A Conversational Data Visualisation Platform for Hierarchical Multivariate Data

E. Kavaz, I. Rodríguez, A. Puig and E. Vives

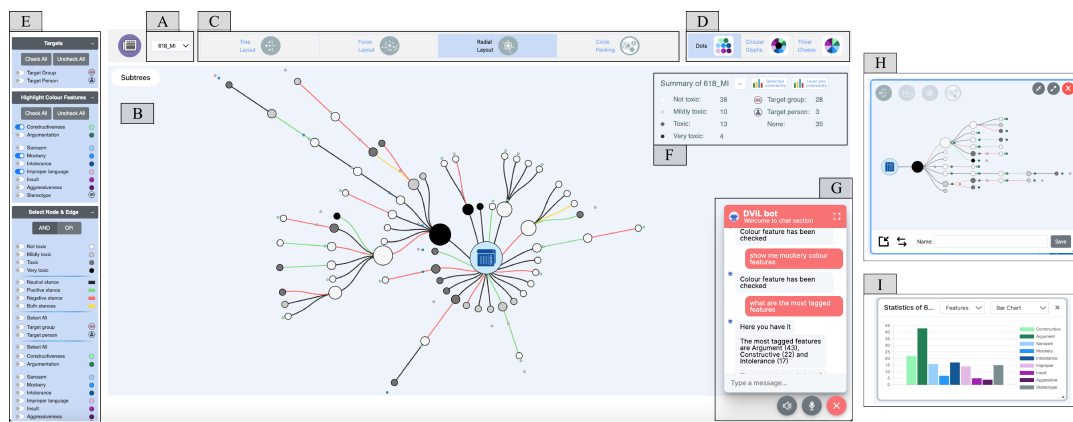Departament de Matemàtiques i Informàtica, Universitat de Barcelona, Spain

Figure 1: *Overview of the platform (A) Datasets menu (B) Main visualisation: hate speech annotation of online news and generated comments, (C) Hierarchical layouts, (D) Glyphs, (E) Filters, (F) Summary of statistics and supplementary charts, (G) Chatbot UI, (H) Pop-up window displaying a region of the dataset, and (I) Supplementary chart illustrating distribution of colour features.*

Abstract
This paper presents a novel data visualisation platform that integrates both direct manipulation and conversational interaction styles for analysing hierarchical multivariate data. The proposed architecture is based on the Rasa conversational AI framework. We show its full potential in a real-life case study for analysing hate speech in online news.

CCS Concepts
*• **Human-centered computing** → **Visualization; Visualization toolkits;** • **Computing methodologies** → **Natural language processing;***

1. Introduction

Traditional data visualisation systems have been primarily designed using direct manipulation interfaces (aka WIMP-Windows Icons Mouse Pointer). Nevertheless, the cluttered design of these interfaces requires from the users significant time and effort to take full profit of them to analyse complex data [LIRC12]. Recent advances in NLP (Natural Language Processing) allow to enhance significantly WIMP interfaces making the visual analysis a seamless user experience through Natural Language conversations [SSL*21]. In this context, we propose a conversational visualisation platform to support the users in analysing hierarchical multivariate data. Our platform includes different visualisation layouts for hierarchical

data (tree, radial, force, circle packing), geometrical elements representing data (dots, points, glyphs, etc.), and complementary basic charts (i.e. histograms, pie, ...). Interacting with both the chatbot and the WIMP, users can filter data, highlight parts of the hierarchy, change visual mapping, and enable/disable glyphs, between others.

In addition, our chatbot - based on Rasa conversational AI [Ras23]- allows the users to perform high-level queries that are not possible using only direct manipulation. Also, it automatically decides visual layouts according to the distribution of nodes resulting from a query [KPR*23].

Although our research can be applied to different scenarios, here

we present the analysis of hate speech in in the comments section of online news. Linguists annotated the toxicity of comments by labelling them with various hate features (e.g., constructiveness, argumentation, intolerance, sarcasm, insult) [TAN*21]. Since comments can generate additional comments, news articles become hierarchical structures with multiple levels. The initial level (level 0) consists of direct comments to the news article, while subsequent levels comprise comments to other comments.

2. Visualisation-oriented Natural Language Interface

Figure 1 overviews the GUI, the multi-views and the chatbot included in our data visualisation platform. Note that the visualisations contained in the blocks B, H and I show different aspects of the data. The chatbot is able to generate all of them as response to user's intents. Specifically, the chatbot, based on the structure of the hierarchy to be displayed, decides automatically the most adequate visual mapping in both the main and the partial visualisation. For example, a radial layout fits better a hierarchy with a root node having lots of children (case of the dataset in block B), and a tree layout is adequate for structures with a root having few but long children (case of visualisation in block H).

Moreover, the chatbot is designed to guide the users during the visual analysis. In its first interaction, the chatbot greets the users and tells them how it can help: i) generating/updating the visualisation in response to users' queries (e.g. *"Please, highlight constructive and argumentative direct comments to the news"*), ii) performing usual operations in the platform (login, logout, open dataset), iii) solving users' doubts about elements of the interface (*"What are the glyphs for?"*), and iv) explaining dataset-related things (*"What are the levels of toxicity?"*). When the chatbot is not able to understand the query, either it shows disambiguation widgets or it prompts the user to rephrase the intent. Additionally, the chatbot provides the users with textual feedback (e.g. *"Here you have the largest thread shown in a new pop-up window."*), and visual feedback since it blinks updated elements of the interface. Another aspect of the chatbot design is the type of supported queries. Low level queries are simple interactions that usually consist of a one turn user-bot exchanges, some of these queries can also be solved by interacting directly with the WIMP. On the other hand, high level queries are more complex ones that can not be performed in the WIMP and usually require several turns (aka follow-up queries). For example, complex queries such as *"Show me the most toxic thread"*, or *"Is the largest thread the most toxic one?"*).

A keypoint of our multimodal platform is that the chatbot and the WIMP interactions are fully synchronised. That is, the action of the chatbot in response to the user query not only produce a new visualisation but also updates the buttons/menus in the interface. For instance, when user writes *"Show me only sarcastic comments"* in the window G in Figure 1, the checkbox *"Sarcasm"* in box E in Figure 1 is checked. And in reverse, the chatbot is aware when the user selects some filters in the WIMP, to be able to follow the conversation about the selected things. For example, when the user asks to the chatbot *"Deselect it"*, the chatbot is able to deactivate the last checked feature. Find a demo of the conversational data visualisation platform in here.

3. Architecture Overview

Our platform is a web-based application as shown in Figure 2 divided in two main modules (Visualisation and the Conversational AI). The Visualisation module is developed using HTML/CSS, JavaScript, D3.js, and jQuery on the frontend, and Python and a PostgreSQL database on the backend. The Conversational AI module is based on the open-source Rasa AI chatbot software.

The Conversational AI module includes two main components: the Rasa Server and the Rasa Action Server. The former is in charge of understanding user's queries (Rasa NLU) and figure out the response (Rasa Core). The Rasa NLU uses DIET (Dual Intent and Entity transformer) which is a multi-task transformer architecture that handles both intent classification and entity recognition. Moreover, the Rasa Core uses TED (Transformer Embedding Dialogue) policies to predict the next action to be performed. When the next action is not a direct response, Rasa Action Server uses custom actions to do operations such as recover information from the conversation tracker, compose responses and perform database queries.

There are three channels to communicate both modules: 1) a websocket to connect the chatbot frontend with the Agent Rasa Server, 2) a rest API to access from the Rasa Action Server to the backend, and 3) a Rasa http API that connect the visualisation frontend with the Rasa Server.
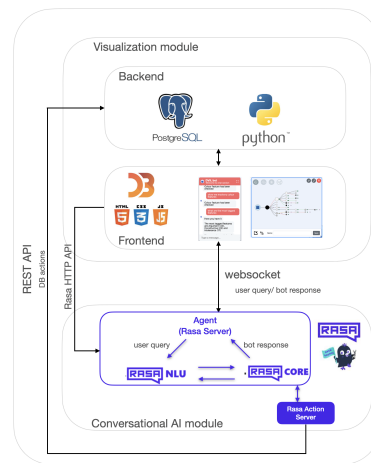

Figure 2: *Visualisation and Conversational AI modules.*

4. Conclusions and Future Work

We have presented a novel data visualisation platform that combines WIMP and natural language interactions to analyse hierarchical multivariate data. We aim to improve our training data, integrate more high-level queries, conduct a user evaluation.

Acknowledgements

References

[KPR*23] KAVAZ E., PUIG A., RODRÍGUEZ I., CHACÓN R., DE-LA-PAZ D., TORRALBA A., NOFRE M., TAULE M.: Visualisation of hierarchical multivariate data: Categorisation and case study on hate speech. *Information Visualization 22*, 1 (2023), 31–51. 1

[LIRC12] LEE B., ISENBERG P., RICHE N. H., CARPENDALE S.: Beyond mouse and keyboard: Expanding design considerations for information visualization interactions. *IEEE Transactions on Visualization and Computer Graphics 18*, 12 (2012), 2689–2698. 1

[Ras23] RASA: Rasa conversational platform, 2023. URL: https://rasa.com/. 1

[SSL*21] SHEN L., SHEN E., LUO Y., YANG X., HU X., ZHANG X., TAI Z., WANG J.: Towards natural language interfaces for data visualization: A survey. *arXiv preprint arXiv:2109.03506* (2021). 1

[TAN*21] TAULÉ M., ARIZA A., NOFRE M., AMIGÓ E., ROSSO P.: Overview of detoxis at iberlef 2021: detection of toxicity in comments in spanish. *Procesamiento del lenguaje natural 67* (2021), 209–221. 2