the deep learning model as a starting point and a proof-of-concept. We note that to maximize the utility of such models, future work should collect a large number of utterances, and additionally rely on data augmentation techniques as we explored.

5.3 Stage 3: Editing Action Synthesis

Given the recognized information, we then map between the independently outputted operations, objects, and parameters, to organize them into a sequence of editing actions. We take a bottom-up approach to traverse the operation list. Because each operation determines its expected parameters and objects, we match the corresponding parameters and objects among the candidates. For example, `setColor` would only accept colors as parameters. Thus, we traverse the list of parameters and objects to search for an entity “blue” labeled as `color`. If there is no parameter or object for an operation, we fill it with a * mark to denote the default state. If there is more than one parameter for an operation, we duplicate the operations into two or more and synthesize more editing actions. We further heuristically rank these actions based on the category they belong to. We prioritize global operations over local operations. The operations in the category of data operations (e.g., sum) and encoding operations (e.g., bind a column to x axis) should be executed before mark operations (e.g., changing the shape of marks), styling operations (e.g., changing the color or the title content), layout operations (e.g., move upward), and annotate operations (e.g., add text labels). The editing actions are then ordered by their appearance in the utterance. The operation mappers can then implemented to map the editing action sequence to application-specific operations.

²<https://github.com/microsoft/VisTalk>

5.4 Guideline for Reusing the Interpreter

Instead of directly translating NL utterances into application-specific commands, our NL interpreter abstracts out the interpretation of editing intents to enable reusability, following the existing NLI framework [39]. The only difference across visualization applications is the implementation of operation mapper, as shown in Figure 1(b). When users’ input is complete and accurate, the interpreter can support easy binding and the editing action can be directly mapped to executable application-wise commands. When users input application- or context-dependent utterances, we expect the application to clarify the vague and relational parameters we define in Section 4.1.3. For example, when users input “change color to cyan” in Figure 7(a), the objects corresponding to the operation `setColor` is underspecified and denoted as “*” by the NL interpreter. Therefore, the system should trace back to find the most recent utterances with the objects where the operation can be applied. Here, we discuss some implications on such handling.

Operation mapper. The visualization applications should implement an operation mapper that maps the interpreted editing actions into tool-specific operations. If the operations are clearly specified, the execution engine blends the objects and calls the related functions. If the operation is not supported, the execution engine can return error messages to notice the users. If the operation is not supported but related operations are supported, the application can further show simple examples to help users understand the functions of the system. When the actions are not fully specified, the application could predefine a set of rules to recommend proper actions for under-specified utterances. Applications could further adopt more advanced chart recommender such as Draco [38] to resolve the underspecified editing actions to improve the editing experience.

Manage Contexts. To better understand users’ editing intents, the applications could manage users’ sessions and maintain a list of objects on canvas with corresponding attributes. If the objects are clearly referred, the application simply looks for the objects and applies operations to the target objects. For objects with selectors, applications can traverse through the parameters of the objects to find out the referred objects. If the objects cannot be resolved, the application could infer the target objects by examining the recent objects being edited, and recommend possible editing. Alternatively, the application could also return error messages to notify users to clarify their intents.

Resolve Ambiguities. The applications should build support and disambiguate between the three types of parameters: exact, vague, and relational ones. For exact parameters, applications can directly map the extracted parameters to the data value and check whether the value can be legally assigned. For example, `10px` for height can be directly assigned to the objects. For vague parameters, applications can develop a set of metadata or rules to explain vague parameters supported by the NL interpreter. For example, when users input a vague parameter, the application looks up the parameter `red` and yields the color code `RGB(255, 0, 0)`. For relational parameters, applications need to parse the parameters into the changing directions of parameter adjustment. For examples, “on the top” for “put it on the top of the *US* bar”, means the object should be placed at a point that has smaller distance to the top border than the *US* bar, but with the same distance to the left/right border as the *US* bar. Therefore, applications should extract the positions for the *US* bar, and calculate a proper position number for the newly added object. To recommend the proper position for newly added objects, advanced chart layout algorithms [69] can be adopted to suggest system generated design under the constraints of users’ intents. Similar design recommendations can also be implemented by the systems to suggest colors, encodings, etc. To enable users to conduct even more fine-tuned operations, other modalities such as mouse and touch should be introduced.

Multimodal Interaction. Obviously, natural language is not always the best choice of input when conducting visualization editing, and we believe it acts as a complementary input modality to traditional WIMP interaction. Other forms of input (e.g., speech, mouse, touch, and pen) could also be combined to support more multimodal interactions. For example, instead of expecting users to grasp the jargon for describing visualization objects (e.g., bar, scatterplot, etc.), they can *select* objects

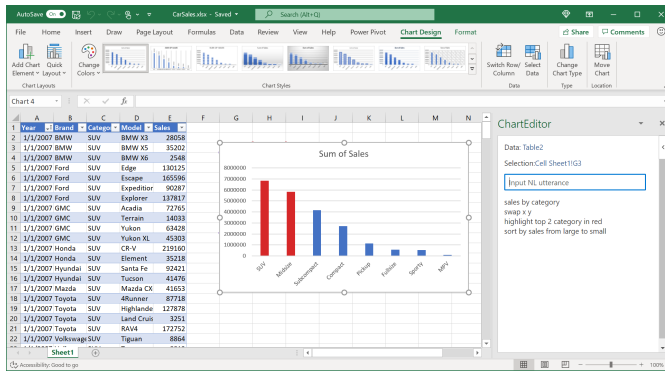


Figure 5. An Excel add-in for chart creation and editing. Users can input NL utterances in the input box. The selected chart updates accordingly.

before voicing natural language commands relevant to the selected object. Furthermore, voice input is also a design choice to combine multiple modalities, where users could save the efforts of typing NL utterances. This could be realized by introducing a speech recognition [45] module at the beginning of our current pipeline to transform voice into text. However, speech-based NLI faces unique challenges (e.g., triggering speech input and transcription errors). The NL input could also be interleaved with input from other modalities. For example, when resolving editing actions that are not partially specified or ambiguous, visualization systems could also prompt other modalities of user interfaces, such as interactive widgets to help users clarify their intents. Multimodal coreference resolution is an important task as users may input NL queries that follow their direct manipulations on the interface [2, 53]. So the design of operator mapper should involve heuristics to handle coreference resolution in multimodal sense based on the design of the NL interactions of the authoring system.

Enhance Discoverability. Users may not be aware of what operations are available to the system and whether there is a preference for a particular language structure in the system. System discoverability is considered an essential factor that improves the user experience [3, 56]. It could happen that the interpreter could parse the utterances into editing actions, but the system is not able to execute the operations. Another possibility is the system has the corresponding operations but the interpreter could not understand the utterance correctly. For the former situation, the system could explain its scope of capability based on the editing actions. For the latter situation, the system could either notice the users to take other modalities to complete their tasks, or educate users on how to phrase queries that can be interpreted correctly by the system. In terms of interaction design, text autocompletion can be leveraged to help users precisely complete NL input; interactive widgets with data/visualization previews can be useful for visualization authoring to enhance discoverability.

6 EXAMPLE APPLICATIONS

As proof-of-concept, we build two example applications of NL-based visualization editing system powered by the NL interpreter in Section 5. The systems are example implementations of the *application* component in Figure 1.

6.1 Excel Chart Editor

Based on our NL interpreter, we build Excel Chart Editor, an Excel add-in to integrate NL-based chart editing into Excel. As shown in Figure 5, we provide a natural language input panel for users to type in their natural language commands. The chart editor creates a new chart when users type in their first sentence and continuously update the chart as they input follow up utterances. Since Excel already offers rich functions of creating and editing charts on canvas, the primary implementation effort is on mapping the editing actions to the predefined executions [1]. This includes resolving object selectors, executing data queries, and converting vague parameters into accurate default values. For example, the editing action `{setColor, mark, color=red}` can directly map

to `setSolidColor(color)` method in `Excel.ChartFill` interfaces. Further, we take a set of simple heuristics to resolve ambiguity. For example, for input utterance “sort”, users do not specify the order and the data field. We by default sort the fields corresponding to the y-axis in descending order. When user input “make the line stroke wider”, we increase the line stroke width by 50%.

Note that the implementation of the Excel add-in does not concern about the textual inputs, they are only parsed by the NL interpreter. It demonstrates how our NL interpreter and editing actions can be plugged into existing tools, and seamlessly augment the larger pipeline of data analysis and presentation workflow.

6.2 VisTalk

We also develop VisTalk, an NL-based standalone chart creation tool. Similar to the Excel add-in, VisTalk takes users’ natural language input through a simple text input box (Figure 6(b)), and automatically re-renders the visualizations to show changes accordingly (Figure 6(c)). But uniquely, while Excel add-in concerns augmenting existing tools, we use VisTalk to show how developers might design their own applications while maximizing the utility of our NL interpreter.

In particular, VisTalk demonstrates how applications can resolve ambiguities in the identified editing actions. As mentioned in Section 6.3, users commonly submit queries that have missing information. Compared to the examples in Figure 2 (“give me a bar chart of Sales by Brand in 2012”), an utterance “Sales by year” misses the `setChartType` operation, and can only be parsed into `{BindX, x axis, data field="Year"}` and `{BindY, y axis, data field="Sales"}`. In response, VisTalk recommends chart types from partial specifications of visual charts, following the research of chart recommendations [35, 52]. In the above example, VisTalk would recommend a bar chart based on the data types. VisTalk also adopts interactive widgets (Figure 6(f)) and utterance auto-completion (Figure 6(g)) to disambiguate parameters (e.g., for elements). For example, users can click on the keywords within the input utterance to specify and refine the choice of icons, colors, chart types, and values in a pop-up window (Figure 6(f)).

Example Gallery: To demonstrate the usability of VisTalk, we produce a variety of charts with example natural language utterances that specify the visualizations as shown in Figure 7. These visualizations cover a wide range of operations that users may take in real world visualization authoring scenarios, and show that VisTalk can enable the creation of basic charts with simple operations such as data operations, encoding operations, and annotate operations with NL utterances.

6.3 Exploratory Study

With VisTalk at hand, we further conduct a user study to understand how users interact with visualization systems to construct visualizations through natural language.

6.3.1 Participants

We recruited 12 participants (three females and nine males, age ranging from 22 to 45) who had normal or corrected-to-normal vision. The participants include undergraduate and graduate students, data analysts, researchers, and software engineers. They are general users who need to create visualizations to present data in daily work. All of the participants had used Excel to create charts before, while five of them had used Tableau or Power BI to create charts.

6.3.2 Study Procedure

Our user study lasted for about one hour. We first surveyed participants’ background information with a questionnaire. Next, we provided a tutorial outlining the features of VisTalk with two examples. As a warm-up exercise, we encouraged them to freely explore the tool with one sample dataset. Then, we provided the participants with **five visualization reproduction tasks**. The participants are asked to reproduce the charts that we provided, ordered by complexity. These tasks included the reconstruction of five charts: (1) one basic bar chart, where participants only need to specify data and encoding; (2) two charts with annotations (i.e., trendline, average line, and annotation band); (3) one pictograph,

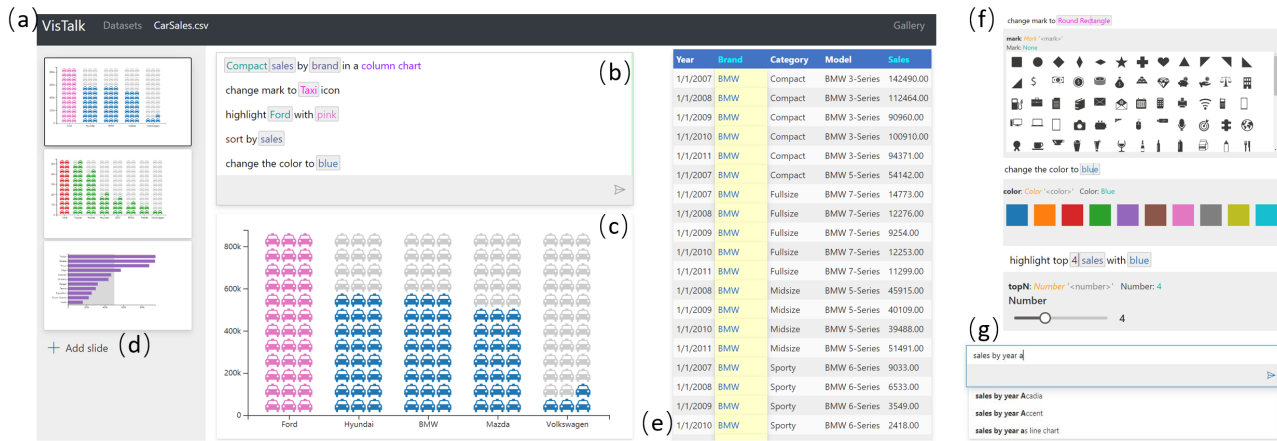


Figure 6. VisTalk interface (a). Users can type in their natural language utterances in the text box (b) to customize their visualization design (c). Users can author different charts at the same time (d). The table view shows the underlying dataset (e). Users can click the words as interactive widgets to resolve ambiguities (f). The auto-completion panel pops out during the input process (g).

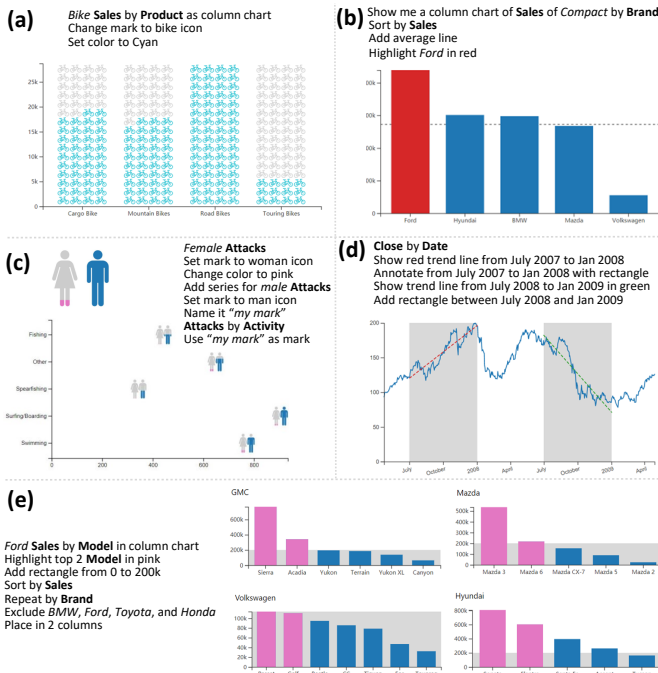


Figure 7. Examples with VisTalk. (a) is based on a product sales dataset, (b) and (e) are based on a car sales dataset, (c) is based on a shark attacks dataset, and (d) is based on a stock dataset. The terms displayed in bold and italic are column names and values in data, respectively.

where participants need to modify the shape of the marks; and (4) one chart with multiple views, where participants need to repeat the chart design on a data column, similar to Figure 7(e). For each task, we displayed the target chart, described the underlying datasets, and asked participants to reproduce the chart in VisTalk using natural language utterances. Afterwards, participants rated their experience with VisTalk in the form of a five-point Likert scale [29]. We further collected their free-form feedback through a semi-structured interview.

6.3.3 Results

Subjective Satisfaction: Overall, the participants are positive about the system. Users highly rated the experience of using VisTalk ($M = 4.36$, $SD = 0.69$). Figure 8 lists the feedback of VisTalk. Eleven of twelve users agree VisTalk is easy to learn and easy to use. Eleven users think VisTalk has improved their productivity. Twelve users felt satisfied. The feedback also reveals space of improvement. Our

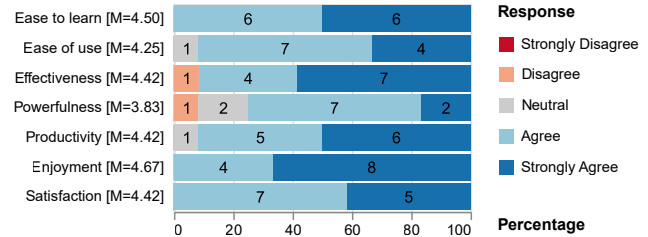


Figure 8. User ratings of VisTalk system with a 5-point Likert scale.

participants rated it relatively low in terms of the powerfulness of the system and the effectiveness of completing their jobs. It may be helpful to design more functions and interactions for VisTalk for future work.

Result Analysis: All the participants have completed the tasks. The participants learned quickly after the training. While they were typing, the system were parsing the query automatically in the meantime. If the utterances were successfully parsed, specific parameters in the utterances were highlighted (Figure 6 (b)), and the visualization was updated. If the system could not parse it, the user could iteratively modify the queries to make them correct.

To understand the complexity of utterances, we use mean length of utterances (MLU) [40]. The number of words used ranged from 1 to 11. The average MLU across 12 users is 3.61 words per utterance ($SD=1.26$). On average, the participants used 2.4 utterances to complete a visualization creation task ($SD=0.72$). The average time of creating one chart was 139s. We found the participants tend to avoid using long sentences, but instead, shorter utterances for incremental configurations. Since the input textbox in VisTalk is directly editable, users tend to correct the NLS inplace instead of appending additional utterances.

Our participants had a positive experience authoring with natural language. For simple tasks, users usually only type in some keywords to see how the systems can parse their intentions. For example, after one user (P2) tried “trend” and then “add trend”, the system automatically added a trend line to the line chart. The participant said, “It worked well. Natural language makes the thing easy.” Participants also mentioned that it is easy to accomplish complex tasks with natural language, because it allows them to combine many operations within a simple sentence in the way that meets their needs. One user (P3) commented, “The experience is very good! In the past I need to search the menus and click the right buttons to perform the tasks. Sometimes it costs a long time to find the right menu item. Now I don’t need to think of where they are.”

The majority of participants (10/12) mentioned that the natural language interface saves their time and efforts. Most participants (8/12) also said they will use natural language as their first choice if their

familiar systems support similar functions. P9 said, *“To design a simple chart, using one simple sentence is enough. It is completely different from using mouse to select menus, select the data field from the dataset, specify chart types, and do further adjustments.”* Impressed by the parsing capability of VisTalk, P1 mentioned *“It is surprising that the system could perform correctly even when the sentence was very short and I misspelled some words and had grammatical errors.”*

The participants (12/12) find the system interactions easy to learn, even for those without much experience in visualization. One participant stated, *“I can quickly understand the effects of the commands.”* One of the participants who was not familiar with visualization creation tools said, *“Now I don’t have motivations to learn to use traditional visualization tools. I will come to VisTalk if I need to create charts. It seems much easier to learn.”* P3 also mentioned, *“It is friendly to non-expert like me. I don’t find difficulties of learning to use this tool. It gives me prompt feedback while I am typing and I feel I can complete the tasks easily. The overall experience is quite smooth.”*

Although the chart recommendation functions took simple strategies, many of our participants (7/12) mentioned the system recommendation is helpful. One participant said, *“The default charts were well-designed and I don’t think I need to do further modifications.”* Three participants mentioned they enjoyed using short phrases to interact with systems’ default recommendations. P9 said, *“I am using these keywords to explore this system, similar to the experience of using search engines. I enjoy the way this system gives me surprises.”*

Participants also found difficulties when using NLIs. For example, when a user (P11) wanted to change the color of the charts, he typed “highlight in blue” and the system did not correctly execute the command because no object to highlight is provided. Some common errors in NLIs also happen in the study, such as synonyms (e.g., “bike” and “bicycle”), ambiguous utterances, and spelling errors. For example, one user (P2) said, “set bike mark to bike”. The system can not parse the query as “bike” is ambiguous (the first is mark and the second is value). One user (P5) also typed some specific configurations (e.g., “arrange by 4 × 2” can not be parsed but “arrange in 2 columns” works) and system control (e.g., “refresh view” and “re-center graph” ask the system to load and position the visualizations) that are not supported by the interpreter. Some analysis-oriented NLI users may pose queries that expresses an intent of data analysis, instead of chart editing (e.g., “What is the relationship between sales and product?”). Some other participants also felt they couldn’t express their requirements especially at the beginning of the tasks. *“Sometimes I forgot the words and felt hard to describe my requirements.”* One user applied to open and read the examples in the tutorial again at the start of the tasks. Another user opened search engine to look for a correct word. Although all of them got used to the natural language experience after a while, we find that cold start can be a problem for users that are not familiar with the system. The reason may be that users are dim about how much the system could tolerate their vague or inaccurate expressions. To solve this problem, we believe more sufficient guidance and feedback can address this issue. The interface should also provide recommendations to give hints about potential choices that they can give commands.

7 DISCUSSION

By taking an one-interpreter-for-all schema, we not only save the efforts of supporting NLIs for visualization authoring applications, but also alleviate the need for users to get familiar with specific concept models. Still, with NLIs freeing users’ mental models from any specific design paradigms [48], the overflexibility can lead to some challenges in practice. Below, we discuss three gaps between what people would say and how a system might (incorrectly) respond, all emerged from NL commands being overly flexible, and discuss potential future work.

Gap 1: Users express commands in diverse patterns, but NL parsers cannot recognize them all. Users that have different backgrounds, or are familiar with different visualization tools, can express the same objective in drastically different ways. For example, Section 4.1 enumerates the different visualization creation types; to create the chart from scratch, the utterances can be as high level as just mentioning the data columns (and the system is expected to automatically

infer the chart type), and it can also go as low-level as binding each specific mark channel (e.g., rectangles) individually. To achieve the seamless switch between these types, the *interpreters* should be able to understand different commands. Following the deep learning approach as in our interpreter (or, even to generalize it with recent pretrained NLP models [13]), a crucial future step would be to collect a diverse set of training utterances that express the same intent in various usage scenarios and levels of details. This can be achieved by feeding crowd labelers with richer *visualization objectives*. We can even diversify utterances by constraining them on *construction workflows*, i.e., to ask expert users to write utterances that can only be parsed into a serial of preset actions. Moreover, it is possible to collect naturally occurring utterances without any predetermined objectives or application workflows. Recently, researchers have asked users to freely submit any possible queries to analyze charts, so to build taxonomies on representative utterances [59], and synthesized natural language to visualization (NL2VIS) benchmarks from NL2SQL benchmarks [33]. Though still targeting at visual analysis, these work shed light on possible designs to collect diverse utterances. We can then further augment the datasets by paraphrasing these commands.

Gap 2: Users have various editing objectives, but the downstream application may not implement them all. While the capability of parsing natural language utterances increases, our applications might be *designed* to only focus on a subset of visualization creation and an editing actions. For example, how should a system react to “change the rectangle to circle”, if it only wants to support analytical interactions, and therefore does not allow mark customization? Would it disappoint people, if by design an application ignores more utterances than another tool, even when the interpreter parses them correctly? User studies on how users react to the boundary of applications are interesting. Alternatively, we can improve command recommendation, make suggestions to users when their utterance does not belong to the supported function type, and help people understand the scope of application functions.

Gap 3: Users use NL for various intentions simultaneously, but the framework does not go beyond authoring. Along with our work, there exist various frameworks that tackle analysis, authoring, data processing, etc. separately. However, users may interchangeably express these needs through natural language all at once, and it is impractical to expect users to swiftly switch between these tools. Here, integrating analysis-oriented and authoring-oriented NLI seems promising. Just as we have hinted in Section 1, users can start with the visualizations produced from analysis-oriented NLIs and make further configurations on top to improve visual interaction with authoring-oriented NLIs. Furthermore, recent NL2SQL [8, 75] advances can also be utilized for low-level data-driven queries. One challenge could be to design additional query type classifier modules that can identify analysis-oriented, authoring-oriented, and database-oriented queries and, more importantly, how these operations should be correctly ranked so that the final visualization reflects all the requirements correctly.

8 CONCLUSION

In this paper, we explore a natural language-based visualization authoring pipeline, which supports the understanding of visualization construction commands. We propose the definition of editing action that describes users’ visualization editing intents, defined as tuples of operations, objects, and parameters. The editing actions bridge the gap between the NL interpreter and visualization applications. We further implement a deep learning-based NL interpreter, to extract operations, objects, and parameters from users’ utterances. From these extracted information, we synthesize the editing actions for visualization applications to handle. Based on the NL interpreter, we demonstrate the utility of our pipeline and NL interpreter with two example applications, an Excel chart editor and a proof-of-concept authoring tool, VisTalk. To assess our approach, we further conduct a user study with VisTalk to understand how users edit visualizations through natural language. Our study is a first step towards NL-based visualization authoring.

REFERENCES

- [1] Microsoft excel javascript api. <https://docs.microsoft.com/en-us/office/dev/add-ins/reference/overview/excel-add-ins-reference-overview>.
- [2] Multimodal Coreference Resolution for Exploratory Data Visualization Dialogue: Context-Based Annotation and Gesture Identification. In *Workshop on the Semantics and Pragmatics of Dialogue*, pp. 41–51. ISCA, 2017.
- [3] Sneak pique: Exploring autocompletion as a data discovery scaffold for supporting visual analysis. In *Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology*, pp. 966–978. ACM, 2020.
- [4] Ibm watson analytics. <http://www.ibm.com/analytics/watson-analytics/>, Mar. 2022.
- [5] Power bi q & a. <https://docs.microsoft.com/en-us/power-bi/natural-language/q-and-a-tooling-intro>, Mar. 2022.
- [6] Sequeval python library. <https://github.com/chakki-works/sequeval>, Mar. 2022.
- [7] Tableau software: Business intelligence and analytics. <https://www.tableau.com/>, Mar. 2022.
- [8] K. Affolter, K. Stockinger, and A. Bernstein. A comparative survey of recent natural language interfaces for databases. *The VLDB Journal*, 28(5):793–819, oct 2019.
- [9] R. Amar, J. Eagan, and J. Stasko. Low-level components of analytic activity in information visualization. In *Proceedings of the 11th IEEE Symposium on Information Visualization*, pp. 111–117. IEEE, 2005.
- [10] I. Androutsopoulos, G. D. Ritchie, and P. Thanisch. Natural language interfaces to databases—an introduction. *Natural language engineering*, 1(1):29–81, 1995.
- [11] J. Aurisano, A. Kumar, A. Gonzalez, J. Leigh, B. Di Eugenio, and A. Johnson. Articulate2: Toward a conversational interface for visual data exploration. In *IEEE VIS*, vol. 16, p. 1, 2016.
- [12] M. Bostock, V. Ogievetsky, and J. Heer. D³ data-driven documents. *IEEE transactions on visualization and computer graphics*, 17(12):2301–2309, 2011.
- [13] 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, NAACL’19*, pp. 4171–4186. ACL, 2019.
- [14] K. Dhamdhere, K. S. McCurley, R. Nahmias, M. Sundararajan, and Q. Yan. Analyza: Exploring data with conversation. In *Proceedings of the 22nd International Conference on Intelligent User Interfaces*, pp. 493–504. ACM, 2017.
- [15] A. Diehl, M. Kraus, A. Abdul-Rahman, M. El-Assady, B. Bach, R. S. Laramée, D. Keim, and M. Chen. Studying Visualization Guidelines According to Grounded Theory. *arxiv*, oct 2020.
- [16] E. Fast, B. Chen, J. Mendelsohn, J. Bassen, and M. S. Bernstein. Iris: A conversational agent for complex tasks. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, pp. 1–12, 2018.
- [17] 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 the 28th Annual ACM Symposium on User Interface Software & Technology*, pp. 489–500, 2015.
- [18] D. Guo, G. Tur, W.-t. Yih, and G. Zweig. Joint semantic utterance classification and slot filling with recursive neural networks. In *2014 IEEE Spoken Language Technology Workshop (SLT)*, pp. 554–559. IEEE, 2014.
- [19] M. Hearst, M. Tory, and V. Setlur. Toward interface defaults for vague modifiers in natural language interfaces for visual analysis. In *2019 IEEE Visualization Conference (VIS)*, pp. 21–25. IEEE, 2019.
- [20] E. Hoque, V. Setlur, M. Tory, and I. Dykeman. Applying pragmatics principles for interaction with visual analytics. *IEEE transactions on visualization and computer graphics*, 24(1):309–318, 2017.
- [21] M. Iyyer, J. Wieting, K. Gimpel, and L. Zettlemoyer. Adversarial example generation with syntactically controlled paraphrase networks. *arXiv preprint arXiv:1804.06059*, 2018.
- [22] E. Kandogan and H. Lee. A Grounded Theory Study on the Language of Data Visualization Principles and Guidelines. *Electronic Imaging*, 28(16):1–9, feb 2016. doi: 10.2352/ISSN.2470-1173.2016.16.HVEI-132
- [23] D. H. Kim, E. Hoque, and M. Agrawala. Answering Questions about Charts and Generating Visual Explanations. In *Proceedings of the 38rd Annual ACM Conference on Human Factors in Computing Systems*, pp. 1–13. ACM, 2020.
- [24] N. W. Kim, E. Schweickart, Z. Liu, M. Dontcheva, W. Li, J. Popovic, and H. Pfister. Data-driven guides: Supporting expressive design for information graphics. *IEEE Transactions on Visualization and Computer Graphics*, 23(1):491–500, Jan 2017. doi: 10.1109/TVCG.2016.2598620
- [25] J. Lafferty, A. McCallum, and F. C. Pereira. Conditional random fields: Probabilistic models for segmenting and labeling sequence data. 2001.
- [26] C. Lai, Z. Lin, R. Jiang, Y. Han, C. Liu, and X. Yuan. Automatic Annotation Synchronizing with Textual Description for Visualization. In *Proceedings of the 38rd Annual Conference on Human Factors in Computing Systems, CHI’20*, pp. 1–13. ACM, 2020.
- [27] G. P. Laput, M. Dontcheva, G. Wilensky, W. Chang, A. Agarwala, J. Linder, and E. Adar. Pixeltone: A multimodal interface for image editing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pp. 2185–2194, 2013.
- [28] Y. Li, H. Yang, and H. Jagadish. Nalix: A generic natural language search environment for xml data. *ACM Transactions on database systems (TODS)*, 32(4):30–es, 2007.
- [29] R. Likert. A technique for the measurement of attitudes. *Archives of psychology*, 1932.
- [30] B. Liu and I. Lane. Attention-based recurrent neural network models for joint intent detection and slot filling. *arXiv preprint arXiv:1609.01454*, 2016.
- [31] C. Liu, Y. Han, R. Jiang, and X. Yuan. ADVISor: Automatic Visualization Answer for Natural-Language Question on Tabular Data. In *Proceedings of the 14th Pacific Visualization Symposium*, pp. 11–20. IEEE, 2021.
- [32] Z. Liu, J. Thompson, A. Wilson, M. Dontcheva, J. Delorey, S. Grigg, B. Kerr, and J. Stasko. Data illustrator: Augmenting vector design tools with lazy data binding for expressive visualization authoring. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI’18*, pp. 123:1–123:13. ACM, New York, NY, USA, 2018. doi: 10.1145/3173574.3173697
- [33] 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 the 2021 International Conference on Management of Data*, pp. 1235–1247. ACM, 2021.
- [34] Y. Luo, N. Tang, G. Li, J. Tang, C. Chai, and X. Qin. Natural Language to Visualization by Neural Machine Translation. *IEEE Transactions on Visualization and Computer Graphics*, 28(1):217–226, 2022.
- [35] J. Mackinlay, P. Hanrahan, and C. Stolte. Show me: Automatic presentation for visual analysis. *IEEE transactions on visualization and computer graphics*, 13(6):1137–1144, 2007.
- [36] N. Malandrakis, M. Shen, A. Goyal, S. Gao, A. Sethi, and A. Metallinou. Controlled text generation for data augmentation in intelligent artificial agents. *arXiv preprint arXiv:1910.03487*, 2019.
- [37] M. Mauri, T. Elli, G. Caviglia, G. Ubaldi, and M. Azzi. RAWGraphs: A visualisation platform to create open outputs. In *Proceedings of the 12th Biannual Conference on Italian SIGCHI Chapter*, pp. 1–5. ACM, Cagliari, Italy, 2017.
- [38] D. Moritz, C. Wang, G. L. Nelson, H. Lin, A. M. Smith, B. Howe, and J. Heer. Formalizing visualization design knowledge as constraints: Actionable and extensible models in draco. *IEEE transactions on visualization and computer graphics*, 25(1):438–448, 2018.
- [39] A. Narechania, A. Srinivasan, and J. Stasko. NL4dv: A toolkit for generating analytic specifications for data visualization from natural language queries. *IEEE Transactions on Visualization and Computer Graphics*, 2020.
- [40] S. Oviatt. Multitmodal interactive maps: Designing for human performance. *Human-Computer Interaction*, 12(1-2):93–129, 1997.
- [41] L. A. Ramshaw and M. P. Marcus. Text chunking using transformation-based learning. In *Natural language processing using very large corpora*, pp. 157–176. Springer, 1999.
- [42] D. Ren, M. Brehmer, B. Lee, T. Höllerer, and E. K. Choe. Chartaccent: Annotation for data-driven storytelling. In *2017 IEEE Pacific Visualization Symposium (PacificVis)*, pp. 230–239. IEEE, 2017.
- [43] D. Ren, T. Höllerer, and X. Yuan. ivisdesigner: Expressive interactive design of information visualizations. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2092–2101, Dec 2014. doi: 10.1109/TVCG.2014.2346291
- [44] D. Ren, B. Lee, and M. Brehmer. Charticulator: Interactive construction of bespoke chart layouts. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):789–799, Jan 2019. doi: 10.1109/TVCG.2018.2865158
- [45] G. Riccardi and D. Hakkani-Tur. Active learning: Theory and applications

- to automatic speech recognition. *IEEE transactions on speech and audio processing*, 13(4):504–511, 2005.
- [46] A. Saktitheswaran, A. Srinivasan, and J. Stasko. Touch? Speech? or Touch and Speech? Investigating Multimodal Interaction for Visual Network Exploration and Analysis. *IEEE Transactions on Visualization and Computer Graphics*, 26(6):2168–2179, 2020.
- [47] A. Satyanarayan and J. Heer. Lyra: An interactive visualization design environment. In *Computer Graphics Forum*, vol. 33, pp. 351–360. Wiley Online Library, 2014.
- [48] A. Satyanarayan, B. Lee, D. Ren, J. Heer, J. Stasko, J. Thompson, M. Brehmer, and Z. Liu. Critical reflections on visualization authoring systems. *IEEE transactions on visualization and computer graphics*, 26(1):461–471, 2019.
- [49] A. Satyanarayan, D. Moritz, K. Wongsuphasawat, and J. Heer. Vega-lite: A grammar of interactive graphics. *IEEE transactions on visualization and computer graphics*, 23(1):341–350, 2016.
- [50] J. Sawicki and M. Burdukiewicz. VisQualdex – the comprehensive guide to good data visualization. *arxiv*, jan 2022.
- [51] V. Setlur, S. E. Battersby, M. Tory, R. Gossweiler, and A. X. Chang. Eviza: A natural language interface for visual analysis. In *Proceedings of the 29th Annual Symposium on User Interface Software and Technology*, pp. 365–377, 2016.
- [52] V. Setlur, M. Tory, and A. Djalali. Inferencing underspecified natural language utterances in visual analysis. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*, pp. 40–51, 2019.
- [53] 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. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–20, 2022.
- [54] L. Shen, E. Shen, Z. Tai, Y. Song, and J. Wang. TaskVis: Task-oriented Visualization Recommendation. In *Proceedings of the 23th Eurographics Conference on Visualization (Short Papers)*, EuroVis’21, pp. 91–95. Eurographics, 2021. doi: 10.2312/evs.20211061
- [55] S. Srinivasa Ragavan, Z. Hou, Y. Wang, A. D. Gordon, H. Zhang, and D. Zhang. GridBook: Natural Language Formulas for the Spreadsheet Grid. In *Proceedings of the 27th International Conference on Intelligent User Interfaces*, pp. 345–368. ACM, 2022.
- [56] A. Srinivasan, M. Dontcheva, E. Adar, and S. Walker. Discovering natural language commands in multimodal interfaces. In *Proceedings of the 24th International Conference on Intelligent User Interfaces*, pp. 661–672, 2019.
- [57] 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 the 2020 CHI Conference on Human Factors in Computing Systems*, pp. 1–13, 2020.
- [58] A. Srinivasan, B. Lee, and J. T. Stasko. Interweaving multimodal interaction with flexible unit visualizations for data exploration. *IEEE Transactions on Visualization and Computer Graphics*, 2020.
- [59] A. Srinivasan, N. Nyopathy, B. Lee, S. M. Drucker, and J. Stasko. Collecting and characterizing natural language utterances for specifying data visualizations. 2021.
- [60] A. Srinivasan and V. Setlur. Snowy: Recommending Utterances for Conversational Visual Analysis. In *Proceedings of the 34th Annual ACM Symposium on User Interface Software and Technology*, pp. 1–17. ACM, 2021.
- [61] A. Srinivasan and J. Stasko. Natural language interfaces for data analysis with visualization: Considering what has and could be asked. In *Proceedings of the Eurographics/IEEE VGTC Conference on Visualization: Short Papers*, pp. 55–59. Eurographics Association, 2017.
- [62] A. Srinivasan and J. Stasko. Orko: Facilitating multimodal interaction for visual exploration and analysis of networks. *IEEE transactions on visualization and computer graphics*, 24(1):511–521, 2017.
- [63] C. Stolte and P. Hanrahan. Polaris: A system for query, analysis and visualization of multi-dimensional relational databases. In *INFOVIS*, 2000.
- [64] ThoughtSpot. <http://www.thoughtspot.com/>, May 2020.
- [65] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- [66] Y. Wang, H. Zhang, H. Huang, X. Chen, Q. Yin, Z. Hou, D. Zhang, Q. Luo, and H. Qu. Infonice: Easy creation of information graphics. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, CHI ’18, pp. 335:1–335:12. ACM, New York, NY, USA, 2018.
- [67] Z. Wen, M. X. Zhou, and V. Aggarwal. An optimization-based approach to dynamic visual context management. In *IEEE Symposium on Information Visualization, 2005. INFOVIS 2005.*, pp. 187–194. IEEE, 2005.
- [68] K. Wongsuphasawat, Z. Qu, D. Moritz, R. Chang, F. Ouk, A. Anand, J. Mackinlay, B. Howe, and J. Heer. Voyager 2: Augmenting visual analysis with partial view specifications. In *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*, pp. 2648–2659. ACM, 2017.
- [69] A. Wu, L. Xie, B. Lee, Y. Wang, W. Cui, and H. Qu. Learning to Automate Chart Layout Configurations Using Crowdsourced Paired Comparison. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI’21, pp. 1–13. ACM, 2021.
- [70] H. Xia, N. Henry Riche, F. Chevalier, B. De Araujo, and D. Wigdor. Dataink: Direct and creative data-oriented drawing. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, p. 223. ACM, 2018.
- [71] P. Xu and R. Sarikaya. Convolutional neural network based triangular crf for joint intent detection and slot filling. In *2013 IEEE Workshop on Automatic Speech Recognition and Understanding*, pp. 78–83. IEEE, 2013.
- [72] B. Yu and C. T. Silva. Flowsense: A natural language interface for visual data exploration within a dataflow system. *IEEE transactions on visualization and computer graphics*, 26(1):1–11, 2019.
- [73] X. Zhang and H. Wang. A joint model of intent determination and slot filling for spoken language understanding. In *IJCAI*, vol. 16, pp. 2993–2999, 2016.
- [74] V. Zhong, C. Xiong, and R. Socher. Seq2sql: Generating structured queries from natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103*, 2017.
- [75] F. Özcan, A. Quamar, J. Sen, C. Lei, and V. Efthymiou. State of the Art and Open Challenges in Natural Language Interfaces to Data. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, SIGMOD’20, pp. 2629–2636. ACM, jun 2020.