Affective graphs: The visual appeal of Linked Data

Suvodeep Mazumdar [a,*], Daniela Petrelli [b], Khadija Elbedweihy [a],
Vitaveska Lanfranchi [a] and Fabio Ciravegna [a]

[a] *OAK Group, Department of Computer Science, University of Sheffield, Regent Court, 211 Portobello Street,
S1 4DP, Sheffield, United Kingdom
Email: {F.Lastname}@sheffield.ac.uk*
[b] *Art & Design Research Centre, Sheffield Hallam University, Sheffield, United Kingdom
E-mail: D.Petrelli@shu.ac.uk*

Abstract. The essence and value of Linked Data lies in the ability of humans and machines to query, access and reason upon highly structured and formalised data. Ontology structures provide an unambiguous description of the structure and content of data. While a multitude of software applications and visualisation systems have been developed over the past years for Linked Data, there is still a significant gap that exists between applications that consume Linked Data and interfaces that have been designed with significant focus on aesthetics. Though the importance of aesthetics in affecting the usability, effectiveness and acceptability of user interfaces have long been recognised, little or no explicit attention has been paid to the aesthetics of Linked Data applications. In this paper, we introduce a formalised approach to developing aesthetically pleasing Semantic Web interfaces by following aesthetic principles and guidelines identified from literature. We apply such principles to design and develop a generic approach of using visualisations to support exploration of Linked Data, in an interface that is pleasing to users. This provides users with means to browse ontology structures, enriched with statistics of the underlying data, facilitating exploratory activities and enabling visual query for highly precise information needs. We evaluated our approach in three ways: an initial objective evaluation comparing our approach with other well-known interfaces for the Semantic Web and two user evaluations with Semantic Web researchers.

Keywords: Linked Data, information visualisation, aesthetics, Visual Analytics, Semantic Web

1. Introduction

The human response to aesthetics has been a subject of study and experimentation for a long time in cognitive psychology, art and industrial design. Aesthetics, or the *"pleasure attained from sensory perception"* [36] plays a significant part in any product design. Norman [19] describes that beautifully designed products make users feel positive and good, thereby putting them in a state of mind that makes them more receptive and open. Semantic Web and Linked Data

Interfaces have traditionally been designed and evaluated performance and reliability, with few evaluations focussing on usability [44,53]. In addition to a greater stress on usability and user experience,[1] a lot of consideration also needs to be paid to aesthetics while designing products [43].

Attention toward how aesthetic pleasure affect perceived usability of interfaces began with the findings

---

[1] A significant lack of attention to usability and user experience has been a concern to the community. Indeed, David R. Karger's talk (http://videolectures.net/eswc2013_karger_semantic/) stressing the need for designing Semantic Web solutions with explicit attention to users bears testimony to the urgent need for such solutions.

---

*Corresponding author. E-mail: S.Mazumdar@sheffield.ac.uk.

of Kurosu and Kashimura [51], where the authors attempted to draw a correlation between users' perceived usability and perceived visual aesthetics. Their results indicated that visually appealing interfaces were perceived to be easier to use. Tractinsky et al. [78] repeated the experiment using a more rigorous approach and proposed the notion *"what is beautiful is usable"* to establish a relation between the perceived usability and aesthetics – showing a strong correlation between the perceived aesthetics and usability. Similar experiments evaluating different versions of websites [50] and designs of DVD players [18] indicated that the perceived quality of information being delivered to the users is influenced by the interaction style of the system.

Considerable research has been conducted in understanding the human perception of aesthetics and identifying principles that can provide a more aesthetically pleasing experience, an explicit focus on aesthetics for Linked Data and Semantic Web applications has been highly limited. Few works note the attention to and the need for aesthetics [20,26,39,68], a more systematic approach toward aesthetically designed interfaces for the Semantic Web is largely missing. The work discussed in this paper is part of our effort aimed at making Semantic Data easily accessible to users. Initially we looked at the use of semantically rich data in the aerospace industry and paired dynamic query with multiple visualisations [66]; through a user-centred design process we identified both the data and the key interaction features that composed the foundation of a highly interactive system for data exploration. The initial successful example pushed us to seek a more general approach that could be applied to any Linked Data set irrespective of the nature of the data. The idea was that any Linked Data could be seen and manipulated through a generic user interface.

A dashboard interface was designed, developed and evaluated with users [57,58]: the focus was on providing multiple, complementary visualisations on the result of retrieved set out of very large linked-data repositories. A number of views widely applicable across a number of different data sets (e.g. time, space, tag-cloud, statistics, etc.) were tested on different data sets (botanic data and DBpedia). The user evaluation clearly showed that the dashboard display was effective and engaging and that domain-specific views would have been useful to domain experts interested in digging into the data set for knowledge discovery. As much as the output was appreciated, the mechanism for providing input was criticised. As the goal was to pro-

vide a generic interface for interacting with the data, the data structure was looked at and presented to the user as a set of features to be composed in a query form. This approach clearly did not work even with users knowledgeable in the domain (botanists, biologists and ecologists). While the visualisation of results was solid, a serious re-thinking of the way linked-data could be presented to users for an initial exploration and query composition was needed. This is the purpose of the work in this paper: it is an attempt to explore alternative ways of looking at Linked Data starting from a visual art approach that favours a neat design and aesthetics to information abundance and bare function. Our interest is not on proving the same framework works for Linked Data, but to see if by combining aesthetic principles and functionalities, a new way of interacting with Linked Data can be proposed that is more engaging and effective. In this paper, we stress the importance in combining aesthetic design and user centered design processes in the development of Semantic Web solutions, and it is the combination of the two that can contribute toward a more pleasant and satisfactory experience. In this paper, we try to address the following challenges to effectively interact with Linked Data by employing Visual Analytics techniques and principles of aesthetic design:

- To develop a generic interface for Linked data
- Support human analysis through visual features
- Visually pleasing to foster engagement
- Support interactivity to foster active exploration

The paper is organised as follows: the next section provides a review of the literature based on Visual Analytics and aesthetics specifically for the Semantic Web; Section 2 discusses the related work on the use of aesthetics for interface design; Section 3 discusses some of the ways data has been visualised in the Semantic Web and how some aesthetic principles can be applicable in Semantic Web; Section 4 lists the principles that have been identified from literature that can help develop aesthetic interfaces for Semantic Web; Section 5 discusses an in-depth investigation into the aesthetic properties of some well-known Semantic Web interfaces; Section 6 then discusses our high-level approach toward generic visualisations for Semantic Web; Section 7 presents our implemented system with a scenario of use; Section 8 discusses the rationale behind our design and the various design decisions which shaped the final solution. Section 9 discusses how the system was implemented, how queries are generated and how are interactions translated into queries. Sec-

tion 10 discusses our evaluations of the system. Section 11 looks back at the Linked Data principles and aligns them with the implemented system while Section 12 concludes the paper with our reflections and and outline of future work.

2. Aesthetics in interface design

Schwarz et al. and Lang [52,73] noted the correlation between the time taken to process an object and human aesthetic response – the perception of beauty can be explained as a function of the fluency in the processing of the object. Two phases of the human cognitive system that come into play are the *preattentive phase* (low level processes before processing the sensory information) and *interpretive phase* (representations that are learned). The perception of aesthetics is therefore based on the *"combination of cognitive and sensory modes of experiences"* [25]. The pre-attentive processing stage exists before conscious processing, and occurs at Norman's visceral level [19,56]. This raises the question – how can information be represented to be quickly processed by our preattentive processes? Very interesting to this context is the work conducted by Healey in [31–33], where the authors investigate visualisation of multivariate data using preattentive processing in a rapid manner (less than 250 ms).

The experiments conducted by Healey drew several interesting conclusions such as:

- Hue can be used as a mechanism to rapidly and accurately determine a target (example, an anomaly);
- Form (shape) can be used to determine targets if hue is not varied; varying hue affects the ability to determine a form-defined target;
- Varying form does not affect the ability to determine a hue-defined target; location is not a deterministic factor in identifying a target.

Several other cognitive aspects have also been proposed elsewhere, such as a minimalistic approach [80], symmetrical layouts, Golden Ratio [21,25] and so on.

Tufte's work, in [80] is also significant for identifying aesthetic principles for information visualisation. He lists several guidelines for building attractive displays of statistical information:

- have a properly chosen format and design
- use words, numbers, and drawing together
- reflect a balance, a proportion, a sense of relevant scale
- display and accessible complexity of detail

- often have a narrative quality, a story to tell about the data
- are drawn in a professional manner, with the technical details of production done with care
- avoid content-free decoration

Bennet [6] discussed Gestalt principles applied to graph drawing from two perspectives – perceptual grouping (the ability to extract low-level primitive visual features from images and formulate a higher-level structure, e.g. grouping simple and stable figures that are similar in shape, located nearby etc.) and perceptual segregation (the ability to separate features from images and grouping them into mutually exclusive areas in order to construct a useful representation of the image, e.g. symmetry, orientation, contours). Several principles were also noted by [48,59] such as *balance* (symmetrical and asymmetrical), *rhythm* (regular, flowing, progressive), *proportion* (proportion of dimension), *dominance* (dominant, sub-dominant, subordinate), *unity* (the relationship between the visual components and elements and the complete visual scene), *emphasis*, *movement*, *pattern*, *harmony* and *variety*.

Beck [5] investigated aesthetic dimensions for dynamic graphs and proposed principles for general aesthetics, dynamic aesthetics and scalability applicable to three types of graph representation techniques – node-link, matrix and list. Among the general aesthetic principles, the authors list principles such as reduce *visual clutter*, *reduce spatial aliases*, *spatial matching of multiple representatives* and *maximise compactness*. Beck also notes while dealing with dynamic data, it is important to preserve user's mental map[2] in order to facilitate graph comparisons and ensure the user requires minimum cognitive load[3] to compare present graphs with previous one to perceive changes. Beck also points out that temporal aliases[4] should be minimised so that a continuity between consecutive frames is established.

3. Data visualiation in the Semantic Web

Semantic Web and Linked Data provides a machine-readable and understandable way of formalising in-

---

[2]Mental Map is the abstract structural information a user gathers by looking at graph layouts.

[3]Cognitive Load is the amount of information needed by the working memory of a user in order to process a visualisation.

[4]Visual elements that are mistaken for one another due to their temporal placement.

formation across different platforms. However, since data is eventually meant for human consumption, there is a need to present such information in an intuitive and meaningful manner. This task is further complicated with the ever-increasing volumes of data continually generated. Extracting actionable information from large volumes of data is a highly complex task for analysts and decision makers. 'Visual Analytics' aims to reduce this complexity by visually representing information to enable users directly interact with the information, gain insight and draw conclusions, thereby aiding decision-making processes [47]. The opportunities that arise from combining Linked Data and Visual Analytics help promote a mutually beneficial research direction: Linked Data can benefit greatly by Visual Analytics – enabling discovery of hidden trends and patterns; Visual Analytics can benefit by the development and evaluation of scalable web-based Visual Analysis techniques for large distributed networks [27].

Several researchers have attempted to support complex querying and/or visualising query results. [35] classifies such attempts in two categories – simple and complex approaches. Complex approaches (such as SPARQLViz[5] and iSPARQL[6]) include advanced user interfaces and query constructs, designed for experienced users and experts. Simple approaches such as mSpace [72], /facet [38] or Parallax [40], on the other hand are designed for casual users, but are limited in answering more complex queries. Such interfaces employ basic visualisations such as lists, tables, maps, matrices or scatter plots to represent information, thereby limiting the analytical dimensions being represented. Data Visualisation in the Semantic Web need a more careful consideration. This is due to more content being added to human readable content to make it machine-processable such as RDFa[7] and microdata.[8] However, the additional information being added is highly structured and well connected – this creates more opportunities to visually represent structured data in a standardised manner.

Green proposed a few guidelines to motivate Visual Analytics research for discovery and knowledge building, based on their human cognition model [28] –

- Provide multiple views (foster discovery of patterns using different views, as humans have different ways of processing information;
- Direct interaction (interaction to be provided without interfering with the user's train of thought);
- Central role of interaction (interaction enables user and machine to share knowledge);
- Insulation of reasoning flow (visualisation should not hamper the rhythm of reasoning);
- Intimate (seamless) interaction.

Most relevant to Semantic Web research is the recent work by Dadzie and Rowe [16], which presents design guidelines for Linked Data by exploring the literature. Starting from high level design guidelines [75], the authors propose the need for Linked Data tools to support multi-dimensional, hierarchical and network data. Additionally, the tools should also provide support for identifying/highlighting relationships within the data and the ability to export data to users/applications for re-use. The authors explored design guidelines for Linked Data interfaces from the point of view of lay-users (regular web users without knowledge of ontologies or RDF) and tech savvy users. We discuss their work and align it with our efforts in more detail in Section 11.

Many Visual Analytics tools such as [2,46,49,55,85, 86] have been built over the years, though not applied to Semantic Web data. However, such tools continue to provide inspiration for Visual Analytics research in the Semantic Web community. Traditional plotting techniques such as Scatter Plots [14], Pie Charts and Parallel Coordinate Plots [41] as well as newer techniques such as Spiral Graphs [12] and Fisheye lenses [71] can be incorporated with different forms of visualisations. A few Visual Analytics systems have also been built specifically for the Semantic Web, such as Stefaner's Elastic Tag Maps,[9] Elastic Lists [77] and the work done by [35,84,88] which have been developed for different application areas such as social network analysis, movement data analysis, bibliographic reference analysis, event detection and so on.

Application of formal aesthetics principles in interface design for the Semantic Web is the purpose of this paper. Research into such areas is needed in Semantic Web as a positive aesthetic experience can greatly

---

[5]SPARQLViz, http://sparqlviz.sourceforge.net/ is a plugin for IsaViz, that enables users to build queries from a SPARQL query interface.

[6]OpenLink's iSPARQL interface graphically renders a user's SPARQL query, showing how query concepts and relations are linked.

[7]http://www.w3.org/TR/xhtml-rdfa-primer/

[8]http://www.w3.org/TR/2011/WD-microdata-20110525/

[9]http://well-formed-data.net/experiments/tag_maps_v5/

influence the acceptability of a solution. As discussed previously, there has been a significant amount of research that investigates the impact of aesthetics on usability and experience with a tool. The nature of Semantic Web data, in addition to being multidimensional and multivariate, is graphical. Several principles of graph visualisation aesthetics are therefore applicable in this context [6,21]. Eichelberger list several aesthetic principles to be followed while drawing class diagrams in UML,[10] which we believe are highly pertinent and can be considered while designing similar LD applications. We list the principles that we believe are most appropriate for data visualisation in the Semantic Web, specifically for graph based visualisations:

– Separate hierarchical and non-hierarchical relations, hierarchy should be clearly visible.
– Centrally position parents or children – this is particularly useful for hierarchical representations such as Semantic Web data. Child nodes to be located at median position of its parents.
– Nodes should be clustered according to semantic reasons – semantically similar nodes should be positioned close.
– Avoid, if possible crossings and overlappings on edges.
– Vicinity of comment nodes – comments connected to other models should be located as close as possible to the connected nodes. In a LD setting, this can be a way of connecting multiple ontologies/datasets/graphs.
– Adornments should be clearly assigned to their model elements. In a LD setting, graphs adornments/additional specifications (e.g. labels, icons etc.) should be standardised based on the same group/ functionality.
– Respect graph drawing constraints – aspect ratio, compact drawing, symmetry, minimisation of edge bends.

Following a survey on aesthetic heuristics, Bennet [6] also proposed similar design principles, classified into syntactic (structural) and semantic (domain-specific) categories.

4. Principles of aesthetics for Linked Data

We studied the literature for Visual Analytics techniques for exploring data, as well as aesthetic princi-

ples for information visualisation and interface design. We also explored the literature for techniques and principles specific to the Semantic Web. Our survey identified several aesthetic design principles that we believe can help designers and developers build interactive and aesthetic user interfaces for exploring Semantic Web and Linked Data. Based on the principles we noted, we propose the following design principles for LD in Tables 1 and 2. We believe these principles are most important for the Semantic Web community and can be used as a set of guidelines while designing and developing interfaces. The guidelines are divided into two sections – general aesthetic principles involve the design and layout of the interface in general (Table 1); node-link representation principles involve the design of node-link graphs for representing Linked Data (Table 2).

5. Evaluating aesthetics of Linked Data applications

The starting point of our research was to understand how well existing Semantic Web tools align with each other with respect to aesthetic properties as aesthetics has not been studied before in the Semantic Web. Our review of the literature explored several evaluations of user interfaces for the Semantic Web and Linked Data, where evaluations have been conducted mostly as user-based studies (understanding usability) or performance analyses (precision, recall, speed etc.). Evaluating the aesthetic properties of Semantic Web-based user interfaces is a step forward in the direction of establishing a research area fundamentally focussed on the development of aesthetically pleasing interfaces for the Semantic Web. We aimed at objectively evaluating interface aesthetic properties as it can provide a simple and inexpensive way of assessing various aesthetic properties of the system. Our experiments were based on the model provided by Ngo [60,61], where different metrics of an interface are computed, on the basis of the layout of visual objects within the visual space. These metrics have also been previously used in identifying most aesthetically pleasing layouts of websites from a set of candidate designs [89,90]. Our survey of the literature did not identify any existing work in Semantic Web research where such metrics have been used in evaluating aesthetic properties of interfaces. Such metrics have also been reported to be highly correlated with subjective scores from users [1,65,67,89]. Following from the work conducted by Purchase [67],

---

[10]UML specification http://www.omg.org/spec/.

Table 1

General Principles for Aesthetic Linked Data Visualisation

Design principle	Description	Proposed by
1. Use words, numbers and drawing to convey information	Graphics alone are not always enough to convey the significance of a piece of information – numbers, narratives, explanations can aid in better communication	[80]
2. Well balanced, proportioned and symmetrical design	Interface should be arranged so that optically larger and smaller objects balance each other in a symmetrical manner; Interface should be well proportioned (i.e. golden ratio, greater horizontal length etc.)	[6,21,25,37,37,48,60,80]
3. Rhythm, unity in design	Interface should be designed from multiple visual elements, coherently constituting a pleasing layout, with regular patterns of visual changes to make the appearance exciting	[59,60]
4. Different weights of lines to represent different information	Contrasting lines indicate different meanings – weights can be associated with values or types of data	[80]
5. Simple, consistent and stable figures	Complex visual representations require greater interpretation and add to the cognitive burden of users; Adhering to Semantic Web principles and standards require a consistent representation of data elements across domains and application areas	[6,10]
6. Using variations of colour, shape, size, intensity to present trends, interesting patterns, anomalies or represent similarity, physical connection	Colour, hue, size, shape etc are visual clues that we can quickly spot, thereby making it easier to observe patterns, anomalies etc.	[6,31,37]
7. Minimalist design, reduce visual clutter	Interface should be minimalistic, and contain as little data-free visual elements as possible	[5,13,37,63,80]
8. Balance in harmony and typicality	Typical solutions require little effort, but at the cost of being a mundane solution – balance in variety and typicality is important	[37,48,50,59]
9. Maintain consistency in visual representations, interaction mechanisms and standards	Visual representations (colour, shape, hue etc) should be consistent across all domains; interaction mechanisms should be familiar to users and standardised (e.g. right-click should present context menus etc.)	[59,63]
10. Follow visual information seeking principles with minimal cognitive burden on users	Provide mechanisms to overview, navigate, filter and access data instances on demand, whilst ensuring minimum cognitive load and changes to the mental map	[70,75,79,80]

Table 2

Principles for Aesthetic Node-Link Representations

Design principle	Description	Proposed by
11. Separate representations of hierarchical and non-hierarchical relations	Differentiating between hierarchical and non-hierarchical relations helps users navigate along or across graphs	[21]
12. Reduce overlapping nodes		[21]
13. Center parents or children	The parent node should be located as close as possible to the median position of the child nodes	[21]
14. Cluster nodes based on semantics	The position of nodes should be based on their semantics so that nodes that are adjacent to each other have a significance	[6,21]
15. Avoid edge crossings or overlaps	Every edge should be as visible and readable as possible and spaced apart from nodes	[6,21]
16. Uniform and minimal edge bends	Minimal angles on the edges to help users follow links	[6]
17. Even distribution of nodes	Well distributed and evenly spaced nodes, but ensuring compactness	[6]
18. Maintain aspect ratio, symmetry	Maintaining symmetry within the layout as well as in the overall interface; aspect ratio of the graph should match the container (interface, screen, page etc.)	[6,21]
19. Minimise total graph area	Compact layout but ensuring readability	[6]

we believe that the placement of visual objects in interfaces can be a strong predictor of aesthetic appeal and perceived usability. It is important to note that this study involved only the interface layout from a general aesthetic point of view and factors such as colour, styling, typography or individual visualisations were
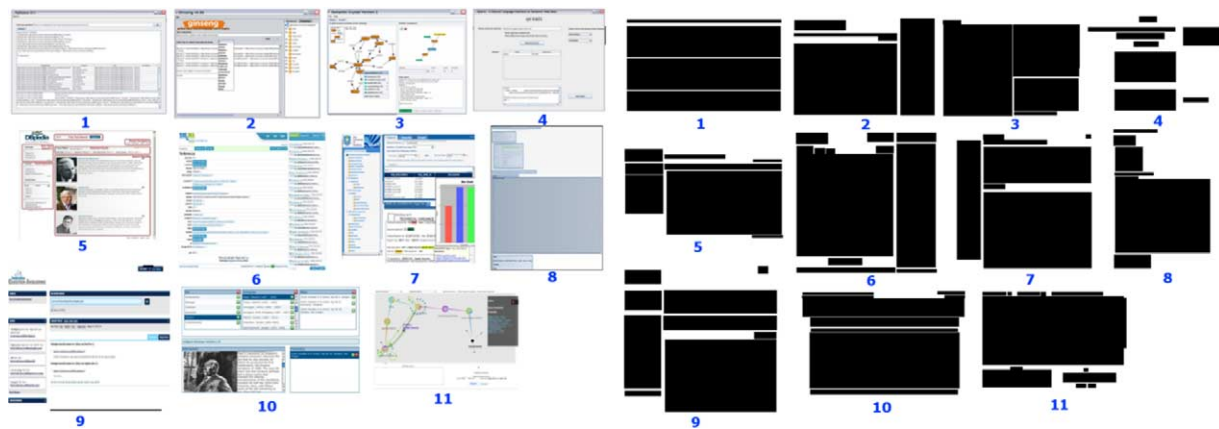
Fig. 1. Layouts used to compute interface aesthetic metrics. The layouts were obtained from the interfaces directly or from the relevant papers, publication materials, websites or user manuals. The layouts shown are of ten well-known Semantic Web interfaces (screenshots of the system are shown on the left, while images on the right show the respective markups): 1. NLP Reduce, 2. Ginseng, 3. Semantic Crystal, 4. Querix, 5. DBpediaFacet, 6. Sig.ma, 7. KSearch, 8. Tabulator, 9. PowerAqua, 10. mSpace and 11. Affective Graphs (later discussed in Section 10).

not a part of this study, but will be explored as part of a future work.

5.1. Experiment design

We based our evaluation on eight of the thirteen metrics provided by Ngo [60,61]. Our starting point is the five most important properties identified by Zain et al. in their similar evaluation exercise with Ngo's measures. [89,90] – **Balance**, **Equilibrium**, **Symmetry**, **Sequence** and **Rhythm**. Balance is a measure of how the visual elements with different optical weights (larger objects are perceived as 'heavier') are distributed in the interface; Equilibrium indicates how centered the layout appears to be; Symmetry indicates how well replicated elements are on either side of the horizontal and vertical axes at the center of the interface; Sequence is a measure of how well objects are arranged in the interface, with respect to the movement of the eye; and Rhythm indicates the variation of visual objects in order to make an interface exciting. We identified two other metrics from the design principles that would be important for our evaluation – **Cohesion** and **Unity**. Cohesion is a measure of the similarity of aspect ratios of the visual elements in an interface; and Unity is a measure of the extent to which the elements seem to belong together.

The final property, Order and Complexity is defined as the sum of all the properties. The metrics not considered in this study are simplicity, regularity, homogeneity, economy and density. Certain visualisations such as graph-based ones can affect how simple or complex a user might interpret the interface. Other visualisa-

tions such as scatterplots or maps can also affect the interpretation of the general economy, regularity, homogeneity and density of an interface. We aim to investigate these metrics in more details in a future study where we explore the implications of complex visualisations and graphs on these factors. All the scores range between 0 and 1, where 0 indicates a highly negative assessment, while 1 indicates a highly positive assessment. These metrics are defined in [60,61] and the relevant formulae are provided in Appendix A. We revisit this evaluation in Section 10, where we investigate how our interface performs as compared to the existing interfaces.

The first step in evaluating the metrics was understanding which tools would be good candidates for our comparative studies. Our survey of the literature identified ten well-known Semantic Web tools that have existed over the past few years: mSpace [72], PowerAqua [54], K-Search [8], Sig.ma [82], Tabulator [7], DBpediaFacet,[11] Semantic Crystal, NLP Reduce, Ginseng and Querix [44,45]. While some of the tools ranging from browsers and search systems (e.g. K-Search, PowerAqua, mSpace, Sig.ma) to standalone interfaces (e.g. Semantic Crystal, NLP Reduce, Querix, Ginseng) are clearly research prototypes, our intention is to also understand how they compare with the rest.

A standalone java-based application was developed, which was fed screen shots of the interface layouts (Fig. 1). The user then manually marked up the areas that contain visual objects by using mouse ges-

---

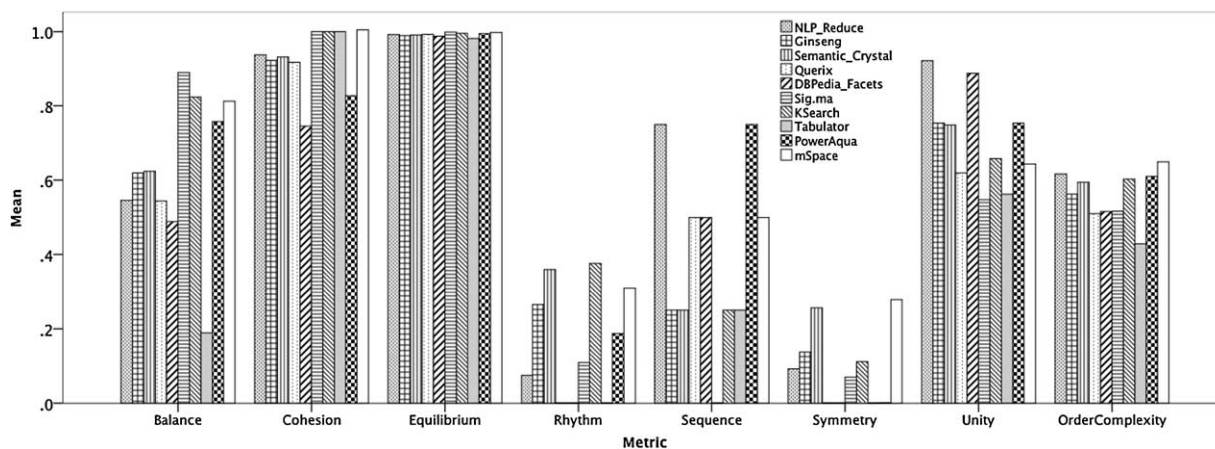[11] DBpedia Faceted Search, http://wiki.dbpedia.org/FacetedSearch.

Fig. 2. Comparative evaluation of seven aesthetic metrics with ten Semantic Web tools. The eighth metric, Order and Complexity is computed as a sum of the others.

tures like click and drag. Each of the manual annotations were then stored locally and their dimensions calculated. The application then aggregates the different measures as the visual objects are marked up, based on the formulae provided by Ngo. The screenshots of the ten systems are either obtained from local installations, publication material, website images or screenshots in user manuals. These measures are then collected and compared against each other. As this process involves human annotations (markups), each interface is annotated three times by one user and the mean is then computed and compared against the others.

5.2. Results

The scores were plotted as shown in the Fig. 2. As can be seen, most of the interfaces performed well on equilibrium. Cohesion scores for DBPedia Facet and PowerAqua were comparitively lower, though all systems scored relatively high. Sig.ma was observed to be the most balanced system, followed by K-Search and mSpace. In general, all systems scored relatively low in Rhythm and Symmetry. Querix, DBpedia Facets and Tabulator scored the least in Rhythm, while Querix, DBpedia Facets, Tabulator and PowerAqua scored the least in Symmetry. Overall, the best performing tool was mSpace, with a Order and Complexity score of 0.65, followed by NLP Reduce and PowerAqua with scores of 0.62 and 0.61. The lowest scoring tool was Tabulator, with a score of 0.43. It can be observed that all the interfaces scored between low to medium, in terms of their overall aesthetic properties.

This provided the starting point for our research, where the first step was to identify the need to inves-

tigate aesthetics and design tools with an explicit attention to aesthetics. It also demonstrated that out of ten well known Semantic Web tools, most interfaces do not score highly on aesthetic properties related to screen design and object positioning. Few other factors such as overall colour palette, typography and iconography also have significant impact on aesthetics [80]. While these factors may be easily evaluated from a subjective point of view, objective analysis in a comparative setting is difficult. This is particularly due to the relative absence of benchmark scores to ascertain the positive or negative scores of interfaces in comparison to other.

We note that other factors such as colour, texture and shape of visual items also need to be considered in order to comprehensively assess aesthetic properties of system. However, as Ngo anticipated, this task is significantly more difficult as it introduces far more variability in order to be considered as an objective analysis. Furthermore, we only consider this evaluation as a preliminary study on the attention to aesthetics in the design of existing Semantic Web interfaces. A more comprehensive study is planned in the near future involving a large number of users to subjectively score available interfaces on several criteria such as colour palette, texture, shapes, typography and so on.

6. Visual analysis of Linked Data – An approach

As it could be expected, representing Linked Data in an abstract way leads to graph representation as Linked Data is essentially multiple data instances connected with links to constitute a highly connected and direc-

tional graph. Hence, our starting point for the design is a node-link graph, where nodes represent concepts and links represent relations the concepts share. Node-link graphs, large ones particularly, are notoriously difficult structures to handle and understand. The challenge here is to present large sets of data, but preserving their links and hierarchical structure.

The most important requirement was to facilitate the user to **explore** unknown (and known) datasets. In addition, we also wanted to enable users **query** for specific information using a visual approach. An important requirement for our design was to put a lot of emphasis on the **aesthetic quality** of our interface. Our approach had to be a **generic** one, in order to ensure any Linked Data set can be consumed. We divided our solution into the following four major functions to address the requirements:

1. **Making the underlying schema of data apparent to users; highlighting further details such as context, relations and statistics**. While ontologies provide a formal specification of the domain, the data itself is what the users are mostly interested in – this generated a need to visualise the ontology as well as the data at the same time.

2. **Support data driven exploration via statistics by making use of standard statistical presentation techniques**. Visualising entire ontologies and data instances can be a useful way of presenting the data as well as the domain, however at the expense of increasing cognitive burden and exhausting screen space – this generates a need for users to access concepts and their data 'on demand', rather than showing everything that is available.

3. **Provide access to individual data instances at all times**. While our initial interest was in presenting statistics with ontological concepts and properties, discussions with users and Semantic Web experts throughout our iterative design process resulted in the need for ways to provide easy access to data

4. **Support highly specific information need by introducing flexible user interactions**. The beauty of Semantic Web is in making available high quality, dynamic and precise information that is highly inter-connected: such information can be exploited from interfaces that can support building complex queries to precisely answer highly specific questions.

We aligned the identified features with the design principles listed in the previous section. Following analysis of the design principles, low-fidelity mock ups were built in order to understand how users would interact with visualisations in an intuitive manner. The following section presents the system that was developed, explaining with an example scenario of use. Section 8 then discusses how different design decisions were taken, aligning the interface with the design principles.

7. Affective graphs – A scenario of use

Figure 3 shows a screenshot of the final implemented system, *Affective Graphs*. The system was built as a result of several re-designs and prototypes, with constant inputs from users and evaluation feedback throughout the implementation. Section marked A shows the interactive node-link representation of the underlying data. The image shows a user exploring the latest DBpedia[12] dataset, presently viewing lakes (the node on focus is at the bottom right corner). Users can gain a large amount of information just by observation and interacting with the visualisation. We list some information that a user can quickly find while interacting with the system:

– The dataset contains information regarding places, persons, work, species, organisations, transportation, films and so on. This is observed by selecting the 'owl:Thing'[13] node and hovering over the pie sections.

– The subclass hierarchy is immediately apparent to the users – lakes are types of bodies of water, which are natural places and persons are types of agents etc. This is observed by following the triangle-shaped hierarchical edges. The colour of the edges indicate the respective pie-sections they originated from.

– There are 41k lakes, 2.5m places, 3.27m persons, 218k natural places in the dataset etc. This can be found from the labels of the nodes as well as tooltips provided on pie sections.

– The amount of information on agents is the most, followed by places (hovering over the pie sections reveal the subclasses as agent, place, work,

---

[12]http://dbpedia.org, as on 20.05.2013.
[13]The first node that users can see is the *owl:Thing* node that encompasses all the data classes described within the dataset.
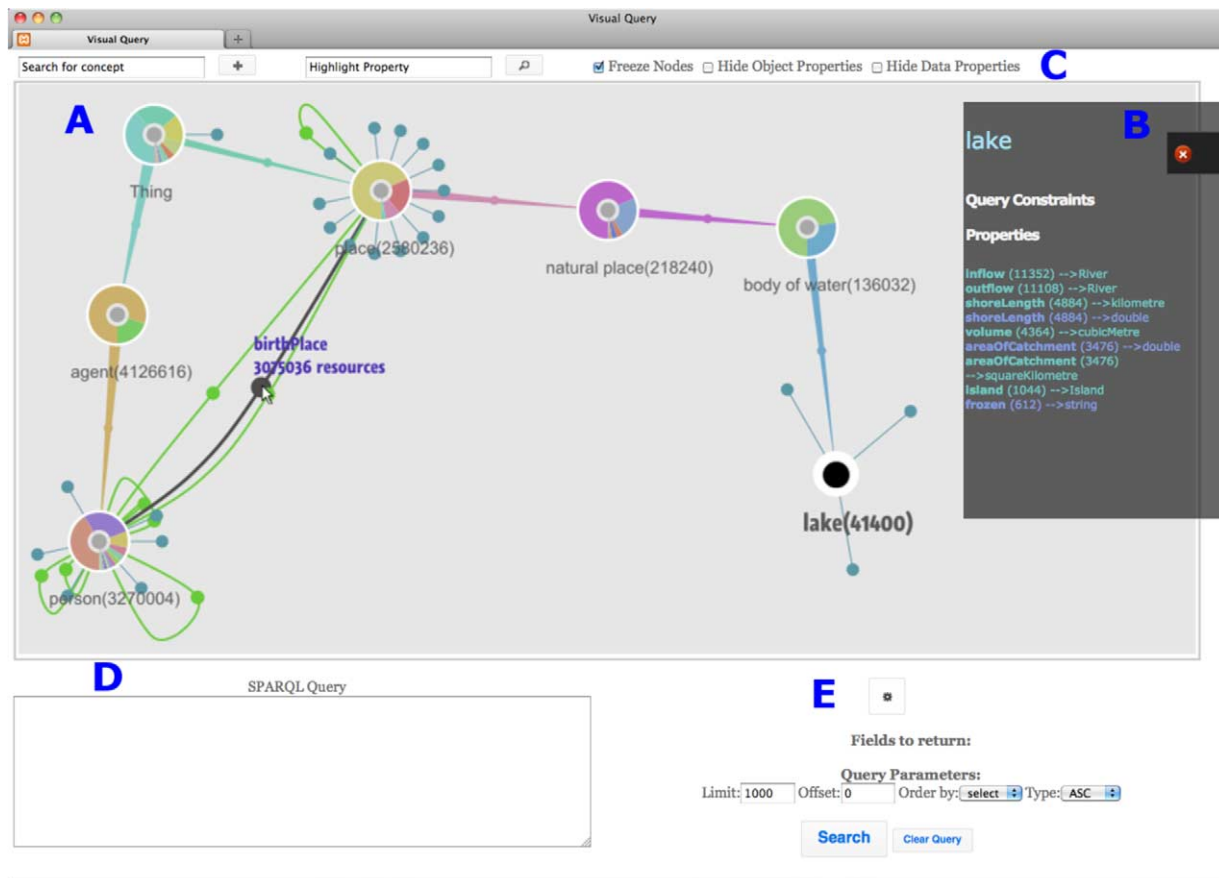
Fig. 3. A screenshot of *AffectiveGraphs*, where the node currently on focus is 'Lake': Section A contains the interactive node-link representation of the data, Section B contains contextual information relevant to the concept currently being explored (here, Lake), Section C contains search elements and controls the visual rendering of the node-link graph, Section D shows the SPARQL query being generated for search, Section E contains advanced features to modify the query.

species etc. in the order of instance counts). This can be observed from the positioning and size of the pie chart sections – the subclasses with the greatest number of instances are the biggest in size as well as positioned at the bottom left of the pie charts.

– The information regarding the birthplace of 3.07m persons are available in the dataset. This is found by hovering over the 'birthplace' relation connecting persons and places.

– There are five relations in persons that are linked to the same concept – *parent* (87k instances), *spouse* (100k instances), *influencedBy* (99k instances), *influenced* (49k instances) and *child* (43k instances). This can be seen as four curved loops originating from and ending in the person node – hovering on the edges reveal the relations and the number of instances.

– Three data properties exist in the lakes concept – *areaOfCatchment* (3.4k instances), *frozen* (612 instances) and *shoreLength* (4.8k instances). This is shown as three edges originating out from the Lake, hovering over the edges show the relation and the number of instances. Similarly, the data properties for person and place can also be investigated.

– Three object properties connect persons with places – *death place* (799k instances), *birth place* (3.07m instances) and *resting place* (52k instances). This is found by hovering over the three curves joining the person node with place node.

– The distribution of all the subclasses of Natural Places (Body of water, Mountain, Mountain range, Lunar crater, Cave), Places (Populated place, Architectural structure, Natural place, Protected area etc.), Body of Water (Stream, Lake),

Agents (Person, Organisation) and Persons (Athlete, Artist, Officeholder, Politician etc.) can also be investigated. This can be done by selecting the respective nodes and hovering over the different pie sections.

In summary, a very high level understanding of 2.5m places, 218k natural places, 136k bodies of water, 4.1m agents and 3.27m persons as well as the broader content of the entire dataset can be very quickly gathered by observing the visualisation and interacting with it. Not only is the hierarchical structure of the dataset made apparent, the links and properties shared by concepts are also made available to the user.

Section B shows the contextual information of the node presently in focus (here, Lake). By investigating the content on the right panel, users can understand that the data properties are shoreLength, areaOfCatchment and frozen (data properties are coloured blue, and object properties are green following the standards set by Protégé). In addition, the number of instances connected with these relations are presented. The object properties are also presented, along with the other concepts the relations are connected to. Clicking on the concepts on the panel trigger new nodes to be formed, and the relations to be visualised.

Section C provides a mechanisms for users to search for a specific concept, if the user prefers to search for any known concepts. Once a required concept has been found, a node is then added to the visualisation (without the user having to manually selecting pie sections). In addition to a concept search box, a property search box is also provided in this section – users can start typing a property name (or URI) that they are interested in, and the system highlights the property of interest in the visualisation. This is an easy way of quickly spotting any property that the users are interested in. There are three other controls provided in the Section C which control the force-directed layout and toggle the visibility of the object and data properties.

Section E provides advanced features for customising the SPARQL query formed while performing a search – such as, selecting the variables to be returned, limiting search results and so on. Section D displays the final SPARQL query formed, in case the user wishes to edit the query before searching. This is a requirement that is most often desired by expert users, who prefer fine-tuning their queries after having created a basic query.

8. From principles to visual design

Directly abstracting Linked Data leads to a node-link representation: this was the starting point of our design. The final interface was developed as a result of a set of several design decisions. The first phase was to understand how to design a node-link representation that can provide additional information about the underlying data as well as adhere to the aesthetic Principles defined in Table 2.

8.1. Consistency in visual representations

The first step was to develop a **consistent** representation for concepts that provided more information about the data in the concept. In our design, concepts are represented as circular nodes, while the properties have been represented as edges. In order to present information about a concept, the circular design of the nodes have been modified to show a pie-chart of the distribution of the subclasses of the concept. The pie sections are sensitive – clicking on each section triggers a new concept node to be created, which contains information about the respective subclass. Pie charts visualise data in a small area and provide a wealth of information that if displayed as a graph would be confusing and difficult to grasp. The new node created is similar where distribution of its subclasses are rendered as another pie chart on the new node (Principle 5, 9 suggest consistent visual representations). If the new concept contains no subclasses, the node would be represented as a circular blank node. The new node is connected with the originating node using a hierarchical edge that signifies a *rdfs:subClassOf* relation.

The design decision was on visualising the distribution of the subclasses – a pie chart was chosen as the most preferred representation due to three main factors: users are familiar with pie charts and can quickly assess the relative amount of data each section contains; pie-charts are able to convey statistical information within a small and regular area better than a table specially when conserving space is important (Principle 19, 7 suggest a minimalist approach); a circular depiction of nodes in a node-link graph is an organic geometric representation that most users are familiar with.

8.2. Representing semantic concepts

The pie-chart **representation** itself provided the next design challenge – our aim was to provide a pie-chart with regions easily distinguishable from one an-
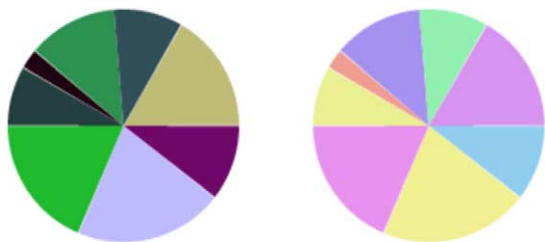
Fig. 4. Comparison of randomly generated pie-chart colours for RGB (left) and HSB (right). Though the colours on the right are completely random, constant values of brightness and saturation generate a more pleasing chart. Images created in http://sketch.processing.org/.
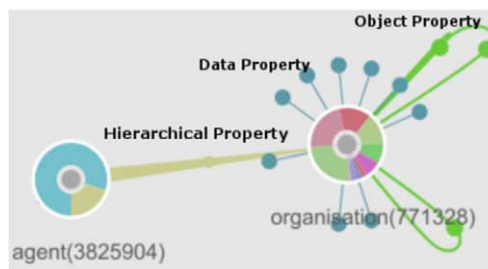

Fig. 5. Different visual representations of edges bear different ontological significance – hierarchical edges are represented as triangles showing the parent and child node, data properties are represented as satellite objects and object properties are represented as curves.

other. The use of colours in order to distinguish pie sections is a standard process – however, the task is further complicated as there are multiple pie charts in the layout. Initial efforts at using a standard bank of 20 unique manually selected colours were unhelpful and caused confusion in the users as standard colours seemed to indicate certain commonality among the similar coloured-sections. Furthermore, a repetitive standard colour palette would reduce the variety in the design (Principle 8 suggests introducing variety in the design). The system was then changed to generate random colours in a HSB (or HSV) scale, varying only the hue values to keep the saturation and brightness consistent as well as provide a greater range of colours. The HSB scale was preferred to the RGB scale as the former provides greater flexibility in varying colours (here, hue) by keeping the saturation and brightness constant, which is not as easily achievable in the latter. Moreover, HSB scale is more natural and user friendly way since it replicates the way we "percieve" colours.[14] Figure 4 explains this with a simple example – the figure on the left was produced from a set of randomly generated HSB colours with brightness and saturation fixed. Though the different regions in the two pie charts are easily distinguishable, the RGB pie chart contains sections of unequal brightness. RGB colours can also achieve the same results as the HSB pie chart, but with more complex methods.

Though randomly generated, there is often a high chance of creating duplicate colours (that may have different colour codes, but are nearly indistinguishable by the human eye). Similar colours can also be wrongly interpreted as the same (or similar) concepts. This was significantly reduced by generating random colours that are different (based on a threshold) from

the set of colours already generated by providing a look-up service for the set of colours already generated. The pie sections are further distinguished with the use of borders and interactive events (such as the respective sections are extended when the user hovers over them)

8.3. Representing semantic relations

The third design challenge was the **representation of relations** (properties) – an important consideration was the different types of relations that may exist within the data set. We identified three types of relations – data properties, object properties and hierarchical properties. Hierarchical properties are a kind of object properties as they essentially describe the relations between a parent object and its child. However, in our representation we have chosen to isolate the hierarchical properties from object properties (Principle 11). Data properties are presented as free edges – where an edge is connected to the node at one end, and a circular object at the other. The positioning of data properties is random, but ensuring that the edges are not overlapped (Principle 15 suggests minimal overlaps of edges). The circular object is the sensitive end of an edge, providing interactions with the users (hovering on the object highlights the edge and provides a tooltip display, clicking activates query mechanisms etc.). We describe the different edge representations in the Fig. 5.

Object properties are represented as curves – this is due to the preference of users to curves as opposed to straight lines [4,9] and also the lack of acute angles and edge bends (Principle 15 suggests avoiding edge bends). Cubic Bézier curves were chosen to render the edges as they provide a greater control and flexibility [87]. Additionally, the presence of straight lines as connections between objects can make it dif-
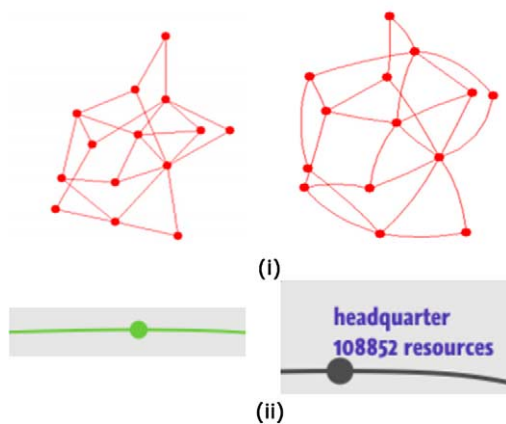
---

[14]http://tools.medialab.sciences-po.fr/iwanthue/theory.php

Fig. 6. (i) Comparison of Straight Edges with curved Edges. Image from [24], (ii) Visual response to hovering over an object property.

ficult to follow links. Figure 6(i) shows an example scenario where the same nodes have been connected using straight edges (left) and curved edges (right). Interactive mechanisms on the curved edges also make it easier for users to follow how the nodes are linked in case of larger graphs. Hovering on the sensitive sections of the edges highlight and increase their thickness, making them easier to spot as shown in Fig. 6(ii) (Principle 4 suggests the use of lines with different thickness to add variety; Principle 6 suggests the use of variations in colour to help users spot elements). The edges of the node currently in focus are also shaded in a darker colour to help users find the edges connected to the current node (Principle 6 suggests using variety in colours).

The representations of hierarchical and non- hierarchical relations are different, as we believe that each can be used as a different interaction mechanism (Principle 11 suggests a different representation of the two). While non-hierarchical relations describe the different properties of objects (such as birthplace, age, date of birth etc.), hierarchical relations contribute more toward explaining how the data is structured. Hence, the design decision was to provide edges with greater weights for representing hierarchical relations. Furthermore, in order to identify the children of a node, the hierarchical edges are represented as triangles signifying directionality, as shown in Fig. 5. Inspired from the design of Protégé,[15] the data properties are coloured blue while the object properties are coloured

green (Principle 5 and 9 suggest maintaining consistency and such colour schemes are well established within the Semantic Web community). This design helps users quickly identify which edges are data and object relations, as well as hierarchical edges (Principle 4 suggests using differently weighted lines to signify different meaning and principle 6 suggests varying visual properties to enable users to quickly spot features).

8.4. Layout

The next design challenge was **layout** of the node-link graph. Initial attempts at automated layouts helped in balancing the representation effectively – nodes were arranged in a force-directed layout, based on the Processing simulation library, developed by Bernstein.[16] However, as the number of nodes increased, the graph grew far more compact than required. Furthermore, the force-directed layout made users disoriented and complicated the interaction as objects kept floating around the graphical space to optimally position them. Two changes were then made for the final design – a change in the layout algorithm and a change in the spring configurations.

The force-directed layout is active only during the initial 5 seconds of a node being created – this allows for enough time for the node to be rendered while positioning itself in an approximate enough position for a good layout. However, the user has the ability to click-and-drag the node to wherever they desire, without being restricted by the force-direction (Principle 9 suggests maintaining consistency in interaction mechanisms – this helps users learn the system more quickly as this is a familiar interaction present in several graph visualisations; Principle 8 suggests a balance in typicality – while representation of nodes as pie charts and introduction of several types of edges with different semantics is a novel addition, familiar interaction mechanisms help users adjust to a new setting). This provides flexibility and freedom to the users as they can layout their graphical exploration in any manner they please, but at the same time provides a rough approximation for a newly created node to ensure readability is maintained. The users however, have an option to disable the feature, which would cause the force-direction to be active at all times.

---

[15]Protégé is one of the most widely used frameworks for modelling ontologies and knowledge systems, with a wide user community and http://protege.stanford.edu/.

[16]Jeffrey Traer Bernstein's physics simulation library is a standard library available for use with Processingjs applications, http://murderandcreate.com/physics/.

The second change involves modifying the contribution of different types of edges toward the force-direction as well as the overall display of the graph. In the final design, the hierarchical edges (subclass relations) are made the only type of edges that contribute directly toward the layout. This makes the layout more balanced and well-spread, instead of a highly compressed layout due to object properties exerting forces between many more nodes (Principle 12 suggests minimal overlapping nodes – the nature of the graph being relatively more spread-out ensures that the minimal number of nodes are overlapped; The balance provided by the hierarchical and non-hierarchical edges in contributing toward the layout ensures that Principle 19 is adhered to, without compromising on readability). The non-hierarchical edges do not contribute directly toward the layout, and are just links that visually connect concepts and interact with users.

8.5. Designing interactions

Another design challenge was to introduce **interactions** in the visualisation. Initial prototypes were designed to make it quicker for the user to perform functions. For example, the nodes supported right-clicks to add concepts to queries, using a control key and left clicking on nodes would hide them and so on. However, this caused a lot of confusion among users, resulting in frustration and needing help constantly. A re-design of the interface introduced similar interaction mechanisms as seen in other graph visualisation tools and interfaces such as Google maps (Principle 8 and 9 recommends using some familiar solutions to reduce the effort required for users in learning the tool). The interface allows drag and drop actions to reposition nodes as users prefer. Right-click on nodes and edges provide users with context menu, enabling them to add concepts or relations to queries, hide nodes and so on.

The final design challenge was to incorporate the visualisation into a complete interface, where all the visual elements are in **unison**. User studies and Visual Analytic principles such as Table 1 showed the need for an interface that provides features for advanced users. In addition, the need for representing the underlying formal query was also raised. In order to balance the layout, the two new visual elements (advanced and formal query display) were positioned below the graph interface. A contextual display that provides more information about the current topic of ex-

ploration is also placed on the right. The positioning of various visual objects have been made to provide a well-balanced and symmetric layout. Consideration was also made to arrange the objects based on a sequence – objects should be positioned in a layout that facilitates movements of the eye (most readers start reading from the top left and move to the bottom right) [60,61].

The final interface was developed as a part of an iterative user-centered design process, with every subsequent iteration resulting in user evaluations or focus groups. An iterative user-centered design process is one where end users are involved in the final development and design of an interface. The whole process of design, implementation and evaluation is carried out several times, where each subsequent iteration results in a refined interface. Several users from different communities such as academia, research, aerospace engineering and knowledge management had been involved in the process, and their comments, feedback and suggestions were extremely helpful in developing the final design of the interface.

9. Logical architecture

The primary goal of Affective Graphs is an automated visualisation of aesthetic graphs from data. This section discusses how we implemented the interface as well as describe how users interact with the system.

The system is composed of two sub systems: the front end (right block, Fig. 7) provides the visualisation, interactions, advanced controls and filters, the backend (left block, Fig. 7) deals solely with querying endpoints using SPARQL generated during the previous processes. Every user-interaction with the various frontend modules results in SPARQL queries being generated in the SPARQL generator, which interprets these actions based on the methods discussed later in this section. The queries are then transferred to the PHP backend. The Query Engine in the backend interpret these queries and make any modifications in the query such as variable names, adding prefixes etc. The Data selector then checks a local cache to see if the same query had been used previously. Due to the unpredictability (and unavailability at times) of SPARQL endpoints, we decided to build a local datastore that stores the responses of queries that have been previously sent to a public endpoint. This was a step
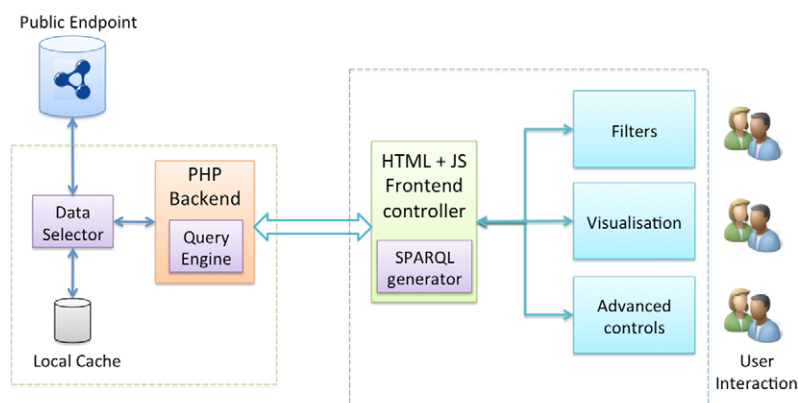
Fig. 7. Architecture.

taken to address an issue highlighted during a previous work where endpoints were found to be unpredictable in terms of their query response times [57,58].

If the query's response had been previously recorded, and if the data provided by the endpoint is not recent, the previous response is gathered from the cache. On the unavailability of any cached result for the query, the public endpoint is queried and the result is stored in the cache to be fetched at a later stage. The result is then converted to a JSON object and returned to the front end. The frontend interprets the object and renders the data into the visualisation as required.

The first step in implementing the solution is to understand how to construct automated queries based on user interactions. Users interact with visual objects in the web interface, triggering calls to the server, passing the URIs of the entities they represent; the server constructs the queries from templates. This process is slightly more complicated, as SPARQL query requires the usage of variables. This is obtained by continuously maintaining a catalogue of the entities being queried for and constructing variables built out of the URIs. For every node being built (as well as during initialisation), the interface sends three requests to the server:

- A *subclass request* for all the subclasses of the concept along with the respective counts of the number of instances within the domain.
- A *domain property request* for all the properties (along with number of instances) that have the concept as its domain.
- A *range property request* for all the properties (along with number of instances) that have the concept as its range.

These requests are translated into formal SPARQL queries using templates such as the following subclass request query:

```
1. SELECT distinct ?subClass count (?x) as ?count
?label
2. WHERE {
3.   ?x a ?subClass.
4.     ?subClass rdfs:subClassOf dbp:Place.
5.     ?subClass rdfs:label ?label.
6.     FILTER langMatches( lang(?label), "EN" )
7. }order by desc(?count)
```

As can be seen from the query, the backend queries the endpoint for all the subclasses of a class that is currently being visualised (the 'Place' concept, as seen on line 4). The filter directs the endpoint to return only labels that are in english. The response from the query would then be converted into JSON in the backend, and then returned to the frontend. A sample response is as follows:

```
[{"subClass type":"uri",
"subClass":"http://dbpedia.org/ontology/PopulatedPlace",
"count type":"literal",
"count":"360296",
"count datatype":"http://www.w3.org/2001/XMLSchema#
integer",
"label type":"literal",
"label":"populated place",
"label lang":"en"},
...,
{"subClass type":"uri",
"subClass":"http://dbpedia.org/ontology/Monument",
"count type":"literal",
"count":"4",
"count datatype":"http://www.w3.org/2001/XMLSchema#
integer",
"label type":"literal",
"label":"monument",
"label lang":"en"}]
```

The response provides the frontend with the subclass and the number of instances that are types of the subclass. This is then parsed by javascript and processing modules to build the pie chart and create the pie sections. Each pie section is built as an interactive el-

ement, listening to mouse gestures and responding accordingly.

Similarly, a sample range property request is as follows:

```
1. PREFIX rdfs: <http://www.w3.org/2000/01/rdf-
schema#>
2. SELECT distinct ?prop count(?instance)
as ?count ?domain
3. WHERE {
4.   ?prop rdfs:range dbp:Stream.
5.      ?prop rdfs:domain ?domain.
6.      ?instance ?prop ?obj.
7. }order by desc (?count)
```

The query retrieves all the properties that have 'Stream' as its range (line 4). Along with the results, the query also requests for the domains of the properties as well as the number of instances. A similar response is returned for the domain and range property requests. The returned objects are properties and the number of instances. The properties are then classified into data type and object type properties.

The graphical space of *AffectiveGraphs* is a particle system, that simulates gravity, drag and apply forces between particles. Particles are objects that can move around within the particle system, based on the forces acting upon them. Spring forces and attractive forces act on the particles – springs ensure the connected particles are always maintained at a minimum distance from each other, while attraction can be positive or negative (repulsion). Unlike non-hierarchical relations, hierarchical ones are used as springs, thereby contributing toward the final layout. The semantic interpretation of a spring is a 'rdfs:subClassOf' (hierarchical) relation in our current implementation of Affective Graphs. However, this can be changed to any other relation that is deemed appropriate for a particular dataset. A non hierarchical edge, on the other hand can have two semantic interpretations: an object edge that connects semantic concepts (like an object property connects two semantic concepts), represented by a green bezier curve between two nodes (edge C in Fig. 8) or a loop that connects one node to itself (edge D in Fig. 8); a data edge that emerges from a node and is represented by a blue straight line (edge B in Fig. 8).

9.1. Query mechanism

Our approach was to exploit the inherent feature in a semantic dataset where concepts are connected to themselves and other concepts with relations. Such a visual approach toward presenting, exploring and querying data stems from our belief that Semantic Web
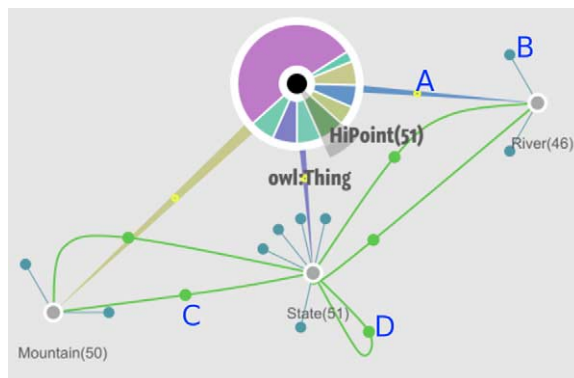

Fig. 8. Different visual representations of nodes and edges bear different ontological significance.

data is fundamentally highly visual and graphical and our approach toward consumption of such data could be more interactive by presenting to users the data as it was conceptualised by data providers at the time of creation. This makes construction of complex queries significantly easier – just by right clicking on nodes and edges and selecting 'Add/Remove Query' from the context menu to set/remove a query term, as shown in Fig. 9.

The figure shows a screenshot of Affective Graphs configured to visualise the geographical dataset from the Mooney Natural Language Learning Data.[17] The left side of the image shows a user right-clicking on a non-weighted object edge (an object property, hasMountain) that connects State and Mountain to load the context-menu. Upon selecting 'Add/Remove Query' from the menu, the concepts State and Mountain are highlighted in blue as well as the edge hasMountain. The highlighting is a visual feedback that communicates that the system has accepted the user's query and has built a corresponding query. Right hand side of the Fig. 9 shows Affective Graphs showing the query that was built.

If a concept was selected as a query item, then the interface highlights only the concept and interprets the action as the user is interested in looking at the instances that are types of the particular concept. Affective Graphs attempts to understand what concept or properties are being selected and associate internal query variables to the selections, partially building/rebuilding a query after every subsequent action. Each partial query is a small fraction of the final SPARQL query that represents what the user actions are inter-

---

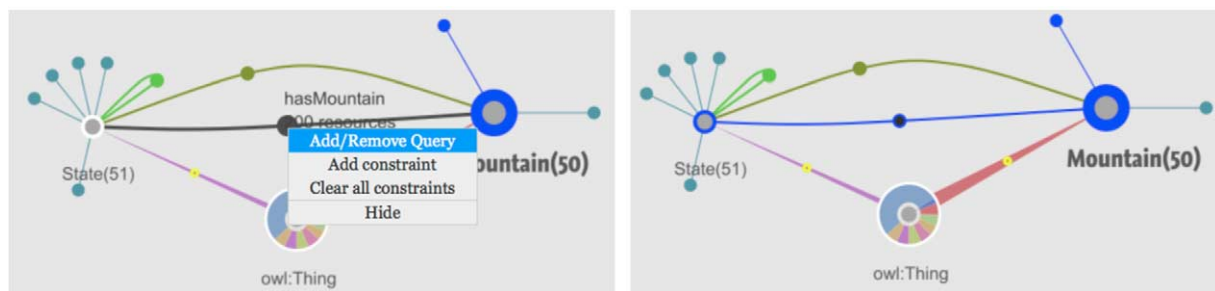[17] http://www.cs.utexas.edu/users/ml/nldata.html

Fig. 9. Visually building queries.

preted as. The following is an example of a SPARQL query generated as a result of the user interactions as shown in Fig. 9

```
@prefix mooney:<http://www.mooney.net/geo#>.
1. SELECT * WHERE {
2.   ?State a mooney:State.
3.   OPTIONAL { ?State rdfs:label ?State_Label}.
4.   ?Mountain a mooney:Mountain.
5.   OPTIONAL { ?Mountain rdfs:label ?Mountain_Label}.
6.   ?State mooney:hasMountain ?Mountain.
7.   ?Mountain mooney:height ?height.
8. }
```

The SPARQL query thus generated consists of two parts: partial query directly built by the user (lines 6 and 7) and partial query prepared by Affective Graphs (lines 2–5). While the user-built partial query is a direct representation of the selection made by the user, the partial queries are built out of the concepts that are selected – the final query looks for instances and their instances that are types of the selected concepts. These instances are 'joined' by the properties that are selected. Labels are used to render the results in a user-friendly way, separating the content from its formalised representation.

Often, users may have highly specific information need that they would like to query for such as the birth date or birth place of Elton John or a generic query such as the height of all mountains and rivers within a state that contains the character sequence *'miss'* within their names. Queries such as these (FILTER queries) can be constructed in Affective Graphs using constraints – users can right click a node or concept that is set as a query and select 'Add Constraint' from the context menu, which would load a dialog that prompts for constraints. Figure 10 shows the user entering a constraint that sets the number of pages in a book authored by a writer to be greater than 300.

Users can also select if this constraint would be set as a negation constraint, as well as an OR query. The data type of property would dictate the type of constraint – if the data type is a string literal, then the fil-
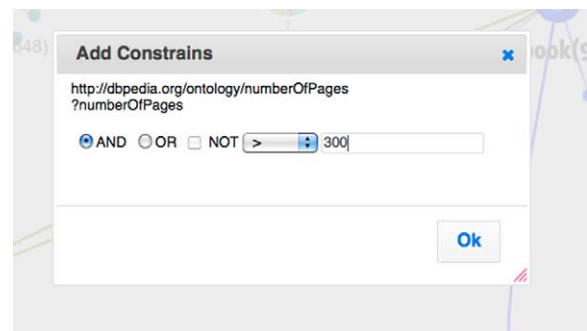

Fig. 10. adding constraints.

ters being applied would be a regular expression filter. Instead, if the data type would be numeric, then comparisons would be possible. After a constraint is set, Affective Graphs would communicate the user of the new action by setting the respective query concept or property in a darker shade.

The following SPARQL query represents the interaction as shown in the figure:

```
@prefix dbont:<http://dbpedia.org/ontology/>.
1. SELECT * WHERE {
2.   ?writer a dbont:Writer.
3.   ?writer rdfs:label ?writer_Label.
4.   ?book a dbont:Book.
5.   ?book rdfs:label ?book_Label.
8.   ?writer dbont:notableWork ?book.
9.   ?writer dbont:birthName ?birthName.
10.  ?book dbont:numberOfPages ?numberOfPages.
11.  FILTER ( (?numberOfPages > 300))
12. }
```

Setting concepts and properties as query items would result in the creation of partial queries on lines 2–10 and setting constraints would result in the creation of partial query on line 11. Multiple constraints would generate multiple lines of FILTER queries. Once the user has completed identifying the concepts and properties of interest, the relevant SPARQL query is displayed at the bottom of the screen (Section D, Fig. 3). It can be often useful for users to configure their queries and only select concepts that they are in-

terested in – e.g. though a query may contain multiple concepts and properties, it could be possible that in the end, a user is only interested in one particular concept and uses the rest as means of constructing logical joins to arrive at the resulting concept. Pagination and limiting result sets could also be a useful feature when dealing with large result sets. Users can click on the 'Search' button (Section E, Fig. 3) to get results, which would be presented at the bottom of the screen.

9.2. User interactions and contextual information presentation

Being a graphical interface, Affective Graphs supports mouse interactions such as left and right clicks, drag and hover as well as pre-configured keyboard short-cuts. Users navigate through the graphs by using conventional techniques such as left click on pie sections to create new nodes, hover on the sections to see concept labels, right click to load a context menu, drag nodes to reposition them to a more convenient location.

The right hand side of the Affective Graphs interface shows contextual information (i.e. Context Section) about the concept currently in focus, as well as the query being built (marked as Section B in Fig. 3). This context section is presented as an overlay on the graphical element, which can be easily hidden if the user wishes to. In addition to the constraints already applied in the query, this section presents a list of all the properties (object and data type) that are associated with the current concept in focus, along with indicating which other concepts are connected to the current concept via the object properties.

9.3. Result presentation

Presentation of results is a challenging process that can have multiple solutions, based on different motivations such as user preference, expertise, application framework, domain and so on. The solutions that we have considered are mostly visualisation of result sets as charts, graphs, maps and so on by incorporating basic visualisations. However, in our current implementation we decided to present results in a sortable table, improved from a standard endpoint presentation as HTML tables.

It should be noted that the presentation of the results is not an integral part of Affective Graphs as the system is to provide users with an interactive and highly visual way of exploring and querying unknown Linked Data.

The final system would be an integrated system, combining Affective Graphs with a dashboard approach for presenting results that had been previously developed [57,58].

9.4. Tools used

Affective Graphs was built using a client-server architecture. A web-based interface was developed using HTML and javascript. The visualisation is built using Processingjs.[18] CSS[19] and jqueryUI[20] are responsible for styling the interface, while jquery[21] handles the interaction with the server. The backend consists of PHP[22] scripts, using ARC2[23] to interact with Linked Data endpoints.

10. Evaluation

As a user-centered development process, several sessions of discussions, focus groups, and evaluations with users shaped the final interface for Affective Graphs. Changes were functional as well as enhancements (such as adding contextual menu items, tooltips, search boxes etc.) after each session with users, hence the interface has significantly evolved since its inception. We discuss three significant evaluations that are most relevant to this paper:

- Evaluation 1: An objective evaluation of aesthetic properties of the system, compared with existing tools.
- Evaluation 2: A user evaluation with experts and casual users to understand how the tool performs compared to other tools.
- Evaluation 3: A user evaluation with Semantic Web experts to understand how well users perceive the system with increased exposure to the tool.

10.1. Evaluation of aesthetic properties

Several rounds of re-design after consecutive user-feedback resulted in a version of Affective Graphs that was relatively mature and ready to be evaluated with

---

[18] http://processingjs.org/js
[19] http://www.w3schools.com/css/
[20] http://jqueryui.com/
[21] http://jquery.com/
[22] http://php.net/
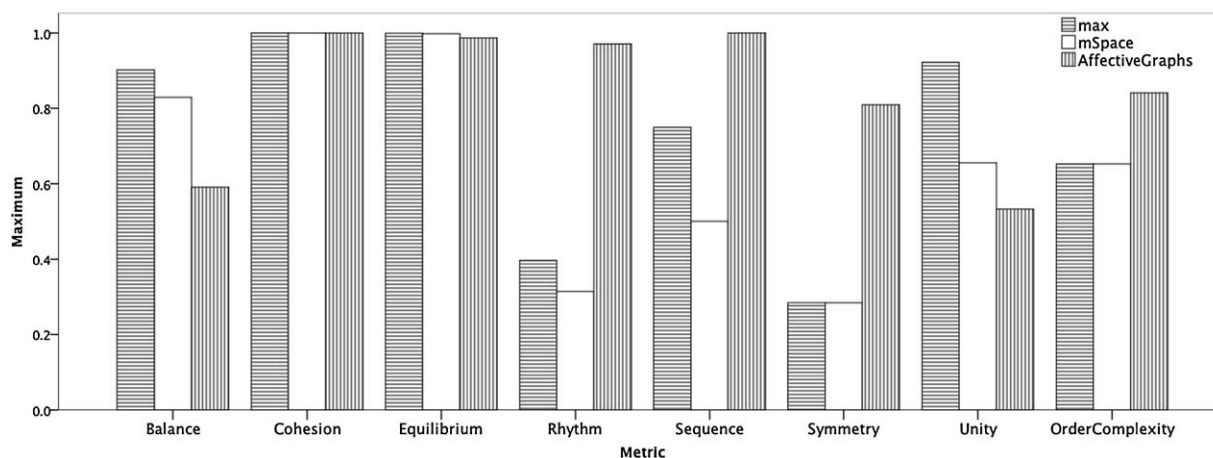[23] https://github.com/semsol/arc2/wiki

Fig. 11. Comparison of Affective Graphs with the highest scoring tool and the maximum scores in each aesthetic metric.

a final set of users. In order to understand the aesthetic properties of the interface, we analysed Affective Graphs using the same metrics as we had previously used to compare existing Semantic Web interfaces (see Section 5). In this evaluation, we wanted to answer two major questions:

– How does the system compare with respect to the system that was judged to be the most aesthetically pleasing tool?
– How does the system compare with the highest score achieved by any tool for the individual categories?

Section 5 showed that the Semantic Web interfaces that we had earlier analysed had relatively low scores, with the highest score obtained by any tool being 0.65 out of a maximum possible score of 1. The eight measures of aesthetic properties were calculated for Affective Graphs and compared with two other sets of scores – the tool which scored the highest in our previous experiment, mSpace. Additionally, we compared Affective Graphs with the highest scores obtained by any tool in each of the individual criteria. The experiment identified the tools which obtained the highest scores for different categories – balance (Sig.ma), Equilibrium (Sig.ma), Symmetry (mSpace), Sequence (PowerAqua, NLP Reduce), Rhythm (K-Search), Cohesion (mSpace) and Unity (NLP Reduce).

Similar to the previous experiment, the scores for each metric are calculated three times. However, in this experiment, we compare the highest scores obtained by each, since we are interested in the best scores. Figure 11 shows a comparative plot of the best scores obtained by any tool, best scoring tool

and Affective Graphs. The figure shows that Affective Graphs scored the highest in four out of the seven categories (Rhythm, Sequence, Symmetry and Cohesion). The system scored lesser than the other two in balance and unity. Equilibrium scores are marginally lower than the other two, with Affective Graphs scoring 0.987 as compared to 0.999 by the highest scoring tool (Sig.ma).

Overall, as can be seen from the graph, Affective Graphs scored the highest, significantly higher than the best scores obtained by the Semantic Web tools (an order and complexity score of 0.8404 as compared to 0.6523). Whilst these scores, are by no means conclusive in deciding the most aesthetically pleasing interface, the positive results serve as a good indication toward developing a more pleasant experience for users. Our intent, in this evaluation was not to judge an interface as the most aesthetically pleasing one, but to explore an alternative way of objectively evaluating Semantic Web interfaces and assessing how the implemented system scores with respect to existing systems. This is important as the system was designed considering aesthetics as one of the most important features.

An objective evaluation as the one explored can only provide an indication of the aesthetic properties of a system. However, the truest reflection on the aesthetic appeal of an interface can only be provided by a subjective judgement of the users. Personal preferences, bias, societal impact, learning experiences and other factors influence a user's judgement and preference for a particular interface. This makes it an extremely difficult task to assess a real user's perception to an interface. While we acknowledge the significant role of users in determining the pleasurable quality of an in-

terface, we believe that early objective analysis of the aesthetic principles is helpful and can provide a starting point for development. As previously discussed, interface layout is only one of the several factors that contribute toward the aesthetic appeal of an interface. However, other factors such as colour and texture are considerably difficult to formulate in order to ascertain an objective value [60].

10.2. User evaluation with casual and expert users

Aligning an interface with aesthetic principles and objective metrics can help estimate the visual pleasure that users may experience while using the system. However, it is important to verify if the approach of the system can be used to perform fact finding tasks, which users of Linked Data engage themselves with. There were two main objectives for this evaluation:

1. How does Affective Graphs perform in comparison to other systems employing different query approaches?
2. How does an aesthetically designed interface affect the user's perception of the system as a whole?

In order to answer the first question, we identified other systems which employ different querying mechanisms (such as natural language, form based and graph based) and evaluated all the tools together with the same questions and dataset. The evaluation also included user satisfaction questionnaires that users were provided with at the end of every session with each tool. The questionnaires were aimed at gathering subjective feedback for each tool, which were then later analysed. The following describe the experiment design and the subsequent analysis.

10.2.1. Experiment design

Five systems (Semantic Crystal, K-Search, Ginseng, Affective Graphs and NLP Reduce) employing different querying techniques (visual, natural language, form based) were evaluated independently in a comparative setting as a part of the second evaluation campaign in the SEALS project.[24] Five questions carefully chosen from the Mooney Natural Language Learning Data[25] were presented to the users, which needed to be answered by interacting with the systems. The questions were of different complexi-

ties: a query involving one concept and one relation (*e.g. All capitals of states in the USA*); a query involving two concepts and two relations (*e.g. Cities in states through which a river Mississippi runs*); a query involving a comparison (*e.g. States that have a city named Columbia with a city population of over 50,000*); a query that involves superlative comparison (*e.g. Lakes that are present in a state with the highest point*) and a negation query (*e.g. Rivers do not traverse the state with the capital Nashville*). Since different query techniques are employed by the systems, some of them are a better natural fit for the type of tasks that were presented to users. For example, a form-based approach would be a better natural fit for answering specific questions compared to a graphical one. While Affective Graphs would fit naturally in an exploratory scenario, the constraints imposed by the SEALS project to comparatively evaluate and benchmark systems limited the types of tasks that could be used for the evaluation. The next evaluation, described in Section 10.3 hence involved exploratory tasks to understand how Affective Graphs could provide support to users in such situations. The users, in either of the evaluations were not aware that an explicit focus on aesthetics was made while designing the system. The first evaluation discussed in this section had the same introduction to all the systems to ensure that the users would not be biased toward any particular system.

Twenty users (10 experts and 10 casual users) aged between 19–46, with a mean age of 30 years were recruited to evaluate the systems via email. Expert users had knowledge and experience with Semantic Web, while casuals had little or no knowledge. The users tried the systems in a random order, to reduce the impact of learning or frustration toward a particular tool.

The evaluation was conducted by a test leader who was independent from the development of any of the tools, to avoid any bias during the evaluation.

For the scope of this paper, we discuss the evaluation from the perspective of Affective Graphs, comparing with the other systems. The results of the comparative study conducted by the SEALS project is discussed in more details by Elbedweihy [23].

10.2.2. Results

We analyse the results for the two user types in two ways: we explore how long it takes for a user to formulate a satisfactory query (query input time) and how many times they executed their queries to perform their tasks (number of attempts). Figure 12 shows three box-
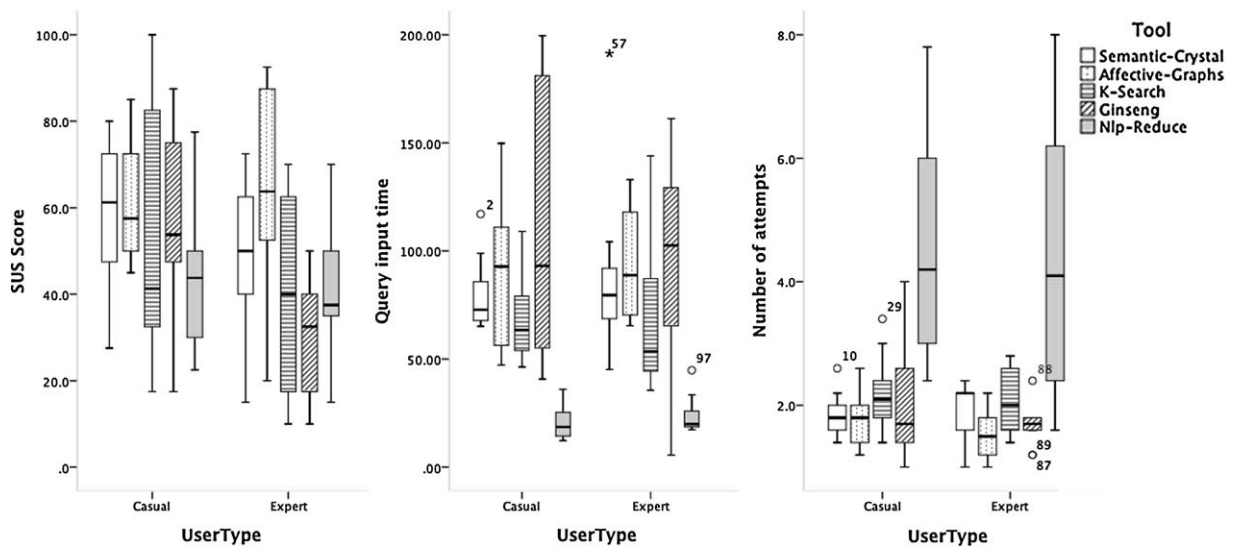

Fig. 12. Evaluation results for SUS scores, query input time and number of attempts for Affective Graphs: the users have been grouped into two types – experts and casual.

plots – SUS score, query input time and number of attempts for each tool clustered by the user type.

Analysing the query input time (middle boxplot) shows a trend that we had expected to observe – NLP Reduce took the least time to formulate queries, while Semantic Crystal and Affective Graphs took considerably longer. The only exception being Ginseng, where users were frustrated by a restrictive natural language interface (this is described in more details in [23]). Users took relatively longer to formulate queries using the graphical approach of Affective Graphs. While conducting the experiment, it was observed that this was mostly due to the fact that users found themselves engaged with the system, and were interested in exploring different features of the interface. The number of attempts also provided an interesting insight – users took the most number of attempts in retrieving satisfactory results using NLP Reduce. Affective Graphs scored among the least for both experts and casuals. However, given their prior knowledge in Semantic Web formalisms, experts took the least number of attempts using Affective Graphs.

Combining the query input time and number of attempts reveals the most interesting observation – casual and expert users took a significantly large amount of time in order to formulate highly precise queries, thereby being able to answer their information need satisfactorily. This finding is key to our efforts, as users can exploit the highly graphical approach to express themselves more precisely to their satisfaction.

The users highly appreciated the system and felt excited about an interactive and intuitive system presenting information in a pleasing way. This was more interesting for the evaluation, as there was no explicit mention of aesthetics playing a critical role in the design of the system. One of the most encouraging comments from one of the participants, *"it is interesting that when you use a colourful and interactive system, you do not mind trying several times to get an answer as it is a playful and enjoyable experience"* clearly identifies with our focus and aim – to help users comprehend, query and explore unknown Linked Data and provide a pleasant, exciting and enjoyable experience. Another interesting comment *"we have been exposed to natural language querying tools like google and yahoo for a long time and hence find ourselves more comfortable with such systems, but had I been introduced to such graphical techniques, I would probably choose them over traditional natural language systems"* shows that a user's pre-disposition and prior experience with natural language interfaces can influence the acceptability of a different solution. However, if the experience of using such an interface is pleasant and enjoyable, there is a greater likelihood of the system being accepted by user communities.

Users were provided with a questionnaire consisting of *System Usability Scale (SUS)* questions, which are standard questions for determining a user's perception toward several aspects of a system. Affective Graphs scored the highest (60.0) in average overall usability

score (SUS) when both the user types are combined (compared to 55 scored by Semantic Crystal, 41.25 by K-Search and 40.0 by Ginseng and NLP Reduce). Interestingly, expert users like the system more than casual users, possibly owing to the prior knowledge of Semantic Web formalisms and graphical approaches to representing data.

While there were users who disliked a completely visual approach toward exploring data, most of the users liked this approach, and would prefer to frequently use the system as a part of their daily analytical tasks. Most users also felt the system was easy to use, though the experts seemed to be more comfortable with the system – we acknowledge this since the experts have prior knowledge of semantic formalisms and have a better understanding of the ontological concepts and visual representations of properties and classes. Similar to ease of use, most often experts found the system easier to learn owing to their knowledge and expertise. Often repeated as comments were that both the casual and expert users found the interface fun, playful and enjoyable to use overall. This is extremely encouraging since this shows that it is possible to interact with Linked Data in a manner which does not involve highly formal and structured ways of querying.

In addition to the questionnaire responses, users were asked qualitative questions that attempt to understand the positive and negative features of the system. The responses of the questionnaires were collated and grouped into different categories:

1. Affective Graphs Interface
2. Visual Query Mechanism
3. Result presentation
4. Others

Our discussion is driven by the positive and negative comments regarding these aspects of the system. The evaluation with the users has been highly satisfying, where users confirmed our approach and appreciated the different features that promise to make interacting with Linked Data an enjoyable, fun and exciting experience for users. Comments such as *"once I got the hang of it, it made much more sense and was easy to use"*, *"Bit of a learning curve but after that it was quite easy and intuitive to use"*, *"Easiest to define queries out of the ones I've used"* and *"this system was simpler to use than I expected"* show that the users had an initial impression of the system to be difficult to use. However, with a little experience and learning, the intuitiveness and ease of use was apparent. Learnability

is a feature that is extremely important specially with new approaches toward consuming Linked Data.

Comments such as *"The graph visualisation worked well graphics were intuitive and easy to use and combine I liked to see the links between the concepts it made it easier to understand"*, *"The nice user interface made for a more pleasant search experience, and the animations made it clear which concepts were connect"*, *"friendly interface, fun"* and *"The graphic interface is really intuitive and easy to use Visual appearance of system was modern and interesting"* reflects the positive feeling that users had after using the system. Upon asking how the system could be improved, several suggestions came up, such as displaying the entire ontology at one go and hide reverse relations (e.g. hasMountain and isMountainOf). While we agree that hiding reverse relations could help reducing the number of edges, it would be beneficial if there are reverse relations existing in the ontologies (a few datasets do not have any reverse relations defined) and the same data is reflected on a reverse relation (it could be possible that the reverse relations are not populated synchronously), thereby possibly increasing the chances of a user missing information. Showing the entire ontology to a user can also have a negative effect of increasing cognitive burden on users by showing them information that is not relevant to their interests. Our current approach has been to present information to the user only if they convey their interest on specific concepts and relations.

Users appreciated the visual query approach, and liked the interactive mechanisms involved in the querying process. Comments like *"The query generation is intuitive and simple to use. It hides all complexity of the underlying query language. You dont need to think in advance the order of the elements that have to be taken into account in the query and add them any time."*, *"I liked to see how the links and circles activated when I added them to my question. That made me realise what I was actually been query, that also gave me and idea of the coverage of my question"* show that the users appreciated the visual communication of the query being applied and how the querying mechanism attempts to 'hide' the complexity involved in building a SPARQL query. However, users felt that several things could be improved in the system in two main areas: automatically linking query concept and properties (where smaller queries can be linked to construct a single larger query, automatically selecting concepts as query elements when a property is selected), a more varied colour palette to prominently highlight filter

constraints. Users also felt that there could be more on-screen help to guide the users in building queries and there seemed to be some cognitive gap among users while converting a natural language question (tasks) into a representative visual query.

Results presentation was by far the weakest aspect of the system – other systems such as K-Search and Semantic Crystal scored relatively higher in this category. We believe there are three reasons for this: a highly visual and interactive mechanism of querying and exploring data generated an expectation of a similar representation for the result sets (comments such as *"better presentation of query results"*, *"perhaps a more graphical approach to the answers like when creating the select statement would help"* indicate that the users were slightly disappointed with a textual result set), lack of enrichment of result sets (comments such as *"i d like to have 'move over' function that brings up a summary of each result"* imply that some level of processing on the results would be helpful) and experimental constraints required users to create queries in a specific manner where they were requested to specify the variable names as a part of the query, thereby resulting in a set of URIs returned rather than labels (*"the results were too SW"*). As previously highlighted, it is also to be noted that Affective Graphs is a tool that was built specifically for querying and exploring Linked Data. A different interface has been built for rendering result sets [57,58], which will be integrated with Affective Graphs at a later stage.

One of the user's comment, *"It would probably look OK as a search system in a Science Fiction B-Movie"* is encouraging and valued to our approach – users seemed to be excited and stimulated by using the system.

Our observations and discussions with the users during the evaluation highlighted several areas and issues that needed more attention. The next iteration of design and implementation addressed such issues and after six months, another set of evaluation was conducted. We discuss the next phase of evaluations in the following subsection.

10.3. Effect of prolonged use (extended learnability)

The motivating factors for conducting this evaluation were two-fold: (1) an evaluation at an earlier development cycle introduced several questions that were highly interesting and required a more in-depth study of how users interact with and respond to the system; (2) as a scheduled evaluation to estimate how newly added features and modifications were perceived by users.

The overall positive results of Affective Graphs was highly encouraging, and user feedback and suggestions were analysed to understand how the system can be improved. The most common feedback was that the users enjoyed interacting with the system and, in spite of being perceived as slightly complex, users found the system highly stimulating and engaging. This was largely credited to the 'playful', attractive and interactive nature of the system. The users, however mentioned that they would like more training and opportunity to learn the system in order to exploit the full potential of the system. This, along with the appreciation of the aesthetic appeal of the tool seeded the next stage of development for Affective Graphs, where our interest was in understanding how learning the system would affect the use of the system. More specifically, we wanted to understand:

1. How easy (in terms of time and effort required) it is to learn how to use Affective Graphs to perform tasks of different complexity and conduct exploratory search tasks?
2. What is the effect of learning on performing tasks?
3. How does learning affect the aesthetic perception of the system?

Learnability, used interchangeably with *ease of learning* is an important criterion of usability that focuses on the ease of learning how to use a system or an interface. [74] describes learnability as the relation of performance and efficiency to training and frequency of use. [62] discusses how learnability can be measured in terms of the time required for a user to be able to perform certain tasks successfully or reach a specified level of proficiency. A similar definition is given by [76] as "the time it takes members of the user community to learn how to use the commands relevant to a set of tasks". [81] argues that measuring usability in a one-time evaluation might be misleading since the use of some applications/systems requires frequency and therefore assessing learnability would be essential.

Learnability has received a significant amount of focus in the literature, some of which focused on assessing learnability as a usability criterion while others investigated how it is affected by different factors (such as interface design). While some of this work focused on *initial learnability* (referring to the initial per-

formance with the system), others looked at *extended learnability* (referring to the change in performance over time) [29]. For example, [34] studied the learnability of two hypermedia authoring tools (HATs) as perceived by academics. Subjects' answers to a set of Likert scale-based questions and their feedback, which was recorded during the sessions, were used to investigate learnability issues. In [64], learnability of two different methods of interaction with databases was compared using similar measures which are based on subjective data (such as questionnaires and users' feedback). [42] assessed the learnability of searching two university Web sites by asking students of the first university to search the other site and vice versa. In contrast to the previous studies, efficiency-based measures, including success rate (number of tasks performed correctly) and the time required to perform the tasks, were used to assess learnability. Additionally, [17,34,69,83] showed that learnability and usability are congruent.

Despite this attention, both IR and Semantic Search evaluations focused either on performance-oriented aspects (such as precision and recall) [3,30] or assessed usability in terms of efficiency, effectiveness and satisfaction, leaving aside learnability and memorability [22,44].

10.3.1. Experiment design

A dataset, consisting of information regarding papers presented in conferences and workshops in the area of Semantic Web[26] was uploaded to a local Virtuoso installation and made available for Affective Graphs to query. There were three main motivations for choosing the dataset, given a group of users – Semantic Web experts are familiar with the dataset; Users have a good understanding of scientific publishing; and Availability of real-world query logs.[27] The logs of the user queries for the dataset were then analysed to understand the different types of requests made by users. Following our analysis, we identified the following four types of queries that are most often used:

1. Simple Task (ST): $C_n A_n F_n$;
   $n = 1$
   Simple queries that comprises only one concept and one attribute but also a filter or a restriction

value applied to the attribute. E.g. *Find the people with first name 'Knud'*

2. Multiple Attributes Task (MAT): $C_n A_m$;
   $n = 1, m \geqslant 1$
   Increased number of attributes without a filter. E.g. *List the name, page and homepage of organisations*

3. Multiple Concepts Task (MCT): $C_n R_m$;
   $n \geqslant 1, m \geqslant 1$
   Searching across multiple concepts, similar to breadth search. E.g. *List all the people who have given keynote talks*

4. Complex Task (CT): $C_n A_m F_o R_p$;
   $n > 1, m, o, p \geqslant 1$
   Include all the four components: concepts with relations linking them, attributes/properties of the concepts as well as filters restricting the values of the attributes. E.g. *Find the page and homepage of each person whose status is 'Academia' and was a chair of a session event and find its location*

where 'C' = Concept, 'A' = Attribute, 'F' = Filter, 'R' = Relation.

Ten expert users (8 men, 2 women) aged between 22–38 (mean of 31 years) were then asked to perform a given set of search tasks using the interface in a controlled laboratory setting over three one-hour sessions. The users were either researchers or software developers highly proficient in Semantic Web technologies as well as conversant with scientific publishing domain. The three sessions were spread over three consecutive days. On the first session, subjects were initially introduced to the experiment and its goals, followed by a 5 minute presentation and explanation of the system. The second session started with a similar 5 minute presentation, with special focus to how the system can be efficiently used and a few shortcuts. Following a 5–10 minutes hands-on practice session, users were then given control of the system and were asked to perform four tasks, one of each type. In addition to the fact-finding tasks, users were also asked to perform an exploratory task, where the real answer to the task was not known, and could depend on the user's interpretation of the question. These exploratory tasks were asked on the first and third sessions. Since we wanted to understand the effect of learning on performing tasks and the exploratory tasks were expected to be time-consuming, we decided to have these tasks as part of the first and third sessions. The set of tasks are provided in Appendix B, where in

---

[26]Semantic Web Dog Food dataset, http://data.semanticweb.org/.
[27]The USEWOD 2012 dataset (http://data.semanticweb.org/usewod/2012/challenge.html) consisted of two years of real-world query logs from 12/2008 till 12/2010.
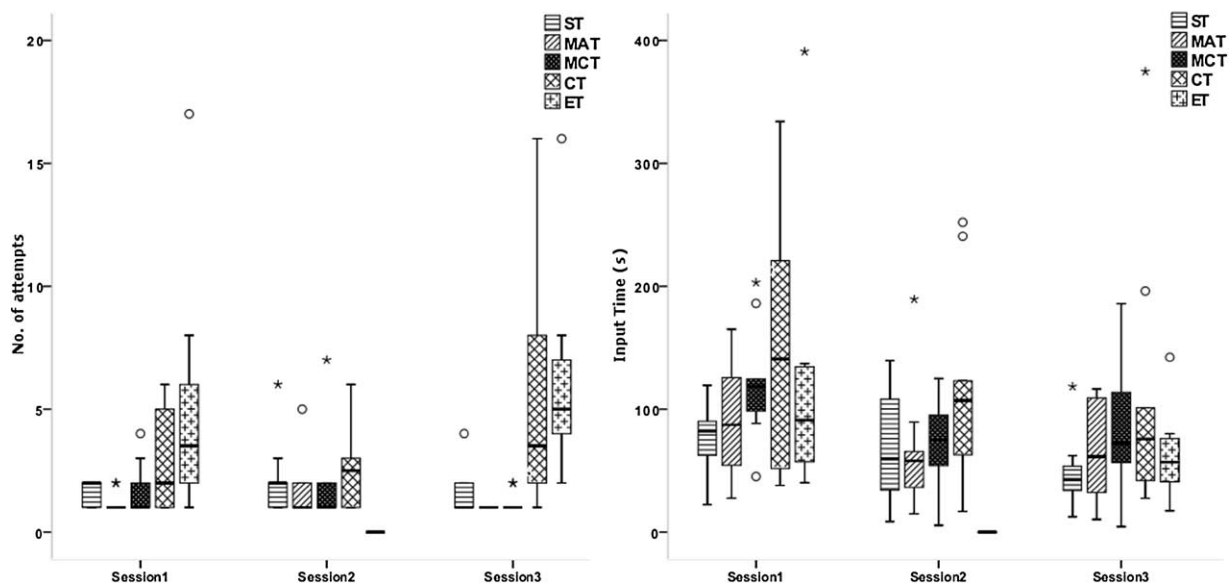
Fig. 13. Analysis of performance of users with five types of tasks: ST: Simple, MAT: Multiple Attribute, MCT: Multiple Concept, CT: Complex, ET: Exploratory. Figure on the left shows a box-plot of the number of attempts required to solve a task satisfactorily, grouped into sessions and task types. Figure on the right shows a box-plot of the input time grouped into sessions and task types.

every session, a random task from each category was selected.

As discussed previously, the most common ways in literature to measure learnability were either based on objective data by comparing users' performance/ efficiency over time or subjectively using learnability questions such as *"I found this interface easy to learn"*. To allow for deeper analysis, we collected both objective and subjective data covering the experiment results. Input time, success rate and number of attempts provided objective data, while responses from three questionnaires (System Usability Scale [11], Extended and Aesthetics) provided the subjective data. Questionnaires were filled in at the end of every session. The system usability scale questionnaire included questions that attempt to understand the affect of learnability on usability. The extended questionnaire included more specific questions on learnability, remembering features and so on. Aesthetics questionnaire included questions on what the user's perception of various aesthetic properties of the system. Additionally, users were also presented with open-ended questions to gather further details about how their experience was.

10.3.2. Data analysis and results

Though a lot of data was collected from our evaluation, our analysis and discussions will focus on two major aspects: how users performed their tasks and

how they perceived the interface. Additionally, we will discuss the feedback users provided in terms of responses to open-ended questions. Following the evaluation, we investigated two main observations: *input time* and *number of attempts*. Here, we define input time as the amount of time taken to compose a satisfying query i.e. time taken from the time the user starts the task till the search is executed. Number of attempts is defined as the number of queries executed by users to complete their tasks. We also analysed these measures to understand how the behaviour of users was affected as they learnt the system.

Objective data, in the form of query logs were collected, which indicated two key features – how quickly do users perform tasks after training, how many attempts at completing the task do the users need before they are satisfied with the results. An evaluation controller collected user interaction events and the logs were later analysed. We grouped the logs into the three sessions and compared the data. This is shown in the Fig. 13.

Efficiency and effectiveness The figure on the left shows a boxplot of the distribution of the number of attempts required to perform the five types of tasks in the different sessions. During the evaluation, we observed two types of behavior among users, based on the type of task. Users behaved similarly when they were faced with simpler tasks (ST, MAT, MCT) and their behavior

changed as the tasks became more complex. Initially, users needed a few attempts at solving simple tasks. Complex tasks required a few more attempts at solving the tasks. The number of attempts required for the simple tasks started reducing over the sessions as users gained more familiarity with the tool. This was an expected result, as users are more comfortable with a new interface with time, and gain more expertise interacting with it. Furthermore, users started trying new techniques and features during the second session, which increased the level of comfort and helped users adapt to Affective Graphs more.

An unexpected observation was the change in behavior while performing slightly complex tasks- users seemed to require more attempts in order to perform the tasks during the latter sessions. This was surprising, as we had expected the users to find such tasks easier with more time and familiarity. The (median) number of attempts for complex task (CT) increased from 2 in session 1 to 2.5 in session 2 and 3.5 in session 3. The (median) number of attempts for exploratory tasks (ET) increased from 3.5 in session 1 to 5 in session 3. This clearly showed a change in the behavior and approach toward solving complex tasks.

The figure on the right shows a boxplot of the distribution of the amount of time required by users to formulate their queries to solve the tasks (input time). In general, we observed a significant decrease of input time from an overall average of 106.48 s in the first session and 72.72 s in the second session to 66.845 s in the final session. All of the types of tasks have shown a steady reduction in the input time. Our observations during the evaluation sessions credited this to the increased comfort and acquaintance with the system and its features with more time and familiarity with the systems. While the relatively simpler tasks (ST, MAT and MCT) have seen a general reduction in the times, the complex and exploratory tasks are of greater interest to our analysis owing to the more complex nature of the tasks.

The user's performance in the complex tasks (CT and ET) appear highly interesting. The reduction in time is significant – from a median of 140.86 s in the first session and 107.01 s in the second to 75.575 s in the final session for the complex tasks (CT) and 91.16 s in the first session to 56.87 s in the final for the exploratory tasks (ET). This explains the earlier observation where the number of attempts increased with more familiarity with the tool – greater familiarity and comfort with Affective Graphs helped users try several things more quickly as users found it easier to formu-

late queries. This was also observed during the evaluations: initially, users were carefully building long queries, connecting multiple concepts. This technique gradually changed to a different one in the second session, where users tried short bursts of queries, gradually building up to form a longer one. Users could use the outcome of the short queries to quickly formulate a longer query, which was well-informed and driven by the results of short queries. Let us consider the second exploratory task in Appendix B: the task requires the user to identify persons who are experts in Knowledge Management. There could be several approaches toward solving this task – users could look at all persons that have organised tutorials that are associated with Knowledge Management, or all persons who have several publications on the topic. The ultimate goal of this task is to connect multiple concepts (either proceedings, tutorials or workshop events) by corresponding relations, and identify people. The approach followed in session 1 was to select all the relevant relations, set constraints and connect the concepts in the very first attempt at one go – this would not give any results for many queries, thereby making the user try re-building the entire query after every attempt. Upon realising this repetitive process, most users gradually shifted their approach toward building short queries (to find all the authors, tutorial authors etc.) and investigate the results to build a final query that was more certain to provide useful answers.

User satisfaction Three questionnaires were presented to users at the end of every session – a general usability questionnaire (ten 5-point Likert scale questions for computing SUS scores [11]), an extended questionnaire (five 5-point Likert scale questions and 2 open questions, related to understanding learnability, ease of use and remembering features) and an aesthetics questionnaire (fifteen 5-point Likert Scale questions, related to understanding various aesthetic features of the interface). Finally, two open-ended questions were asked at the end of the third session, which aimed at understanding what users liked/disliked in the system, and how that changed over the sessions (as a result of learning and increased familiarity).

Overall, the system was appreciated by users, and the high SUS scores indicate the positive response. The high SUS score of 76.25 in the initial session indicated the first reactions of the users to the system – users seemed to like the interactive and 'playful' environment, which evoked an overall positive response. The second session noted a significant improvement in the

SUS scores (82.5) – this was more due to the users being excited and satisfied about successfully trying several features. The final session shows a marginally decreased SUS score (81.25), but still significantly higher than the first. We believe this is due to the users being more accustomed to the system, and having learnt new features over the sessions, the initial excitement of discovering new features had normalised by the third session. Users appreciated the visual approach adopted by Affective Graphs, as shown in the answers when asked what the users liked about the system:

- *"The visual aspect of the system makes it less tedious to perform searches. I also found it easier to define the relations and constraints on the relations than in other systems"*
- *"interface is visually appealing: responsive, colourful, professional feedback is good nice to explore the structure of the underlying data without clutter"*
- *"the graphical interface makes it very easy to use. the exploratory nature is very good, it is easy to learn by exploring"*
- *"This interface eases a lot the use of SPARQL and makes simple querying an RDF dataset"*
- *"Overall though, I liked how queries could be built using the system to retrieve precise facts. It is much easier to use a system like this than having to build queries by hand"*

Affective Graphs was judged by users using 5-point Likert scales, values ranging from 1 to 5, 1 being the most positive, while 5 being the most negative. The system was judged moderately easy to use and understand (Fig. 14, bottom right) initially with a median score of 2 in the initial two sessions. The system was also initially judged to be moderately straightforward. With more familiarity and experience with the tool, users felt more confident and the median scores reduced to 1 from 2 in the first two sessions. Users gradually felt more comfortable exploring new features by trial and error, and this also reflected in the user's subjective scores – A median of 2 in the first two sessions reduced to 1 in the third. Overall, users also did not face any issues remembering features throughout the experiment, and by the third session, they appeared to remember features without any effort. A few comments also noted how the users perception and experience of the system changed over time:

- *"Ability to use the system effectively and my confidence/speed in using the system grew over time."*

- *"The system became easier to use and understand over time."*
- *"With practise, the tool became easier to use. I found satisfactory solutions to the questions much quicker in the second and third sessions than I did in the first."*
- *"At first the system seems complicated to use. After using it a few times I found it a lot easier to construct the queries."*
- *"I also got more confidence with using the system to progressively build queries instead of trying to get everything in place at once."*

Users were also initially confused by few features and visual elements on the screen – for example, the context window on the right (Section B, Fig. 3) was initially perceived as an element which was not helpful, but with more familiarity the use of the window was more evident and users could eventually perform their tasks much better. User comments such as *"(I) thought having to use the right hand context box to find out possible onward links to other concepts was a pain. got used to this and in the last session was using this box all the time to find suitable associated concepts."* and *"With time I learned how to use the pane on the right more effectively to identify potential relationships I could explore between an object already on the canvas and other objects."* show that there is much scope for improvement in such areas, but the value of the content on the window is highly useful if the users are familiar with the tool.

Our continual insistence on an aesthetically pleasing experience as one of the key design requirements had significant benefits on the users. Several users appreciated the interface, and during informal discussions showed a lot of interest in how the tool was built. Comments such as:

- *"Nice UI, clear design to see results"*
- *"easy to use; intuitive; friendly good way to visualise the structure and data. Good to see possible links from selected concepts"*
- *"interface is visually appealing: responsive, colourful, professional feedback is good nice to explore the structure of the underlying data without clutter"*
- *"Its a great search tool!"*
- *"UI nice and is good to use, faster to use than typing SPARQL."*
- *"(Liked) The layout and the connections between the different sets. As it was easy to see where the connection between sets were."*
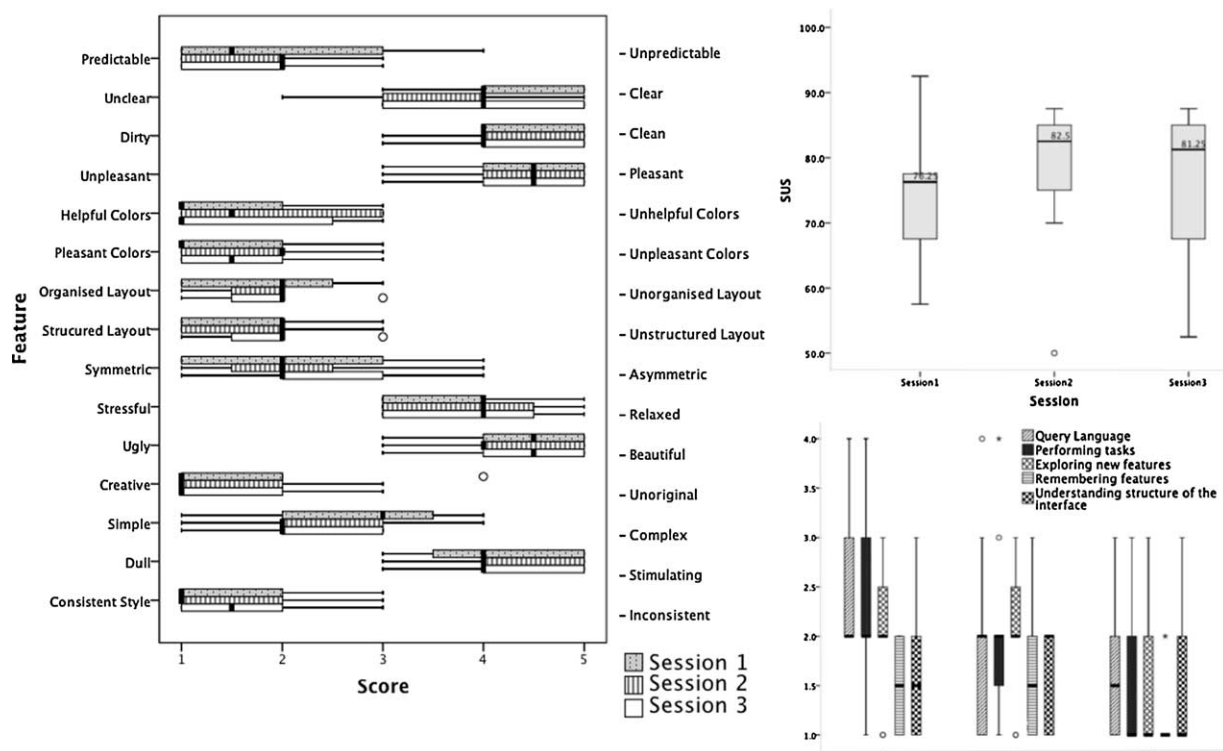
Fig. 14. Results of the User Evaluations. Left – Aesthetics questionnaire responses, Right Top – SUS scores, Right Bottom – Extended questionnaire responses.

– "*(Liked) the highlighting tool when searching for particular entities was well designed, and helped with finding the correct query*"

show the overall pleasant experience users had while interacting with Affective Graphs. We believe that the response to the interface is directly an outcome of our design decisions while developing the system. The Fig. 14 shows how the users perceived different aesthetic properties of the interface over time. We did not notice significant changes in the user's perception of aesthetic the properties. The median scores of 9 of the 15 properties remained constant throughout the sessions. 4 of the properties had a minor variance in the sessions. Most of the users found the interface to be creative, beautiful and stimulating. The colours presented to the users were pleasant and helpful in general. Users initially found the system relatively complex, but with more learning, the system appeared to be simpler to use.

10.3.3. Discussions

While the overall reception of the tool was positive, there were three main aspects which were highlighted during the evaluation. One aspect was the placement of edges in the layout. Recalling from the design decisions, the layout was designed in such a way that executed the force direction algorithm for a few seconds once a user interaction triggers a new node to be rendered. This stopped nodes from overlapping, but there were no measures for removing overlapping properties. User comments such as *"The relational arms in the user interface sometimes overlapped and made it difficult to trace the connections and read the labels"*, *"Arms of the relations overlap sometimes, especially when multiple concepts/relations are on the screen"* and *"Reading the names of certain properties can be difficult when they overlap with the concepts behind them- perhaps the option to drag individual property nodes rather than just the concept nodes would be a way to overcome this?"* highlighted the issue. Increasing the amount of time allotted to execute the force-direction algorithm would help position the new node in a more optimal position, but at the cost of more browser processing. A different suggestion that came from the mentioned user comments was be to enable click-and-drag events to allow users to customise the position of relations. We are also investigating machine learning techniques that model a user's

Table 3

Design guidelines proposed by Dadzie and Rowe [16] visual information presentation catering to two types of users – lay and tech

Tech-users	Lay-users
T1. Intuitive Navigation through LD structures	L1. Intuitive navigation through the large amounts of complex, multi-dimensional data
T2. Data exploration to understand content and structure	L2. Exploratory knowledge discovery
T3. Data exploration to identify links across and within datasets	L3. Support for basic to advanced querying, to support filtering and IR in order to cater to experts as well as casual users
T4. Data exploration to identify errors, noise and anomalies in content and syntax	L4. Detailed analysis of ROIs
T5. Advanced querying using formal query syntax	L5. Publication/syndication
T6. Publication/syndication, verification and validation of new data and links	L6. Data extraction for reuse
T7. Data extraction for reuse	L7. Presentation of the results of analysis to different audiences

re-positioning events on the interface to identify best possible locations for a new node, based on the user's previous behavior or a common behavior identified by analysing multiple users.

Another issue highlighted was the context window on the right of the interface – users had clearly not noticed the usefulness of this section at the start of the experiment. It is important to understand the reasons for this – while, a better and more elaborate explanation of the section would certainly increase its visibility, a change in the design of the section is worth investigating. Users could realise the value of the section only during the final stages of the second session and the final session when they started trying different features and exploring different elements of the interface – a comment by a user on the final session, *"I just noticed today there's a box on the right with information about the concepts and the relations that can be defined with other concepts. I think this information is essential and maybe it's should be made more evident to the user"* sums up the need to re-think the design of the section. Some solutions were also proposed by the users – comments such as *"The help on the side about related concepts is not straight-forward in the first session, perhaps a different view would be more helpful (maybe a tree view?? maybe a small concept graph??)"* provide some suggestions to adhere to a highly visual approach toward presentation of contextual information.

The third issue highlighted was the presentation of the nodes themselves. While users were satisfied with the presentation of the information, a few users noted that hovering over pie sections to find out subclasses of interest could be an intensive process and may be negatively affecting a positive user experience. A few proposed solutions have risen as a result of informal discussions with users – showing labels of all pie sections of the node in focus, showing labels of the most significant sections in the pie chart and showing a legend or list of pie sections in a placeholder to name a few. This, however needs further investigation as providing more labels do not necessarily contribute to a greater visual experience. A minimalist approach has been suggested by many in our literature survey, and multiple labels visible on the work-space would clutter the screen and add to the user's information load.

11. Discussions

Our initial experiments with Linked Data interfaces and aesthetic measures has highlighted the need for explicit attention to aesthetics while designing interfaces for the Semantic Web. Starting from principles developed by the HCI and the Visual Analytics communities, we proposed principles that could be used to develop Semantic Web applications. We used these principles to design and build an interface that facilitates exploratory browsing of Linked Data. [16] provided a comprehensive review of Linked Data visualisation approaches and categorised their design guidelines based on the perspectives of a tech user and a lay-user. We look back at these guidelines and align them with the features provided in the Affective Graphs interface in order to understand how it fits with the greater expectations of the Semantic Web community. The guidelines proposed by the authors are shown in Table 3.

One of the main guidelines proposed by the authors for both lay and tech users is the need for an intuitive interface that facilitates browsing of large complex multi-dimensional data (L1, T1). User evaluations and focus groups have been highly positive and indicate that users had a good experience with the tool while exploring data. The exploratory tasks fo-

cussed on understanding how well users can explore an unknown dataset to find relevant information. Users found the interactive and visual approach stimulating and were willing to explore data in a playful manner, thereby finding answers to their tasks (L2, T2). We also noted a change in behaviour as users started becoming more confident with the system with increasing familiarity and practice. Affective Graphs also makes it apparent to the user how different concepts are hierarchically related and what are the common relationships they share (L3,T3). Exploring nodes in focus provided users with more information about the node, which helped users gather an understanding of a concept without actively searching for it's content (L4).

Although not discussed within the scope of this paper, Affective Graphs also has features for exporting data from query results (T7, L6, L7). The SPARQL queries are stored in the system, and if a user is interested in the result sets, they can export the results in a file. This feature was disabled during the evaluations as it was not a focus of our experiments. Since Affective Graphs is not meant to be a standalone system, and will be integrated with another visualisation framework, which makes it possible to simultaneously visualise result sets in multiple facets (L7). One of the most useful features identified by Dadzie and Rowe is the ability to query for specific instances of data within a dataset. Affective Graphs helps users build queries in a highly visual an interactive manner. We believe this would be of immense help to lay users as they would not be trained in formal query syntax. Their interactions with visual elements would generate queries which would enable them to answer specific information needs (L4, L3). The interactions can also serve as a starting point for advanced users, who can then directly edit the SPARQL query thereby generated or modify various query parameters such as limit the number of results, order results and so on (T5).

12. Conclusions

The starting point of Affective Graphs was understanding how to approach exploration of semantic data in a highly visual approach. In addition to satisfying a user's information need and providing an intuitive and interactive experience, an aesthetically pleasing interface is paramount to a user's positive impression on a system. The recent movement toward the release of thousands of open datasets as Linked Data by governments and organisations have fostered an environment

where Semantic Web practitioners, enthusiasts, developers and researchers can develop highly useful applications, integration services and mashups.[28] However, greater stress on aesthetic appeal of interfaces for the Semantic Web is essential. This is more urgent as Semantic Web technologies and Linked Data grows to a more established and mainstream research direction – standards such as HTML5, Schema.org and movements like big data, knowledge graph,[29] semantic search[30] have already been employed to great success.[31] We acknowledge the assertion of Cruesen [15] as aesthetics being one of the most important factors that influence product choice. Tracktinsky's [78] notion of "what is beautiful is usable" stresses the importance of an aesthetically pleasing design in influencing perceived usability – aesthetically pleasing solutions are perceived to be more usable. Our experience with the user evaluations also noted a similar response from users – even when users were unable to perform tasks, they enjoyed the experience and were willing to try several times even after many unsuccessful attempts. While most solution developers attempt to answer the user's information need at the first few attempts, it is often possible that the user would need to re-attempt several times. Our observations noted that an aesthetic interface helped reduce frustration among users if they failed to perform tasks after several attempts. Discussions with users also showed that users tend to remember their experience when an aesthetically pleasing tool is used. Several participants had shown keen interest in the functioning of the system, as well as requested for copies of the application deployed on their own datasets. While one might argue that a positive experience arises out of a combination of several factors such as functionality, effectiveness, intuitiveness and so on, we believe that aesthetics has helped in influencing how users perceived the system.

As described previously, the prime focus of the approach has always been motivated toward an aesthetically pleasing experience. Our investigation into the literature highlighted several recommendations and principles that are relevant to Linked Data explo-

---

[28]As of March 2013, CKAN (http://thedatahub.org/) registered 5216 open datasets including data about railways, census reports, emission data, meteorology and so on.

[29]http://www.google.com/insidesearch/features/search/knowledge.html

[30]https://www.facebook.com/about/graphsearch

[31]http://semanticweb.com/the-semantic-web-has-gone-mainstream-wanna-bet_b27329

ration, and the Semantic Web community in general. We distinguish the design principles into two: general principles and node-link principles. The most common general recommendation that was identified from our literature survey suggests using interfaces that are well-balanced, proportional and symmetric. Another common suggestion is to use a minimalist approach that can aid in reducing visual clutter. Suggestions to provide mechanisms to support standard visual information seeking tasks have also been recommended. Principles such as these can be quickly referenced while developing prototypes and solutions in order to ensure an aesthetically pleasing experience. Affective Graphs was developed starting from the aesthetic principles, and design decisions were taken to ensure the tool follows such recommendations.

Following an iterative user-centered development process, Affective Graphs has been re-designed several times, every time as a result of an evaluation or focus group. After several re-designs, the final version of the system was objectively evaluated on the basis of the layout of visual elements on the screen. This layout was compared with 10 other well known interfaces and research prototypes. The results from this evaluation was highly encouraging, with Affective Graphs scoring the highest among all the Semantic Web tools. Such a study had not been conducted in the Semantic Web community. It is important to understand how visual elements are arranged in an interface, specifically because human response to visual objects is dictated by several properties such as shape, order, symmetry, balance and so on. We followed such an approach as it provides a relatively inexpensive way of initially validating the layout of an interface. The final evaluation was conducted with 10 expert users where users evaluated the system in three sessions, over a period of three consecutive days. The aim of this evaluation was to understand how well users perceive the system and also how their perception changes over sessions. We analysed objective and subjective data to gather an understanding of how users change their behavior while performing tasks as well as how their judgement of the system changes as they gain more familiarity with the system.

Overall, the positive results of the user evaluations are highly promising and we believe that more focus should be stressed on aesthetics while developing tools for the Semantic Web. This is particularly important at this critical juncture where thousands of Linked Datasets are being released to the public, to be exploited. While more work is needed in establishing the role of aesthetics in the Semantic Web community, we believe this is a step forward and the positive evaluations with the users and their feedback is highly encouraging. Future work includes modifying the context window to provide a greater visibility to users. Changes in the layout algorithm would also be investigated. In the future, another interesting study could be to understand the impact of learnability sessions on users in a much larger scale – to initially evaluate the system with users and then conduct a multi-session evaluation for learning, and finally re-doing initial evaluation to understand how the sessions have impacted use. The most important follow-up work would be in integrating a dashboard visualisation interface [57,58] with Affective Graphs, where users can query using the visual approach of Affective Graphs, and visualise the results in multiple visual perspectives of the dashboard.

Acknowledgements

Appendix

A. Aesthetic metrics

The following provides a description and the formulae to calculate the metrics as used in the evaluation. These definitions and formulae were proposed by Ngo.

Balance is defined as the distribution of optical weight in a picture, where optical weight is the perception that some objects appear heavier than others. Larger objects are heavier while smaller objects are lighter.

Balance

$$= 1 - \frac{|BM_{vertical}| + |BM_{horizontal}|}{2} \in [0, 1] \tag{1}$$

Equilibrium represents stabilisation, a midway center of suspension. It can be defined as equal balance between opposing force, various visual objects are centers of forces. A layout is in equilibrium when its center coincides with the center of the frame.

$$Equilibrium = 1 - \frac{|EM_x| + |EM_y|}{2} \in [0, 1] \quad (2)$$

Symmetry denotes the balanced distribution of equivalent elements about a common line. Essentially representing axial duplication, symmetry defines how well a unit on one side of the center is replicated on the other side.

$$Symmetry$$
$$= 1 - \frac{|SYM_v| + |SYM_h| + |SYM_r|}{3} \in [0, 1] \quad (3)$$

The measure of **sequence** relates to the way that visual objects are positioned in a layout with respect to the movement of the eye – heavier objects being on the top left, while lighter and smaller objects at the bottom right.

$$Sequence$$
$$= 1 - \frac{\sum_{j=UL,UR,LL,LR} |q_j - v_j|}{8} \in [0, 1] \quad (4)$$

Rhythm relates to understand the variety in the arrangement, dimension, number and form of visual objects within a layout.

$$Rhythm$$
$$= 1 - \frac{|Rhythm_x| + |Rhythm_y + |Rhythm_{Area}|}{2}$$
$$\in [0, 1] \quad (5)$$

Cohesion denotes how the aspect ratios of each visual element relates to the screen's width and height.

$$Cohesion = \frac{|CM_{fl}| + |CM_{lo}|}{2} \in [0, 1] \quad (6)$$

Unity signifies coherence, where visual elements appear to belong together, seen together as one thing –

similar sized objects, using less space between elements, larger margins and so on.

$$Unity = \frac{|UM_{form}| + |UM_{space}|}{2} \in [0, 1] \quad (7)$$

Order and Complexity is defined as the sum of all the above measures for a layout.

$$Order\_Complexity = \frac{\sum_{i=1}^{7} M_i}{7} \in [0, 1] \quad (8)$$

B. Tasks for user evaluation

B.1. Simple Tasks (ST)

1. Find the people with first name 'Knud'.
2. Find the "inproceedings" whose title contains 'Semantic Search'.
3. Find the organisations whose name contains 'Karlsruhe'.

B.2. Multiple Attributes Task (MAT)

1. List the name, page and homepage of organisations.
2. List the name, familyName and status of all people.
3. List the location, homepage and summary of all tutorial events.

B.3. Multiple concepts Task (MCT)

1. List all the conference venues and their meeting rooms.
2. List the programme committee members and the conference events they participated at.
3. List all the people who have given keynote talks.

B.4. Complex Task (CT)

1. Find the description and summary of keynote talks which took place at 'WWW' conferences and the name of the presenter.
2. Find the name, homepage and page of people who were workshop organisers for a workshop about 'Ontology Matching'.
3. Find the page and homepage of each person whose status is 'Academia' and was a chair of a session event and find its location.

B.5. Exploratory Task (ET)

1. Imagine you are a young researcher, starting your career in 'Ontology alignment'. Since the organisation you are affiliated to conducts research in a different area, you do not have direct access to experts in your area of research. The only access to information is via Affective Graphs, which provides visual means to look for information. Using such systems, can you identify a few researchers in the area of your interest and why have you chosen them?

2. Imagine you are organising a day-long workshop on knowledge management in business at an organisation. As a part of the workshop, there would be two tutorials from experts. Who are the experts you would choose and why?

References

[1] A. Altaboli and Y. Lin, Objective and subjective measures of visual aesthetics of website interface design: The two sides of the coin, in: *Human-Computer Interaction. Design and Development Approaches*, J. Jacko, ed., Lecture Notes in Computer Science, Vol. 6761, Springer, Berlin Heidelberg, 2011, pp. 35–44, ISBN 978-3-642-21601-5.

[2] G. Andrienko, N. Andrienko and S. Wrobel, Visual analytics tools for analysis of movement data, *SIGKDD Explor. Newsl.* **9**(2) (December 2007), 38–46, ISSN 1931-0145.

[3] K. Balog, P. Serdyukov and A.P. de Vries, Overview of the TREC 2010 Entity Track, in: *TREC 2010 Working Notes*, NIST, November 2010.

[4] M. Bar and M. Neta, Humans prefer curved visual objects, *Psychological Science: A Journal of the American Psychological Society/APS* **17**(8) (August 2006), 645–648, ISSN 0956-7976, URL http://www.ncbi.nlm.nih.gov/pubmed/16913943.

[5] F. Beck, M. Burch and S. Diehl, Towards an aesthetic dimensions framework for dynamic graph visualisations, in: *Proc. of the 2009 13th International Conference Information Visualisation, IV '09*, IEEE Computer Society, Washington, DC, USA, 2009, pp. 592–597, ISBN 978-0-7695-3733-7, URL 10.1109/IV.2009.42.

[6] C. Bennett, J. Ryall, L. Spalteholz and A. Gooch, The aesthetics of graph visualization, in: *Proc. of the Third Eurographics Conference on Computational Aesthetics in Graphics, Visualization and Imaging, Computational Aesthetics'07*, Aire-la-Ville, Switzerland, 2007, pp. 57–64, ISBN 978-3-905673-43-2, Eurographics Association, URL http://dx.doi.org/10.2312/COMPAESTH/COMPAESTH07/057-064.

[7] T. Berners-Lee, Y. Chen, L. Chilton, D. Connolly, R. Dhanaraj, J. Hollenbach, A. Lerer and D. Sheets, Tabulator: Exploring and analyzing linked data on the semantic web, in: *Proc. of the 3rd International Semantic Web User Interaction*, Athens, Georgia, USA, 2006.

[8] R. Bhagdev, S. Chapman, F. Ciravegna, V. Lanfranchi and D. Petrelli, Hybrid search: Effectively combining keywords and ontology-based searches, in: *Proc. of the 5th European Semantic Web Conference*, M. Hauswirth, M. Koubarakis and S. Bechhofer, eds, LNCS, Springer Verlag, Berlin, Heidelberg, June 2008.

[9] U. Brandes and D. Wagner, Using graph layout to visualize train interconnection data, in: *Proc. of the 6th International Symposium on Graph Drawing, GD '98*, Springer-Verlag, London, UK, 1998, pp. 44–56, ISBN 3-540-65473-9, URL http://dl.acm.org/citation.cfm?id=647550.728937.

[10] R. Brath, M. Peters and R. Senior, Visualization for communication: The importance of aesthetic sizzle, in: *Proc. of the Ninth International Conference on Information Visualisation, IV '05*, IEEE Computer Society, Washington, DC, USA, 2005, pp. 724–729, ISBN 0-7695-2397-8.

[11] J. Brooke, SUS: A quick and dirty usability scale, in: *Usability Evaluation in Industry*, P.W. Jordan, B. Weerdmeester, A. Thomas and I.L. Mclelland, eds, Taylor and Francis, London, UK, 1996.

[12] J.V. Carlis and J.A. Konstan, Interactive visualization of serial periodic data, in: *Proc. of the 11th Annual ACM Symposium on User Interface Software and Technology, UIST '98*, ACM, New York, NY, USA, 1998, pp. 29–38, ISBN 1-58113-034-1.

[13] P. Chandler and J. Sweller, Cognitive load while learning to use a computer program, *Applied Cognitive Psychology* **10**(2) (1996), 151–170, ISSN 1099-0720.

[14] W.C. Cleveland and M.E. McGill, *Dynamic Graphics for Statistics*, 1st edn, CRC Press, Inc., Boca Raton, FL, USA, 1988, ISBN 053409144X.

[15] M.E.H. Creusen and J.P.L. Schoormans, The different roles of product appearance in consumer choice, *Journal of Product Innovation Management* **22**(1) (2005), 63–81, ISSN 1540-5885, URL http://dx.doi.org/10.1111/j.0737-6782.2005.00103.x.

[16] A.-S. Dadzie and M. Rowe, Approaches to visualising linked data: A survey, *Semant. Web* **2**(2) (April 2011), 89–124, ISSN 1570-0844, URL http://dx.doi.org/10.3233/SW-2011-0037.

[17] F.D. Davis, R.P. Bagozzi and P.R. Warshaw, User acceptance of computer technology: A comparison of two theoretical models, *Manage. Sci.* **35**(8) (August 1989), 982–1003, ISSN 0025-1909, URL http://dx.doi.org/10.1287/mnsc.35.8.982.

[18] A. De Angeli, A. Sutcliffe and J. Hartmann, Interaction, usability and aesthetics: What influences users' preferences? in: *Proc. of the 6th Conference on Designing Interactive Systems, DIS '06*, ACM, New York, NY, USA, 2006, pp. 271–280, ISBN 1-59593-367-0.

[19] A.N. Donald, *Emotional Design: Why We Love (Or Hate) Everyday Things*, Basic Books, January 2004.

[20] M. Dontcheva, S.M. Drucker, D. Salesin and M.F. Cohen, Relations, cards, and search templates: User-guided web data integration and layout, in: *Proc. of the 20th Annual ACM Symposium on User Interface Software and Technology, UIST '07*, ACM, New York, NY, USA, 2007, pp. 61–70, ISBN 978-1-59593-679-0.

[21] H. Eichelberger, Nice class diagrams admit good design? in: *Proc. of the 2003 ACM Symposium on Software Visualization, SoftVis '03*, ACM, New York, NY, USA, 2003, pp. 159–ff, ISBN 1-58113-642-0.

[22] K. Elbedweihy, S.N. Wrigley, F. Ciravegna, D. Reinhard and A. Bernstein, Evaluating semantic search systems to identify future directions of research, in: *Second International Workshop on Evaluation of Semantic Technologies*, CEUR Workshop Proceedings, Vol. 843, May 2012, pp. 25–36.

[23] K. Elbedweihy, S.N. Wrigley and F. Ciravegna, Evaluating semantic search query approaches with expert and casual users, in: *Proc. of the 11th International Conference on the Semantic Web – Volume Part II, ISWC'12*, Springer-Verlag, Berlin, Heidelberg, 2012, pp. 274–286, ISBN 978-3-642-35172-3, URL http://dx.doi.org/10.1007/978-3-642-35173-0_18.

[24] B. Finkel and R. Tamassia, Curvilinear graph drawing using the force-directed method, in: *Proc. of the 12th International Conference on Graph Drawing, GD'04*, Springer-Verlag, Berlin, Heidelberg, 2004, pp. 448–453, ISBN 3-540-24528-6, 978-3-540-24528-5.

[25] P. Fishwick, S. Diehl, J. Prophet and J. Löwgren, Perspectives on aesthetic computing, *Leonardo* **38**(2) (2005), 133–141, ISSN 0024094X.

[26] R. García, J. Gimeno, F. Perdrix, R. Gil, M. Oliva, J. López, A. Pascual and M. Sendín, Building a usable and accessible semantic web interaction platform, *World Wide Web* **13**(1–2) (2010), 143–167, ISSN 1386-145X.

[27] M. Granitzer, V. Sabol, W. Kienreich, D. Lukose and K. Weng Onn, Visual analyses on Linked Data – An opportunity for both fields, Technical report, STI International, Riga, Latvia, 2010.

[28] T.M. Green, W. Ribarsky and B. Fisher, Building and applying a human cognition model for visual analytics, *Information Visualization* **8**(1) (January 2009), 1–13, ISSN 1473-8716, URL http://dx.doi.org/10.1057/palgrave.ivs.2008.28.

[29] T. Grossman, G. Fitzmaurice and R. Attar, A survey of software learnability: Metrics, methodologies and guidelines, in: *Proc. of the SIGCHI Conference on Human Factors in Computing Systems, CHI '09*, ACM, New York, NY, USA, 2009, pp. 649–658, ISBN 978-1-60558-246-7, URL http://doi.acm.org/10.1145/1518701.1518803.

[30] H. Halpin, D.M. Herzig, P. Mika, R. Blanco, J. Pound, H.S. Thompson and T. Tran, Duc. Evaluating Ad-Hoc Object Retrieval, in: *Proc. of the 1st International Workshop on Evaluation of Semantic Technologies (IWEST)*, 2010.

[31] C.G. Healey, K.S. Booth and J.T. Enns, Harnessing preattentive processes for multivariate data visualization, in: *Proc. of Graphics Interface '93*, Vancouver, BC, Canada, Canada, 1992, University of British, Columbia.

[32] C.G. Healey, K.S. Booth and J.T. Enns, Visualizing real-time multivariate data using preattentive processing, *ACM Trans. Model. Comput. Simul.* **5**(3) (July 1995), 190–221, ISSN 1049-3301.

[33] C.G. Healey, K.S. Booth and J.T. Enns, High-speed visual estimation using preattentive processing, *ACM Trans. Comput.-Hum. Interact.* **3**(2) (June 1996), 107–135, ISSN 1073-0516.

[34] M. Hearst, A. Elliott, J. English, R. Sinha, K. Swearingen and K.-P. Yee, Finding the flow in web site search, *Commun. ACM* **45**(9) (September 2002), 42–49, ISSN 0001-0782, URL http://doi.acm.org/10.1145/567498.567525.

[35] P. Heim, S. Lohmann, D. Tsendragchaa and T. Ertl, Semlens: Visual analysis of semantic data with scatter plots and semantic lenses, in: *Proc. of the 7th International Conference on Semantic Systems (I-SEMANTICS 2011)*, ACM, New York, NY, USA, 2011, pp. 175–178, ISBN 978-1-4503-0621-8.

[36] P. Hekkert, Design aesthetics: Principles of pleasure in design, *Psychology Science* **48**(2) (2006), 157–172, ISSN 1614-9947.

[37] P. Hekkert and H. Leder, *Product Experience, Chapter Product Aesthetics*, Elsevier Science, San Diego, USA, 2008, pp. 259–285, ISBN 9780080450896.

[38] M. Hildebrand, J. van Ossenbruggen and L. Hardman, /facet: A browser for heterogeneous semantic web repositories, in: *Proc. of the 5th International Conference on The Semantic Web, ISWC'06*, Springer-Verlag, Berlin, Heidelberg, 2006, pp. 272–285, ISBN 3-540-49029-9, 978-3-540-49029-6, URL http://dx.doi.org/10.1007/11926078_20.

[39] A. Hogan, A. Harth, J. Umbrich, S. Kinsella, A. Polleres and S. Decker, Searching and browsing linked data with swse: The semantic web search engine, *Web Semantics* **9**(4) (December 2011), 365–401, ISSN 1570-8268, URL http://dx.doi.org/10.1016/j.websem.2011.06.004.

[40] D.F. Huynh and D. Karger, Parallax and companion: Set-based browsing for the data web, 2009, http://davidhuynh.net/media/papers/2009/www2009-parallax.pdf, Online; accessed 28-Oct-2014.

[41] A. Inselberg, *Parallel Coordinates: Visual Multidimensional Geometry and Its Applications*, Springer-Verlag, New York, Inc., Secaucus, NJ, USA, 2009, ISBN 0387215077.

[42] J. Jeng, Usability assessment of academic digital libraries: Effectiveness, efficiency, satisfaction, and learnability, *LIBRI: International Journal of Libraries and Information Services* **55** (2005), 96–121.

[43] P.W. Jordan, Human factors for pleasure in product use, *Applied Ergonomics* **29**(1) (February 1998), 25–33, ISSN 00036870.

[44] E. Kaufmann and A. Bernstein, How useful are natural language interfaces to the semantic web for casual end-users? in: *Proc. of the 6th International the Semantic Web and 2Nd Asian Conference on Asian Semantic Web Conference, ISWC'07/ASWC'07*, Springer-Verlag, Berlin, Heidelberg, 2007, pp. 281–294, 3-540-76297-3, 978-3-540-76297-3, http://dl.acm.org/citation.cfm?id=1785162.1785184.

[45] E. Kaufmann and A. Bernstein, Evaluating the usability of natural language query languages and interfaces to semantic web knowledge bases, *Web Semant.* **8**(4) (November 2010), 377–393, ISSN 1570-8268, URL http://dx.doi.org/10.1016/j.websem.2010.06.001.

[46] D.A. Keim, T. Nietzschmann, N. Schelwies, J. Schneidewind, T. Schreck and H. Ziegler, A spectral visualization system for analyzing financial time series data, in: *EuroVis*, B. Sousa Santos, T. Ertl and K.I. Joy, eds, Eurographics Association, 2006, pp. 195–202, ISBN 3-905673-31-2.

[47] D.A. Keim, F. Mansmann, J. Schneidewind, J. Thomas and H. Ziegler, Visual data mining, in: *Visual Data Mining, Chapter Visual Analytics: Scope and Challenges*, S.J. Simoff, M.H. Böhlen and A. Mazeika, eds, Springer-Verlag, Berlin, Heidelberg, 2008, pp. 76–90, ISBN 978-3-540-71079-0, URL http://dx.doi.org/10.1007/978-3-540-71080-6_6.

[48] N. Kim, A history of design theory in art education, *Journal of Aesthetic Education* **40**(2) (2006), 12–28, ISSN 00218510. URL, http://www.jstor.org/stable/4140227.

[49] R. Kosara, F. Bendix and H. Hauser, Parallel sets: Interactive exploration and visual analysis of categorical data, *IEEE Transactions on Visualization and Computer Graphics* **12**(4) (July 2006), 558–568, ISSN 1077-2626, http://dx.doi.org/10.1109/TVCG.2006.76.

[50] M. Kumar and N. Garg, Aesthetic principles and cognitive emotion appraisals: How much of the beauty lies in the eye of the beholder? *Journal of Consumer Psychology* **20**(4) (October 2010), 485–494, ISSN 10577408.

[51] M. Kurosu and K. Kashimura, Apparent usability vs. inherent usability: Experimental analysis on the determinants of the apparent usability, in: *Conference Companion on Human Factors in Computing Systems, CHI '95*, ACM, New York, NY, USA, 1995, pp. 292–293, ISBN 0-89791-755-3.

[52] A. Lang, Aesthetics in information visualization, Technical report, University of Munich, Department of Computer Science, 2010.

[53] H. Lausen, Y. Ding, M. Stollberg, D. Fensel, R.L. Hernández and S.-K. Han, Semantic web portals: State-of-the-art survey, *J. Knowledge Management* **9**(5) (2005), 40–49.

[54] V. Lopez, E. Motta and V. Uren, Poweraqua: Fishing the semantic web, in: *Proc. of the 3rd European Conference on The Semantic Web: Research and Applications, ESWC'06*, Springer-Verlag, Berlin, Heidelberg, 2006, pp. 393–410, ISBN 3-540-34544-2, 978-3-540-34544-2, URL http://dx.doi.org/10.1007/11762256_30.

[55] A. Malik, R. Maciejewski, B. Maule and D.S. Ebert, A visual analytics process for maritime resource allocation and risk assessment, in: *IEEE VAST*, IEEE, 2011, pp. 221–230, ISBN 978-1-4673-0015-5.

[56] K. Marriott, H. Purchase, M. Wybrow and C. Goncu, Memorability of visual features in network diagrams, *IEEE Transactions on Visualization and Computer Graphics* **18**(12) (2012), 2477–2485, ISSN 1077-2626.

[57] S. Mazumdar, A. Varga, V. Lanfranchi, D. Petrelli and F. Ciravegna, A knowledge dashboard for manufacturing industries, in: *The Semantic Web: ESWC 2011 Workshops*, R. García-Castro, D. Fensel and G. Antoniou, eds, Lecture Notes in Computer Science, Vol. 7117, Springer, Berlin Heidelberg, 2011, pp. 112–124, ISBN 978-3-642-25952-4.

[58] S. Mazumdar, D. Petrelli and F. Ciravegna, Exploring user and system requirements of linked data visualization through a visual dashboard approach, *Semantic Web – Interoperability, Usability, Applicability* **5** (2014), 203–220.

[59] J.D. McClurg-Gevenese, *The Principles of Design. Design in Theory and Practice*, 2005, http://www.digital-web.com/articles/principles_of_design.

[60] D.C.L. Ngo, L.S. Teo and J.G. Byrne, Modelling interface aesthetics, *Information Sciences* **152**(1) (June 2003), 25–46, ISSN 0020-0255, URL http://dx.doi.org/10.1016/S0020-0255(02)00404-8.

[61] D.C.L. Ngo, A. Samsudin and R. Abdullah, Aesthetic measures for assessing graphic screens, *Journal of Information Science and Engineering* **16**(1) (2000), 97–116.

[62] J. Nielsen and T.K. Landauer, A mathematical model of the finding of usability problems, in: *Proc. of the INTERACT '93 and CHI '93 Conference on Human Factors in Computing Systems, CHI '93*, ACM, New York, NY, USA, 1993, pp. 206–213, ISBN 0-89791-575-5.

[63] J. Nielsen and R. Molich, Heuristic evaluation of user interfaces, in: *Proc. of the SIGCHI Conference on Human Factors in Computing Systems, CHI '90*, ACM, New York, NY, USA, 1990, pp. 249–256, ISBN 0-201-50932-6, URL http://doi.acm.org/10.1145/97243.97281.

[64] S. Park, Usability, user preferences, effectiveness, and user behaviors when searching individual and integrated full-text databases: Implications for digital libraries, *J. Am. Soc. Inf. Sci.* **51**(5) (March 2000), 456–468, ISSN 0002-8231, URL http://dx.doi.org/10.1002/(SICI)1097-4571(2000)51:5<456::AID-ASI6>3.3.CO;2-F.

[65] S.-e. Park, D. Choi and J. Kim, Visualizing e-brand personality: Exploratory studies on visual attributes and e-brand personalities in Korea, *International Journal of Human-Computer Interaction* **19**(1) (2005), 7–34.

[66] D. Petrelli, S. Mazumdar, A.-S. Dadzie and F. Ciravegna, Multi visualization and dynamic query for effective exploration of semantic data, in: *Proc. of the 8th International Semantic Web Conference, ISWC '09*, Springer-Verlag, Berlin, Heidelberg, 2009, pp. 505–520, ISBN 978-3-642-04929-3, URL http://dx.doi.org/10.1007/978-3-642-04930-9_32.

[67] H.C. Purchase, J. Hamer, A. Jamieson and O. Ryan, Investigating objective measures of web page aesthetics and usability, in: *Proc. of the Twelfth Australasian User Interface Conference – Volume 117, AUIC '11*, Australian Computer Society, Inc., Darlinghurst, Australia, Australia, 2011, pp. 19–28, ISBN 978-1-920682-97-2, URL, http://dl.acm.org/citation.cfm?id=2460616.2460619.

[68] M. Rico and P. Castells, Jacinta: A mediator agent for human-agent interaction in semantic web services, in: *Proc. of The Semantic Web-International Semantic Web Conference (ISWC)*, 2004.

[69] T.L. Roberts and T.P. Moran, The evaluation of text editors: Methodology and empirical results, *Commun. ACM* **26**(4) (April 1983), 265–283, ISSN 0001-0782, URL http://doi.acm.org/10.1145/2163.2164.

[70] R. Rosenholtz, Y. Li and L. Nakano, Measuring visual clutter, *Journal of Vision* **7**(2) (2007), 1–22.

[71] M. Sarkar and M.H. Brown, Graphical fisheye views of graphs, in: *Proc. of the SIGCHI Conference on Human Factors in Computing Systems, CHI '92*, ACM, New York, NY, USA, 1992, pp. 83–91, ISBN 0-89791-513-5.

[72] M.C. Schraefel, M. Wilson, A. Russell and D.A. Smith, mspace: Improving information access to multimedia domains with multimodal exploratory search, *Commun. ACM* **49**(4) (April 2006), 47–49, ISSN 0001-0782, URL http://doi.acm.org/10.1145/1121949.1121980.

[73] N. Schwarz, P. Winkielman, B. Morton, J. Skowronski and B. Whittlesea, Processing fluency and aesthetic pleasure: Is beauty in the perceiver's processing experience, *Social Psychology Review* **8** (2004), 364–382.

[74] B. Shackel, Ergonomics in design for usability, in: *Proc. of the Second Conference of the British Computer Society, Human Computer Interaction Specialist Group on People and Computers: Designing for Usability*, Cambridge University Press, New York, NY, USA, 1986, pp. 44–64, ISBN 0-521-33259-1, URL, http://dl.acm.org/citation.cfm?id=17324.24077.

[75] B. Shneiderman, The eyes have it: A task by data type taxonomy for information visualizations, in: *VL '96: Proc. of the 1996 IEEE Symposium on Visual Languages*, IEEE Computer Society, 1996.

[76] B. Shneiderman and C. Plaisant, *Designing the User Interface: Strategies for Effective Human-Computer Interaction*, Addison-Wesley Longman Publishing Co., Inc., 1986.

[77] M. Stefaner and B. Muller, Elastic lists for facet browsers, in: *Proc. of the 18th International Conference on Database and Expert Systems Applications, DEXA '07*, IEEE Computer Society, Washington, DC, USA, 2007, pp. 217–221, ISBN 0-7695-2932-1, URL http://dx.doi.org/10.1109/DEXA.2007.63.

[78] N. Tractinsky, A.S. Katz and D. Ikar, What is beautiful is usable, *Interacting with Computers* **13** (December 2000), 127–145, ISSN 09535438.

[79] E. Tufte, *Envisioning Information*, Graphics Press, Cheshire, CT, USA, 1990, ISBN 0-9613921-1-8.

[80] E.R. Tufte, *The Visual Display of Quantitative Information*, Graphics Press, Cheshire, CT, USA, 1986, ISBN 0-9613921-0-X.

[81] T. Tullis and W. Albert, *Measuring the User Experience: Collecting, Analyzing, and Presenting Usability Metrics*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2008, ISBN 0123735580, 9780123735584.

[82] G. Tummarello, R. Cyganiak, M. Catasta, S. Danielczyk, R. Delbru and S. Decker, Sig.ma: Live views on the web of data, *Web Semantics: Science, Services and Agents on the World Wide Web* **8**(4) (2010), 355–364, ISSN 1570-8268.

[83] J. Whiteside, S. Jones, P.S. Levy and D. Wixon, User performance with command, menu, and iconic interfaces, *SIGCHI Bull.* **16**(4) (April 1985), 185–191, ISSN 0736-6906, URL, http://doi.acm.org/10.1145/1165385.317490.

[84] N. Willems, W.R. van Hage, G. de Vries, J.H.M. Janssens and V. Malaise, An integrated approach for visual analysis of a multisource moving objects knowledge base, *Int. J. Geogr. Inf. Sci.* **24**(10) (October 2010), 1543–1558, ISSN 1365-8816.

[85] P.C. Wong, B. Hetzler, C. Posse, M. Whiting, S. Havre, N. Cramer, A. Shah, M. Singhal, A. Turner and J. Thomas, Inspire infovis 2004 contest entry, in: *Proc. of the IEEE Sympo-sium on Information Visualization, INFOVIS '04*, IEEE Computer Society, Washington, DC, USA, 2004, p. 216.2, ISBN 0-7803-8779-3, URL http://dx.doi.org/10.1109/INFOVIS.2004.37.

[86] P.C. Wong, G. Chin Jr., H. Foote, P. Mackey and J. Thomas, Have green – A visual analytics framework for large semantic graphs, in: *IEEE VAST*, P.C. Wong and D.A. Keim, eds, IEEE, 2006, pp. 67–74, ISBN 1-4244-0591-2.

[87] F. Yamaguchi, *Curves and Surfaces in Computer Aided Geometric Design*, 1st edn, Springer-Verlag, Berlin, Nov 1988, ISBN 978-3540174493.

[88] X. Yang, S. Asur, S. Parthasarathy and S. Mehta, A visual-analytic toolkit for dynamic interaction graphs, in: *Proc. of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '08*, ACM, New York, NY, USA, 2008, pp. 1016–1024, ISBN 978-1-60558-193-4.

[89] J.M. Zain, M. Tey and Y. Goh, Does aesthetics of web page interface matters to mandarin learning? *CoRR*, abs/1101.1608 (2011).

[90] J.M. Zain, M. Tey and Y. Goh, Probing a self-developed aesthetics measurement application (sda) in measuring aesthetics of mandarin learning web page interfaces, *CoRR*, abs/1101.1606 (2011).