

A B A K Ó S

Instituto de Ciências Exatas e Informática



Licença Creative Commons Attribution 4.0 International

DocGenealogy – Visualizing the doctoral advisors and mentors genealogic tree*

DocGenealogy - Visualizando os conselheiros de doutorado e mentores da árvore genealógica

David Paiva Fernandes¹ Elizabeth Simão Carvalho²

Resumo

DocGenealogy é um projeto de mineração de dados que permite a visualização e a exploração do relacionamento de assessoria / orientação de doutorados por meio de um grafo interativo. A Wikipedia é conhecida por sua natureza colaborativa informal, com base em uma forma livre e não estrutural de produzir informações. Esse é um dos seus pontos fortes, na verdade, a abertura a todo tipo de colaborações. A DocGenealogy utiliza os dados disponíveis da Wikipedia sobre assessoria e orientação para doutorados para descobrir e rastrear as relações existentes entre conselheiros ou mentores e seus alunos. O ponto de partida do processo de mineração de dados é sempre um cientista bem conhecido. Esse processo continua, até encontrar os conselheiros e mentores relacionados com renome co-relacionados, sendo construído iterativamente o grafo de relacionamento. Além disso, o DocGenealogy também mostra visualmente como as pessoas estão distribuídas por alma mater e área de conhecimento. O objetivo principal do trabalho é avaliar a eficácia de técnicas orientadas ao grafo na interpretação dos dados das infoboxes da Wikipedia.

Palavras-chave: Visualização da Informação. Visualização Gráfica. Recuperação da Informação. Mineração visual de dados. Wikipedia.

^{*}Submetido em 18/07/2017 - Aceito em 15/02/2018

¹Mestrando em Tecnologias e Sistemas Informáticos, pela Universidade Aberta, Lisboa, Portugal – david.fernandes.paiva@gmail.com

²Professora Auxiliar na Universidade Aberta, investigadora no pólo CIAC/UAb, Lisboa, Portugal – elizabeth.carvalho@uab.pt

Abstract

DocGenealogy is a data mining project that allows the visualization and exploration of the doctoral advisement/mentoring relationship through an interactive graph chart. Wikipedia is known for its informal collaborative nature, based on a somewhat loose and non-structural way of producing information, and one of its strengths is, in fact, the openness to all sorts of collaborations. DocGenealogy uses the Wikipedia available data on doctoral advisement and mentoring to find out and track the existing relationships between advisors or mentors and their students. The data mining process starting point is always some well-known scientist. This process continues until finding all the ultimate reputed co-related advisors and mentors, and builds iteratively the relationship graph. Besides this, DocGenealogy also visually shows how people are distributed by alma mater and field of knowledge. The main objective of the work is to evaluate the effectiveness of the graph-oriented techniques in the interpretation of data from Wikipedia infoboxes.

Keywords: Information visualization. Graph visualization. Information retrieval. Visual data mining. Wikipedia.

1 INTRODUÇÃO

Information visualization (InfoVis), the study of transforming data, information, and knowledge into interactive visual representations, is very important to users because it provides mental models of information. The boom in big data analytics has triggered a broad use of InfoVis in a variety of domains, ranging from finance to sports or politics (LIU et al., 2014). There are different kinds and types of visualization techniques that can be used to visually map the data.

The study of complex activities such as scientific production and software development often requires modeling connections among heterogeneous entities including people, institutions, and artifacts (HEER; PERER, 2014). Graphs have been widely used to visualize complex relationships such as the trade relationships among cities and the friend relationships in social networks. A graph consists of nodes and edges where the nodes represent entities and the edges represent the relationships among the entities/nodes. For example, graphs can show vivid trade routes by representing cities as nodes and trade relationships as edges. Large data sizes often bring out the visual clutter problem in graphs (ZHOU et al., 2013). A node-link diagram visualization technique is frequently used to map graphs.

DocGenealogy is a web-based application that easily allows discovering the existing relationship between doctoral advisers and mentors along the time. It offers an interactive dynamic force-direct node-link diagram to assist end-users on this task, besides other charts to support a better information insight. Although DocGenealogy prototype is used in the context of extracting the correlation between advisor and mentor, the main goal behind it is to evaluate the effectiveness of using graph-based techniques to properly map Wikipedia's infobox information. We actually want to know if the end-user can easily understand the existing interdependency behind Wikipedia's infoboxes, when its information is represented in the form of a visual graph.

The rest of the article is structured as follows: we will first point out the scientific background to this work, with a special focus on visual techniques used to visually map graphs, and then we will describe the DocGenealogy conceptual model and its prototype, including its data model, source, design, and technology. Finally, we will discuss the main conclusions and outputs, besides future works to be developed.

2 BACKGROUND

Information visualization is concerned with exploiting the cognitive capabilities of human visual perception in order to convey meaningful patterns and trends hidden in abstract datasets (MOERE et al., 2012). As data has steadily become more complex in terms of its size, dimensionality, and time-variance, the field has been challenged to create new techniques that are more sophisticated and to develop objective evaluation methods that are able to benchmark

these different techniques against each other.

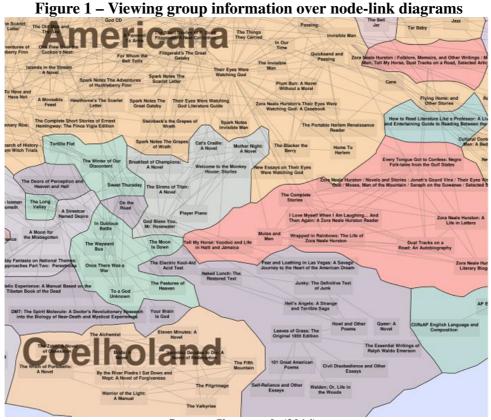
Visualization is suitable when there is a need to augment human capabilities rather than replace people with computational decision-making methods (MUNZNER, 2014). The design space of possible Vis idioms is huge and includes the considerations of both how to create and how to interact with visual representations. Vis design is full of trade-offs, and most possibilities in the design space are ineffective for a particular task, so validating the effectiveness of a design is both necessary and difficult.

Graphs in general form one of the most important data models in computer science because many problems and domains can be modeled as graph structures. Just to name a few, there are automata in theoretical computer science, flow networks such as pipes and roads, digital and non-digital social networks, computer networks such as the Internet, networks of companies and financial transactions, chemical reaction chains and molecular interactions, epidemic spreads of diseases in communities, or correlations of controlled variables in experiments (BECK et al., 2014).

Researchers have created a diversity of visualization techniques for networks and graphs. Two common representations used in social network analysis are node-link (JIANU et al., 2014) diagrams and adjacency matrix views (KANG et al., 2014). Hybrids of the two have also been proposed (RUFIANGE et al., 2012). These approaches organize elements according to the linkage structure of the graph creating a combination of these visual techniques and even others.

An alternative approach is to plot network data according to the attributes of the nodes, as in a scatter plot or so-called semantic substrates. Shneiderman and Aris (2006) proposed a network visualization layout based on a user-defined semantic substrate with a node-links diagram as an underlying visualization. Semantic substrates are spatially non-overlapping regions that are built to hold nodes based on some category present in the dataset. This approach is well suited for assessing potential correlations between node attributes and network structure (SHNEIDERMAN; ARIS, 2006).

Others have researched means of dealing with large graphs in excess of tens of thousands nodes. Common strategies include filtering and aggregation. Dunne and Schneiderman (2013) in order to help address this problem introduced a technique called motif simplification, in which common patterns of nodes and links are replaced with compact and meaningful glyphs. Zinsmaier, Brandes, Deussen, & Strobelt (2012) proposed a technique that allows straight-line graph drawings to be rendered interactively with an adjustable level of detail. The approach consists of a novel combination of edge cumulation with density-based node aggregation and is designed to exploit common graphics hardware for speed. Figure 1 illustrates some node-link diagrams. (DUNNE; SHNEIDERMAN, 2013) (ZINSMAIER et al., 2012)



Source: Jianu et al. (2014).

Data Mining (DM) aims to extract useful knowledge from raw data (CORTEZ; EMBRECHTS, 2013). Interest in this field arose due to the advances in Information Technology and rapid growth of business and scientific databases. These data hold valuable information such as trends and patterns, which can be used to improve decision making. Two important DM tasks are classification and regression. Both tasks use a supervised learning paradigm, where the intention is to build a data-driven model that learns an unknown underlying function that maps several input variables to one output target.

Liao, Chu and Hsiao (2012) reviewed the data mining techniques developed recently and several major kinds of data mining methods, including generalization, characterization, classification, clustering, association, evolution, pattern matching, data visualization and meta-rule guided mining. The techniques for mining knowledge from different kinds of databases, including relational, transactional, object-oriented, spatial and active databases, as well as global information systems were also examined by them in addition to potential data mining applications and some research issues. (LIAO et al., 2012)

Recently, applications of evolutionary algorithms have been found to be particularly useful for automatic processing of large quantities of raw noisy data for optimal parameter setting and to discover significant and meaningful information (MUKHOPADHYAY et al., 2014). Many real-life data mining problems involve multiple conflicting measures of performance, or objectives, which need to be optimized simultaneously. Under this context, multi-objective evolutionary algorithms are gradually finding more and more applications in the domain of data mining since the beginning of the last decade.

3 DOCGENEALOGY PROTOTYPE

It is frequently interesting to find out the relationship between people. To achieve a master or doctoral degree, a student must have an advisor, who guides him through the research process and stands up for him when the defense of his thesis takes place. However, it is very easy to lose track of this relationship. Many of the students that were advised do not pursue an academic career or develop relevant and notorious research.

The Wikipedia is a free encyclopedia, written collaboratively by the people who use it. It is a rich repository of information, not exhaustive, but holds a huge quantity of useful data. It contains, for instance, many articles and details about scientists, including the name of their notorious advisors or students. However, it is not easy to interpret the relationship between them.

Node-link diagram is one of the most suitable information visualization techniques to represent the relationship between data. Scientific mentors and advisors relationships can be easily interpreted using this kind of visual approach. A node-link diagram can be visually explored in many ways, as well as interacted using different methods.

The next sections describe what is behind the DocGenealogy prototype. We will refer to its data model, source and mining process. We will also describe its design, interface and technology.

3.1 Data

An infobox of a Wikipedia article generally contains key facts in the article and is organized in attribute-value pairs (TRAN; CAO, 2013). Infoboxes not only allow readers to rapidly gather the most important information about some aspects of the articles in which they appear, but also provide a source for many knowledge bases derived from Wikipedia. However, not all the values of the attributes of an infobox are updated frequently and accurately. Although not mandatory, the use of infoboxes is quite common and about 33% of Wikipedia articles contained an infobox (Figure 2 illustrates an infobox inside a Wikipedia page).

For this work we are interested in the Infobox scientist which has the following definition template described in Figure 3. Highlighted in red, we show the attributes that were necessary to be extracted from each page, such as: the name, the birth and death dates, the nationality and citizenship, the alma mater, the advisors of the scientist and his notorious students.

Alan Turing WikipédiA A enciclopédia livre Origem: Wikipédia, a enciclopédia livre. Alan Mathison Turing OBE Página principal Alan Turing Alan Mathison Turing Conteúdo destacado (Paddington, Londres, 23 de junho Eventos atuals de 1912 - Cheshire East, Esplanada Cheshire, 7 de junho de 1954) foi Página aleatória um matemático, lógico, Portals criptoanalista e cientista da Informar um erro computação britânico. Foi influente Loja da Wikipédia no desenvolvimento da ciência da Colaboração computação e na formalização do Boas-vindas conceito de algoritmo e AJuda computação com a máquina de Página de testes Turing, desempenhando um papel Portal comunitário importante na criação do Mudanças recentes computador moderno.[1][2][3] Foi Manutenção também pioneiro na inteligência Criar página Páginas novas artificial e na ciência da Contato computação.[4] É conhecido como Donativos Conhecido(a) Máquina de Turing o pai da computação. Teste de Turing, Prémio Turing Imprimir/exportar Durante a Segunda Guerra Nascimento 23 de Junho de 1912 Mundial, Turing trabalhou para a Criar um Ilvro gton, Londres, Inglaterra. inteligência britânica em Bletchley Descarregar como PDF Versão para Impressão 7 de Junho de 1954 (41 anos) Wilmslow, Cheshire, inglaterra, Park, num centro especializado em Morte quebra de códigos. Por um tempo Reino Unido Noutros proletos ele foi chefe do Hut 8, a seção Nacionalidade pritánica Wikimedia Commons responsável pela criptoanálise da Universidade de Cambridge Alma mater frota naval alemã. Planejou uma Ferramentas Universidade de Princeto série de técnicas para quebrar os Prêmios Officer of the Order of the British Páginas afluentes códigos alemães, incluindo o Alterações relacionadas Fellow of the Royal Society Carregar ficheiro método da bomba eletromecânica, Causa da Suposto sulcídio por Ingestão de Páginas especiais uma máquina eletromecânica que morte Hiperligação permanente poderia encontrar definições para a Orientador(es) Alonzo Church Informações da página máquina Enigma. Orientado(s) Robin Gandy Elemento Wikidata Instituições Universidade de Manchester Após a guerra, trabalhou no Citar esta página National Physical Laboratory Universidade de Cambridge Laboratório Nacional de Física do ٥ Noutros Idlomas Reino Unido, onde criou um dos Campo(s) matemática, lógica e অসমীয়া primeiros projetos para um

Figure 2 – Infobox (inside the red rectangle) within a Wikipedia page example

Source: Wikipedia (2017).

Figure 3 – Wikipédia infobox scientist attributes

```
{{Infobox scientist
  honorific_prefix :
  honorific_suffix =
  native_name =
native_name_lang =
                                 <!--(filename only, i.e. without "File:" prefix)-->
  image_size =
  caption
                                 <!--{{birth date | YYYY | MM | 00}}-->
   birth place =
                                 <!--{{death date and age | YYYY|MM|DD | YYYY|MM|DD}} (death date then birth date)-->
  death_place =
  death_cause =
  resting_place = resting_place_coordinates = <!--{(coord|LAT|LONG|type:landmark|display=inline,title}}-->
  other names =
   vorkplaces
   education
                                 <!--(or | thesis1_title = and | thesis2_title = )-->
<!--(or | thesis1_url = and | thesis2_url = )-->
<!--(or | thesis1_year = and | thesis2_year = )-->
<!--(or | doctoral_advisors = )-->
  thesis_title =
  thesis_url =
thesis_year =
doctoral_adviso
  influences
  influenced
awards
  author_abbrev_bot = author_abbrev_zoo =
  spouse
partner
children
                                 <!--(or | spouses = )-->
<!--(or | partners = )-->
  signature = signature_alt =
                                 <!--(filename only)-->
                                 <!--{{URL|www.example.com}}-->
   website
```

Source: Wikipedia (2017).

Table 1 describes the data that is used by DocGenealogy. Most of it is mined. We have a total of seven dimensions and three measures (the totals and counting). The dimensions are discrete independent variables describing the data while the measures are dependent variables holding data values that can be aggregated. The totals of advisors or mentors per alma mater and knowledge fields are calculated. The counting of direct relationships from and to an advisor or mentor is also computed.

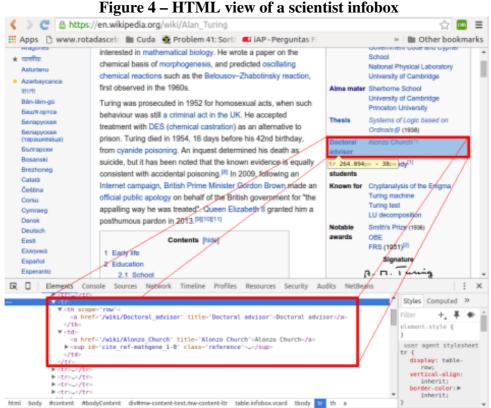
Table 1 – Data dimension/measures and classification

Dimensions/measures	Classification	
Name	Nominal	
Birth date	Numeric and temporal	
Death date	Numeric and temporal	
Citizenship	Nominal and geo-referenced	
Nationality	Nominal and geo-referenced	
Alma mater	Nominal	
Doctoral advisor	Nominal	
Totals per alma mater	Numeric, and quantitative	
Totals per knowledge field	Numeric, and quantitative	
Count of relationships	Numeric, and quantitative	

The structure showed in Figure 3 is used in the Wikipedia Editor as a template for article editing. The result, and the information that we have access to, as in any other type of dynamic web page, is pure HTML, which obeys another type of structure. The data-mining tool that we implemented, which may be considered a simplified web crawler, reads the HTML of a page, extracts the relevant information needed and, making use of the URL's of the doctoral advisors and students, recursively obtains more information about the graph.

The recursive mining uses a depth-first search strategy, starting on a given scientist's Wikipedia page (Alan Turing was chosen for this work) and crawling the Wikipedia pages of its doctoral students first, downstream, and doctoral advisors afterwards, upstream (LIU; MOTODA, 2012). In order to respect the best Wikipedia practices for automated content crawling (WIKIPEDIA, 2017), we limited the request frequency to 1 page request per second.

An example for Alan Turing's infobox HTML regarding his Advisors and Students (all the other attributes were extracted in a similar way) is showed in Figure 4. In order to put some limit on the time and space needed to gather the data, we established two limiting values: one on depth, for the maximum number of levels to crawl, up and downstream, and total elements for the maximum number of scientists to crawl. We started defining only a maximum of 100 for both, but of course, this number will have to be further tested and tuned.



Source: Wikipedia (2017).

3.2 Technology

As for the data mining/web crawling process we have implemented a specific java application that executes direct http requests. The HTML content of each scientist page is then explored using XML XPath queries (W3C, 2016), extracting the relevant data and organizing it in a JSON data structure suitable for direct consumption by the visualization framework.

As for the visualization and graphic interaction we used D3.js (BOSTOCK, 2016a), such as its histogram functions and force-layout graph representation (BOSTOCK, 2016b). D3.js is a JavaScript library for manipulating documents based on data. D3 helps to bring data to life using HTML, SVG, and CSS. D3's emphasis on web standards gives the full capabilities of modern browsers without tying yourself to a proprietary framework, combining powerful visualization components and a data-driven approach to DOM manipulation.

Finally, a website was created to show case the visualization and interaction features using plain HTML/CSS/JS. This first version of the prototype was used to run some initial evaluation tests.

3.3 Visualization

Any interactive multimedia project must follow certain steps in order to facilitate their development and implementation. It is also important to consider the contents and the interface to be chosen to provide an interactive structure that meets the original objectives of data exploration.

Because we wanted to represent the existing relationship between doctoral advisers and mentors along the time, we assumed that a dynamic node-link diagram visual representation should be the most suitable. We also wanted to give some flexibility to navigate and explore the graph. Another sensitive issue is how the graph will behave while being explored (like a zoom or selection directly on it) and its usability aspects.

When designing an algorithm for drawing any node-link diagram (not just trees), one must consider three categories of often-contradictory guidelines: drawing conventions, constraints, and aesthetics (WARD et al., 2010). Conventions may include restricting edges to be either a single straight line, a series of rectilinear lines, polygonal lines, or curves. Other conventions might be to place nodes on a fixed grid, or to have all sibling nodes share the same vertical position. Constraints may include requiring a particular node to be at the center of the display, or that a group of nodes are located close to each other, or that certain links must either go from top to bottom or left to right. Each of the above guidelines can be used to drive the algorithm design. Aesthetics, however, often have a significant impact on the interpretability of a tree or graph drawing, yet often result in conflicting guidelines. Some typical aesthetic rules include: to minimize line crossings, to maintain a pleasing aspect ratio and to minimize the total area of the drawing.

Force-directed algorithms are among the most flexible methods for calculating layouts of simple undirected graphs (KOBOUROV, 2012). Also known as spring embedders, such algorithms calculate the layout of a graph using only information contained within the structure of the graph itself, rather than relying on domain-specific knowledge. Graphs drawn with these algorithms tend to be aesthetically pleasing, exhibit symmetries, and tend to produce crossing-free layouts for planar graphs. We considered this approach ideal to represent the existing relationship between doctoral advisors and mentor along the time.

The graph nodes can also offer additional information about the age of the advisor or mentor, or even if he is still alive or not. The color attribute is used to map this information. The end-user can drag, filter, select (by name, alma mater institution or knowledge field) and zoom the graph to have a better information insight. Annotation is also available.

Besides the relationship, we also assumed that the distribution per alma mater or field of knowledge would be also helpful to the end-user. The graph itself is not the best way of showing this information. Ordered bar charts would easily show these distributions.

Although most of the data is directly extracted from the infobox, some of it has to be computed, like the number of students or the total per alma mater. The visual encodings where chosen taking into consideration the order of their effectiveness according to the data attributes (MUNZNER, 2014). Figure 5 illustrates this order. Table 2 summarizes the visual encodings that were considered for the existing dimensions and measurements in the data.

Channels: Expressiveness Types and Effectiveness Ranks Identity Channels: Categorical Attributes Magnitude Channels: Ordered Attributes Position on common scale Spatial region Position on unaligned scale Color hue Length (1D size) Motion Tilt/angle Shape Area (2D size) Depth (3D position) Color luminance Color saturation Curvature Volume (3D size)

Figure 5 – Visual encodings

Source: Munzner (2014).

Table 2 – Chosen visual encodings

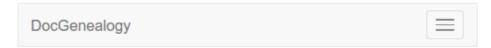
Dimensions/measures	Visual encoding	Type
Name	Text	Nominal
Birth date	Node colour	Quantitative
Death date	Node colour	Quantitative
Citizenship	Text	Nominal
Nationality	Text	Nominal
Alma mater	Text	Nominal
Doctoral advisor	Text	Nominal
Number of students	Node diameter	Quantitative
Number of mentors/advisors	Node border width	Quantitative
Totals per alma mater	Bar size and gradient color	Quantitative
Totals per knowledge field	Bar size and gradient color	Quantitative
Count of relationships	Node size	Quantitative

Hue color associations have been the topic of significant study in psychology literature and recent supports the idea that the formation and triggering of color associations can be understood through models of semantic memory such as associative network (LABRECQUE; MILNE, 2012). Although these studies are restricted in the number of colors and types of emotions and associations they test, the effects of colors remain relatively consistent across studies, which provides some empirical evidence of a systematic relationship between color and emotions and psychological functioning. Color is a relevant visual attribute that we use to map some of our quantitative data, such as the birth and death date or the alma mater or knowledge field distributions values. Finally, nominal data is mapped to text while size denotes the count of relationships or number of students.

The end-user interface is composed of two different types of panels. The main panel depicted in Figure 6 shows the interactive graph and presents filtering and thematic map selection tools. Each one of its input boxes gives auto-completion aid (Figure 7) as soon as a first letter is typed. It allows filtering by a combination of:

- Name: which allows the highlight of a particular scientist giving its name;
- Alma mater: allows the filter of scientists of a particular university;
- Fields: that allows the filtering of scientists by field of knowledge.

Figure 6 - Graph visualization, filters and thematic map selection



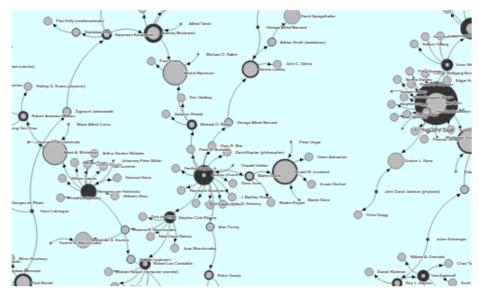
Network

Start typing on one of the textboxes shows an autocompletion list of matches. Selecting one of the elements of that list and pressing the ">" button, will apply the

selection as a filter on the nodes of the graph. Name: > Alma mater: > Fields: > Color by allows you to color the nodes on two diferent dimensions: age and dead/alive of the scientist. Color by: None

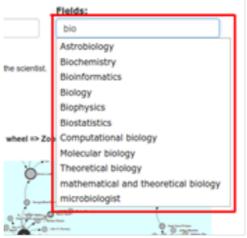
Actions:

Background Click and Drag => PAN || Node Click and Drag => Change graph || mouse wheel => Zoom In/Out || Shift+Click => Node Information || Dblclick => Node (un)select || CTRL+Dblclick => Node cumulative (un)select



Source: Elaborated by authors.

Figure 7 – Auto-completion help



Source: Elaborated by authors.

A combo box labeled "Color by" is presented and allows the selection of two different dimensions of the data: age and dead/alive condition, both reflecting in the coloring of the nodes as seen in Figure 8. The second type of panels shows static frequency histograms on two attributes of the data: Alma mater and field of knowledge. Figure 9 shows the one regarding Alma mater.

Color by allows you to color the nodes on two different dimensions: age and dead/alive of the scientist.

Age

Actions:

Background Click and Drag => PAN || Node Click and Drag => Change graph || mouse wheel => Zoom In/Out || Shift+Click => Node Informat CTRL+Dblclick => Node cumulative (un)select

Legend:

Legend:

Legend:

John David Jackson Ishington

Managed Pastin

Scoth, Sherker

Daniel Frank Walts

Gregorio Ricci-Curbastro

Figure 8 – Thematic map on age

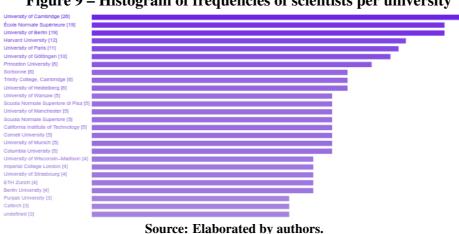


Figure 9 – Histogram of frequencies of scientists per university

Although we were able to find a reasonable visual representation for the data, some challenges do exist and had to be overcame. The quantity of nodes have to be well controlled, otherwise the graph may become unreadable. This is a limitation, because the user may want to have a deeper view of the graph. Filtering can help to see some portions, but on the other hand,

there is a lack of the overall graph.

4 MAIN CONCLUSIONS AND FUTURE WORK

We adopted the guidelines proposed by Lam et al. (2012) to define how the visualization process was going to be designed and conducted. Because our main goal is to assess a visualization tool's (or more specific, its data representations) ability to support visual analysis and reasoning about data to the end-users, we adopted the Evaluating Visual Data Analysis and Reasoning (VDAR) scenario (LAM et al., 2012).

In this case, outputs are both quantifiable metrics such as the number of insights obtained during analysis or subjective feedback such as opinions on the quality of the data analysis experience. The potential end-users are identified as general public who wants to find out the relationship between doctoral advisors and mentors.

Because the study of how a visualization tool may support analysis and reasoning is problematic and the products of an analysis are hard to homogenize and quantify, we assume that a case study should be further designed to properly evaluate the effectiveness of DocGenealogy.

This paper introduced the DocGenealogy web-application. It is still an on-going work. It will be improved and extended to visualize the result of the mining of other types of Wikipedia's infoboxes contents, and thus support visually its interpretation and insight.

We have already conducted preliminary tests to evaluate its performance and its overall visual quality both in data mapping and aesthetics level as well as its interface usability. The main goal of this work is to evaluate the effectiveness of using graph-based techniques to properly translate the existing interdependency and interconnectivity pointed out in the Wikipedia's

infoboxes. Based on the output that we received from preliminary tests, we may say that graph-based techniques should always be considered as an effective visual technique to map Wikipedia's infoboxes data, and thus, be included in any tool that aims to establish correlations between any Wikipedia's pages (through its infoboxes). The end-users were able to easily discover how different contents are interconnected and infer existing hierarchical aspects of the relationship, using the graphs visualization offered by DocGenealogy prototype.

DocGenealogy is able to be applied in other data contexts. The underlying structure of the prototype can be easily extended and used in other fields and data categories, such as biology, ecology, chemistry etc. This makes it also interesting to run tests against other data scenarios, validating its usefulness much more broadly.

Besides including fish-eye visual technique, another enhancement will be to perform numerical analysis on the selected quantitative data of the graph. The development of such functionality in our prototype will allow the end-user to easily discover more refined statistics on the filtered data, as well as what is already possible.

REFERÊNCIAS

BECK, Fabian et al. The state of the art in visualizing dynamic graphs. **EuroVis STAR**, v. 2, 2014.

BOSTOCK, Mike. **Data-driven documents.** 2016. Disponível em: https://d3js.org/>. Accessed: 2016-1-15.

BOSTOCK, Mike. **Force layout.** 2016. Disponível em: https://github.com/d3/d3/wiki. Accessed: 2016-1-15.

CORTEZ, Paulo; EMBRECHTS, Mark J. Using sensitivity analysis and visualization techniques to open black box data mining models. **Information Sciences**, Elsevier, v. 225, p. 1–17, 2013.

DUNNE, Cody; SHNEIDERMAN, Ben. Motif simplification: improving network visualization readability with fan, connector, and clique glyphs. In: **Proceedings of the SIGCHI Conference on Human Factors in Computing Systems**. New York: ACM, 2013. p. 3247–3256.

HEER, Jeffrey; PERER, Adam. Orion: A system for modeling, transformation and visualization of multidimensional heterogeneous networks. **Information Visualization**, Sage Publications Sage UK: London, England, v. 13, n. 2, p. 111–133, 2014.

JIANU, Radu et al. How to display group information on node-link diagrams: an evaluation. **IEEE Transactions on Visualization and Computer Graphics**, IEEE, v. 20, n. 11, p. 1530–1541, 2014.

KANG, U et al. Net-ray: Visualizing and mining billion-scale graphs. In: **Pacific-Asia Conference on Knowledge Discovery and Data Mining**. Switzerland: Springer, 2014. p. 348–361.

KOBOUROV, Stephen G. Spring embedders and force directed graph drawing algorithms. **arXiv preprint arXiv:1201.3011**, 2012.

LABRECQUE, Lauren I; MILNE, George R. Exciting red and competent blue: the importance of color in marketing. **Journal of the Academy of Marketing Science**, Springer, v. 40, n. 5, p. 711–727, 2012.

LAM, Heidi et al. Empirical studies in information visualization: Seven scenarios. **IEEE transactions on visualization and computer graphics**, IEEE, v. 18, n. 9, p. 1520–1536, 2012.

LIAO, Shu-Hsien; CHU, Pei-Hui; HSIAO, Pei-Yuan. Data mining techniques and applications—a decade review from 2000 to 2011. **Expert systems with applications**, Elsevier, v. 39, n. 12, p. 11303–11311, 2012.

LIU, Huan; MOTODA, Hiroshi. **Feature selection for knowledge discovery and data mining**. New York: Springer Science & Business Media, 2012. v. 454.

LIU, Shixia et al. A survey on information visualization: recent advances and challenges. **The Visual Computer**, Springer, v. 30, n. 12, p. 1373–1393, 2014.

MOERE, Andrew Vande et al. Evaluating the effect of style in information visualization. **IEEE** transactions on visualization and computer graphics, IEEE, v. 18, n. 12, p. 2739–2748, 2012.

MUKHOPADHYAY, Anirban et al. A survey of multiobjective evolutionary algorithms for data mining: Part 1. **IEEE Transactions on Evolutionary Computation**, IEEE, v. 18, n. 1, p. 4–19, 2014.

MUNZNER, Tamara. Visualization analysis and design. New York: CRC press, 2014.

RUFIANGE, Sébastien; MCGUFFIN, Michael J; FUHRMAN, Christopher P. Treematrix: A hybrid visualization of compound graphs. In: **Computer Graphics Forum**. New Jersey: Blackwell Publishing Ltd, 2012. v. 31, n. 1, p. 89–101.

SHNEIDERMAN, Ben; ARIS, Aleks. Network visualization by semantic substrates. **IEEE Transactions on Visualization and Computer Graphics**, IEEE, v. 12, n. 5, p. 733–740, 2006.

TRAN, Thong; CAO, Tru H. Automatic detection of outdated information in wikipedia infoboxes. **Research in Computing Science**, v. 70, p. 211–222, 2013.

W3C. **Xml path language (xpath).** 2016. Disponível em: https://www.w3.org/TR/xpath/. Online; accessed 30-April-2017.

WARD, Matthew O; GRINSTEIN, Georges; KEIM, Daniel. Interactive data visualization: foundations, techniques, and applications. New York: CRC Press, 2010.

WIKIPEDIA. **Wikipedia robots.txt**. 2017. Disponível em: https://en.wikipedia.org/robots.txt. Online; accessed 24-March-2017.

ZHOU, Hong et al. Edge bundling in information visualization. **Tsinghua Science and Technology**, TUP, v. 18, n. 2, p. 145–156, 2013.

ZINSMAIER, Michael et al. Interactive level-of-detail rendering of large graphs. **IEEE Transactions on Visualization and Computer Graphics**, IEEE, v. 18, n. 12, p. 2486–2495, 2012.