

Chapter 11

Articulate: Creating Meaningful Visualizations from Natural Language

Yiwen Sun

University of Illinois at Chicago, USA

Jason Leigh

University of Illinois at Chicago, USA

Andrew Johnson

University of Illinois at Chicago, USA

Barbara Di Eugenio

University of Illinois at Chicago, USA

ABSTRACT

This chapter presents an approach to enable non-visualization experts to craft advanced visualizations through the use of natural language as the primary interface. The main challenge in this research is in determining how to translate imprecise verbal queries into effective visualizations. To demonstrate the viability of the concept, the authors developed and evaluated a prototype, Articulate, which allows users to simply ask the computer for questions about their data and have it automatically generate visualizations that answer these questions. The authors discovered that by relieving the user of the burden of learning how to use a complex interface, they enable them to focus on articulating better scientific questions and wasting less time in producing unintended visualizations.

INTRODUCTION

Nearly one third of the human brain is devoted to processing visual information. Vision is the dominant sense for the acquisition of information from our everyday world. It is therefore no surprise that visualization, even in its simplest form,

remains the most effective means for converting large volumes of raw data into insight. Over the past three decades, much has been investigated in the design of sophisticated visualization tools in a variety of disciplines. However, the effort end-users have to make to craft a meaningful visualization using these tools has been mostly

DOI: 10.4018/978-1-4666-4309-3.ch011

overlooked. The users of such tools are usually lay people or domain experts with marginal knowledge of visualization techniques. When exploring data, they typically know what questions they want to ask, but often do not know, or do not have the time to learn, how to express these questions into a series of interactions that produce an effective visualization.

A 2008 National Science Foundation report “Enabling Science Discoveries through Visual Exploration” (Ebert, 2008) also noted that one of the main barriers hindering the adoption of advanced visualization tools is the steep learning curve associated with them. 2010 findings by Grammel (2010) showed that novices to Information Visualization still tended to use traditional bar, line and pie charts over other chart types by more than 70% because of their familiarity with them. Modern visualization tools offer such an expansive array of capabilities that they can only be wielded by an expert trained in visualization. In some ways it is like expecting someone to know how to build a house by simply sending them to Home Depot.

Meanwhile, the 2008 NSF report noted “there is a strong desire for conversational interfaces that facilitate a more natural means of interacting with science.” In other words, scientists “simply” want to tell the computer what they want to see and have it just create it. They do not want to have to become visualization experts. Even a decade ago this would have seemed far-fetched, but today we are seeing renewed interest in the use of natural language as an interface to computing. For example, survey results according to search engines like Ask.com show that a third of search queries are entered as natural language questions rather than keywords. Siri, the intelligent personal assistant on iPhone 4S, allows users to send messages, schedule meetings, and place phone calls by directly speaking into their smartphones. The field of natural language processing has made great strides in the last decades, with a variety of

models that are able to understand the meaning of sentences in recommender systems, educational technology and health applications.

This inspired us to consider the use of a conversational interface for the automatic generation of visualizations. A system such as this would allow an end-user to pose natural language inquiries, and then let the system assume the burden of creating the most appropriate visual representation of the inquiry. It is hoped that such a capability can potentially reduce the learning curve necessary for effective use of visualization tools, and thereby expand the population of users who can successfully conduct visual analysis. Note however that in this work we are not simply translating explicit visualization commands such as “*make a plot of temperature vs pressure*”- though it is certainly possible within the context of this research. Instead the expectation is that our approach will enable a user to ask deeper questions about data, such as “*what is the correlation between temperature and depth with temperature below zero*”, without having to follow or memorize a strict grammar or command structure, as has been in the past. Furthermore users will be able to ask follow-up questions where the computer has some knowledge of what has already been asked and visualized. Therefore the fundamental value of this approach is that it enables the end-users to focus on the scientific question they are trying to ask rather than the specific visualization task they are trying to perform.

The remainder of this chapter is organized as follows. We first describe prior work in related fields. Then we explain in detail our methodology for translating conversation into a precise visualization. Next, we present details of our user studies. While the initial prototype produces information visualizations, we will also explain and show through a case study how the approach is conceptually extensible to scientific visualizations as well. Lastly we outline directions for future work.

16 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/articulate/78721

Related Content

Big Data Analytics and IoT in Smart City Applications

Mamata Rath (2022). *Research Anthology on Big Data Analytics, Architectures, and Applications* (pp. 364-380).

www.irma-international.org/chapter/big-data-analytics-and-iot-in-smart-city-applications/290991

A Predictive Analytics Framework for Blood Donor Classification

Kavita Pabrejaand Akanksha Bhasin (2021). *International Journal of Big Data and Analytics in Healthcare* (pp. 1-14).

www.irma-international.org/article/a-predictive-analytics-framework-for-blood-donor-classification/277644

A Brief Survey on Big Data in Healthcare

Ebru Aydindag Bayrakand Pinar Kirci (2020). *International Journal of Big Data and Analytics in Healthcare* (pp. 1-18).

www.irma-international.org/article/a-brief-survey-on-big-data-in-healthcare/253842

I'd Like to Teach the World to Think: Mind Genomics, Big Mind, and Encouraging Youth

Howard Moskowitz, Jyotsna, Andi Sciaccaand Andrew Lester (2018). *Harnessing Human Capital Analytics for Competitive Advantage* (pp. 55-90).

www.irma-international.org/chapter/id-like-to-teach-the-world-to-think/199991

GPS Travel Diaries in Rural Transportation Research: A Focus on Older Drivers

Trevor Hansonand Eric Hildebrand (2018). *Intelligent Transportation and Planning: Breakthroughs in Research and Practice* (pp. 609-625).

www.irma-international.org/chapter/gps-travel-diaries-in-rural-transportation-research/197153