Common Sense Design: Aided Visualisation of the Semantic Web

Gus Gollings

Globalism Institute,
School of Global Studies, Social Science and Planning
RMIT University, Melbourne, Australia
gus.gollings@rmit.edu.au
http://globalism.rmit.edu.au/

Abstract. This paper provides an outline for a general purpose Semantic Web browser which aims to provide a sensible visualisation for any RDF data. Sketched here is a web proxy for common browsers which interlopes on RDF/XML transmissions, transforming them at 'diplaytime' to provide the viewer with a meaningful view on the intercepted model. It suggests using the Wikipedia-style *million minds* approach to developing two large repositories, one of common sense knowledge about design principles, and another of design templates using the Fresnel pattern. An 'intelligent' agent could make use of these two repositories, as well as personalisation information and RDF introspection, in the rendering of useful interpretations of raw RDF material.

1 Introduction

In the near term, before the rise of agents and pervasive RDF documentation of the modern world (such is the vision for the Semantic Web [1]), the kind of user interaction experience the Semantic Web might offer can be seen in several traditional web-applications—Google's venerable 'Page Rank', Amazon's gratifying book recommendations and the tag-based selection of blog entries offered by most blog aggregators, to name just a few obvious examples. There are also early examples of 'real' Semantic Web user interaction, the 'first' Semantic Web site, Mindswap [2], the lesser known FOAFnaut [3], and moreover, there are hybrid systems like the breakthrough CiteSeer [4, 5], which has all the hallmarks of a Semantic Web system without actually being one [6]. The examples of 'real' Semantic Web applications, as well as the traditional web-applications employing 'semantic' techniques, all utilise handcrafted and pre-prepared visual presentations from centralised and domain specific positions—and this will continue be a common approach in designing for the Semantic Web, particularly for popular or specific field needs, for example, a popular need is a photo album and a specific need is medical proteins investigation interface. It is expected, however, that there will also be a need in between the popular and the specific to interact with Semantic Web data which is of a 'general interest' nature, has not been pre-arranged, as it were, and where ideal data representations have not been selected on behalf of the person engaged in viewing. This paper proposes the design of a *Common Sense Design Agent* with particular features for making it possible to meaningfully explore this latter area of 'uncharted' Semantic Web data.

This paper is arranged in two parts, with book-ends of this Introduction and a *conclusion*. The first part, *background*, sets the context for this paper and provides a brief literature review. The second and major part, *project overview*, introduces the idea of using a 'common sense' knowledge repository to drive dynamic presentations of RDF models at 'displaytime' to facilitate the visual investigation of new associations between uncommon class combinations.

2 Background

There is a rich history behind today's traditions of modern Western human communication, from the times of primary oral cultures, to the development of alphabetic writing, to the radicalisation of the written form in printed typographic material, through to our present information age. The most striking observation about this history of communication is that alongside each epoch in the progression from oral, to written, to printed, to 'networked' society, are identifiable transformations in the characteristics of social thought and expression [7].

At a personal level, in the recent past we would characterised an author as deciding how to present their ideas and data on the page. They may have selected from any one of many organising metaphors, from prose paragraphs, to unordered lists to scatter plots, to best represent their data and make a communicable meaning for an intended audience. The audience would receive in an essentially fixed way the material given by the author.

The classic Web, to some degree, and more so the Semantic Web now bring a new dimension to reading, and that is that active construction of meaning from 'meta-analysing', which is in some senses a more agile approach to data mining than we are used to with today's brute force text mining tools, like Google's search engine. The active construction of meaning (or paths to meanings) through reasoning about Semantic Web data associations is notably different to the historical generation and distribution of knowledge. It makes abstract the serialisation process of the 'raw data' from any point of authorship, and in this new arrangement, each person viewing the information becomes the author, or arranger of data being interrogated. But to the point, a person needs tools to do this work or arranging. The tools designed for this task to date have been quite literal in their approach, revolving around what is possible in terms of past graph visualisation practice on the one hand, and the possible flat formats (RDF itself, and plain text lists and tables) on the other—neither of which necessarily capture the spirit of the data as arranged by the modern day 'viewer-author'.

The ways of viewing information are themselves standards, socials contracts, cultural imperatives, dialectic, and part of the broader set of habits surrounding language and its communicable meanings. There is a distinction between an artful visualisation of information that is ordinarily 'made to order', and the canonical, established, archetypal representation for a given type of information that is a recognisable form. The former is viewed with the loose embrace of a light conversation, its meaning doesn't need to be apparent. The latter takes on the form of an insult when poorly executed, or deliberately interupted—for example it makes one feel awkward to navigate on a map of the globe which has been turned upside down, books without paragraph breaks disturb the reader's concentration, and so on.

Writing, and especially typographic (printed) writing, has produced an exhaustable list of basic information encoding forms which visually aid the distribution of meaning contained withing certain common data structures:

- Editions (version control)
- Leaves (page numbers)
- Sections (chapters)
- Lists (trees)
- Charts (tables)
- Figures (plots)
- Dictionaries (indexes, table of contents)
- Maps (cartography, schemata)

Printing has particularised the collection and organisation of texts, so that institutions and conventions inside and outside the console of the text have been developed. Inside the book we have tables of contents, publication information and topic indexes. Outside the book we have location systems like the Dewey classification, which problematically doubles as a subject classification, so the second part of a three volume series might not be stored alongside the first and third volumes on the shelf. The world of text is a world of standards. In some ways what we are engaged in designing when we talk about designing for the Semantic Web, is a series of standards.

This paper understands that the predominant organising metaphor for user interfaces design on the World Wide Web has been formed in the opinion of Marshall McLuhan's idea of media as an extension of man [8], and in that conflation, every user interface is built as an individual, and an extension of the corporate or individual persona. Friedrich Kittler, the German media theorist famous for his work on *technischen medien* (new media, or technical media) debunks this idea. 'Media are not pseudopods for extending the human body. They follow the logic of escalation that leaves us and written history behind it.' [9]. Instead of different interfaces on essentially the same information, there is evidence that what is desired is different information through the same interface (aggregators, newsreaders, search over browse, and so on).

As such, the information age is not a moment of extending mans ulterior connection to media, or that through our use of computers we were extending ourselves, but it is a transitionary moment in the way society expresses itself which is part of a longer tradition of communication development from oral to chriographic, from chriographic to typographic, and this latest moment moving from typographic to 'network' (a base also referred to as *'information'*, or *'mulitliteracies'* [10, 11]).

2.1 Brief literature review

To note, there are already systems for general purpose RDF browsing and/or viewing: the MIT Simile project with its Longwell (facetted) and Welkin (arc-node) viewers [12], Aduna's Sesame explore mode feature (tabular) [13] and the World Wide Web Consortium's IsaViz (dynamic graph) [14]—which in turn is using ATT's GraphViz [15] to create plots of arc-node graphs. Like the Horus project at the Freie Universität Berlin [16], some of the above browsers are so-called 'new breed' RDF visualisation tools, which can take advantage of the innovative browser-independant Fresnel specification [17] for applying 'lenses' and 'formats' to nodes of a graph with the effect of styling them for display.

This specification is a welcome addition to the Semantic Web family as it provides a means to share common visualisation configurations between rendering agents and across media. This adds value to the work people put toward crafting views of RDF classes and encourages sharing and reuse. More subtly, the Fresnel specification provides a basis from which the dynamic generation of 'lenses' and 'formats' can be made at 'displaytime' for the visualisation of previously unseen models (commonly encountered in the compilation of new associations between graphs).

```
:Chris rdf:type foaf:Person ;
   foaf:name "Chris Bizer" ;
   foaf:mbox <mailto:chris@bizer.de> ;
   foaf:mbox <mailto:bizer@gmx.de> ;
   foaf:depiction <http://www.../Chris.jpg> .
```

A.

B.

Fig. 1. Fresnel: the best of both worlds, computer readable markup with a human understandable presentation. From A. Conceptual Space (Notation 3 syntax for RDF of `foaf:Person`) to B. Visual Space (possible output using Fresnel). *From* [17].

Vladimir Geroimenko and Chaomei Chen's edited collection 'Visualizing the Semantic Web' [18] demonstrates the wide range of possibilities for visualising Semantic Web data due to its highly uniform nature, and highlight some of the cross-overs and divergence between the information visualisation and information management disciplines. David Duke refined much of the same argument in his earlier work documenting the relationship between physical space visualisation metaphors and abstract space visualisation metaphors [19]. Much work has also been done on the variety of approaches to RDF representations, for example [20, 21].

Suffice to say that it is possible to imagine mapping the several most useful visualisation techniques to certain kinds of RDF models, or even describing the conditions under which certain visualisations would be useful. Moreover, it is also conceivable that each visualisation, or approach to visualising, could be enshrined in a Fresnel 'lens' and 'format' pair with accompanying CSS and XSLT scaffolding. These two ideas form the foundation of the following section which describes the outline of a default Semantic Web browser.

3 Project Overview

Following is a description for an ambitious set of parts making up an envisioned general purpose Semantic Web browser experience. Its overriding goal is to produce an RDF/XML viewer which enables the visual tuning of new associations between unlike sources. Mediated by a Semantic Web-style agent accessed through any one of today's existing web browsers, this system is proposed to have the capability of making visual meaning out of the purely conceptual relations found in RDF documents, from any source. It suggests using an agent to marshal between an artificial intelligence repository of common sense knowledge about design principles coupled with a large repository of canonical Fresnel templates used for transforming RDF statements about particular classes of things into display data. It is envisaged that the agent take cues from personalisation preferences when compiling display information from, for example, the results of searching a hybrid RDF/HTML web search engine.

3.1 A Common Sense for Design

John McCarthy's now classic 1959 paper 'Programs with Common Sense' [22] discussed the design of a common sense reasoning program called the *advice taker*. This research marked the beginning of the the idea that heuristics for logical responses to 'common sense' issues could be stored outside of the program itself, and it produced a view of common sense that it was knowledge that was able to be learned and once learned, held captive for later re-use. This was an important point and it developed the understanding of using common sense for computing applications as the antithesis of the more general artificial intelligence challenge typified by expert systems where a lot is know about a small topic:

> Common Sense: The mental skills that most people share. Common sense thinking is actually more complex than many of the intellectual accomplishments that attract more attention and respect, because the mental skills we call "expertise" often engage large amounts of knowledge but usually employ only a few types of representations. In contrast, common sense involves many kinds of representations and thus requires a larger range of different skills. [23]

Judea Pearl's insightful 1988 book 'Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference' [24], began a new specialisation in the use of probability and utility theory to improve software design, an important consideration for common sense systems. By far the largest common sense repository to date is Douglas Lenat's 'CYC' [25] which captures common sense knowledge in a rigourous way, with skilled authors of such statements employed to convert knowledge into the 'CYC' frame of reference, in much the same way language dictionaries and encyclopædias are typically made. A more recent approach is Push Singh's 'Open Mind Common Sense' [26] repository, which collects common sense statements from the general public in typed natural language. This has the effect of allowing anyone to contribute to the repository, but means that the interpretive work of figuring strong semantics from the repository is loaded onto heuristic and algorithmic text mining processes behind the scenes.

Design by numbers. There is an interesting tradition in graphic design of a belief that design is principled—in that a design is a solution to a communication problem, and in that way it is formulaic. This belief in the formula of design is mainly embodied in a purely rhetorical way within the design communities, through any number of common design lessons like 'scale/tension/dynamics'.

In recent times this belief has also been explored in a more literal and radical way by practitioners using computers to solve design problems. Muriel Cooper, John Maeda and Joshua Davis have all pursued this idea with various programs. Under Cooper's direction, the MIT Media Lab Visible Language Workshop [27] created dynamic magazine page layout routines, Davis has for several years worked on 'generative' patterns to create prints [28] and user-interface layouts in a similar vein to the concepts behind Maeda's *Design by Numbers* work [29]). The experience of electronic typesetting has also borne the message that design can be reduced to a set of principles, for example, *the first line of a body paragraph should be indented when it follows another body paragraph*, or similar. And so it stands that this kind of procedural design knowledge is not simply creative and irreducible. It is therefore feasible as a design community to encode these design principles into an artificial intelligence repository for the *Common Sense of Design*, in a similar fashion to the way we have developed the knowledge contained in Wikipedia.

We also want to build in the opportunity to use and experiment with some of the wide variety of canonical and emergent graphical representations of RDF models.

Rise of common sense applications. Massachusetts Institute of Technology's Media Laboratory has over the last few years developed some compelling examples of common sense technology, including video editing recommendation systems and photographic library mining to augment the context of an author's current text, useful for writing emails.

One of the key researchers from this school, Henry Lieberman, has made the observation that common sense artificial intelligence can be used effectively today, even with relatively small reservoirs of knowledge, if it is used as an assistant to a primary task that could anyway be carried without any intervention. In that way, if you have a fallback position, the inadequacy of the artificial intelligence is not a blockage to progress, but any contributions it may make are felt as welcome help.

The architecture of common sense programs can take several forms, from neural networks to highly procedural approaches. The concept of using so called critics to weigh in on the reasoning about a situation enable the creation of self reflective machinery, where one critic will take actions, where another critic will reflect on the outcome of those actions and modify the decision chain so that subsequent action can take advantage of previous experiences known as stories. This very structured approach in making sensible automation is ideal for margin tasks like preparing interfaces for RDF models as the consequences are always low for bad decisions, and there is a great opportunity for human guidance in the analysis of past actions and development of experience stories simply through the observation of natural browsing tendencies, and iterative refinement of initial results—the so called *million minds* approach of continuous improvement.

Common sense of design knowledge. A knowledge repository of common sense can be successfully built by a lay community of interest. Such was the development of the Open Mind [26] catalogue of everyday relationships between things, e.g. *shampoo is found in the shower*. By using natural language as the input, there is a very low barrier to community contribution and the knowledge base can grow quickly. The natural language is then processed with an arsenal of text mining tools to arrive at a computer readable network of connections.

Creating a repository of common sense design knowledge in the way of the Open Mind model is an interesting possibility to bridge the gap in the human comprehension of RDF data. The analogy being that a human trying to understand RDF data is like a computer trying to understand HTML data: understanding is possible to some degree in both cases, but some translation needs to be done first. In the case of computers understanding HTML some guesswork based on the common sense of HTML markup must first take place, *an `H1` tag signals an important concept on the page*. Similarly, for a human to understand RDF, some guesswork must take place based on the common sense understanding of what the RDF is representing, *a `foaf:Person` class is a concept that looks like an address book entry*.

The pragmatic answer for the person with an interest in viewing, or comprehending, RDF data is to apply their own or another pre-prepared Fresnel accoutrement (remembering Fresnel provides a means to modify and arrange RDF data in preparation for visual display suitable for humans). We will still encounter a display hurdle in the situation where there is no pre-prepared Fresnel 'lens' or 'format'. And even where there is a pre-existing Fresnel view design, it is likely that there will be small changes desired from person to person and from node-set to node-set [30]. It is from this characterisation of the viewing requirements for RDF data that the suggestion for using common sense computing in the display of RDF is made.

In the case of computers comprehending HTML, there is a translation process from *visual* to *semantic*. In the case of humans comprehending RDF, there is translation process from *semantic* to *visual*. For the case of a human trying to understand RDF data, there is another aspect beyond the literal comprehension of what the RDF itself is representing, and that is the notion of what relevant relationships there are to the data under scrutiny. For example, if we have a series of geographical latitude and longitude co-ordinates there is a relevant relationship between that information and map image information, and the kind of relationship, if you will, is the plotting of the geographic co-ordinates on the map, along with labels for each pair of co-odinates. The obscure relationships between things may often only become relevant in light of the particular goals of the viewer. Relationships themselves may also be goals, where open-ended reasoning is required to 'look' for the the possible connections between disparate sources from various places. All of these aspects of a meaningful design would need to be considered, as not all requirements could captured in a repository of common sense design knowledge.

Canonical Fresnel repository. The Fresnel specification [17] makes great strides in the basic issues surrounding the presentation of Semantic Web data. Fresnel is an important development for two reasons. First, it standardises the way we describe the display of RDF across different browsers (and its output is not limited to text-based media, it could also form the basis of an interactive SVG animation, for example). This is similar to the enabling common ground of HTML for the classic Web. Second, it allows us to make declarations about how to design instances of abstract and unstructured concepts, including how to design for the relationships between such concepts, which as an example might translate to a set of dates being plotted on a timeline.

To complement the repository of common sense design knowledge, a dictionary of reusable Fresnel code could be developed, again as a community endeavour. This repository would contain both canonical Fresnel treatments of particular RDF classes (like the `foaf:Person` example) and also templates that could be used and combined in different situations for a given design goal. For example, a 'lens' describing the pertinent parts of a timeline and a corresponding 'format' to compile a display, could handle any number of different classes needing to be represented on the timeline, as might be determined by the common sense design knowledge repository described above.

The Fresnel specification defines a vocabulary of core RDF display concepts, based on two metaphors: that a 'lens' is used to select from a complete graph only those nodes and properties that are of interest to a viewer, and a corresponding 'format' for that selection which describes how to display that set of nodes and properties. A 'lens' and 'format' together create a more-or-less static view, or serialisation, of a particular node class in an RDF graph.

It is also possible to target untyped nodes by selecting them directly with an FSL [31] statement (FSL is a query language for graphs similar to XPath for XML [32]), or with a single class SPARQL [33] query. A strategy to customise the display from a 'lens'/'format' pair is provided using special Fresnel declarations: sub-lenses (`fresnel:sublens`), alternative formats (`fresnel:alternateProperties`) and merging properties (`fresnel:mergeProperties`), but these approaches are still not entirely suitable to cater for various interests and needs requiring particular layouts at different times for different reasons.

Fresnel remains a strategic design solution that comprehensively answers to its aims [17]:

- useful for rendering different output formats like HTML, SVG, PDF, plain text, and others,
- applicable across different RDF display paradigms, (nested box-based textual representation à la XHTML+CSS, node-link diagram, etc.),
- built on existing Web technology,
- extensible for more specialized needs,
- easy to learn and use.

While Fresnel doesn't to attempt to answering the final challenge of how a Semantic Web browser might look, it does provide the means to share view-decriptions between rendering agents. So it is plausible to invest in a shared, centrally located, database of Fresnel 'lens' and 'filter' code which covers many of the display strategies and techniques available for many of the common classes.

These Fresnal templates, if you will, act as the starting point for the *Common Sense Design Agent*.

3.2 Common Sense Design Agent

The Common Sense Design Agent aims to enable the automatic composition of visual representations of Semantic Web data (RDF/XML) and display them with an ordinary classic Web browser (Firefox, Internet Explorer and their like). The *intelligence* of the Common Sense Design Agent is modelled on a simplified version of Push Singh's Emotional Machine, or EM-ONE, and its architecture for reflective common sense thinking [34]. It could also take other forms of artificial intelligence—without making a firm claim on the implementation of such a system, it suffices to say that the idea is to provide some intelligence in the representational design of RDF/XML, as the connections between ontologies will grow and emerge asychronously to the development of Fresnel definitions, or any other visualisation to display them.

Practically, the agent acts as a transparent *web proxy* for classic web data—HTML and its supports—and where it finds RDF/XML content it analyses the document(s) and transforms it for display in the browser (commonly as HTML, but possibly as SVG or any other format). The analysis will try to incorporate any of the following aspects if they are available: preference settings given by the user, common sense design knowledge and Fresnel 'lenses', 'formats' and associated stylesheets which directly target classes from the RDF to be viewed. The

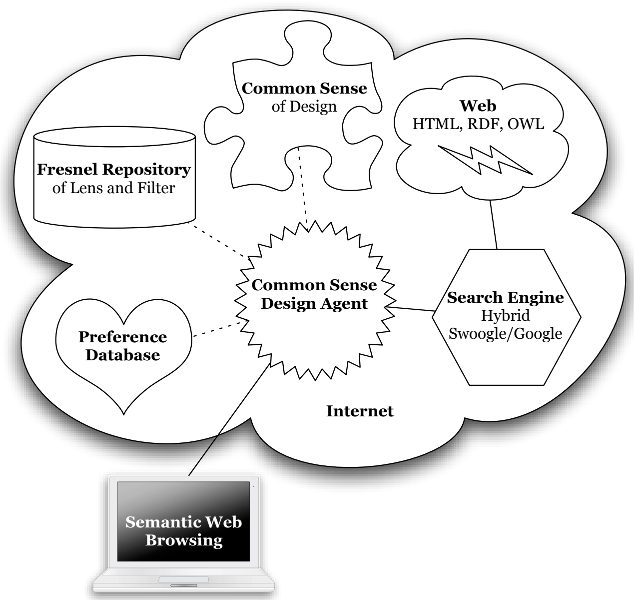

Fig. 2. Schematic of the Common Sense Design Agent

common sense design knowledge will match with properties of the RDF being transformed and the users preferences or goals, and enable the dynamic creation of Fresnel lenses and formats to manipulate the RDF/XML into meaningful assemblies brought to life with dynamically generated stylesheets.

Knowing how to look. Because the RDF data model is abstract and unstructured, RDF instance data can take wild forms, which don't naturally lend themselves to sensible serialisations. As an investigator, you will always see in a scene what you want to see, so your interest can help to guide the display to be more specific about what to show you. Furthermore, the available data (set of classes, and the sets of each class's properties—as well as untyped subjects) can provide context for the selection of an appropriate presentation action. Other vectors for the selection of a graphic design principle are:

- *Inferred:* 'Kinds' of subjects, class membership, domain and range of particular properties.
- *Declared Interest:* property weights, literal search topics, Haystack-style 'recommendations' [35]
- *Common Sense:* Mapping a central repository of Fresnel-style layouts information to common design situations

Ordinarily, faced with a task of comparing the frequency of publications on a particular topic over time between three cities would require finding source data for all publications on the queried topic, filtering a set of data for each city of origin and then plotting the filtered sets on a timeline for comparison. It was common sense that told us that a comparison could take the form of a composite rendering of several sets of data on the one plane, and that mapping time-based data could be achieved on a timeline, and that timelines have a domain and range and labels.

It is suggested that this kind of task could be achieved by a computer in two composite passes. First, data selection is made in the fashion of Noadster's global view [36], via a query which may be successively refined based on initial results, or reasoning based in personal preferences. Second, once the data is isolated, it can be modified to provide a common sense representation based on the content of the isolated data-set and user goals, again, this representation of the data can be iterative, based on successive refinements according to a person's taste and need.

Hybrid search engine. How do you find RDF files to view in the first place? One answer among several is the suggestion that a hybrid search engines returning pointers to Semantic Web code as well as classic Web pages could be a simple way to provide access to the growing wealth of RDF resources and ontology descriptions becoming available. Of course, the RDF could be found and investigated by other means as well.

It is expected that the classic Web will continue to operate the way it does, and that people using it will continue using it the way they do. The development of the Semantic Web will mean that more and more Semantic Web code is available on the classic Web. While the Semantic Web code will be used behind the scenes, so to speak, to enable a whole new class of highly inferential applications, that same code may be viewable by people, and with the aided visualisation described in this paper, might provide an important *default* Semantic Web browser in its own right, notwithstanding the many excellent domain specialisation browsers.

Presentation personalisation: honing common sense. There are many ways to view RDF data, but the most meaningful will always be subjective and informed by personal preferences and goals.

In terms of display strategies and techniques, there is a growing number of demonstrations of specialised approaches to the visualisation of Semantic Web data (several of which have been summarised by Ivan Herman, Guy Melançon and Scott Marshall [21]). Sometimes these are old techniques brought to bear on new requirements, and sometimes they are fresh approaches, ranging from sorted serialisations of basic RDF statements, to the display of RDF objects on a geographical map with named subjects, to the automatic generation of a calendar of events, to interactive 'fish-eyed' graphical plots of entire RDF stores.

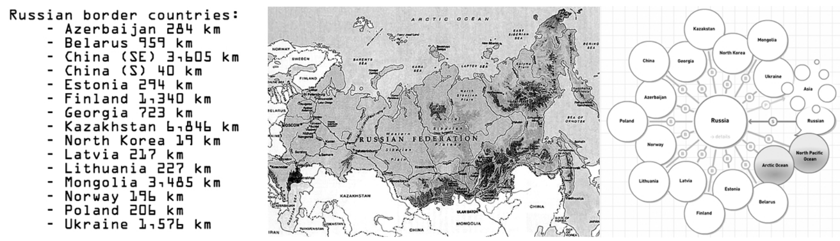

Fig. 3. "How would you like your Russian borders today? As a list, a map, a plot. . . "

Individually, each of these approaches is useful as a display generalisation, in so far as each approach can take as input cross-domain RDF data. However, the sense-making validity of each approach is contingent upon what it is that the person wants to see—their personal interest or preference. For example, if you want to know where each attendee of a particular conference came from, then it would make sense to view this information on a map, but if you wanted to know who attended a particular conference, it would make sense to see a list of attendees sorted by surname displaying first name, surname, affiliation. On the other hand, if you were interested in the connectivity between several major genetics databases, then a comprehensive view of several ontology domains possibly with some reasoning to simplify connections prior to display might be most useful. So a person's interest is a determining factor in the most appropriate display of Semantic Web data.

The model for personalisation used in the design of the Common Sense Design Agent draws on the work done for the The Personal Publication Reader [37] and CHIP [38] projects, countenancing several factors of personal preference among previous browsing context, query, check-box preference setting and property weighting.

Significant numbers of research projects framing approaches to the personalisation of presentation and visualisation aspects of data retrieval have been undertaken over the last decade (see, for example, [39, 38]). This research commonly relies on either of, or combinations of, *context of behaviour analysis and inference* or *domain programming and user preference declaration*. Context of behaviour analysis alone is sometimes less effective compared with layered approaches, because once a context is identified, the system will often still rely on a decision of the system's programmer to arrange for the identified context. If the programmer tries to assuage this concern by not taking steps to modify the layout substantially for a given context, then there is a reduction in the general applicability of the system to display a variety of different information spaces. Moreover, user preference declaration leaves the important and sometimes very detailed requirements of property weighting or the difficult demands of layout trajectories to the person using the system, someone who is possibly disinterested and unskilled in the these tasks.

Nonetheless, the examples of the Flamenco [40], Topia [41], mSpace [42] and Noadster [43] programs show how close analysis of the RDF store along multiple assessment criteria can significantly aid structuring of tables of contents and prepare layout of domain independent material. Another approach to the customisation of the display for RDF data achieved through a means of pre-existing, hand-crafted, ground-rules to inform the display logic is shown in the proposal for the Semantic Channel Specification Language [44] and more recently in Dennis Quan and David Krager's discussion of the Haystack system [35].

Based on analysis of the alternative ways in which Semantic Web data can be represented graphically and in serialisation, it is apparent that data style and the interests of the person making the inquiry of the data determine to a large degree the pertinent display of that information. For example, Figure 3 shows three reasonable ways to present information about the countries sharing borders with Russia. The most useful display out of the reasonable presentations of the information is left to the discretion of the person at the console.

4 Conclusion

The Semantic Web should not be a friend who endlessly puts new indexes, dictionaries and thesauri on your desk. It should instead be a friend who quietly makes intelligent suggestions to you, which you can follow in a pinch or leave aside. The ideals of user interaction for the Semantic Web are characterised by the type of data at hand (scientific, educational, informational, recreational, etc.), the associations available to that data or the associations being made (sometimes data has very little *natural* interconnectedness) and the person's goal in using the Semantic Web (comprehension of an entire field of data or answer to a particular question).

Looking at information structures that we are not accustomed to retards the perception of the information being served. The consequence of this visual discomfort can be seen on today's web, where communities of users have naturally tended toward the development of presentation standards. Web pages have become more and more the same in terms of their basic information delivery characteristics. For example, we will commonly find a web page to have discreet parts for a heading, navigation, content and footer. This consistency enables people to confidently read across unlike sources. Similarly, through repeated usability refinement, competing search engine interfaces and functions have become more similar, allowing users to reuse their existing understanding of operation and navigation metaphors. Semantic Web data, and in particular the dynamic creation of new associations from disparate data sources, presents a dilemma for interface and interaction metaphors.

The challenge then remains to adapt to a general pattern for 'information based knowledge discovery'. One of the hypotheses of the Semantic Web is that there is an enormous amount of new information, or information waiting to be inferred, sitting just below the surface of traditional web querying capacities. Existing examples of the benefit of these inference techniques are apparent in our earlier examples of Google's 'page rank', in NEC and Penn State's CiteSeer, and in Amazon's product recommendations.

To build an application level common sense of how to design for any scenario—how to design in principle—we endow the Semantic Web with the ability to grow without having to wait for new data presentations, to sensibly display new associations between data that have not been pre-arranged and to co-exist with the much vaunted classic Web. This confirms a principle of being able to browse any data anywhere.

Manuel Castells concludes his three volume treatise concerning 'The Information Age, Economy and Society':

> The most fundamental political liberation is for people to free themselves from uncritical adherence to theoretical or ideological schemes, to construct their practice on the basis of their experience while using whatever information or analysis is available to them, from a variety of sources. [45]

Castells' presage to the information age is that the individual's liberation is predicated on their access and control over primary sources of information. In a practical way, the Semantic Web offers just such an opportunity.

References

1. Berners-Lee, T., Hendler, J., Ora, L.: The semantic web. Scientific American (2001)
2. Mindswap: http://www.mindswap.org (1996–2006) Maryland Information and Network Dynamics Lab Semantic Web Agents Project. Accessed August 2006.
3. FOAFnaut: http://www.foafnaut.org (2002–2006) The SVG-based FOAF user interface. Accessed August 2006.
4. CiteSeer: http://citeseer.ist.psu.edu (1999–2006) Scientific Literature Digital Library. Accessed August 2006.
5. Lawrence, S., Giles, C.L., Bollacker, K.: Digital libraries and autonomous citation indexing. IEEE Computer **32**(6) (1999) 67–71
6. Petinot, Y., Giles, C.L., Bhatnagar, V., Teregowda, P.B., Han, H., Councill, I.G.: A service-oriented architecture for digital libraries. In: International Conference on Service-Oriented Computing. (2004)
7. Eisenstein, E.L.: The Printing Press as an Agent of Change: Communications and Cultural Transformation in Early-Modern Europe. Cambridge University Press, Cambridge (1979)
8. McLuhan, M.: Understanding Media: The Extensions of Man. Routledge, London (1964 (2001))
9. Kittler, F.A.: Geschichte der kommunikationsmedien. In Huber, J., Müller, A.M., eds.: Interventionen 2: Raum und Verfahren. Museum für Gestaltung, Zürich (2002)
10. Cope, B., Kalantzis, M., eds.: Multiliteracies: Literacy Learning and the Design of Social Futures. Routledge, London (2000)
11. Castells, M.: The Power of Identity. Volume 2 of The Information Age: Economy, Society and Culture. Blackwell, Cambridge, MA (1997)

12. SIMILE: http://simile.mit.edu (2004–2006) Semantic Interoperability of Metadata and Information in unLike Environments. Accessed August 2006.
13. Broekstra, J., Kampman, A., van Harmelen, F.: Sesame: A generic architecture for storing and querying RDF and RDF schema. In: The International Semantic Web Conference. Volume 2342 of Lecture Notes in Computer Science., Springer (2002)
14. IsaViz: http://www.w3.org/2001/11/isaviz/ (2001–2006) A Visual Authoring Tool for RDF. Accessed August 2006.
15. Graphviz: http://www.graphviz.org (–2006) Graph Visualization Software. Accessed August 2006.
16. Horus: http://www.wiwiss.fu-berlin.de/suhl/bizer/rdfapi/tutorial/horus (2005–2006) A simple RDF browser which supports the Fresnel display vocabulary. Accessed August 2006.
17. Fresnel: http://www.w3.org/2005/04/fresnel-info (2005–2006) Display Vocabulary for RDF. Accessed August 2006.
18. Chen, C.: Top 10 unsolved information visualization problems. IEEE Computer Graphics and Applications **25**(4) (2005) 12–16
19. Duke, D.: Modular techniques in information visualization. In Eades, P., Pattison, T., eds.: Australian Symposium on Information Visualisation, Sydney, Australia, ACS (2001)
20. Sayers, C.: Node-centric RDF graph visualization. Mobile and Media Systems Laboratory, HP Labs (2004)
21. Herman, I., Melançon, G., Marshall, M.S.: Graph visualization and navigation in information visualization: A survey. IEEE Transactions on Visualization and Computer Graphics **6**(1) (2000) 24–43
22. McCarthy, J.: Programs with common sense. In: Proceedings of the Teddington Conference on the Mechanization of Thought Processes, London, Her Majesty's Stationary Office (1959) 75–91
23. Minsky, M.L.: Society of Mind. Simon and Schuster, Inc., New York, NY (1985)
24. Pearl, J.: Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1988)
25. Lenat, D.B.: CYC: A large-scale investment in knowledge infrastructure. Communications of the ACM **38**(11) (1995) 33–38
26. Singh, P., Lin, T., Mueller, E., Lim, G., Perkins, T., Zhu, W.: Open mind common sense: Knowledge acquisition from the general public. Proceedings of the First International Conference on Ontologies, Databases, and Applications of Semantics for Large Scale Information Systems **2519** (2002)
27. Abrams, J.: http://www.aiga.org/content.cfm?contentid=655 (1997) Muriel Cooper's Visible Wisdom. Accessed August 2006.
28. Maeda, J.: Creative Code: Aesthetics + Computation. Thames and Hudson, New York (2004)
29. Maeda, J., Antonelli, P.: Design By Numbers. The MIT Press, Boston (2001)
30. Quan, D., Huynh, D., Karger, D.R.: Haystack: A platform for authoring end user semantic web applications. In: 2nd International Semantic Web Conference (ISWC). (2003) 738–753
31. Pietriga, E.: http://www.w3.org/2005/04/fresnel-info/fsl-20050726 (2005) Fresnel Selector Language for RDF.
32. Clark, J., DeRose, S.: http://www.w3.org/tr/xpath (1999) XML Path Language (XPath) version 1.0.
33. Prud'hommeaux, E., Seaborne, A.: http://www.w3.org/tr/rdf-sparql-query (2005) SPARQL Query Language for RDF.

34. Singh, P.: EM-ONE: An Architecture for Reflective Commonsense Thinking. PhD in Computer Science, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Massachusetts (2005)
35. Quan, D., Karger, D.R.: How to make a semantic web browser. In: WWW2004. (2004)
36. Noadster: http://homepages.cwi.nl/ media/demo/noadster (2004–2006) Domain-Independent Generation of Full Navigation and Document Structure from Semantics. Accessed August 2006.
37. Reader, T.P.P.: http://www.personal-reader.de (2004–2006) Research Center L3S and Information Systems Institute. Accessed August 2006.
38. CHIP: http://www.chip-project.org (2005–2006) Cultural Heritage Information Personalization. Accessed August 2006.
39. Staab, S., Angele, J., Decker, S., Erdmann, M., Hotho, A., Maedche, A., Schnurr, H.P., Studer, R., Sure, Y.: Semantic community web portals. Computer Networks **33**(1–6) (2000) 473–491
40. Flamenco: http://flamenco.berkeley.edu (2004–2006) FLexible information Access using MEtadata in Novel COmbinations. Accessed August 2006.
41. Rutledge, L., Alberink, M., Brussee, R., Pokraev, S., van Dieten, W., Veenstra, M.: Finding the story: Broader applicability of semantics and discourse for hypermedia generation. In: Hypertext. (2003) 67–76
42. Gibbins, N., Harris, S., Schraefel, M.: Applying mspace interfaces to the semantic web. Intelligence, Agents, Multimedia (2003)
43. Rutledge, L., van Ossenbruggen, J., Hardman, L.: Making RDF presentable: Integrated global and local semantic web browsing. In: The 14th International World Wide Web Conference, Japan (2005) 199–206
44. Spyratos, N., Tzitzikas, Y., Christophides, V.: On personalizing the catalogs of web portals. In: FLAIRS'02, Pensacola, Florida (2002)
45. Castells, M.: The End of the Millenium. Volume 3 of The Information Age: Economy, Society and Culture. Blackwell, Cambridge, MA (1998)