Information Realisation: Textual, Graphical and Audial Representations of the Semantic Web

Owen Gilson[1], **Nuno Silva**[2], **Phil W. Grant**[1], **Min Chen**[1], **João Rocha**[2]
[1](Department of Computer Science, University of Wales Swansea, UK
{csowen, P.W.Grant, M.Chen}@swan.ac.uk)

[2](ISEP Instituto Superior de Engenharia, Instituto Politécnico do Porto, Portugal
{Nuno.Silva, Joao.Rocha}@dei.isep.ipp.pt)

Abstract: Information Realisation is the process of presenting data as Textual, Graphical or Audial information to a human user. In this paper, we discuss the importance of this concept with respect to the accessibility of Semantic Web data to a diverse target audience. We provide an ontological point of view, defining the expressive characteristics and application domain of representation formats, thus presenting a system which produces representations customised to the user environment and the nature of the source data. Our approach considers the semantics of the data, not just the structure, and aims to present the information in the most semantically appropriate manner for the given target environment. We provide examples of a simple data set being realised as popular target representation formats: textual (XHTML, RSS); graphical (SVG, X3D); and audial (SoundML, VoiceXML).

Keywords: visualisation, information realisation, accessibility, ontology, semantic web, XML
Categories: H.3.1, H.3.2, H.3.3, H.3.5, H.5.1, H.5.2, H.5.4, H.5.5

1 Introduction

As we approach the vision of the Semantic Web [Berners-Lee and Fischetti, 1999], an increasing amount of data is being provided on the web in machine readable formats such as XML and RDF. This allows machines to reason about data thus providing additional services. However, in this process, we must not forget about the needs of human users. There is still the need for humans to be able to view this data as well as for machines to process it. Human users should be able to access data at every stage that a machine processes information on the Semantic Web. This provides users with an additional level of confidence about the validity of the machine processing task. This is especially important as a complement to the upper levels of the Semantic Web stack (i.e. trust and proof).

Information visualisation is focussed on the visual (i.e. using the eyes) aspect of humans accessing information sources. We believe that a more general approach should be taken which considers other formats which target different sensory organs. For example, blind or partially sighted people should be able to “see” data as an audio representation of data. Additionally, there is the need to have data realised in different formats depending on the environment and circumstance of the user: interactive or passive use; background or foreground activity. A user who is performing a task which requires most of their attention, such as driving a car, would not benefit from a detailed, interactive, graphical visualisation of a weather forecast. A passive summary which is presented audially would be more appropriate. A further example is a blind user who wishes to know statistics about a sports team's performance over a season. This could be presented as a summary audio stream.

This process considers multiple output types (textual, graphical and audial) as well as multiple audience environments. The environment consists of user factors, technological factors and data factors. We term this technique *Information Realisation* and believe that it is an important part of future work in the Semantic Web. This technique is similar to the framework outlined in [Jung and Sato, 2005]. We summarise Information Realisation in figure 1.

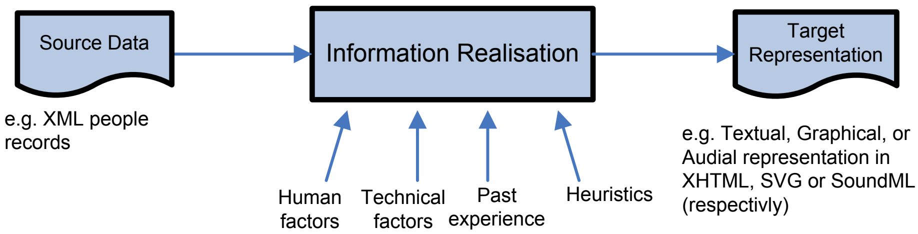

Figure 1: A general model of Information Realisation.

In Section 2 we provide an overview of related work. In Section 3 we give a full description of the Information Realisation process. We then describe in Section 4 the three main ontologies used in the process and provide a worked example in Section 5. In Section 6 we give concluding remarks and our thoughts on future work.

2 Related work

Much related work can be derived from the information visualisation community. Most of this work is domain specific in that it concentrates on the data and user characteristics of a specific subject area. For example, medical visualisation is concerned with viewing large datasets of body tissue [Ackerman, 1994]. The datasets are high-resolution, but low in semantics. In volume graphics, each element of a 3D volume is represented as a voxel (a 3D pixel). The voxel usually has a single value such as a Hounsfield unit in CT scans, but this is the limit of the semantical information. However, there are millions of voxels in a typical data set which means that there is a great deal of data to manipulate and analyse. In contrast, Semantic Web data is usually lower in volume, but higher in semantics. We thus need different approaches for visualising (or realising) semantically rich data compared to high-volume, high resolution, low semantics data sets.

Much of the work on visualising the semantic web has been concerned with RDF data. This data is in the form of a labelled-directed graph structure with visualisations being concerned with examining the structure of this data [Geroimenko and Chen, 2006]. The user is then able to analyse the data to find clusters, patterns and trends [Thomas and Cook, 2006]. While this can provide useful information about the data, we believe that there is also a benefit from visualisations which consider additional semantics of the data. For example, in a dataset concerning records of people, the age of a person should ideally be visualised differently from the number of children the person has. In this way, the visualisation becomes richer and more meaningful, also allowing other representation formats (textual and audial) to be considered.

We therefore view Information Realisation as the representation of data which has a high level of semantics, where we attempt to model the data according to these semantics.

3 The Information Realisation Process

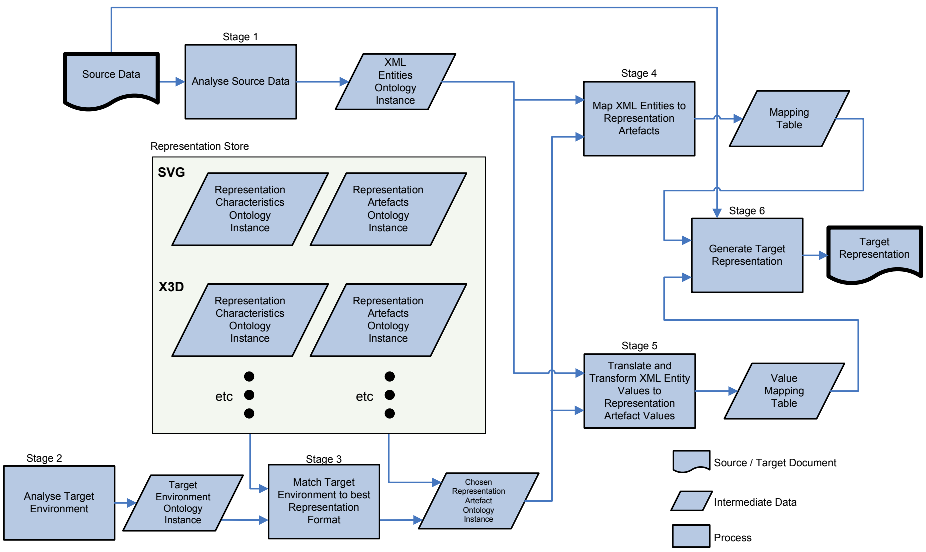

Figure 2 : The Information Realisation Process

The process is shown in figure 2, and each stage is described below:

1. Analyse Source Data - In order to create an XML Entities ontology instance (see Section 4.1).
2. Analyse Target Environment – The information about the nature of the target environment. The output is a Target Environment ontology instance (see Section 4.3).
3. Match Target Environment to Target Representation Format - The system compares instances of both ontologies for similarity. The output is the Representation Artefacts Ontology Instance (see Section 4.2) of the chosen representation format.
4. Map XML Entities to Representation Artefacts - The inputs are: the Representation Artefacts Ontology Instance from the previous stage and the XML Entities Ontology Instance. We use target environment factors, heuristics and past experience to create mappings between XML Entities and Representation Artefacts. The output is a mapping table between XML Entities and Representation Artefacts.

5. Translate and Transform XML Entity Values to Representation Artefact Values - In many cases the values of the source data can not be used directly in the Target Representation. Instead, a value mapping or translation process must occur. This is described in more detail in the worked example (Section 5).
6. Generate Target Representation - We generate the Target Representation by creating Representation Artefacts (in the target representation format) with the values supplied by the Value Mapping Table. The output of this stage is the final Target Representation.

4 Ontologies

In this section we present the three ontologies which are used during the Information Realisation process.

4.1 XML Entities Ontology

The ontology shown in figure 3 considers XML elements and attributes as a generalised concept called an Entity. This is similar to the "Layered Normal Form" which is one of the XML Normal Forms [Thompson, 2001]. An Entity has a value, a name, and an XPath. It also has a parent entity, and in some cases it has child entities. In this way, we can concentrate on the values and structure of the data rather than how it is represented.

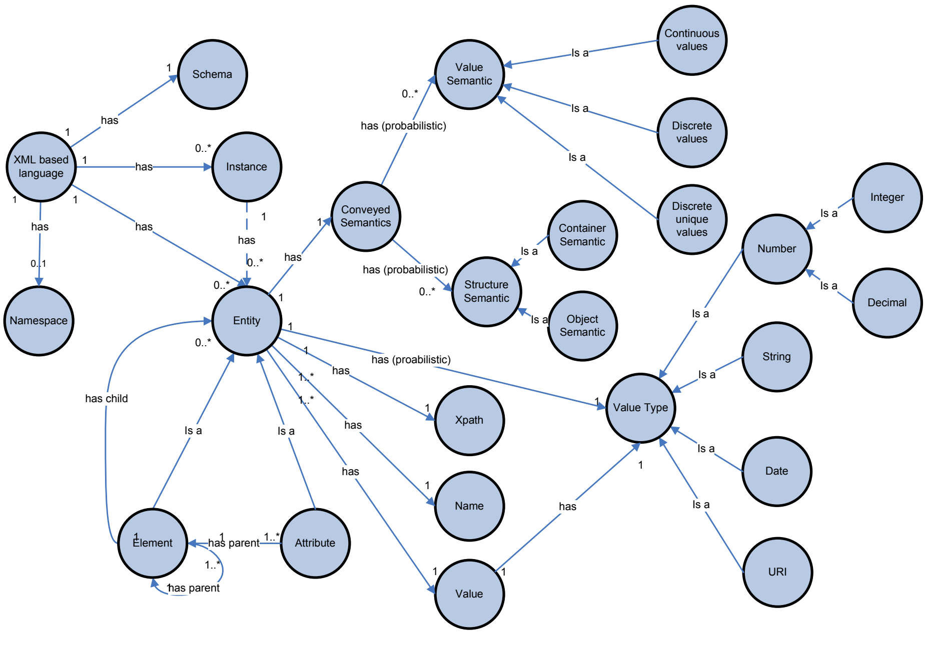

Figure 3 : The XML Entities Ontology

Value Semantics of the Entity are either: Continuous Values (e.g. age), Discrete Values (e.g. nationality) or Discreet Unique Values (e.g. name). Value Semantics are deemed to be probabilistic because the system may not be able to give a definitive Value Semantic categorisation without user intervention. A Structure Semantic is a probabilistic value that the Entity is a Container or an Object.

- **Container** is an Element which has child elements. It may have attributes, but usually has no value.
- **Object** is an Element which has no child elements. It may have attributes and/or a value.

4.2 Representation Artefacts Ontology

Each target representation language will have its features categorised as Representation Artefacts. In figure 4, we give examples of each Representation Artefact. White rectangles represents Graphical artefacts, shaded rectangles represent Audial artefacts.

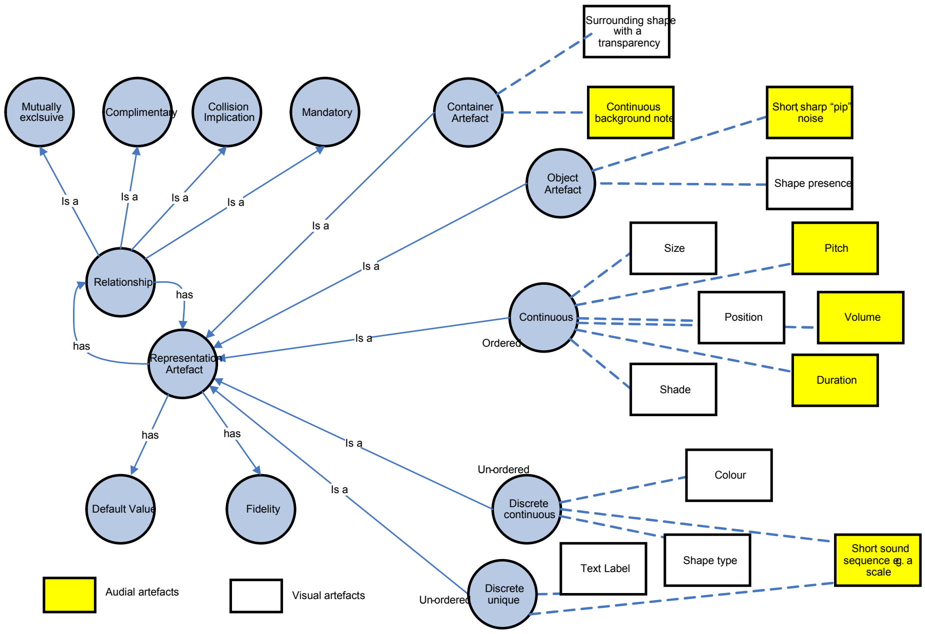

Figure 4 : The Representation Artefacts Ontology

4.3 Target Environment Ontology

This ontology is made up of:

- **User Abilities**. The sensory abilities of the user (sight, hearing, cognitive, motor control, etc).
- **User Situation**. The situation that the user will be in when presented with this target representation. For example, driving a car (cognitively engaged,

physically engaged), or using a computer terminal (cognitively available, physically available).

- **Technology Characteristics**. This includes details of the capabilities of the target medium. For example, a computer terminal is interactive and can be provided with a visual representation. However, a car's sound system is non-interactive and must be provided with an audio representation.

Further information about this ontology can be found within the appendix [WebAppendix, 2006].

5 Worked Example

The example concerns a set of sports fans represented in an XML file. The source data consists of 26 people (only 3 shown here) with their name, age, tallness ("tallness" is used rather than "height" to avoid confusion with the SVG attribute of the same name), nationality, scarves (number of scarves they own) and games (number of games they have been to). The source data is detailed below:

```
<fans>
    <person name="alice" age="28" tallness="1.41" nationality="welsh" scarves="2" games="19" />
    <person name="bob" age="37" tallness="1.02" nationality="scotish" scarves="4" games="33" />
    <person name="colin" age="16" tallness="1.84" nationality="irish" scarves="6" games="8" />
    ...
</fans>
```

The process begins by analysing the Source Data (section 3, stage 1). This results in an instantiated XML Entities ontology. The data is summarised in Table 2.

Entity Name	XPath	Element / Attribute	Value type	Value Semantic	Structure Semantic
fans	fans	element			Container (root)
person	fans / person	element			Object
name	fans / person / @name	attribute	text	Discrete unique	
age	fans / person / @age	attribute	numeric	Continuous	
tallness	fans / person / @tallness	attribute	numeric	Continuous	
nationality	fans / person / @nationality	attribute	text	Discrete	
scarves	fans / person / @scarves	attribute	numeric	Continuous	
games	fans / person / @games	attribute	numeric	Continuous	

Table 2: A summary of the information held in the XML Entities Ontology Instance

We then perform the Analyse Target Environment stage (Section 3, stage 2). In this example, the user has full sight, hearing, cognitive and motor control abilities. The user is at a computer terminal which has full interactive and display capabilities (800 by 600 pixels screen size). The system therefore matches the Target Environment to the SVG Target Representation Format (Section 3, stage 3). The output is a Representation Artefact Ontology Instance (Section 4.2). The data held in the SVG ontology instance is summarised in Table 3.

Entity Name	XPath	Element / Attribute	Value type	Value Semantic	Structure Semantic
svg	svg	element			Container (root)
rect	svg / rect	element			Object
x	svg / rect / @x	attribute	nums	Continuous values	
y	svg / rect / @y	attribute	nums	Continuous values	
width	svg / rect / @width	attribute	nums	Continuous values	

height	svg / rect / @height	attribute	nums	Continuous values	
fill	svg / rect / @fill	attribute	text	Discrete values	
title	svg / rect / @title	attribute	text	Discrete unique values	

Table 3 : A summary of the information held in the Representation Artefact Ontology Instance for SVG

The system then maps the sports fans XML Entities to the SVG Representation Artefacts (Section 3, stage 4). The mapping process is based on matching Value Semantics and Structure Semantics and is described further in future work. In this case, the mapping is relatively simple and is detailed in Table 4.

Sports fans XML Entity	SVG Representation Artefact
fans	svg
person	rect
age	x
tallness	y
scarves	width
games	height
nationality	fill
name	title

Table 4 : Mapping table between sports fans XML Entities and SVG Representation Artefacts

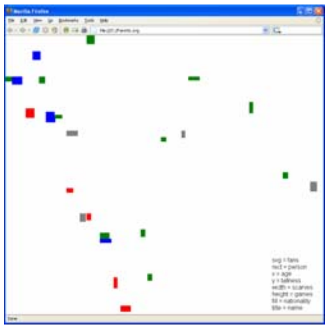

Figure 5 : The SVG target representation

The next stage is to Translate and Transform XML Entity Values to Representation Artefact Values (Section 3, item 5). This stage involves the processing of value mappings. For example, the system can not merely assign the *fill* attribute as the *nationality* of the person (e.g. Welsh). Instead there must be a mapping to the available values for the *fill* attribute. This is an example of an attribute which the user may decide to adjust if the assumption made by the system is incorrect. For example, if the system has an available choice of 10 primary fill colours, it may assign the *fill* value "blue" to the *nationality* value of "Welsh". The user would probably want to change this to "red" to more accurately reflect the traditional nationality colour of Wales.

Also, the values of *age, tallness, scarves* and *games* would need translating before being set to *x, y, width* and *height* respectively. These values would need to be scaled to the dimensions of the screen (in this case 800 by 600 pixels). Again, these are settings which the user may wish to adjust after examining the output.

The final stage is to generate the Target Representation (Section 3, item 6). This takes the Mapping Table, together with the original Source Data (the sports fans data) and generates the Target Representation in SVG.

The SVG target representation is shown in figure 5 and the code shown below:

```
<svg xmlns="http://www.w3.org/2000/svg">
  <rect title="alice" x="213" y="471" fill="red" width="12" height="18" />
  <rect title="bob" x="373" y="800" fill="blue" width="16" height="25" />
  <rect title="colin" x="0" y="109" fill="green" width="21" height="13" />
  ...
</svg>
```

A proof of concept system is detailed within the appendix [WebAppendix, 2006].

6 Conclusions

In this paper we have described a broader concept of information visualisation which we have termed Information Realisation. This encompasses textual, graphical and audial representation of data and can also scale to accept other representation formats should they become practical and common place. Our technique also considers a broader view of the target users, encompassing the environment of the user and the capabilities of the available technology.

At every stage of the Information Realisation Process, the system will have made assumptions about the nature of the data based on probabilistic decision making. As such, there are likely to be some choices which are not what the user would have chosen. To address this, the user has the ability to tweak settings and change any of the assumptions and decisions made by the system. In this way, the system learns from the user's intervention, thus creating a more accurate process for future realisations. As part of future work, we intend to investigate the use of data mining and machine learning tools, such as YALE [Ritthoff et al., 2001] for this process.

Acknowledgements

This work took place during a research fellowship at the Knowledge Engineering and Decision Support Research Group (GECAD) at the Polytechnic Institute of Porto, Portugal. This research is funded by a University of Wales Swansea Postgraduate Research Studentship under the supervision of Dr. Phil Grant and Prof. Min Chen.

The authors would also like to acknowledge FCT, FEDER, POCTI, POSI, POCI and POSC for their support to R&D projects and the GECAD unit.

References

[Ackerman, 1994] Ackerman, M.J.: "The Visible Human Project"; Proc. of the IEEE , 86, 3 (1998), 504-511.

[Berners-Lee and Fischetti, 1999] Berners-Lee, T. and Fischetti, M.: "Weaving the Web The Original Design and Ultimate Destiny of the World Wide Web"; Harper, San Francisco (1999).

[Geroimenko and Chen, 2006] Geroimenko, V., Chen, C.: "Visualising the Semantic Web"; Springer-Verlag, London (2006).

[Jung and Sato, 2005] Jung, E., Sato, K.: "A Framework of Context-Sensitive Visualization for User-Centered Interactive Systems"; Proc. of 10th International Conference on User Modeling, Springer-Verlag Berlin Heidelberg, Edinburgh (2005), 423-427.

[Ritthoff et al., 2001] Ritthoff, O., Klinkenberg, R., Fischer, S., Mierswa, I., and Felske, S. "YALE: Yet Another Machine Learning Environment"; In LLWA'01/FGML-2001 – Proc. GIWorkshop- Woche Lernen – Lehren – Wissen – Adaptivität (2001), 84–92.

[Thomas and Cook, 2006] Thomas, J.J, Cook, K.A.: "A Visual Analytics Agenda"; IEEE Computer Graphics and Applications, 26, 1 (2006), 10-13.

[Thompson , 2001] Thompson, H.: "Normal Form Conventions for XML Representations of Structured Data"; Proc. of XML 2001, IDEAlliance, Alexandria, VA USA.

[WebAppendix, 2006] http://www.cs.swan.ac.uk/~csowen/InformationRealisationAppendix/.