



City Research Online

City St George's, University of London

Citation: Turkey, C. & Henkin, R. (2018). Towards Natural Language Empowered Interactive Data Analysis. Poster presented at the EuroVis 2018 20th EG/VGTC Conference on Visualization, 04 - 08 June 2018, Brno, Czech Republic. doi: 10.2312/eurp20181128

This is the accepted version of the paper.

This version of the publication may differ from the final published version. To cite this item please consult the publisher's version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/20110/>

Link to published version: <https://doi.org/10.2312/eurp20181128>

Copyright and Reuse: Copyright and Moral Rights remain with the author(s) and/or copyright holders. Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge, unless otherwise indicated, provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way. For full details of reuse please refer to [City Research Online policy](#).

Towards Natural Language Empowered Interactive Data Analysis

C. Turkey¹ and R. Henkin¹

¹giCentre, Department of Computer Science, City, University of London, UK

Abstract

The recent advances in natural language based interaction methodologies offer promising avenues to enhance the interactive processes within the human-machine dialogue of visual analytics. We envisage Multimodal Data Analytics as a novel approach for conducting data analysis that builds on the strengths of visual analytics and natural language as an expressive interaction channel. We investigate the potential enhancements from such a multimodal approach and discusses the preliminary outline for a structured methodology to study the role of natural language in data analytics. Our approach builds on a simple model of human machine dialogue for interactive data analysis which we then propose to instantiate as visual analytics workflows – representations to study and operationalise interactive data analysis routines empowered by natural language interaction.

CCS Concepts

•Human-centered computing → Visual analytics;

1. Introduction

The success of visual analytics systems relies on an iterative human-computer dialogue where the knowledge and the creativity of the human, and the power of computation operate together to deal with the complexities in the data sources [SSS*14, ALA*]. In such human-in-the-loop data analysis approaches, interactive visualisation methods are core facilitators of this dialogue [TKBH17a] and have proven their merit in several ways. However, despite the wide literature and established products in visualisation [ERT*17], such iterative visual analytics processes mostly rely on conventional, not often intuitive interaction mechanisms that can introduce unnecessary complexities into the process [Lam08].

The recent advances in natural language based interaction methodologies offer promising avenues to rethink the interaction processes within visual analytics [BHCDV99] and natural language is getting increasing interest from the visual analytics community as a new channel to conduct and support visual analysis [HSTD18, SS17a, SS17b]. We envisage *Multimodal Data Analytics* as a novel potential approach for conducting data analysis that combines the strengths of visual analytics techniques and natural language – as an expressive channel – for interaction. In this poster, we discuss preliminary outcomes from an ongoing 18 month long research project and present initial considerations towards this analysis approach, and present the preliminary attempts to build a structured approach in exploring the open questions within this emerging approach to data analysis. In the context of this work, we consider settings where a computational method is embedded within an interactive visual analytics framework, with the algorithm having a means to responding to user inputs.

2. Enhancing Data Analysis with Natural Language Interaction

This new paradigm of multimodal analytics offers a number of potential enhancements for interactive data analysis processes. In the following, we investigate them under a number of headings and then provide a few illustrative and speculative scenarios.

E1: Enhanced expressiveness for guidance: Active learning is an established technique through which users can provide examples of correct results to algorithms [CGJ96]. Users, however, often have more than only the correct results to provide to the algorithms [SRL*09] and natural language has the potential to facilitate users in providing guidance to algorithms through complex utterances that explain their way of thinking more effectively to reveal their tacit knowledge of the domain.

E2: Enhanced model steering: Natural language inputs can enable users to *steer* algorithms by targeting these inputs in the data, in a mixed role of enriching and guiding the system. A key aspect for this is presenting users with incomplete results and enabling the algorithm to be steered [TKBH17b].

E3: Enhanced comprehension and transparency: Expressiveness of the natural language opens up new possibilities to improve the communication of model results and the execution logic of algorithms. Greater transparency has been shown to lead to increased satisfaction with results [KSBK12] with also increased trust in computational models. Natural language offers an effective channel in this regard to embellish and enrich visual representations, and draw attention to uncertainties.

E4: Enhanced knowledge elicitation: A more natural flow of the




Figure 1: A two-way dialogue model where a number of interaction facets are listed as points where useful interactions can occur between the user and an algorithm.

interaction dialogue can foster more effective knowledge elicitation. In particular in cases where the analytics system supports several workflows, language provides an effective channel to learn and adapt to users' tasks and expectations within an analysis session.

Speculative Examples:

Ex-1: In response to a projection of high-dimensional points, an utterance "... the observations on the left do not share the same income characteristics, so this looks wrong ..." from the user could be translated into updates of the weights of income related variables in the computation (**E1**, **E2**) and further guidance can be gathered through a question "... there seems to be redundancy within income related features as seen above (referring to a new visualisation of income related features), consider excluding some? ... (**E1**, **E4**).

Ex-2: In the presentation of clustering results of the same points and following the guidance from the user above, a visual representation of the clusters could be supported by an utterance of the system – ... here, more importance is given to income, so the two groups differ significantly in their income characteristics but notice now the weaker discrimination overall ... – not only enhances comprehension but also highlights the issues in the results (**E3**).

3. Investigating the Role for Natural Language Interaction

In order to investigate the role of natural language interaction, we firstly investigate the theoretical models of visual analytics [ERT*17, ALA*] and also in other disciplines such as intelligent systems [MKF*15] and robotics [ACKK14]. Building on these works, a user can: **R1 - Adjust/enrich** by modifying the *inputs* used by algorithm, such as the data items or classes used in classification, **R2 - Assist** by mimicking the algorithm through correcting the *outputs* to be fed back to the system, **R3 - Guide**: by modifying the *data features* used by algorithms, **R4 - Tune** by modifying the *parameters* of algorithms, such as the number of clusters.

In addition to these roles, we also consider a number of 'interaction facets' where the interaction between the user and the computer could conceptually take place over (see Figure 1). Here, we list six interaction facets as points where useful interactions can occur between the user and an algorithm: Data Inputs, Algorithmic Outputs, Data Features, Algorithmic Parameters, Algorithmic Logic, Progress State of the algorithm.

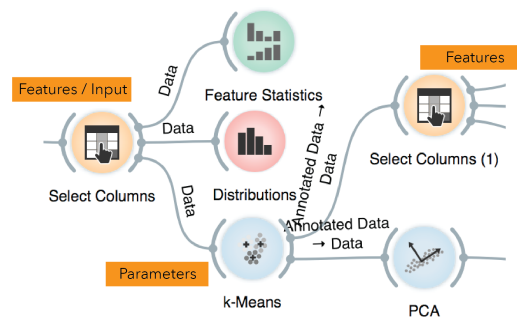


Figure 2: An analytical workflow illustration (built here using the Orange software <https://orange.biolab.si/>) that serves as an instantiation of the high-level communication model presented in Figure 1 for a clustering algorithm with facets indicated.

3.1. Analytical Workflows

The above general communication model provides the high-level framework to study in which of the *interaction facets*, natural language could benefit most in coordination with visualisation methods. As instantiations of this high-level framework, we are developing analytical workflows for a wide range of algorithmic methods, identify the interaction facets within these interactive processes, and use these systematic approaches both for empirically studying the benefits of using different modalities and also as analysis templates for potential users of such approaches. One example of this is the case of using the k-means algorithm to cluster a high-dimensional data set. For such an analytical task, we build analytical workflows (Figure 2) where a number of the interaction facets are explicitly denoted. In Figure 2, for instance, *Selecting Columns* for k-means (**R3**) to consider is an instance of the interaction facet *Features* that can be captured in a similar fashion as **Ex-1**, where as the *Selection of k value* is an instance of the interaction facet *Parameters* which is of role type **R4**. Future work here is to cover a wide range of algorithms [ERT*17] and provide structured workflows where most of the interaction facets are explicitly denoted.

4. Conclusion

With the increasing adoption of complex algorithmic approaches in data analysis, methods to enable analysts to effectively utilise algorithms and action on their results are of critical importance. In this poster, we present initial discussions and the sketch of a systematic way to approach the study of multimodal analytics. To advance the field, further research is needed to: gather empirical information on the language used by analysts both at a generic and domain specific level, analyse the effectiveness of mediums and techniques for the different interaction facets, and develop a wide range of algorithm agnostic visual analytics workflows.

Acknowledgements

This work is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) [nlvis: *Natural Language Interaction for Visual Data Analysis*, grant number: EP/P025501/1]

References

- [ACKK14] AMERSHI S., CAKMAK M., KNOX W. B., KULESZA T.: Power to the People: The Role of Humans in Interactive Machine Learning. *AI Magazine* 35, 4 (dec 2014), 105. 2
- [ALA*] ANDRIENKO N., LAMMARSCH T., ANDRIENKO G., FUCHS G., KEIM D., MIKSCH S., RIND A.: Viewing visual analytics as model building. *Computer Graphics Forum* 0, 0. 1, 2
- [BHCDV99] BRADEN-HARDER L., CORSTON S. H., DOLAN W. B., VANDERWENDE L. H.: Apparatus and methods for an information retrieval system that employs natural language processing of search results to improve overall precision, Aug. 3 1999. US Patent 5,933,822. 1
- [CGJ96] COHN D. A., GHAHRAMANI Z., JORDAN M. I.: Active learning with statistical models. *Journal of artificial intelligence research* (1996). 1
- [ERT*17] ENDERT A., RIBARSKY W., TURKAY C., WONG B., NABNEY I., BLANCO I. D., ROSSI F.: The state of the art in integrating machine learning into visual analytics. In *Computer Graphics Forum* (2017), vol. 36, Wiley Online Library, pp. 458–486. 1, 2
- [HSTD18] HOQUE E., SETLUR V., TORY M., DYKEMAN I.: Applying Pragmatics Principles for Interaction with Visual Analytics. *IEEE Transactions on Visualization and Computer Graphics* 24, 1 (jan 2018), 309–318. 1
- [KSBK12] KULESZA T., STUMPF S., BURNETT M., KWAN I.: Tell me more? The Effects of Mental Model Soundness on Personalizing an Intelligent Agent. *Proceedings of the 2012 ACM annual conference on Human Factors in Computing Systems - CHI '12* (2012), 1. 1
- [Lam08] LAM H.: A framework of interaction costs in information visualization. *IEEE transactions on visualization and computer graphics* 14, 6 (2008). 1
- [MKF*15] MEIGNAN D., KNUST S., FRAYRET J.-M., PESANT G., GAUD N.: A Review and Taxonomy of Interactive Optimization Methods in Operations Research. *ACM Transactions on Interactive Intelligent Systems* 5, 3 (2015), 1–43. 2
- [RD10] ROSENTHAL S. L., DEY A. K.: Towards maximizing the accuracy of human-labeled sensor data. *International Conference on Intelligent User Interfaces, Proceedings IUI* (2010), 259–268.
- [SRL*09] STUMPF S., RAJARAM V., LI L., WONG W. K., BURNETT M., DIETTERICH T., SULLIVAN E., HERLOCKER J.: Interacting meaningfully with machine learning systems: Three experiments. *International Journal of Human Computer Studies* 67, 8 (2009), 639–662. 1
- [SS17a] SRINIVASAN A., STASKO J.: Natural Language Interfaces for Data Analysis with Visualization: Considering What Has and Could Be Asked. In *EuroVis 2017 - Short Papers* (2017), Kozlikova B., Schreck T., Wischgoll T., (Eds.), The Eurographics Association. 1
- [SS17b] SRINIVASAN A., STASKO J.: Orko : Facilitating Multimodal Interaction for Visual Exploration and Analysis of Networks. *InfoVis TVCG* 24, SI (2017), 11. 1
- [SSS*14] SACHA D., STOFFEL A., STOFFEL F., KWON B. C., ELLIS G., KEIM D. A.: Knowledge generation model for visual analytics. *IEEE transactions on visualization and computer graphics* 20, 12 (2014), 1604–1613. 1
- [SWM17] SAMEK W., WIEGAND T., MÜLLER K.-R.: Explainable Artificial Intelligence: Understanding, Visualizing and Interpreting Deep Learning Models.
- [TKBH17a] TURKAY C., KAYA E., BALCISOY S., HAUSER H.: Designing progressive and interactive analytics processes for high-dimensional data analysis. *IEEE transactions on visualization and computer graphics* 23, 1 (2017), 131–140. 1
- [TKBH17b] TURKAY C., KAYA E., BALCISOY S., HAUSER H.: Designing Progressive and Interactive Analytics Processes for High-Dimensional Data Analysis. *IEEE Transactions on Visualization and Computer Graphics* 23, 1 (jan 2017), 131–140. 1