

From Web Data to Visualization via Ontology Mapping

O. Gilson¹, N. Silva², P. W. Grant¹ and M. Chen¹

¹Department of Computer Science, Swansea University, UK

²School of Engineering - Polytechnic of Porto, Portugal

Abstract

In this paper, we propose a novel approach for automatic generation of visualizations from domain-specific data available on the web. We describe a general system pipeline that combines ontology mapping and probabilistic reasoning techniques. With this approach, a web page is first mapped to a Domain Ontology, which stores the semantics of a specific subject domain (e.g., music charts). The Domain Ontology is then mapped to one or more Visual Representation Ontologies, each of which captures the semantics of a visualization style (e.g., tree maps). To enable the mapping between these two ontologies, we establish a Semantic Bridging Ontology, which specifies the appropriateness of each semantic bridge. Finally each Visual Representation Ontology is mapped to a visualization using an external visualization toolkit. Using this approach, we have developed a prototype software tool, SemViz, as a realisation of this approach. By interfacing its Visual Representation Ontologies with public domain software such as ILOG Discovery and Prefuse, SemViz is able to generate appropriate visualizations automatically from a large collection of popular web pages for music charts without prior knowledge of these web pages.

Categories and Subject Descriptors (according to ACM CCS): I.3.0 [Computer Graphics]: General

1. Introduction

Visualization is one of the indispensable means for addressing the rapid explosion of data and information. Although a large collection of visualization techniques have been developed over the past three decades, the majority of ordinary users, who handle data and information everyday, have little knowledge about these techniques. Despite there being many interactive visualization tools (e.g., ILOG Discovery [BHS04], Prefuse [HCL05], Spotfire [Ah96]) available in the public domain or commercially, producing visualizations remains a skilled and time-consuming task.

One approach for cost-effective dissemination of visualization techniques is to use captured expert knowledge for helping ordinary users generate visualizations automatically. To some users, this approach may serve as an introduction to new visualization techniques or an initial overview of possible styles of visualizations, which is followed by a more intensive interaction to create finely-tuned and customised visualizations. To others, this approach may provide an adequate visualization service without consuming excessive effort to learn and utilise the various visualization tools directly.

The method of “design galleries” [MAB*97], placed the first footprint in this direction. However, for visualizations with a very large parameter space, the method generally requires a long, iterative process involving the user before the search converges on a satisfactory result. In this work, we propose to use *ontologies*, which represent captured expert knowledge, to reduce the parameter space, providing a more effective automated solution to the dissemination of visualization techniques to ordinary users. As an example, we consider the visualization of music chart data on the web, and aim to generate visualizations automatically from the data. We present an ontology-based pipeline to map tabular data to geometrical data, and to select appropriate visualization tools, styles and parameters, producing formatted data that can be fed to the visualization tools automatically to generate visualizations.

The novel design of this pipeline features three ontologies, namely a *domain ontology* (DO) for storing domain knowledge about the source data (i.e., music charts in this work), a *visual representation ontology* (VRO) for storing the knowledge about visualization tools, styles and parameter space, and a *semantic bridging ontology* (SBO) for storing the knowledge about the mapping from DO to VRO. We use a

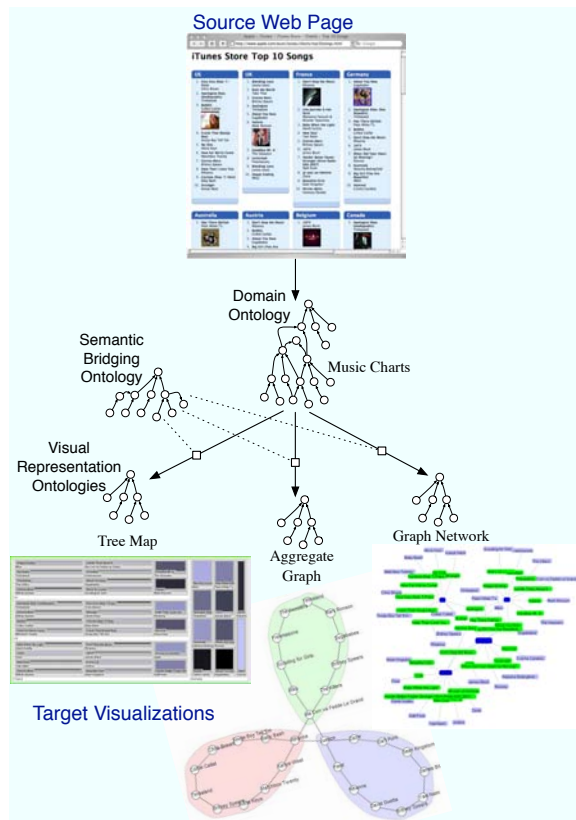


Figure 1: SemViz pipeline showing: Domain Ontology (DO); Semantic Bridging Ontology (SBO); and Visual Representation Ontology (VRO).

probabilistic ontology mapping technique loosely based on OMEN (Ontology Mapping Enhancer) [MNJ05] to realise the automatic data mapping within the pipeline, and create interfaces between the pipeline and two popular visualization tools, ILOG Discovery [BHS04] and Prefuse [HCL05]. Figure 1 shows an example web page containing the iTunes Store song chart, and visualizations generated by ILOG Discovery and Prefuse automatically via our pipeline.

In the remainder of the paper, we give an overview of related work in Section 2. In Section 3, we present an overview of the main components of our pipeline. In Section 4, we detail the three main ontologies, followed by a description of the ontology mapping algorithm in Section 5. We present and discuss our results in Section 6, and provide our concluding remarks in Section 7.

2. Related Work

There are four main methodologies commonly deployed in the visualization process, the first three were discussed in [PLB*01].

(a) The *trial and error* methodology [YaKSJ07] relies on the interaction between users and the visualization system to derive satisfactory results with minimum assistance from the computer. A large collection of visualization tools (e.g., SpotFire [Ahl96] and ILOG Discovery [BHS04]) support this approach by providing fast rendering and effective exploration of the visual space.

(b) The *design galleries* methodology [MAB*97] is a data-centric approach that relies on limited knowledge of any underlying data model. With some basic knowledge of the application domain and visualization tool (i.e., volume visualization in [MAB*97]), the visualization system automatically selects parameters and generates a set of visualizations, from which users select the most relevant and useful visualizations. This process is repeated until satisfactory visualizations are obtained in a manner resembling the semi-automatic genetic algorithm.

(c) The *information-assisted* methodology relies on some understanding of the underlying model of the data. It extracts more abstract information from the data (e.g., histogram [YMC05], cluster [GDGL07] and topology [WBP07]), and uses it to guide users in their interactive visualization process. Methodologies (b) and (c) involves partial automation, but users' interaction is an essential part of the process.

(d) The *automatic visualization* methodology attempts to generate visual representations from data automatically. [Fei85] and [Mac86] first set the agenda for this research direction. [MHS07] presented a set of user interface commands, "Show Me", as part of the user interface of Tableau, providing a number of automated functions in user interaction. In comparison with the other three methodologies, this approach is least studied.

Music chart data is a relatively simple form of data, but varies greatly in data organisation, terminology used, and levels of details. Many general purpose tools for non-spatial data (or information) visualization can be used to create appropriate visualizations by hand. There are also visualization tools which are tailored towards music. CoMIRVA [SKSP07] contains a powerful set of functionalities that are able to produce music visualizations in a variety of styles. Through its user interface, a skilled user can create impressive visualizations from music meta-data stores and digital audio files as long as they conform to accepted formats. There are some 15 other music visualization tools for performance, tune, pitch etc. All these are based primarily on methodology (a).

Ontologies are used in domains such as knowledge engineering, enterprise data sharing and the semantic web. Due to the heterogeneity of these systems, translating between ontologies is of great importance and is known as Ontology Mapping, Matching or Alignment [ES07]. The large size of a typical ontology means that automatic ontology mapping is essential. Techniques use a combination of similarity, distance, structure and external semantics. Our work looks at

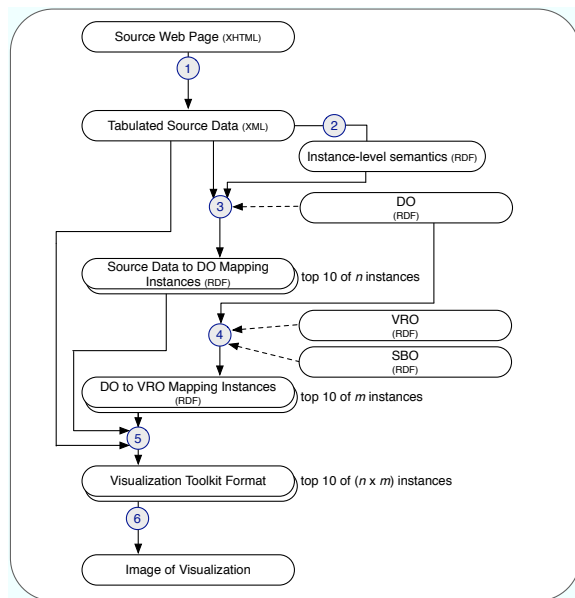


Figure 2: The SemViz technology pipeline, from Source web page to Target visualizations

automatic ontology mapping from the perspective of probabilistic techniques [MNJ05], [TLL*06].

Ontologies have been considered in visualization. [DBDH05] suggested the need to build an ontology of visualization to capture the concepts and characteristics of visualization. [RKR06] developed a web application for categorising and storing information about systems for software visualization. There are also studies on visualization techniques for displaying ontological structures (e.g. [FSvH04], [BBP05], [KHL*07]), which is not the focus of this work. Our work falls into the same scope of [DBDH05] and [RKR06], but we have gone a step further by proposing to use concepts and knowledge stored in ontologies to facilitate automatic generation of visualizations. Hence, we are pursuing the same goals as the above-mentioned previous work in the scope of methodology (d), but by employing more powerful and systematic tools, that is, ontologies and ontology mapping.

3. System Overview

We have developed a prototype, SemViz which is able to produce an end-to-end automatic visualization of tabulated data from a selection of music chart web pages. SemViz allows the mapping algorithm's parameters to be adjusted and includes custom code for interfacing to the visualization toolkits. We choose to output visualizations using either the ILOG Discovery or Prefuse visualization toolkits. The pipeline stages (figure 2) are as follows:

- 1. Extract Tabulated Source Data from Web Page** If an XML or CSV link to the tabulated data is not provided, a screen-scraping/data extractor such as Solvent and Piggy Bank [HMK07] can be applied.
- 2. Perform Instance-level Data Analysis on Source Data** This stage is optional, but can be used to augment the Domain Ontology, particularly if there is a large amount of data where valuable semantics can be usefully extracted.
- 3. Create the Source Data to Domain Ontology Mappings** This component uses string similarity measures of the data column and domain ontology concept names and also the instance data to probabilistically reason on the most likely mappings. Each mapping permutation is scored and the top n of the possible permutations are stored. This is a schema mapping process.
- 4. Create the Domain Ontology to VRO Mappings** Depending on which concepts in the Domain Ontology have been stimulated by the Source Data, the Mapper uses the rules stored in the Semantic Bridging Ontology to create mappings which aim to result in useful visualizations. Each mapping permutation is scored, and the top m of the permutations are stored. This is a schema mapping process. The Ontology Mapping algorithm is described in Section 5.
- 5. Execute the Mappings** With the top n permutations from stage 3 and m permutations from stage 4, this results in $n \times m$ possible mapping permutations. Each permutation is given a score, they are ranked, and the highest 10 scoring permutations are combined with the original tabulated source data to form 10 VRO instances. In SemViz, 10 are chosen as a good trade-off. The 10 VRO instances are converted into the specific files necessary for each Visualization toolkit supported by the system.
- 6. Generate Visualizations** The toolkits are invoked and the visualizations are generated before being presented to the user.

4. Ontologies

An ontology provides an explicit conceptualisation (i.e., meta-information) that describes the semantics of data [Fen01]. A language for defining ontologies is syntactically and semantically richer than other common approaches (e.g. databases). An ontology consists of concepts, relationships and attributes. This can be seen in Figure 3. Concepts (circles) are related via relations (arrows). For example, an "Artist" concept is related to a "Song" concept via a "has" relation. A concept can also have attributes (rectangles in the diagram). For example, an "Artist" concept has an "is-PrimaryKey" attribute.

The ontologies used in SemViz were developed using Stanford's Protege tool [NSD*01] and are expressed in RDF/OWL [LS98]. In the following worked example, we use the BBC's top 40 web page visualized as a 2D Graph. We restrict the diagrams to show only stimulated concepts and also the strongly weighted relationships.

DO	VRO	Relationship / Attribute
synonyms	-	A
instanceHistory	-	A
has	contains	R
complements	complements	R
priorityWrt	priorityWrt	R
isQualitative	isQualitative	A
isQuantitative	isQuantitative	A
isPrimaryKey	isInformational	A
-	isMandatory	A

Table 1: Semantic Equivalence of relationships and attributes as used in the DO and VRO.

4.1. Domain Ontology (DO)

The purpose of the Domain Ontology is to store the semantics of the subject area which the source web page covers. The semantics are derived in such a way that they can be easily mapped to artefacts in the Visual Representation Ontology (VRO) (see section 4.2). This is done by defining a controlled set of relationships and attributes for use in both the DO and VRO. Some of the relationships and attributes have semantic equivalence. This forms the basis of our ability to map between DO concepts (i.e. data entities) and VRO concepts (i.e. visual artefacts) and therefore produce cognitively useful visualizations. Table 1 lists the relationships and attributes used in the DO and VRO, together with their semantic equivalence (if applicable). Note that these relationships and attributes are general in that they can be applied to any new DO instance (e.g. car records) or VRO instance (e.g. 3D Graph) which may be added to the SemViz system.

The DO for the Music Charts area is shown in Figure 3. Each relationship and attribute has a *strength value* which is a real number between 0 (weakest) and 1 (strongest). The only exception to this is the “priorityWrt” (priority with respect to) relationship which is 0.5 if the two linked concepts are of equal priority, or > 0.5 if the source concept has a higher priority than the target concept. The system records all relationships and attributes, no matter how strong or weak. In fact, there are relationships between every concept. These are present because the pipeline is based on probabilistic reasoning where we score permutations in order to decide on the best mapping. In general, a DO is initially created by a domain expert who “primes” the ontology with appropriate values for the relationship strengths.

The first mapping stage of SemViz is between the Tabulated Source Data and the DO (stage 3 in Figure 2). This process uses string similarity (Levenshtein distance) to measure the likelihood of a column in the source data having a match with a concept in the DO. String similarity is performed on both the column/concept names and the instance data. The DO keeps a record of concept name history (*synonyms* attribute) and instance value history (*instanceHistory* attribute). The score for the mapping of a column to a con-

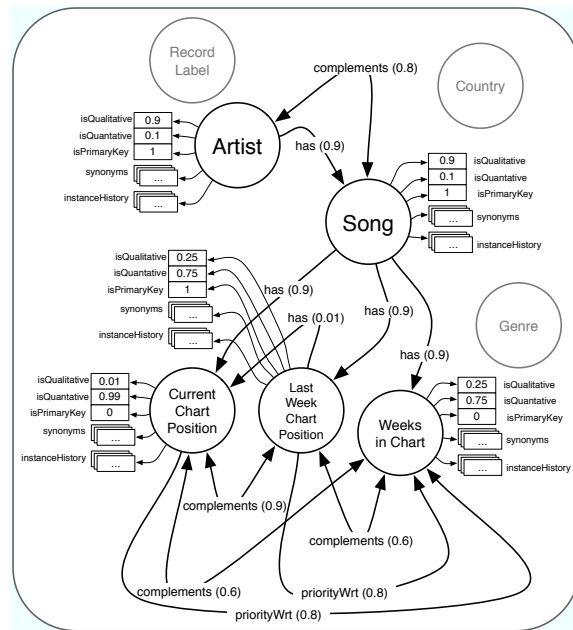


Figure 3: The Domain Ontology instance for music charts subject area (as mapped to the BBC top 40 charts web page).

cept is based on the top similarities of the concept name synonyms plus the proportion of historical instances which are the same as the instances in the source. The second mapping stage between the DO and VRO is more involved and is discussed in Section 5.

In Figure 3 it can be seen that the BBC Top 40 charts web page has been mapped to the “Artist”, “Song”, “Current Chart Position”, “Last Week Chart Position”, and “Weeks In Chart” concepts in the DO. These 5 concepts are known as the *stimulated* concepts of the DO.

4.2. Visual Representation Ontology (VRO)

The VRO captures the semantics of a particular visual representation (e.g. 2D Graph). It does this by modelling visual artefacts (e.g. X coordinate, Y coordinate, Colour, etc.) as concepts and the relationships between them. In this way, we can match relationships in the DO with relationships which have semantic equivalence in the VRO. We can also perform a similar task with semantically equivalent attributes.

We have built VRO’s for 2D graphs (see Figure 4), TreeMaps, Parallel Coordinates and Graph Networks. The major source of information during the domain modelling exercise was ILOG Discovery. The user interface has a Projection Inspector which allows users to control the mappings between source data entities (e.g. “Current Chart Position”) and target visualization artefacts (e.g. “X coordinate”). ILOG Discovery’s Projection Inspector therefore provides a good source of *executable* and *pragmatic* seman-

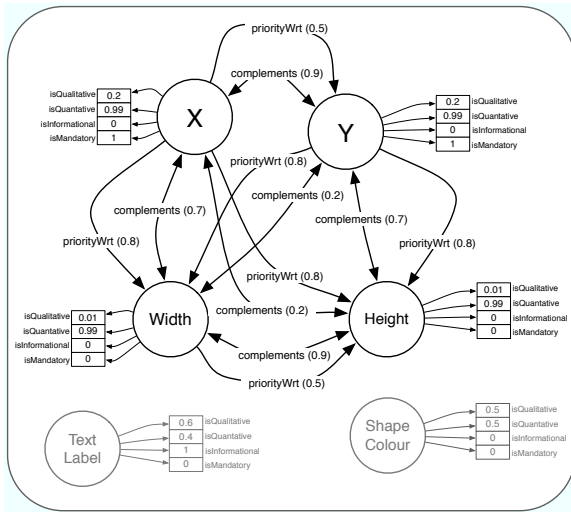


Figure 4: The Visual Representation Ontology instance for a 2D Graph (as mapped to the DO instance in Figure 3).

tics covering different visual representations. The 2D Graph VRO is also able to capture the semantics used by 2D Graphs in other visualization toolkits such as Prefuse [HCL05].

The VRO in Figure 4 shows the 2D graph concepts (visual artefacts) which have been mapped to the concepts (data entities) in the Music Charts DO from Figure 3. The “Text Label” and “Shape Colour” concepts have no visible relationships since they are all weak relationships. Note that the *isMandatory* attribute in the VRO has no semantic equivalence to any attribute in the DO. However, it is used as a control feature to ensure that visualizations are valid through having all mandatory VRO concepts (i.e. visual artefacts) mapped.

4.3. Semantic Bridging Ontology (SBO)

The purpose of the SBO is to capture and store the available expert knowledge about how various subject domains can be usefully visualized by different visual representations. This allows the complexity of the number of mapping permutations to be reduced. It also allows the accuracy of the scoring algorithm to be increased (see Section 5). The SBO is made up of Semantic Bridge concepts (or “semantic bridges”). Each semantic bridge records a single mapping between a DO concept (data entity) and a VRO concept (visual artefact), together with its appropriateness value. In this way, the SBO is a fully-connected graph of all possible permutations between the DO(s) and the VRO(s) in the system. By default, the appropriateness given to each semantic bridge is 100. However, this value can be increased or decreased to reflect specific expert knowledge. The SBO shown in Figure 5 highlights the semantic bridges which have non-default appropriateness values.

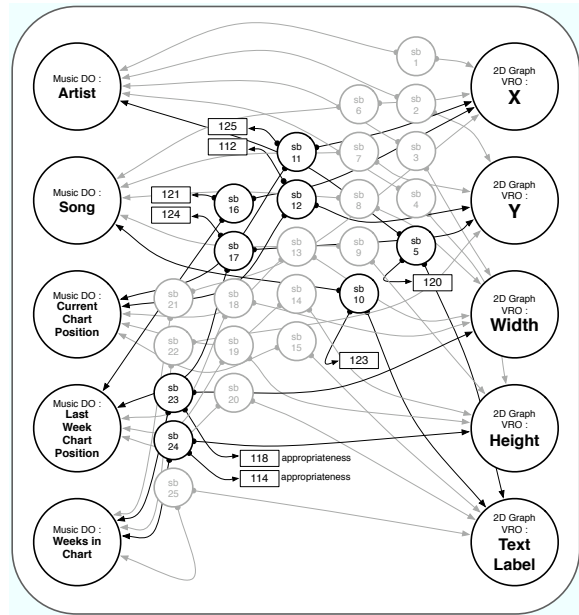


Figure 5: The Semantic Bridging Ontology containing the domain and visualization knowledge for mappings between the music DO and the 2D Graph VRO.

5. Ontology Mapping

The algorithm we employ to score the mapping permutations from the DO to VRO is loosely based on a version of the OMEN (Ontology Mapping ENhancer) algorithm described in [MNJ05]. OMEN uses a set of meta-rules that capture the influence of the ontology structure and the semantics of ontology relations to match nodes that are neighbours of already matched nodes in two ontologies. Instead of a Bayesian network (which cannot easily be defined by experts), we use the SBO to manage the complexity and scalability of the mapping process.

It is possible to consider all permutations between concepts from DO to VRO. However, this leads to an algorithm with a factorial computational complexity. Therefore, for non-trivial examples, the number of permutations to check quickly becomes unwieldy. To reduce the number of permutations, we use the SBO to ensure that only a subset of the permutations will be considered - those with semantic bridge appropriateness values over a pre-determined threshold. This expert knowledge can also be used by the scoring algorithm.

With respect to the base example in Figure 6, let θ be the mapping from DO to VRO, so:

$$V = \theta(D)$$

$$V' = \theta(D')$$

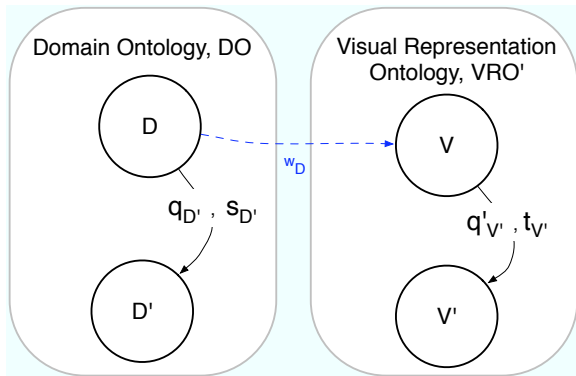


Figure 6: The DO and VRO of a base example.

If we wish to find w_D (the weighting of the concept pair mapping in this permutation in θ), we know that D to D' has a relationship of type $q_{D'}$. We also know that V to V' has a relationship of type $q'_{V'}$. If $q_{D'}$ and $q'_{V'}$ have semantic equivalence, $q_{D'} \sim q'_{V'}$ (see table 1) then we can compare the strength values: $s_{D'}$ and $t_{V'}$. The closer these strength values are to each other, the higher the probability of them being equivalent. In order to get w_D , we apply a “fitness function” which takes the two strength values as parameters (s and t). In the example visualizations in this paper, we use the first fitness function (f_1).

The overall score given to the whole permutation, $totalw_\theta$ (indicating the calculated cognitive value of the visualization) is the sum of all concept pair weight values. This is formalised as:

$$w_D = \sum_{D' \neq D} f_1(s_{D'}, t_{V'})$$

$$totalw_\theta = \sum_{D \in DO} w_D$$

where

$$f_1(s, t) := 1 - |s - t|$$

Other fitness functions are possible, such as f_2 , which takes into account the size of the values of s and t .

$$f_2(s, t) := (1 - |s - t|) \cdot \frac{s + t}{2}$$

An alternative version of $totalw_\theta$ takes into account the appropriateness value ($a_{D'}$) stored in the SBO.

$$totalw_\theta = \sum_{D \in DO} w_D \cdot \frac{a_{D'(D)}}{100}$$

The approach of using a SBO allows us to reduce the permutation search space while utilising existing domain and visualization knowledge.

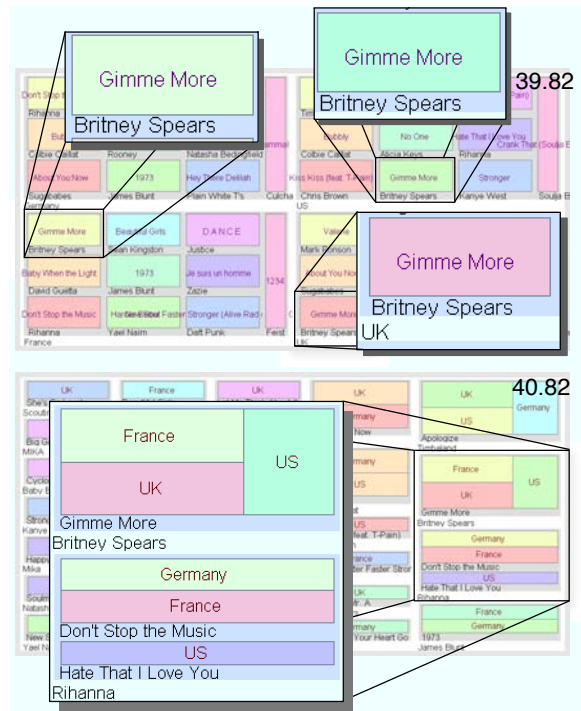


Figure 7: Top: Schema-based semantics deduced from the Music Chart DO. Bottom: Instance-based semantics derived from the source data by the Data Analysis module.

6. Results and Remarks

6.1. Instance vs. Schema-level Categorisation

There are two methods of deducing the semantics of the web page source data:

1. Schema-level Categorisation. The semantics of each concept in the Domain Ontology are pre-defined. We use these semantics to calculate an effective mapping between the DO and VRO.
2. Instance-level Categorisation. An analysis of the actual values of the source data provides semantics. This analysis can optionally be performed by the Data Analysis module (see stage 2 in Figure 2) during a first pass through the visualization pipeline.

In Figure 7 we see the effects of the two methods on the visualization of the iTunes music chart using a TreeMap. The top visualization uses the DO as defined (schema level). This shows a Country \rightarrow Artist \rightarrow Song hierarchy. The second uses an instance level analysis to augment and override the DO. As such, we have an Artist \rightarrow Song \rightarrow Country hierarchy. This is because the Data Analyser deduces that a Song “has” a Country, rather than a Country “has” an Artist. For the iTunes Store music chart, the second method produces a cognitively more valuable visualization compared to the first method which doesn’t provide much more insight over the original web page’s table of data.

6.2. The Gallery Selection Methodology

The gallery interaction methodology [MAB*97] presents the user with multiple visualisations for one data set. This is based on the principle of the user being able to choose which visualization is most applicable for their needs. The ontology-based pipeline we present in this paper lends itself to this style of interaction. Since the result of the pipeline is a probabilistic score for each possible mapping permutation, we can present the user with a manageable set of the best visualizations. This is shown as scored thumbnail visualizations in Figure 8. The benefits of the gallery selection methodology are:

- the user gets “something” to see, even if the certainty of its appropriateness is low.
- the visualization thumbnails which the user selects provides the system with feedback on the mapping decisions which were made. This provides the basis for a learning system based on users’ interactions.

6.3. Comprehending Automatically Created Visualizations

At the top of Figure 8, we show the highest scoring visualization for the BBC Top 40 webpage. Nearest to the origin, we can clearly see a cluster of shapes just below the $X=Y$ line, representing those songs which have fallen least since the previous week. Shapes along the X-axis represent new song entries since they have no value (zero) for “Last Week”.

Notice that the six visualizations with the highest scores have a diagonal line ($X = Y$) overlaid on the visualization. SemViz has automatically instructed ILOG Discovery to draw this line to assist the end-user with observing trends. A rule exists in the system which states that when a mapping permutation’s concept pairs pertaining to the the X coordinate and the Y coordinate have a *complements* value greater than a certain threshold, then there is a benefit in drawing the user’s attention to the placement of shapes relative to the $X = Y$ line. Therefore, for this particular permutation, an assistance line is drawn.

7. Conclusions

We have described a pragmatic method of producing automatic visualizations using domain knowledge captured in ontologies. The Domain Ontology (DO) captures knowledge about the source web pages’ subject domain; the Visual Representation Ontologies (VRO) capture the semantics of popular visual representations/styles; and the Semantic Bridging Ontology (SBO) holds key knowledge about the relationships between data entities of the source subject domain and the visual artefacts of the target visualizations. We have rationalised the relationships between concepts in the DO and VRO into a core set of *semantic equivalences* which form the basis of the scoring algorithm.

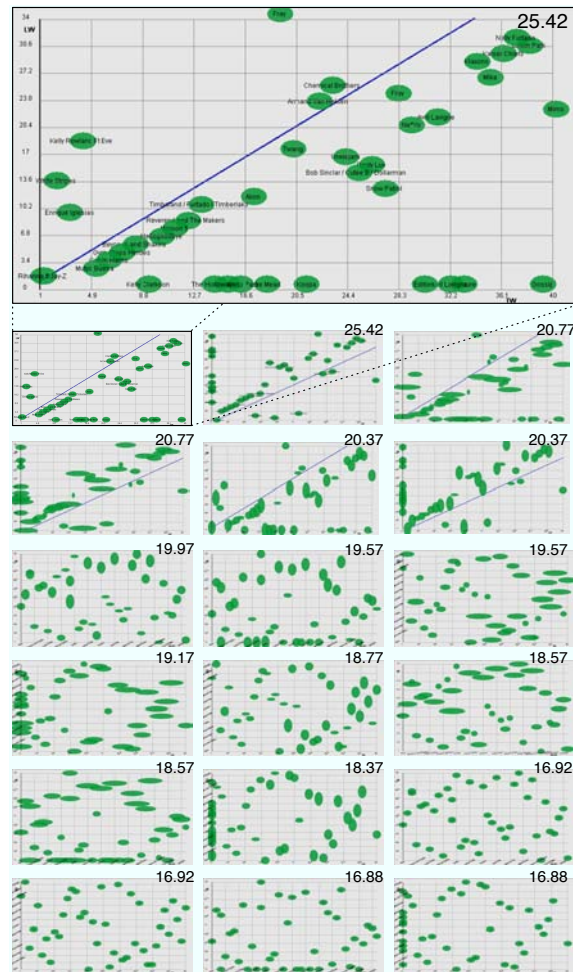


Figure 8: Top: The highest scoring visualization. Thumbnails: Images showing all (usable) permutations of the BBC Top 40 web page to 2D Graph using ILOG Discovery.

We have adapted an existing ontology mapping algorithm to encompass probabilistic relationships. This algorithm has a good trade-off between computational cost and ability to produce high quality automatic visualizations. We have implemented the visualization pipeline in a prototype, SemViz which functions end-to-end from source web page to target visualization. SemViz interfaces with two public-domain visualization frameworks.

We have shown demonstrable results by taking music chart web pages and using SemViz to interface with the ILOG Discovery and Prefuse visualization toolkits to produce examples in a variety of popular visualization styles. The visualization pipeline and supporting data-structures provide a good framework on which to extend and refine the current ontology mapping algorithm.

8. Acknowledgements

The authors wish to acknowledge the support of Swansea University and the GECAD research unit at ISEP. This work has been part-funded by the European Union through the Welsh Assembly Government.

References

- [Ahl96] AHLBERG C.: Spotfire: an information exploration environment. *SIGMOD Rec.* 25, 4 (1996), 25–29.
- [BBP05] BOSCA A., BONINO D., PELLEGRINO P.: Ontosphere: more than a 3d ontology visualization tool. In *Proceedings of the 2nd Italian Semantic Web Workshop* (2005).
- [BHS04] BAUDEL T., HAIBLE B., SANDER G.: *Visual Data Mining with ILOG Discovery*. 2004.
- [DBDH05] DUKE D. J., BRODLIE K. W., DUCE D. A., HERMAN I.: Do you see what i mean? *IEEE Computer Graphics and Applications* 25, 3 (2005), 6–9.
- [ES07] EUZENAT J., SHVAIKO P.: *Ontology matching*. Springer-Verlag, Heidelberg (DE), 2007.
- [Fei85] FEINER S.: Apex: An experiment in the automated creation of pictorial explanations. *IEEE Computer Graphics and Applications* 5, 11 (1985), 29–37.
- [Fen01] FENSEL D.: *Ontologies: A Silver Bullet for Knowledge Management and Electronic Commerce*. Springer-Verlag, Heidelberg, Germany, 2001.
- [FSvH04] FLUIT C., SABOU M., VAN HARMELEN F.: Supporting user tasks through visualisation of lightweight ontologies. In *Handbook on Ontologies*. Springer, 2004, pp. 415–434.
- [GDGL07] GIACOMO E. D., DIDIMO W., GRILLI L., LIOTTA G.: Graph visualization techniques for web clustering engines. *IEEE Transactions on Visualization and Computer Graphics* 13, 2 (2007), 294–304.
- [HCL05] HEER J., CARD S. K., LANDAY J. A.: Prefuse: a toolkit for interactive information visualization. In *CHI '05* (2005), ACM Press, pp. 421–430.
- [HMK07] HUYNH D., MAZZOCCHI S., KARGER D.: Piggy bank: Experience the semantic web inside your web browser. *Web Semant.* 5, 1 (2007), 16–27.
- [KHL*07] KATIFORI A., HALATSIS C., LEPOURAS G., VASSILAKIS C., GIANNPOULOU E.: Ontology visualization methods—a survey. *ACM Comput. Surv.* 39, 4 (2007), 10.
- [LS98] LASSILA O., SWICK R.: *Resource Description Framework (RDF) model and syntax specification*. Tech. rep., 1998.
- [MAB*97] MARKS J., ANDALMAN B., BEARDSLEY P. A., FREEMAN W., GIBSON S., HODGINS J., KANG T., MIRTICH B., PFISTER H., RUMMLER W., RYALL K., SEIMS J., SHIEBER S.: Design galleries: a general approach to setting parameters for computer graphics and animation. In *SIGGRAPH '97* (NY, 1997), pp. 389–400.
- [Mac86] MACKINLAY J. D.: Automating the design of graphical presentations of relational information. *ACM Trans. Graph.* 5, 2 (1986), 110–141.
- [MHS07] MACKINLAY J., HANRAHAN P., STOLTE C.: Show me: Automatic presentation for visual analysis. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1137–1144.
- [MNJ05] MITRA P., NOY N., JAISWAL A.: Ontology mapping discovery with uncertainty. In *Proceedings of the 4th International Semantic Web Conference (ISWC)* (Galway (IE), 2005), vol. 3729 of *LNCS*, pp. 537–547.
- [NSD*01] NOY N. F., SINTEK M., DECKER S., CRUBEZY M., FERGERSON R. W., MUSEN M. A.: Creating semantic web contents with protege-2000. *IEEE Intelligent Systems* 2, 16 (2001), 60–71.
- [PLB*01] PFISTER H., LORENSEN B., BAJAJ C., KINDLMANN G., SCHROEDER W., AVILA L. S., MARTIN K., MACHIRAJU R., LEE J.: The transfer function bake-off. *IEEE Computer Graphics and Applications* 21, 3 (2001), 16–22.
- [RKR06] RHODES P., KRAEMER E., REED B.: Vision: an interactive visualization ontology. In *ACM-SE 44: Proceedings of the 44th annual Southeast regional conference* (New York, NY, USA, 2006), ACM Press, pp. 405–410.
- [SKSP07] SCHEDL M., KNEES P., SEYERLEHNER K., POHLE T.: The comirva toolkit for visualizing music-related data. In *Proc. 9th Eurographics/IEEE VGTC Symposium on Visualization (EuroVis '07)* (Norrköping, Sweden, May 2007).
- [TLL*06] TANG J., LI J., LIANG B., HUANG X., LI Y., WANG K.: Using bayesian decision for ontology mapping. *Web Semant.* 4, 4 (2006), 243–262.
- [WBP07] WEBER G., BREMER P.-T., PASCUCCI V.: Topological landscapes: A terrain metaphor for scientific data. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1416–1423.
- [YaKSJ07] YI J. S., AH KANG Y., STASKO J., JACKO J.: Toward a deeper understanding of the role of interaction in information visualization. *IEEE Transactions on Visualization and Computer Graphics* 13, 6 (2007), 1224–1231.
- [YMC05] YOUNESY J., MOLLER T., CARR H.: Visualization of time-varying volumetric data using differential time-histogram table. *Fourth International Workshop on Volume Graphics, 2005* (2005), 21–224.