
Towards Text Search for Information Visualization Retrieval

Romain Vuillemot

Harvard University
79 JFK Street Cambridge, MA
United States
Romain.Vuillemot@hks.harvard.edu

Mehmet Akmanalp

Harvard University
79 JFK Street Cambridge, MA
United States
Mali.Akmanalp@hks.harvard.edu

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

Copyright is held by the owner/author(s).
CHI'15 Extended Abstracts, Apr 18-23, 2015, Seoul, Republic of Korea
ACM 978-1-4503-3146-3/15/04.
<http://dx.doi.org/10.1145/2702613.2732753>

Abstract

We investigate the use of text search to retrieve information visualizations. This is important as the body of available visualizations on the web is growing. And in the meantime, they are difficult to find as they don't have an immediate textual description. Our approach is first to create a taxonomy of textual terms to describe visualizations, both in a general and in a specific way using economics visualization as application domain. Then we designed and implemented a search engine to query this vocabulary on a real website, The Atlas of Economic Complexity, using economics data. Results from an exploratory study informed us on the types of tasks to support and on the visual design of the search box, such as showing recommendations similar to Google results. Those early and promising results pave the way for a more diverse and complex vocabulary that exploits the full wealth of information that graphical elements can contain. Such a search feature has wide ranging applicability, from making visualizations more accessible for less technical users, to using other input modalities such as voice.

Author Keywords

Information Visualization, Information Retrieval.



Figure 1: Example of Treemap visualization from The Atlas of Economic Complexity. Each visualization is described with a title (in this case *"What did Germany export in 2012?"*) that we use as vocabulary for the search feature.

ACM Classification Keywords

H.5.m [Information interfaces and presentation (e.g., HCI)]: Miscellaneous.

INTRODUCTION

Information visualizations are becoming increasingly popular, and the Web makes them theoretically accessible to everyone who is online. However, finding a specific visualization remains difficult. While users are just a handful of keywords away from historical visualizations, e. g., Napoleon's March to Moscow by Minard, visualizations that are not one of a kind are more complex to find. Searching for popular systems such as Tableau, Many Eyes or toolkits like D3 may return interesting user-created examples, but it casts a net too broad to find a specific visualization. To match the current way users find information online, one would search for visualizations using queries such as *"Change of child mortality across time for Africa"*, and once the result is found, text search would also allow to refine the original query to change the visualization. So far, such a query eventually leads to visualizations indexed by image search engines, where keywords are from text surrounding visualizations, and not from their labels or data sources. The reason for this is because visualizations are primarily composed of graphical elements, which are not understood by traditional search engines. Prior research focusing on textual description of visualizations focused on their construction [7], finding semantic icon sets [8] or adding contextual annotations [6]. But as far as we know, no attempt has been made to apply textual description to index and retrieve visualizations.

Text search falls in the category of text-based interfaces which have numerous advantages. They empower

experts with a mechanism to quickly access specific commands using a keyboard. They allow users to browse a large and deep hierarchy of terms and suggestions, and they can dynamically assist them when typing through auto-complete. Such interfaces are now ubiquitous in our environment, not only on the Web but also on the desktop, such as Apple Spotlight [2]. However, text commands are not easily discoverable and difficult to learn and remember for novice users [1]. The main alternative to text-based interfaces are widget-based interfaces, e. g., interfaces using buttons, check boxes or sliders. Widgets have the benefit of being visible and discoverable commands, allowing for trial and error, but are limited when building and repeating complex queries. To that end, [9] concludes that the two paradigms of text-based interfaces and widgets are complementary and not mutually exclusive.

A VOCABULARY FOR VISUALIZATION

Our first step was to focus on creating the textual descriptions for visualization, to be described and retrieved using text-commands.

Methodology

We first used our prior knowledge regarding information visualization to textually describe charts, and then we collaborated with practitioners in macroeconomics to broaden and refine this vocabulary using a specific website containing standard visualizations. This website is The Atlas of Economic Complexity [5], which is publicly available at <http://atlas.cid.harvard.edu/>. It implements interactive standard visualizations like tree-maps (Figure 1), stacked graphs, geo maps, scatterplots and node-link graphs for macroeconomics data to answer elementary level questions [10, 3] such as *What did Germany Export to Italy?* or *Who Exported*

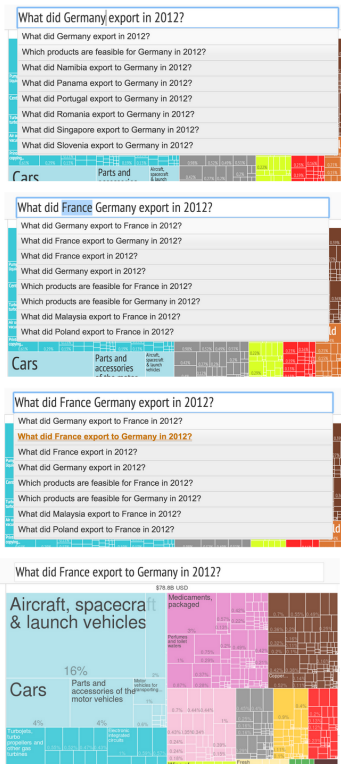


Figure 2: Workflow for the user (from top to bottom): the user clicks on the visualization's title which shows a recommendation of similar visualizations; then the user adds a new keyword ("France") which updates the list of recommendations; the user selects one of the results; finally, a new page showing the recommended visualization appears.

Cars in 2012? which are visible as titles. It can be considered among the state-of-the-art systems for this specific application domain.

General Vocabulary

Figure 3 shows the categories of textual description of graphical elements we identified from our experience designing visualizations. The **data model** contains the major choices regarding the data being used, and the **chart** is composed of several layers and elements, as described by Bertin [3]. **Interactions with the view** are the main changes (e. g., highlight, zoom and details) that do not require additional data to be loaded. We also identified textual descriptions outside the chart itself, with **metadata** and textual context from the **document** embedding the visualization.

Use of Vocabulary by Domain Experts

We collected and reviewed textual material from our experts classes and seminars using The Atlas (e. g., hand outs, Power Points, white-board notes) and had informal discussions with them. The materials were issued for a broad audience ranging from students to groups of executives for the past two years and written by about 20 different researchers.

Our main observation was the heavy use of the visible text descriptions related to **metadata** surrounding visualizations to describe them. In particular, visualizations' **titles** are probably the most crucial reference as they are concise and the most visually salient text element on the website. Such titles are automatically generated by permuting interrogative pronoun, country, trade flow and year. This extensive reliance on titles is surprising as 1) some pages may have the same title and show different visualizations, 2) titles do not capture the complexity of the visualization's

parameters. Nonetheless, we considered those titles as the main vocabulary experts refer to, in order to find a specific visualization and will be the core of the vocabulary to retrieve visualizations.

PROTOTYPE IMPLEMENTATION

We built a search engine prototype leveraging both the vocabulary from experts, as well as the general one we identified. We used an iterative approach which consisted by first generating all the **Titles** only, and storing them in a database (Figure 4). We disambiguated them by including more vocabulary such as **Template** as keywords (e. g., using *scatterplot* on figure 6). We also allowed **Visual Mapping** interactions to reach specific views, such as specifying a custom Y-axis for a chart. We allowed congruency for keywords to query countries by continent, products by alternative names (often simpler), products by categories (e. g., *Cars* are in the *Machinery* category) and both countries and products by shortcut (e. g., USA for United States as on figure 5). Regarding the user interface (UI) design, we implemented a Google Suggestion-like approach which triggers, after a 3 character minimum input, a ranked list of 8 suggestions. The UI highlights the matching input keywords in the list of results and shows suggestions when clicked (Figure 2). It should theoretically be equivalent to use either widgets or the text search to interact with the visualization. We implemented the system using Elasticsearch (<http://www.elasticsearch.org/>) to provide fast error-tolerant (fuzzy) matching for titles. Elasticsearch's built-in ranking uses a mixture of TF-IDF and the vector space model that tends to work well for our needs. In terms of performance, for an index of 615,859 pages (94.1MB compressed), the average response time was below 15ms. The code is available as an open source

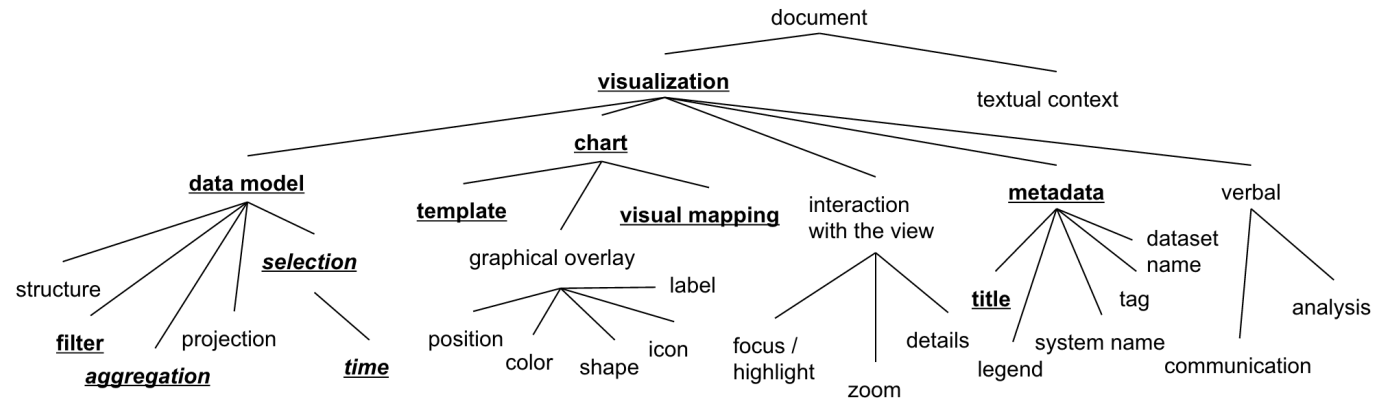


Figure 3: Example of categories of textual terms to describe a visualization and how they relate to each others. Our prototype focused on the underlined ones.

project (<https://github.com/cid-harvard/>) with a permissive license.

EXPLORATORY STUDY

We conducted an exploratory study of The Atlas augmented with the search feature prototype with experts not involved in the design process, but familiar with the dataset and visualizations. We provided a quick demo of the search feature to each of our four participants (2min); followed by an *adoption* phase where they explored the parameters of their choice (e.g. country); followed by a *reproduction* phase where they solved a problem they already experienced; and finally an *exploration* phase where they focused on a real-world question they have and not solved yet. We used a think aloud technique where participants gave a live commentary of their thoughts. Sessions lasted 30 minutes on average.

Participant #1 "I know my countries, but not the products"
 Our first expert frequently uses The Atlas. He typed in a country he recently worked on, *Colombia*, and observed the suggestions. He found them relevant, but would have liked them ranked by geographical locations to explore neighboring countries. He particularly appreciated being recommended results that are usually difficult to find through the traditional UI such as the *Product Feasibility* one. This visualization shows complex measures, and despite being very useful it is often missed when exploring a country. Then he said "*I usually know countries but usually have a hard time finding products*". He started picking up products by name, but struggled to find their names in the dataset. Thus, congruency was helpful as it provided a tolerant way to retrieve *Cocoa* by typing in *Chocolate*. He suggested using an even higher level of description or user-created sets of products for more natural queries.

France Italy Wine

What did France import from Italy in 2012?

What did France export to Italy in 2012?

What did Italy export to France in 2012?

What did Italy import from France in 2012?

Where did France export Wine of fresh grapes to in 2012?

Where did France export Wine lees, argol to in 2012?

Where did France import Wine of fresh grapes from in 2012?

Where did France import Wine lees, argol from in 2012?

Figure 4: Data-related query.

usa 8702

Where did United States import Motor vehicles for the transport of > 10 persons to in 2012?

Where did United States export Motor vehicles for the transport of > 10 persons to in 2012?

8702 export

Where did Cape Verde export Motor vehicles for the transport of > 10 persons to in 2012?

Where did Burundi export Motor vehicles for the transport of > 10 persons to in 2012?

Where did Guinea export Motor vehicles for the transport of > 10 persons to in 2012?

Where did Gambia export Motor vehicles for the transport of > 10 persons to in 2012?

Figure 5: Shortcuts-related query.

France scatterplot complexity

What did France import in 2012?

Which products are feasible for France in 2012?

What did France export in 2012?

Where did France import from in 2012?

Where did France export to in 2012?

What did Botswana import from France in 2012?

What did Cote d'Ivoire import from France in 2012?

What did Ethiopia import from France in 2012?

Figure 6: Specific chart-related query.

Participant #2 "It better supports my reporting process"

Our second expert was about to conduct a macroeconomics analysis on *Jamaica*, a country he doesn't know in particular but wanted to investigate the impact of quotas from the U.S. in the late 1990s. As he already used The Atlas for the same type of analysis, but with a different country, he did not particularly pay attention to the recommendations, and only found it useful to auto-complete his query. To navigate using the website, he carefully edited titles to switch the visualization, such as to change country or year. He also edited the URL a couple of times, as this his usual way of changing the page as opposed to using the UI. During the session he only operated on the widget to explore data, such as time ranges or filters. According to him, the feature is *"fantastic as it saved me plenty of time and did not require any learning effort"*.

Participant #3 "I want something like Google"

Our third expert is familiar enough with The Atlas that he started by saying that he was always struggling to find a specific visualization which was *"To which countries France exports cars"*. The parameter to reach this query is indeed hidden behind a tab. Finding this kind of visualization was immediate using the text search, even though his query did not perfectly match the title. Then he started exploring the range of the vocabulary by typing many queries, and immediately grasped everything he could do as he knew the structure with which titles were systematically generated. He attempted querying for several types of metals not in the list, to end with *"Who buys Copper?"*. This annoyed him a bit as he wanted *"something like Google"* which is tolerant and always provides him with some results, even if they might not be totally related.

Participant #4 "Search for the words I don't know"

Our last expert only used The Atlas a couple of times, beforehand. He started by exploring his home country and its connections with neighboring ones and browsed the standard visualizations. When investigating visualization on complex economic indicators, he typed those in the text field hoping for a definition and eventually some results. One of the terms used was *"Complexity"* to find something related to it, while not knowing this is an axis for a scatterplot that is on the website. He then looked up for its definition in the glossary and used several terms from this glossary hoping that the system would *"let me search for the words I don't know"*. He would have wanted the text search as a unique entry point to browse everything on the website, visual or not.

FINDINGS AND DESIGN IMPLICATIONS

The results from our study suggest commonalities in how experts interact with the text search and widgets. We enumerate them as recommendations for the design or improvement of this kind of feature.

- **Text search, in general, is familiar to users** as they have already been trained using it with major similar systems and may anticipate similar behavior.
- **Users are influenced by textual terms** from the charts. For example, titles, buttons or axis labels will be used in the text search and shape users' verbal communication.
- **Follow-up queries are often done using widgets** once a specific chart has been found, as users want to explore the view. We did not observe the use of relative terms (e.g., *"find previous year"*) to navigate between charts [7].

- **Suggesting relevant charts** theoretically allows better performance and satisfaction, but no consensus emerged from participants on how to rank results. The diversity of results was seen as a positive reminder of existing features.

CONCLUSION AND FUTURE WORK

We released the text search feature on The Atlas of Economic Complexity website and started recording user interactions to better understand users' behaviors from a quantitative perspective. Thus, despite considered as an hard problem [4], our main perspective is to further evaluate the current text search prototype but also against its variations, such as the ones using a larger inventory of terms and synonyms. We are also interested in extending text description to descriptive statistics (e. g., asking *What are the top-5 exporters of oil?* or *Who are the main trade partners of Europe?*) to gradually go towards intermediate and high level questions [10, 3] (while remaining within range of tractable natural language processing analysis). Such questions require, for instance, to automatically identify patterns or outliers, which are heavily domain dependent [6] and are often only detectable by users. On the long run we sought to observe how text search and visualizations will co-evolve and benefit from each other, e. g., how user queries may provide better guidelines for visualizations titles or labels creation.

References

[1] Ahlberg, C., Williamson, C., and Shneiderman, B. Dynamic queries for information exploration: An implementation and evaluation. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '92, ACM (New York, NY, USA, 1992), 619–626.

[2] Apple, M. A. C. *OS X spotlight*. 2009.

[3] BERTIN, J. *Semiologie graphique: les cartes, les diagrammes, les reseaux*. Pars-La Haia: Mouton-Gautier Villars,, 1973.

[4] Google. Search evaluation at google. <http://googleblog.blogspot.com/2008/09/search-evaluation-at-google.html> (2008).

[5] Hausmann, R., and Hidalgo, C. A. *The atlas of economic complexity: Mapping paths to prosperity*. MIT Press, 2014.

[6] Hullman, J., Diakopoulos, N., and Adar, E. Contextifier: Automatic generation of annotated stock visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, ACM (New York, NY, USA, 2013), 2707–2716.

[7] Metoyer, R., Lee, B., Henry Riche, N., and Czerwinski, M. Understanding the verbal language and structure of end-user descriptions of data visualizations. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, ACM (New York, NY, USA, 2012), 1659–1662.

[8] Setlur, V., and Mackinlay, J. D. Automatic generation of semantic icon encodings for visualizations. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, ACM (2014), 541–550.

[9] Shneiderman, B., and Plaisant, C. *Designing the User Interface: Strategies for Effective Human-Computer Interaction (4th Edition)*. Pearson Addison Wesley, 2004.

[10] Wainer, H. Understanding graphs and tables. *Educational Researcher* 21, 1 (1992), 14–23.