

University of Groningen

Adapting Graph Visualization Techniques for the Visualization of RDF Data

Frasincar, Flavius; Telea, Alexandru; Houben, Geert-Jan

Published in:
EPRINTS-BOOK-TITLE

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2006

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Frasincar, F., Telea, A., & Houben, G.-J. (2006). Adapting Graph Visualization Techniques for the Visualization of RDF Data. In *EPRINTS-BOOK-TITLE* University of Groningen, Johann Bernoulli Institute for Mathematics and Computer Science.

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

9.1 Introduction

The foundation language for the Semantic Web is the Resource Description Framework (RDF). RDF is intended to describe the Web metadata so that the Web content is not only machine readable but also machine understandable. In this way one can better support the interoperability of Web applications. RDF Schema (RDFS) is used to describe different RDF vocabularies (schemas), that is, the classes and properties associated to a particular application domain. An instantiation of these classes and properties form an RDF instance. It is important to note that both an RDF schema and an RDF instance have RDF graph representations.

Realizing the advantages that RDF offers, in the last couple of years, many tools were built in order to support the browsing and editing of RDF data. Among these tools we mention Protégé (Noy et al., 2001), OntoEdit (Sure et al., 2003), and RDF Instance Creator (RIC) (Grove, 2002). Most of the text-based environments are unable to cope with large amounts of data in the sense of presenting them in a way that is easy to understand and navigate (Card et al., 1999). The RDF data we have to deal with describes a large number of Web resources, and can thus easily reach tens of thousands of instances and attributes. We advocate the use of visual tools for browsing RDF data, as visual presentation and navigation enables users to effectively understand the complex structure and interrelationships of such data. Existing visualization tools for RDF data are: IsaViz (Pietriga, 2002), OntoRAMA (Eklund et al., 2002), and the Protégé visualization plug-ins like OntoViz (Sintek, 2004) and Jambalaya (Storey et al., 2001).

The most popular textual RDF browser/editor is Protégé (Noy et al., 2001). The generic modeling primitives of Protégé enable the export of the built model in different data formats, among which is also RDF/XML. Protégé distinguishes between schema and instance information, allowing for an incremental view of the instances based on the selected schema elements. One of the disadvantages of Protégé is that it displays the information in a hierarchical way, that is, using a tree layout (Sugiyama et al., 1981), which makes it difficult to grasp the inherent graph structure of RDF data.

In this chapter, we advocate the use of a highly customizable, interactive visualization system for the understanding of different RDF data structures. We implemented an RDF data format plug-in for GViz (Telea et al., 2002), a general-purpose visual environment for browsing and editing graph data. The largest advantage that GViz provides in comparison with other RDF visualization tools is the fact that it is easily

and fully customizable. GViz was architected with the specific goal in mind of allowing users to define new operations for data processing, visualization, and interaction to support application-specific scenarios. GViz also integrates a number of standard operations for manipulation and visualization of relational data, such as data viewers, graph layout tools, and data format support. This combination of features has enabled us to produce, in a short time, customized visualization scenarios for answering several questions about RDF data. We demonstrate our approach to RDF data visualization by using a real dataset example of considerable size.

In the next section, we describe the real-world dataset we use, and show the results obtained when visualizing it with several existing RDF tools. Our visualization tool, GViz, is presented in Section 9.3. Section 9.4 presents several visualization scenarios we built with GViz for the RDF dataset used, and details of various lessons learned when building and using such visualizations. Finally, Section 9.5 concludes the chapter, proposing future directions for visualizing RDF information.

9.2 Background

Throughout this chapter, we will use an example based on real data made available by the Rijksmuseum in Amsterdam, the largest art and history museum in the Netherlands. In the example there is a museum schema used to classify different artists and their artefacts. The museum instance describes more than 1,000 artists and artefacts. For comparison purposes, we chose to represent the same museum RDFS schema in several browsing tools.

Figure 9.1 depicts the museum schema in Protégé. As can be noticed from this figure such a text-based representation cannot nicely depict the structure of a large amount

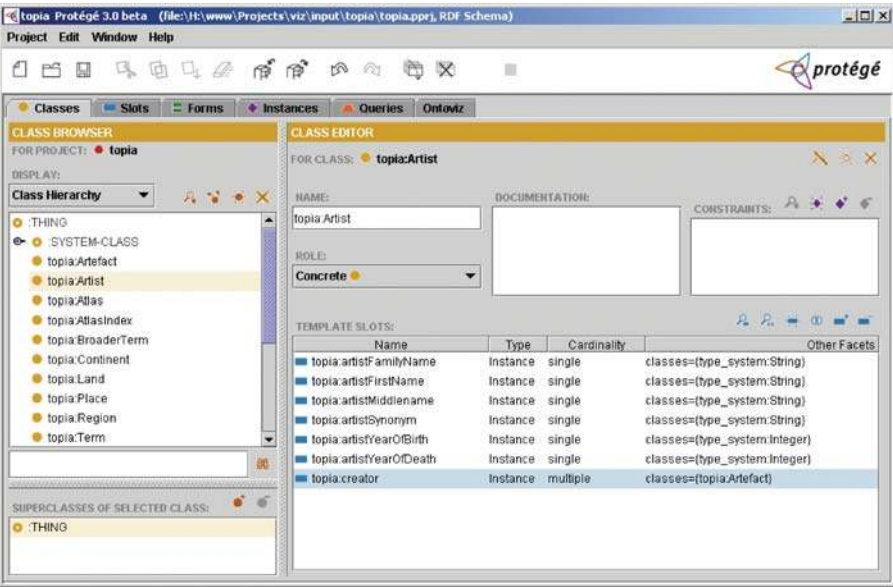


Figure 9.1 Museum schema in Protégé (text-based).

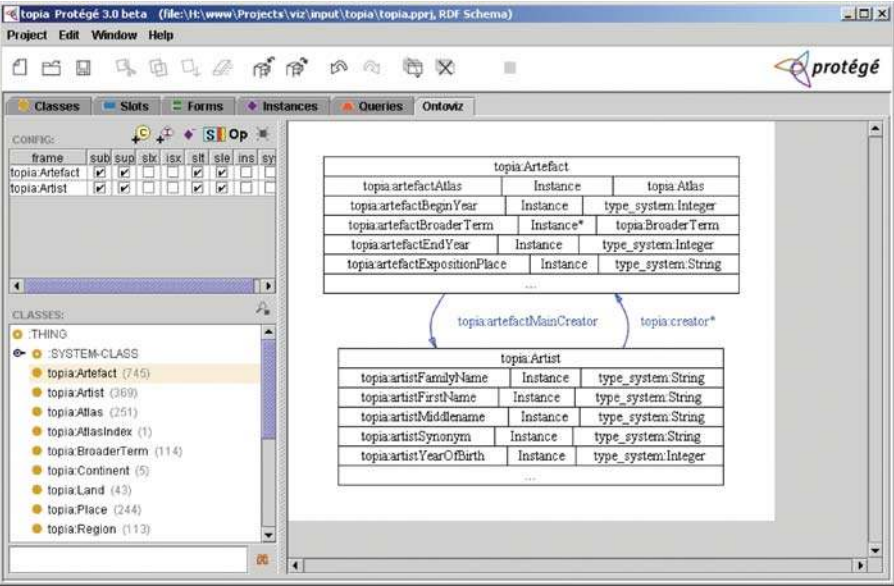


Figure 9.2 Museum schema in Protégé (with OntoViz plug-in).

of data. More exactly, a text-based display is very effective for *data mining*, that is, posing targeted queries on a dataset once one knows what structure one is looking for. However, text-based displays are not effective for *data understanding*, that is, making sense of a given (large) dataset of which the global structure is unknown to the user.

In order to alleviate the above limitation, Protégé offers a number of built-in visualization plug-ins. Figure 9.2 shows the graph representation generated by the OntoViz plug-in for two classes from the museum schema. The weak point of OntoViz is the fact that it is not able to produce good (understandable) layouts for graphs that have more than 10 nodes.

IsaViz (Pietriga, 2002) is a visual tool for browsing/editing RDF models. IsaViz uses AT&T's GraphViz package (North and Koutsosios, 1996) for the graph layout.

Figure 9.3 shows the same museum schema using IsaViz. The layout produced by the tool is much better than the one generated with OntoViz. However, the directed acyclic graph layout used (Sugiyama et al., 1981) becomes ineffective when the dataset at hand has roughly more than a hundred nodes, as can be seen from Figure 9.3. IsaViz has a 2.5D GUI with zooming capabilities and provides numerous operations like text-based search, copy-and-paste, editing of the geometry of nodes and arcs, textual attribute browsing, and graph navigation.

For all these reasons, we believe that IsaViz is a state-of-the-art tool for browsing/editing RDF models. However, its rigid architecture makes it difficult to define application-dependent operations other than the standard ones currently provided by the tool. Experience in several communities interested in visualizing relational data in general, such as software engineering and Web engineering, and our own experience with RDF data in particular, has shown that tool customization is extremely important. Indeed, there is no “silver bullet” or best way to visualize large graph-like datasets. The questions to be answered, the data structure and size, and the user

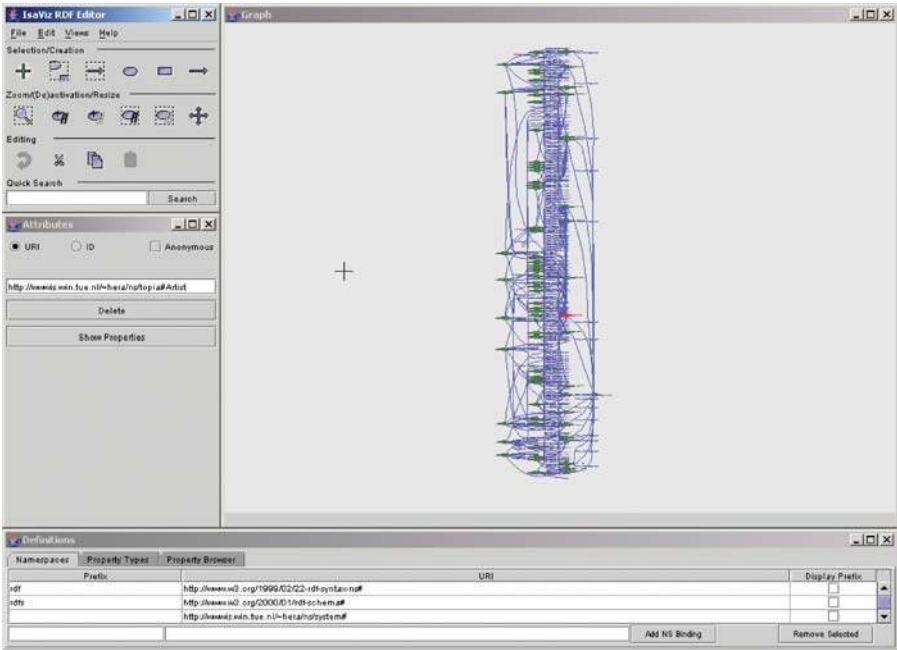


Figure 9.3 Museum schema in IsaViz.

preferences all determine the “visualization scenario,” that is, the kind of (interactive) operations the users may want to perform to get insight into the data and answers to their questions. It is not that each separate application domain demands a specific visualization scenario. Users of the same domain and/or even the same dataset within the same domain may require different scenarios. Building such scenarios often is responsible for a large part of the complete time spent in understanding a given dataset (Telea, 2004). This clearly requires the visualization tool in use to be highly (and easily) customizable.

9.3 GViz

In our attempt to understand RDF data through visual representations, an existing tool was used. We implemented an RDF data format plug-in for GViz (Telea et al., 2002), a general-purpose visual environment for browsing and editing graph data. The largest advantage that GViz provides in comparison with other RDF visualization tools is the fact that it is easily and quickly customizable. One can seamlessly define new operations to support application-specific scenarios, making the tool more amenable for the user needs. In the past, GViz was successfully used in the reverse engineering domain, in order to define application-specific visualization scenarios.

Figure 9.4 presents the architecture of GViz based on four components: selection, mapping, editing, and visualization. In the next section we describe the data model used in GViz. Next, we outline the operation model describing the tasks that can be defined on the graph data. We finish the description of the GViz architecture

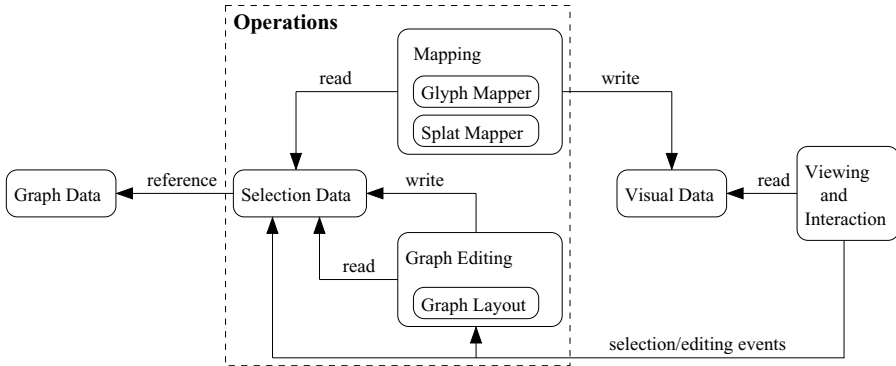


Figure 9.4 GViz architecture.

with the visualization component, which we illustrate using the museum schema dataset.

The GViz core implementation is done in C++ while the user interface and scripting layer were implemented in Tcl (Raines, 1998) to take advantage of the runtime scripting and weak typing flexibility that this language provides. All the GViz customization code developed for the RDF visualization scenarios presented in this chapter was done in Tcl.

9.3.1 Data Model

The data model of GViz consists of three elements:

1. *Graph data.* This is the RDF graph model, that is, a labeled directed multigraph in which no edges between the same two nodes are allowed to share the same label. Nodes stand for RDF resources/literals and edges denote properties. Each node has a `type` attribute that states if the node is a `NResource` (named resource), an `AResource` (anonymous resource), or a `Literal`. The label associated to a node/edge is given by the `value` attribute. The labels for `NResource` nodes and edges are URIs. The label for `Literals` is their associated string. The value of an `AResource` is an internal identifier with no RDF semantics. Note that the `type` and `value` attributes are GViz-specific attributes that should not be confused with their RDF counterparts. Since GViz's standard data model is an arbitrary attributed graph, with any number of (name, value) type of attributes per node and edge, the RDF data model is directly accommodated by the tool.
2. *Selection data.* Selections are subsets of nodes and edges in the graph data. Selections are used in GViz to specify the inputs and outputs of its operations; their use is detailed in Section 9.3.2.
3. *Visual data.* This is the information that GViz ultimately displays and allows the user to interact with. Since GViz allows customizing the mapping operation (i.e., the way graph data are used to produce visual data), the latter may assume different look-and-feel appearances. Section 9.4 illustrates this in the context of our application.

9.3.2 Operation Model

As shown in Figure 9.4, the operation model of GViz has three operation types: selection, graph editing, and mapping. Selection operations allow users to specify subsets of interest from the whole input graph. In the RDF visualization scenarios that we built with GViz, we defined different complex selections based on the attributes of the input model. These selections can perform tasks like: “extract the schema from an input set of RDF(S) data (which mixes schema and instance elements).” Custom selections are almost always needed when visualizing relational data, since (1) the user doesn’t usually want to look at too many data elements at the same time, and (2) different subsets of the input data may have different semantics, and thus have to be visualized in different ways. A basic example of the latter assertion is the schema extraction selection mentioned earlier.

Graph editing operations enable the modification, creation, and deletion of nodes/edges and/or their attributes. For our RDF visualization scenarios, we did not create or delete nodes or edges. However, we did create new data attributes, as follows. One of the key features of GViz is that it separates the graph layout, that is, computing 2D or 3D geometrical positions that specify *where* to draw nodes and edges, from the graph mapping, that is, specifying *how* to draw nodes and edges. The graph layout is defined as a graph editing operation that computes position attributes. Among the different layouts that GViz supports we mention the spring embedder, the directed tree, the 3D stacked layout, and the nested layout (Telea et al., 2002). Although based on the same GraphViz package as IsaViz, the layouts of GViz are relatively more effective, as the user can customize their behavior in detail via several parameters.

Mapping operations, or briefly mappers, associate nodes/edges (containing also their layout information) to visual data. The latter is implemented using the Open Inventor 3D toolkit, which delivers high-quality, efficient rendering and interaction with large 2D and 3D geometric datasets (Werneck, 1993). GViz implements two mappers: the glyph mapper and the splat mapper. The glyph mapper associates to every node/edge in the input selection a graphical icon (the glyph) and positions the glyphs based on the corresponding node/edge layout attributes. Essentially, the glyph mapper produces the “classical” kind of graph drawings (e.g., similar to those output by IsaViz). However, in contrast to many graph visualization tools, the glyph mapper in GViz allows full customization of the way the nodes and edges are drawn. The user can specify, for example, shapes, sizes, and colors for every separate node and edge, if desired, by writing a small Tcl script of 10 to 20 lines of code on the average. We used this feature extensively to produce our RDF visualizations described in Section 9.4. The splat mapper produces a continuous two-dimensional splat field for the input selection. For every 2D point, the field value is proportional to the density of nodes per unit area at that point. Essentially, the splat mapper shows high values where the graph layout used has placed many nodes, and low values where there are few nodes. Given that a reasonably good layout will cluster highly interconnected nodes together, the splat mapper offers a quick and easy way to visually find the clusters in the input graph (see Figure 9.9, Section 9.4). For more details on this layout, see Van Liere and De Leeuw (2003).

A final way to customize the visualizations in GViz is to associate custom *interaction* to the mappers. These are provided in the form of Tcl callback scripts that are called by the tool whenever the user interactively selects some node or edge glyph with the mouse, in the respective mapper windows. These scripts can initiate any desired operation using the selected elements as arguments, for example, showing some

attributes of the selected arguments. Examples of this mechanism are discussed in Section 9.4.

As explained above, GViz allows users to easily define new operations. For the incremental view of RDF(S) data, we defined operations as: extract schema, select classes and their corresponding instances, select instances and their attributes. As for the glyph mappers, these operations have been implemented as Tcl scripts of 10 to 25 lines of code. The usage of the custom selection, layout, and mapping operations for visualizing RDF(S) data is detailed in the remainder of this chapter.

9.3.3 Visualization

Figure 9.5 presents the museum data schema in GViz. We use here a radial tree layout, also available in the GraphViz package, instead of the directed tree layout illustrated

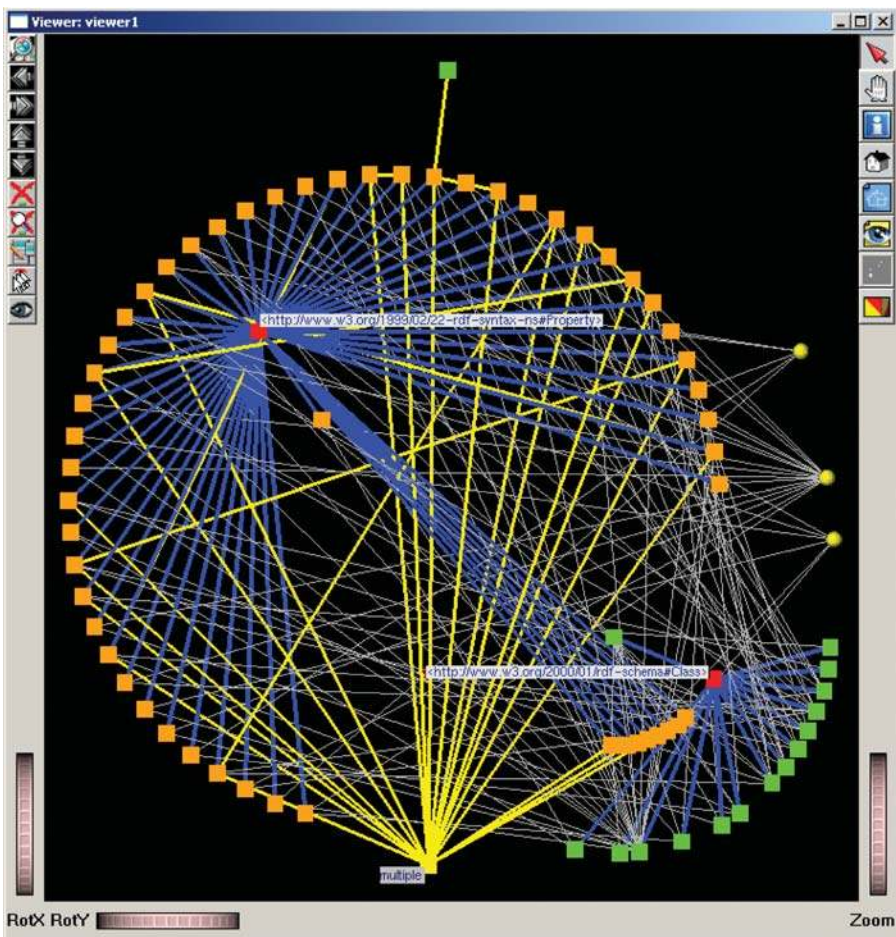


Figure 9.5 Museum schema in GViz (2D).

in Figure 9.3 for IsaViz. As a consequence, the structure of the schema is easier to understand now.

In the picture in Figure 9.5 the edges with the label `rdf:type` are depicted in blue. There are two red nodes to which these blue edges connect, one with the label `rdfs:Class` and the other with the label `rdf:Property`, shown near the nodes as balloon pop-up texts. We chose to depict the property nodes (laid out in a large circular arc around the upper-left red node) in orange and the class nodes (laid out in a smaller circle arc around the lower-right red node) in green. As can be noticed from the picture there are a lot of orange nodes, which is in accordance with the property-centric approach for defining RDFS schemas. In order to express richer domain models we extended the RDFS primitives with the cardinality of properties, the inverse of properties, and a media type system. These extensions are shown in yellow edges (see also below) and yellow spheres (positioned at the right end of the image). The yellow edges that connect to orange nodes represent the inverse of a property. The yellow edges that connect an orange node with the yellow rectangle labeled “multiple” (positioned at the middle of the figure bottom) state that this property has cardinality one-to-many. The default cardinality is one-to-one. Note that there are no many-to-many properties as we had previously decomposed these properties in two one-to-many properties. The three yellow spheres represent the media types: `String`, `Integer`, and `Image`. The light-gray thin edges denote the domain and the range of properties. Note that only range edges can have a media node at one of its ends. As these edges are (1) not so important for the user and (2) quite numerous and quite hard to lay out without many overlaps, we chose to represent them in a visually inconspicuous way (i.e., make them thin and use a background-like, light-gray color).

The tailoring of the graph visualization presented above is only one example. One can define some other visualizations depending on one’s needs. Figure 9.6 presents a 3D view of the same museum schema example. Here, we used a spring embedder layout, also available from the GraphViz package, to position all schema nodes in a 2D plane. Next, we designed a custom operation that selects the two `rdfs:Class` and `rdf:Property` nodes and offsets them away from the 2D layout plane, in opposite directions. This creates a 3D layout, which allows the user to better distinguish the different kinds of edges. For example, the edges labeled `rdf:type` (colored in blue) are now clearly separated, as they reach out of the 2D plane to the offset nodes.

9.4 Applications

In order to better understand the context in which we developed our visualization applications we now briefly describe the Hera project (Vdovjak et al., 2003). Hera is a methodology for designing and developing Web Information Systems (WISs) on the Semantic Web. All the system specifications are represented in RDFS. For the scope of this chapter it is important to look at two of these specifications: the conceptual model (domain model) and the application model (navigation model).

The conceptual model describes the types of the instances that need to be presented. An example of the conceptual model we already saw in Figure 9.5. A conceptual model is composed of concepts and concept properties. There are two kinds of concept

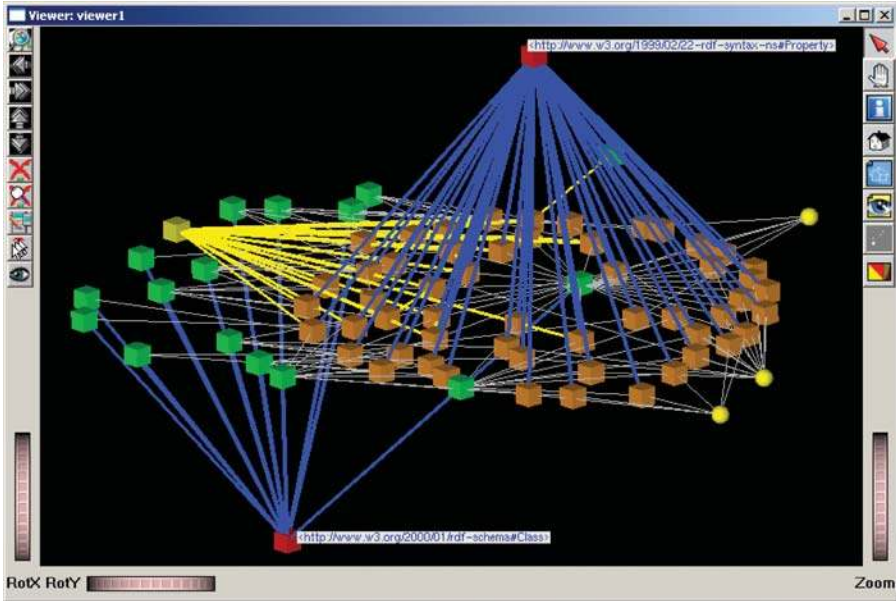


Figure 9.6 Museum schema in GViz (3D).

properties: relationships (properties between concepts) and attributes (properties that refer to media types).

The application model defines the navigation over the data, that is, a view on the conceptual model that contains appropriate navigation links. The application model is an overlay model over the conceptual model, a feature exploited in the definition of the transformations of the conceptual model instances into application model instances. An application model is composed of slices and slice properties. A slice contains concept attributes (not necessarily from the same concept) as well as other slices. There are two kinds of slice properties: compositional properties (aggregations) and navigational properties (links). The owner relationship is used to associate a slice to a concept. Each slice has a title attribute related to it.

A conceptual model instance and an application model instance are represented in RDF (which should be valid according to the corresponding RDFS specifications, i.e., the conceptual model and the application model, respectively). In the WIS application it is only the application model instance that will be visible to the user. We consider four types of RDF(S)-related visualization scenarios that are relevant in the support of the WIS application designer:

1. Conceptual model visualization
2. Conceptual model instance visualization
3. Application model visualization
4. Application model instance visualization

In Section 9.3.3 we already showed how one can visualize conceptual models. A second similar scenario for the conceptual model visualization is described next.

9.4.1 Conceptual Model Visualization

The conceptual model visualization enables one to better understand the structure of the application's domain. It answers questions such as:

What are the concepts?

What are the properties?

What are the relationships between concepts and properties?

What are the most referenced concepts?

What are the most referenced media types?

Figure 9.7 shows the extracted schema from an RDF file that contains both the schema and its associated instance. The extraction is done by a custom selection operation, as described in Section 9.3.2. The picture is very similar to the one from Figure 9.5. However, there are two differences between this picture and the one from Figure 9.5. First, we now use a different layout, that is, a spring embedder instead

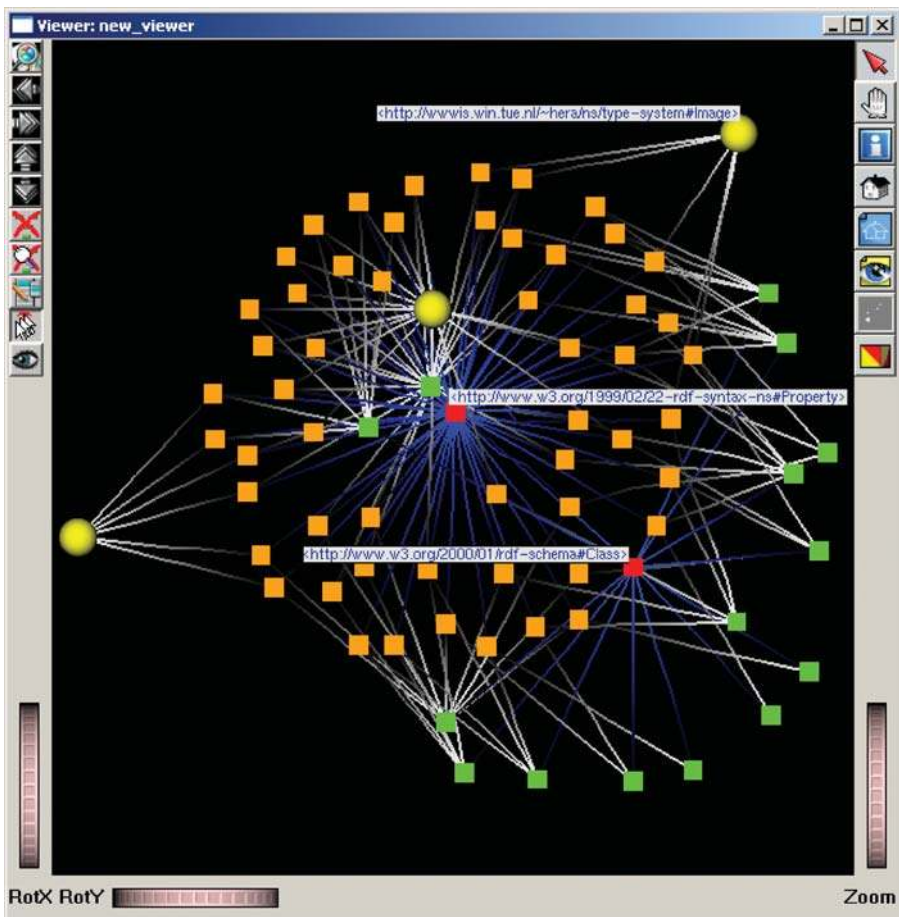


Figure 9.7 Museum extracted conceptual model.

of a radial tree. Second, we now depict also the *direction* of the edges. The edges are fading out toward the start node. A direction effect is created: the edges get brighter as they approach the end node. We found this representation of the edge direction much more effective than the arrow representation when visualizing large graphs, as the drawing of arrows produces too much visual clutter in this case. Moreover, the edge fading glyph is faster to render than an arrow glyph, as it involves a single (shaded) line primitive.

From Figure 9.7 we can deduce that the most used media type is `String` (the text-based descriptions are the most popular for this domain specification) and the most referenced concept is the `Artefact` (it has the most relationships). Each artefact is classified by some museum terms (e.g., `Self Portraits`). There is a hierarchy of museum terms; terms are grouped in broader terms (e.g., `Portraits`), and broader terms are grouped in top terms (e.g., `Paintings`).

9.4.2 Conceptual Model Instance Visualization

The conceptual model instance visualization answers questions such as:

What are the instances of a certain concept?

What are the relations between two selected concept instances?

What are the most referenced instances?

What are the attributes of a selected instance?

In most of the encountered situations, there are (much) more concept instances than concepts. For example, our museum dataset contains tens of thousands of instances. It is easy to imagine other applications where this number goes up to hundreds of thousands, or even more. Drawing *all* these instances simultaneously is neither efficient nor effective. Indeed, no graph layout we were able to test could produce an *understandable* image of an arbitrary, relatively tightly connected graph with tens of thousands of nodes in a *reasonable* amount of time (e.g., tens of seconds). In order to keep the instance visualization manageable, we decided for an incremental view scenario on the RDF(S) data. First, the user selects the subpart of the schema for which he wants the corresponding instances to be visualized. Next, we use a custom interaction script (Section 9.3.2) of about 10 lines of code to separately visualize the instances of the selected items. For example, when the user selects the `Artist` and `Artefact` concepts from Figure 9.7, the GViz tool automatically shows the instances of these concepts and their relations in another window, using a spring embedder layout (Figure 9.8). In Figure 9.8 we used a custom glyph mapper to depict the artefacts with blue rectangles and the artists with green rectangles. The relations between these instances are represented by fading white edges. One can note that there are more artefacts than artists, as expected.

Figure 9.9 shows the same selected data (artists and artefacts) but using a splat mapper instead of the classical glyph mapper. The scalar density function (splat field) is constructed as outlined in Section 9.3.2. We visualize the splat field using a blue-to-red colormap that associates cold hues (blue) to low values and warm hues (yellow, red) to medium and high values. Figure 9.9 (left) shows the splat field as seen from above. Figure 9.9 (right) shows the same splat field, this time visualized using an elevation plot that shows high-density areas also by offsetting these points in the Z (vertical) direction. A red/yellow color (Figure 9.9 left and right) or a high elevation point (Figure 9.9 right) indicates that there are a lot of relations for a particular instance

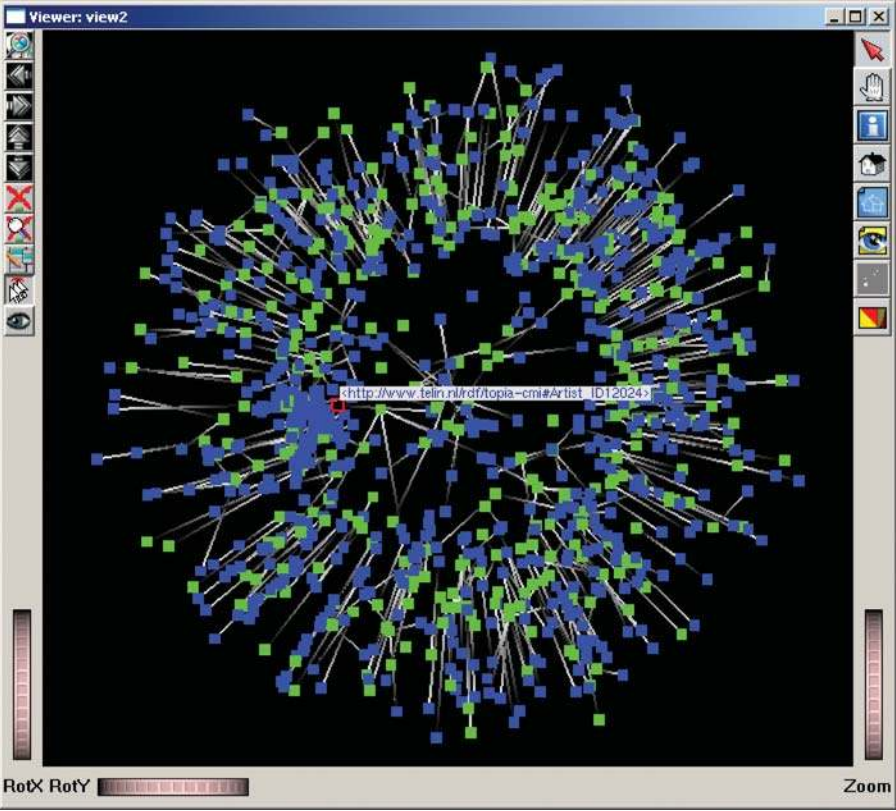


Figure 9.8 Artists/artefacts properties in the conceptual model instance.

or group of instances. In this way one can notice from Figure 9.9 which are the artists with the most artefacts. The artists with the most artefacts are the unknown artists (potter, goldsmith, bronzesmith, etc.) that show up as the singular peak in the left of Figure 9.9 (left). On the average, these artists have several tens (up to 60) of artefacts. They are followed by Rembrandt and the unknown painters, who show up as the other two higher peaks to the right of Figure 9.9 (left). This can be explained by the fact that in the seventeenth century, for which the Rijksmuseum has a special focus, there were a lot of artefacts done by unknown artists.

We have further customized our visualization scenario as follows. When the user selects one instance of Figure 9.8, we use a custom interaction script on the mapper of Figure 9.8 to pop up another window to display the instance attributes.

The selected instance is shown as having the balloon pop-up label in Figure 9.8. Figure 9.10 shows the attributes of the selected instance, in this case Rembrandt: the painter's year of birth, year of death, first name, etc.

9.4.3 Application Model Visualization

The application model visualization enables one to better understand the navigation structure of a hypermedia presentation. It answers questions such as:

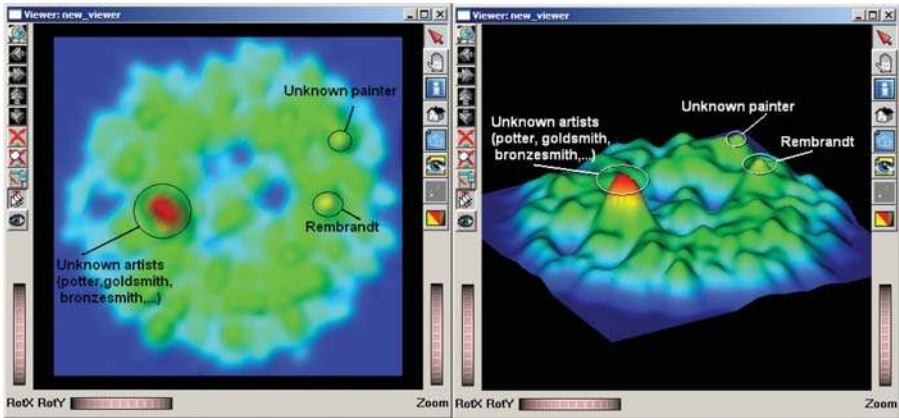


Figure 9.9 Conceptual model instance splatting (left: 2D; right: elevation plot).

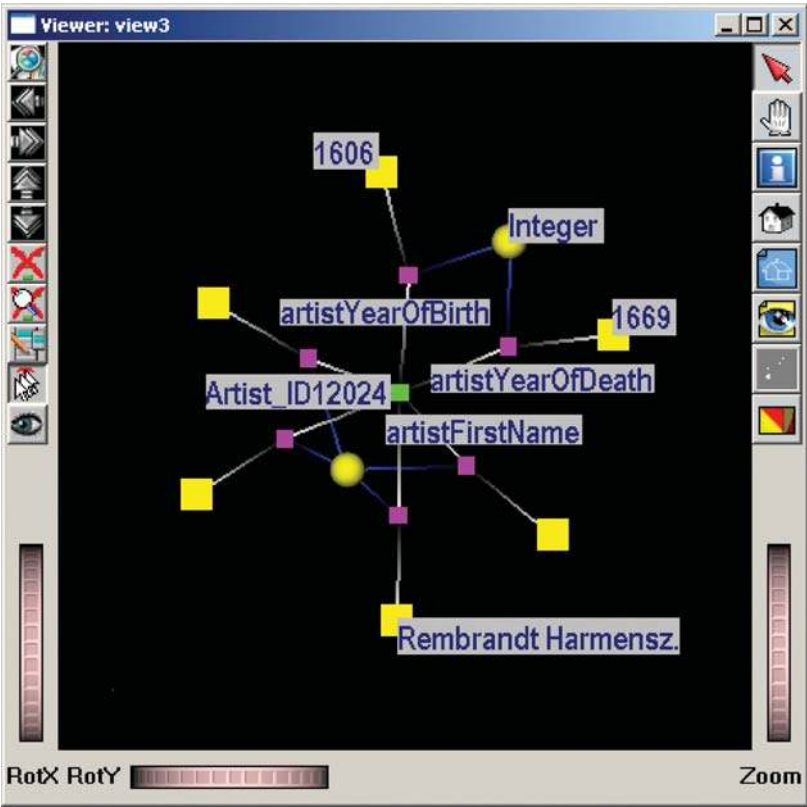


Figure 9.10 Attributes of a selected concept instance.

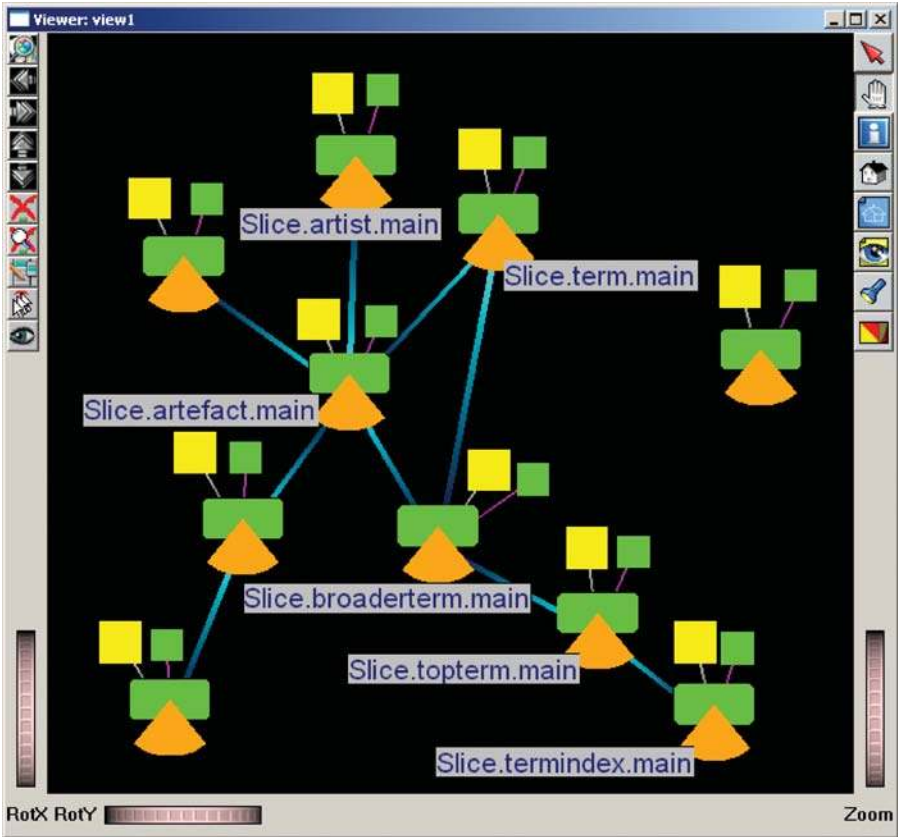


Figure 9.11 Museum application model.

- What are the application model slices?
- What are their links?
- What are the slice owners?
- What are the slice titles?
- What slices are navigation hubs?
- What are possible navigation paths from a certain slice?

Figure 9.11 depicts the application model for the museum example. We chose to present here the top-level slices (slices that correspond to Web pages) and the links between them in order to decrease the complexity of the picture. A new glyph shape was designed in order to represent the pizza-slice shape for slices (as defined in the application model’s graphical representation language). The blue thick edges represent links between slices. Each slice has associated with it two attributes. We use a custom layout to place these nodes right above the top of the slice node. The slice nodes themselves are laid out using the spring embedder already discussed before. The two attributes of each slice are visualized by using two custom square glyphs, as follows: the yellow glyph (left) stands for the name of the slice and the green glyph (right) denotes the concept owner of the slice (remember that the concept owner is a concept from

the conceptual model). In the center of the picture is the `Slice.artefact.main` slice, which has the most links associated with it (i.e., it is a navigation hub). The figure also shows the designer's choice to present the museum information based on the terms hierarchy: top terms, broader terms, and terms.

9.4.4 Application Model Instance Visualization

The application model instance visualization answers questions like:

What are the instances of a certain slice?

What are the slice instances reachable from a certain slice instance?

What are the most referenced slice instances?

What are the attributes of a selected slice instance?

As there are more slice instances than slices, in order to keep the visualization manageable we used the same visualization scenario as for conceptual model instances (i.e., to use incremental views). The user can select from the mapper in Figure 9.11 the slices for which she wants the corresponding instances to be visualized. For example, after selecting the `Slice.topterm.main`, `Slice.broaderterm.main`, and `Slice.term.main` slices from Figure 9.11, we use the same mechanism of a custom interaction script (Section 9.3.2) to pop up another window that shows the instances of these slices and their associated links. Figures 9.12 and 9.13 show the corresponding slice instances, as described below.

For the visualizations in Figures 9.12 and 9.13, we use yellow sphere glyphs for nodes labeled `Slice.topterm.main`, green sphere glyphs for nodes labeled `Slice.broaderterm.main`, and blue rectangle glyphs for nodes labeled `Slice.term.main`. The chosen colors and shapes are motivated by the need to produce an expressive, easy-to-understand picture when presenting a large number of instances coming from three slices linked in a hierarchical way, as follows. We did give up the pizza-slice glyph for these visualizations as we found out that this glyph produces too much visual clutter for large graphs. Next, we chose colors of increasing brightness (blue, green, and yellow) to display items of increasing importance (terms, broader terms, and top terms, respectively). The size of the glyphs used for these items also reflects their importance (the top term glyphs are the largest, whereas the term glyphs are the smallest). A final significant cue is the shape: the more important top and broader terms are drawn as *3D shaded* spheres, whereas the less important terms are drawn as *2D flat* squares. For the edges connecting these glyphs in the visualization, we used a varying color and size scheme that varies both line color and line thickness along the edge between the end nodes' colors and sizes respectively. Summing up, the combination of above choices produces a visualization where the overall structure of top terms and broader terms “pops” into the foreground, whereas the less important terms and their links “fade” into the background. As a comparison, we were unable to get the same clear view of the structure by just varying the layout parameters and using the same glyph for all nodes.

After selecting the slice instance corresponding to the `Paintings` top term, we obtain in Figure 9.12 the broader term slice instances accessible *after one step*, shown in red. By this, we mean the terms that a user of the Web site (whose design our dataset captures) can access after one navigation step. This translates to nodes that are directly connected (via an edge) to the selected slice instance in our RDF dataset. In Figure 9.13 we visualize the term slice instances accessible from the same `Paintings` top




Figure 9.12 Broader term slice instances accessible from the *Paintings* slice instance (one step).

term instance slice *after two steps*, also drawn in red. These correspond to Web pages that the user of the Web site can access after two navigation steps. An example for the second step is the navigation from the broader term *Portraits*. Using such a visualization scenario one can view which are the slice instances reachable from a selected slice instance after a certain number of navigation steps.

9.5 Future Work

In the future, we would like to explore the GViz 3D visualization capabilities for RDF data, possibly getting an even better insight into the data structure. Another research direction would be to use GViz in conjunction with a popular RDF query language (like RQL, for example). Our purpose is here twofold: to use the RDF query language as a selection operation implementation for GViz when visualizing RDF data and to support the RDF query language with the visualization of the input and resulting set of RDF data. Finally, as it is planned in the Hera project to use OWL instead of RDF for the future input data/design specifications we would like also to conduct visualization experiments on the more semantically rich OWL data.

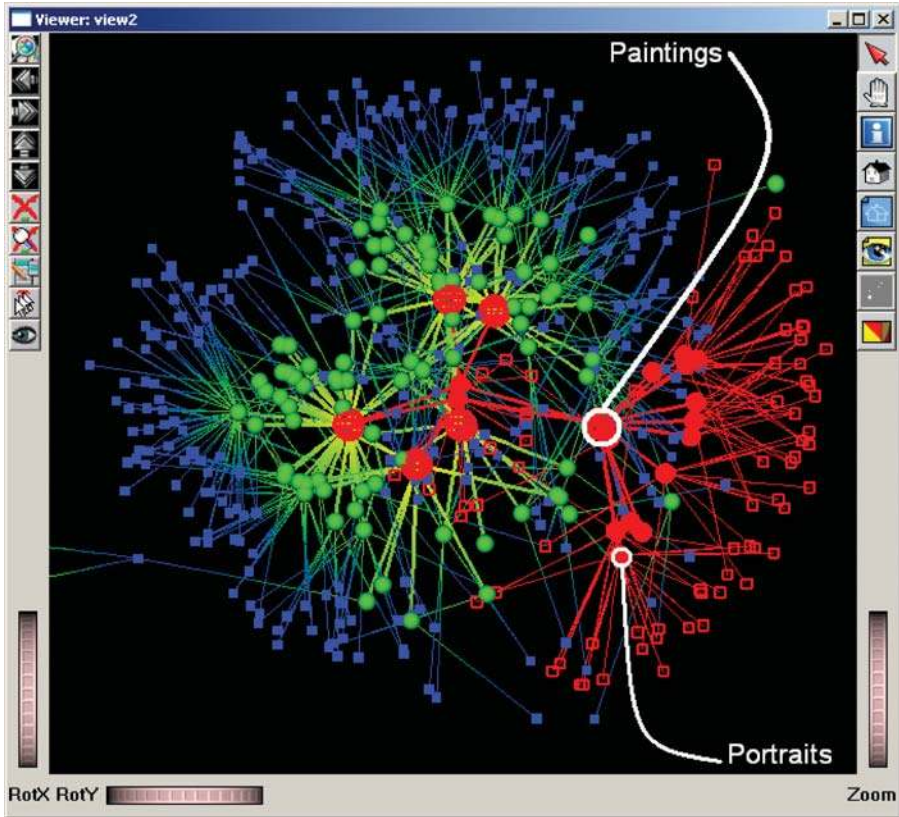


Figure 9.13 Term slice instances accessible from the Paintings slice instance (two steps).

9.6 Summary

In this chapter we have shown how a general-purpose graph visualization tool, GViz, can be used for the visualization of large RDF graphs produced from real-world data. All experiments were performed in the context of the Hera project, a project that investigates the designing and developing of Web Information Systems on the Semantic Web. The visualization of large amounts of RDF input data and RDF design specifications enabled us to answer complex questions about these data and to give an effective insight into their structure.

Several ingredients were crucial for obtaining these results. First, the amount of customizability of the GViz tool (layouts, selections, node and edge drawing, choice between glyph and splat mappers, and custom user interaction) was absolutely necessary to produce the desired visualization scenarios. We found all these elements to be necessary to create the desired results. We have actually experimented with customizing just the layout but not the glyphs and/or the interaction. In all cases, the results were not flexible enough to give the users the desired look-and-feel that would make the scenario effective for answering the relevant questions. Second, the script-based customization mechanism of GViz allowed a user experienced with Tcl scripting to

produce the scenarios described here (which were imagined by a different user, inexperienced with Tcl) in a matter of minutes. Third, we found that using several visual cues (shape, color, size, and shading) together to enhance a single attribute, as for example described in Section 9.4.4, is much more effective than using a single cue. Finally, we mention that none of the investigated RDF visualization tools (Section 9.2) showed the high degree of customization of GViz needed for our scenarios.

9.7 Acknowledgments

The authors would like to thank the Rijksmuseum in Amsterdam for their kind permission to use copyrighted data in our running example. We also wish to thank our colleagues Lloyd Rutledge and Lynda Hardman, from CWI, Amsterdam, for helpful discussions.

9.8 References

- Card, S., Mackinlay, J., Shneiderman, B. (1999). *Readings in Information Visualization*. Morgan Kaufmann.
- Eklund, P.W., Roberts, N., Green, S.P. (2002). OntoRama: Browsing RDF ontologies using a hyperbolic-style browser. *The First International Symposium on CyberWorlds (CW2002), Theory and Practices*, IEEE CS Press, pp. 405–411.
- Grove, M. (2002). *RDF Instance Creator (RIC) for the MIND-SWAP Project*. Available: <http://www.mindswap.org/~mhgrove/RIC/RIC.shtml>.
- North, S.C., Koutsofios, E. (1996). *DOT and NEATO User's Guide*. AT&T Bell Labs Reports.
- Noy, N.F., Sintek, M., Decker, S., Crubezy, M., Ferguson, R.W., Musen, M.A. (2001). Creating Semantic Web contents with Protege-2000. *IEEE Intelligent Systems*, 16(2):60–71.
- Pietriga, E. (2002). IsaViz: A visual environment for browsing and authoring RDF models. *The Eleventh International World Wide Web Conference (WWW2002), Developer's Day*.
- Raines, P. (1998). *Tcl/Tk Pocket Reference*. O'Reilly & Associates.
- Sintek, M. (2004). *OntoViz Tab: Visualizing Protégé Ontologies*. Available: <http://protege.stanford.edu/plugins/ontoviz/ontoviz.html>.
- Storey, M.-A.D., Musen, M., Silva, J., Best, C., Ernst, N., Ferguson, R., Noy, N. (2001). Jambalaya: Interactive visualization to enhance ontology authoring and knowledge acquisition in Protégé. *Workshop on Interactive Tools for Knowledge Capture (K-CAP-2001)*.
- Sugiyama, C., Tagawa, S., Toda, M. (1981). Methods for visual understanding of hierarchical systems. *IEEE Transactions of System, Man, and Cybernetics*, 11(2):109–125.
- Sure, Y., Angele, J., Staab, S. (2003). OntoEdit: Multifaceted inferencing for ontology engineering. *Journal on Data Semantics*, LNCS 2800, pp. 128–152, Springer.
- Telea, A., Maccari, A., Riva, C. (2002). An Open Toolkit for Prototyping Reverse Engineering Visualization. *IEEE EG Symposium on Visualization (VisSym'02)*, pp. 241–250. The Eurographics Association.
- Telea, A. (2004). An open architecture for visual reverse engineering. *Managing Corporate Information Systems Evolution and Maintenance*, pp. 211–227, Idea Group Inc.
- Van Liere, R., De Leeuw, W. (2003). GraphSplatting: Visualizing graphs as continuous fields. *IEEE Transactions on Visualization and Computer Graphics*, 9(2):206–212, IEEE CS Press.
- Vdovjak, R., Frasinca, F., Houben, G.J., Barna, P. (2003). Engineering Semantic Web information systems in Hera. *Journal of Web Engineering (JWE)*, 2(1–2):3–26, Rinton Press.
- Wernecke, J. (1993). *The Inventor Mentor: Programming Object-Oriented 3D Graphics*. Addison Wesley.