Multimed Tools Appl (2012) 59:523–556
DOI 10.1007/s11042-011-0762-9

Linked Data and multimedia: the state of affairs

Bernhard Schandl · Bernhard Haslhofer ·
Tobias Bürger · Andreas Langegger · Wolfgang Halb

Published online: 22 February 2011

Abstract Linked Data is a way of exposing and sharing data as resources on the Web and interlinking them with semantically related resources. In the last three years significant amounts of data have been generated, increasingly forming a globally connected, distributed data space. For multimedia content, metadata are a key factor for efficient management, organization, and retrieval. However, the relationship between multimedia and Linked Data has been rarely studied, leading to a lack of mutual awareness and, as a consequence thereof, technological deficiencies. This article introduces the basic concepts of Linked Data in the context of multimedia metadata, and discusses techniques to generate, expose, discover, and consume Linked Data. It shows that a large amount of data sources exist, which are ready to be exploited by multimedia applications. The benefit of Linked Data in two multimedia-related applications is discussed and open research issues are outlined with the goal

B. Schandl (✉) · B. Haslhofer
Department of Distributed and Multimedia Systems, University of Vienna,
Liebiggasse 4/3-4, 1010 Wien, Austria
e-mail: bernhard.schandl@univie.ac.at

B. Haslhofer
e-mail: bernhard.haslhofer@univie.ac.at

T. Bürger
Salzburg Research Forschungsgesellschaft mbH, Jakob Haringer-Straße 5/3,
5020 Salzburg, Austria
e-mail: tobias.buerger@salzburgresearch.at

A. Langegger
Institute for Application-Oriented Knowledge Processing, Johann Kepler University Linz,
Altenberger Straße 69, 4040 Linz, Austria
e-mail: al@jku.at

W. Halb
DIGITAL—Institute for Information and Communication Technologies,
JOANNEUM RESEARCH Forschungsgesellschaft mbH, Steyrergasse 17, 8010 Graz, Austria
e-mail: wolfgang.halb@joanneum.at

of bringing the research fields of multimedia and Linked Data closer together in order to facilitate mutual benefit.

Keywords Linked Data · Semantic Web · Multimedia semantics · Multimedia metadata

1 Introduction

Over the recent years, the Web has evolved from a collection of mostly text-based content to a giant multimedia database. Today we are able to access all kinds of multimedia material using the stack of Web technologies. This includes images, video and audio clips, live TV streams, but also animated and interactive content. These vast amounts of multimedia data are provided by different sources, including professional media producers and end users alike.

Metadata are key for the efficient management of multimedia content in institutional media repositories [62]. Metadata cannot only be used to describe low-level technical attributes of multimedia documents such as their length, resolution, or color depth, but even more importantly, to describe high-level semantic features such as a genre classification or information about depicted persons. It is widely accepted that the availability and quality of metadata is critical for effective and efficient search and retrieval. Content providers can achieve this by creating metadata descriptions, enriching them with structured and well-defined specifications of background knowledge (thesauri, classification schemes, and ontologies [29]), and introducing links to semantically related resources.

The Web is a giant source of such potentially relevant background information. Consider, for example, data sources such as Wikipedia[1] or WordNet[2], which provide the user with detailed background information about terms, their meanings, and the relationships between them. This knowledge can be exploited to enhance existing metadata descriptions with high-quality semantic information. However, the current World Wide Web has been designed for human consumption, but not for automatic processing by machines. If structured data exists, it is either hidden in databases (the so-called "deep web") or, in the best case, is provided for bulk download in proprietary formats, which are difficult to process.

1.1 Motivating example

If we search for a certain resource such as Stanley Kubrick's movie *The Shining*, we will find out that there is a lot of information about this movie available on the Web. Among the top hits will most likely be the Wikipedia article and the Internet Movie Database (IMDB)[3] entry about *The Shining*. From these sites we can learn a

[1] Wikipedia: http://www.wikipedia.org

[2] WordNet: http://wordnet.princeton.edu

[3] Internet Movie Data Base: http://www.imdb.com

lot about that movie: we see that *Jack Nicholson* is starring, that the movie's genre is *Horror/Thriller*, that the movie is distributed by *Warner Bros*, and so forth.

The problem is that this information is available in a human-readable representation only. Applications that need to further process these data encounter technical problems because they need to parse the data out of HTML documents, which is usually an imprecise and error-prone task. Alternatively they can use data source-specific APIs to retrieve the raw data about the movie *The Shining*; however this requires customization of applications for each specific data source they want to access (Fig. 1).

In the "classic" human-consumable Web it is unimaginable that every web page would require the client to adapt to their specifics—would anyone accept the requirement to install a specific browser plugin for each distinct web site? The World Wide Web works because it is built upon technical standards that are understood

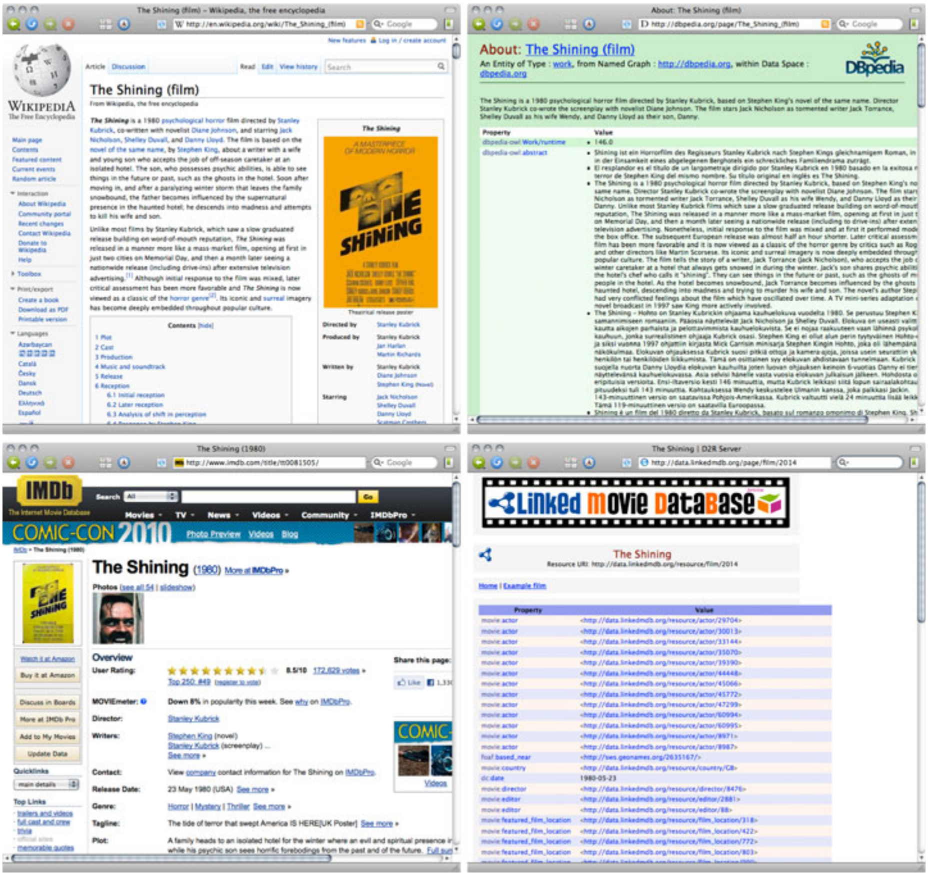

Fig. 1 Web data sources about the movie "The Shining": human-readable (Wikipedia and IMDB, *left*) and machine-readable (DBpedia and LinkedMDB, *right*)

by all involved actors, ranging from identification through transport to information representation. For machine-processable raw data, this is not yet the case.

1.2 Multimedia and Linked Data

The vision of the *Linked Data* movement is to provide a uniform access infrastructure for *data* on a global scale, just like the World Wide Web is for *documents*. The goal of this paradigm is to make data available for consumption by humans *and* machines using standardized formats and access mechanisms, thus enabling developers of applications to re-use data easily and in a unified manner. One part of this goal is to *connect* different sources that provide similar data, thus forming a global graph that can be traversed by clients in order to discover new information. This global data graph is also denoted as the *Web of Data*.

DBpedia [16] is an example for such a source in the Web of Data. It exposes all the information available in Wikipedia in a structured form and provides links to related information in other data sources such as the *Linked Movie Data Base* (LinkedMDB, cf. Section 4.2). As of November 2009, the DBpedia knowledge base describes more than 2.9 million things such as persons, music albums, or films in 91 different languages. It provides a user-generated knowledge organization system comprising approximately 415,000 categories and millions of links to semantically related resources on the Web.

In the case of multimedia objects, we are typically talking about *metadata describing multimedia objects*. A metadata description about the movie *The Shining* in one data source (e.g., LinkedMDB), for instance, should contain links to other data sources (e.g., DBpedia) that maintain information about the same movie. This gives direct access to a magnitude of potentially relevant data to metadata-centric applications. What we, however, are currently missing with respect to multimedia in the Web of Data is the support for more fine-grained semantic descriptions and look-up.

We believe that content providers can make use of the Web of Data in two ways. First, they can consume and re-use the high-quality and often multi-lingual information provided by publicly available knowledge bases, such as DBpedia, to semantically enrich their metadata descriptions. This significantly reduces the required effort for creating and maintaining proprietary knowledge bases and increases the quality of organization, search, and retrieval of multimedia content. Second, they can publish their metadata descriptions as Linked Data on the Web, which will increase their visibility and the expand the coverage of their content to a Web scale.

To reach this goal, it is important to bring together knowledge from the two fields of Web Science and multimedia research. Therefore, this article gives an overview on the current state of the art of Linked Data research, with a special focus on work that is relevant in the context of multimedia. The core set of underlying technologies (Section 2) are introduced, and methods how Linked Data can be produced (Section 3) are described. Further, multimedia-related data sets that are already existing on the Web and are ready to be used by applications are discussed (Section 4), and methods to consume these data are described (Section 5). The practical benefit of Linked Data is illustrated by the means of two concrete, multimedia-related applications; namely, Web-based multimedia annotations and personal semantic multimedia (Section 6).

2 Technical foundations

The technical foundations of Linked Data are similar to those of the traditional, human-centered World Wide Web, and extend them with technologies for the representation and interpretation of machine-processable data. In the following we give a brief introduction to the fundamental design principles and technologies that underly the Web of Data by the means of the motivating example described in the previous section.

2.1 Linked Data principles

In 2006 Tim Berners-Lee postulated the so called *Linked Data principles* [10] as a recommended best practice to share and connect structured data on the Web. These are:

1. *Use URIs to identify things.*
2. *Use HTTP URIs so that people can look up those names.*
3. *When someone looks up a URI, provide useful information, using Semantic Web standards such as RDF and SPARQL.*
4. *Include links to other URIs, so that they can discover more things.*

These principles have been taken on by a constantly growing community; a comprehensive overview of recent research and industry activity related to Linked Data has been given by Bizer et al. [15].

The Linked Data principles resemble the *Resource-Oriented Architecture* of RESTful Web Services [27]. In REST applications, "*everything that is important enough to be referenced as a thing in itself*" should be a resource and "*have at least one URI*" [53]. HTTP serves as a *uniform interface* for looking up and manipulating remote resources and their representations. For retrieving *useful information*, that is for deciding between alternative resource representations, REST proposes to use *content negotiation*. RESTful services should also include links to other resources in the data they serve.

Linked Data goes beyond the RESTful architecture and also lifts the resource representations to the level of the Web. A RESTful service defines how to access and exchange data *through the Web* but does not provide any recommendation on the technical characteristics of resource representations (the exchanged data). Linked Data proposes to use RDF, a model based on the Architecture of the World Wide Web [41], for representing resources *on the Web*. Therefore, we can regard the Linked Data principles as a Web-centric specialization of the Resource-Oriented Architecture defined in REST.

2.2 Underlying technologies

In the following we discuss how the various technologies provided by the Semantic Web can be used to implement our *The Shining* movie example according to the Linked Data principles.

2.2.1 Uniform Resource Identifier (URI)

The *Uniform Resource Identifier (URI)* [50] is a fundamental concept in the Web architecture. It defines a generic identifier syntax for various types of resources, including ones with digital manifestations (e.g., videos, Web sites, and pictures) and ones without (e.g., people, books, and locations). Because URIs are a simple way of identifying things, they displaced many other identification schemes and fundamentally contributed to the success of the World Wide Web.

The Linked Data principles *(1)* and *(2)* demand identifiers to be dereferenceable HTTP URIs. In DBpedia, for instance, the movie "The Shining" is identified by http://dbpedia.org/resource/The_Shining_(film). According to principle *(3)*, when the URI is dereferenced via HTTP, the server should deliver an HTML representation of the metadata describing this resource, if the requesting client is an ordinary Web browser, or an RDF representation, if it is an RDF-aware application client. In both cases it is recommended to distinguish between the *thing itself* (i.e., the actual movie), and any resource that conveys *information about* the thing, like an HTML page. It is a best practice (cf. [57]) to also assign URIs to the different representations of this entity; for instance, http://dbpedia.org/page/The_Shining_(film) for a human-readable web page (HTML) about the movie, and http://dbpedia.org/data/The_Shining_(film) for a machine-readable RDF document containing metadata about the movie.

2.2.2 Resource Description Framework (RDF)

In the Web of Data, the *Resource Description Framework (RDF)* [70] is used as the model for representing data and metadata *about* resources. It allows us to formulate *statements*, each of which consists of a subject, a predicate, and an object. The subject and predicate in a statement must always be resources, the object can either be a resource or a literal node.[4] A statement is represented as a *triple*; triples can be grouped to form a *graph*. Figure 2 shows how we can use RDF to describe the movie "The Shining".

To be exchanged between systems, RDF graphs must be encoded into a concrete serialization format. The original RDF specification proposes the *RDF/XML Syntax* [69]. Other, less verbose encoding formats such as Notation3 (N3),[5] Turtle [9], or N-Triples[6] are also widely used.

2.2.3 RDFS, OWL, and SKOS

The *RDF Vocabulary Description Language*, also called *RDF Schema (RDFS)* [68], and the *Web Ontology Language (OWL)* [72] are means to describe the vocabulary terms used in an RDF model. RDFS provides the basic constructs for describing classes and properties and allows to arrange them in subsumption hierarchies.

OWL extends RDFS and introduces a distinction between attribute-like (`owl:DatatypeProperty`) and relationship-like (`owl:ObjectProperty`) properties. Additionally it provides more expressive modeling primitives (e.g., class union

[4]The RDF standard also contains the concept of *blank nodes*, i.e., resources that are not assigned a URI. However, since these resources cannot be dereferenced using HTTP they are less relevant in the context of Linked Data.

[5]Notation3 (http://www.w3.org/DesignIssues/Notation3.html) is not a "real" serialization format because it encodes only a superset of the RDF data model.

[6]N-Triples: http://www.w3.org/2001/sw/RDFCore/ntriples/

and intersections, cardinality restrictions on properties, etc.). Further it provides the `owl:sameAs` property, which indicates that two URIs identify the same real-world entity (as used in Fig. 2), which is important for the purpose of linking data sets that contain information about the same entities.

The *Simple Knowledge Organization System (SKOS)* [74] is a model for expressing the structure and constituents of concept schemes (thesauri, controlled vocabularies, taxonomies, etc.) in RDF. With SKOS one can attach multi-lingual labels to concepts and arrange them using two kinds of semantic relationships: *broader* and *narrower* relationships for constructing concept hierarchies, and *associative* relationships for linking semantically related concepts. Figure 3 shows how these technologies are used to classify the movie *The Shining* as a film and, in particular, as *1980s Horror Films*.

In the Web of Data it is important that the terms and concepts defined using one of these languages are again dereference-able HTTP URIs. This allows data providers to reuse vocabulary definitions by integrating them into their metadata by linking to the URIs. The *Best Practice Recipes for Publishing Vocabularies* [12] gives detailed guidelines for that.

2.2.4 SPARQL

Many data providers offer means to execute complex queries over their published data. The *SPARQL Query Language for RDF* [71] is what SQL is to databases: an expressive query language for formulating query patterns over RDF graphs. Additionally it defines a protocol for sending queries from clients to a SPARQL endpoint and for retrieving the retrieved results via the Web. Currently, the abstract protocol specification has bindings for HTTP and SOAP. The important distinction between SPARQL and other languages such as SQL is that it operates entirely through the Web: a client executes a query against a given endpoint and retrieves the result set through common Web transport protocols.

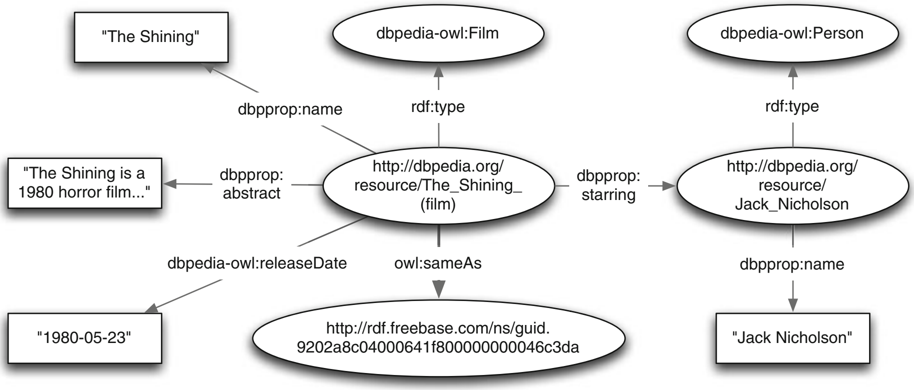

Fig. 2 An RDF graph representing data about the movie "The Shining". The name (`dbpprop:name`) of the movie is represented as a literal node, the starring Jack Nicholson as a resource node, which is further described by RDF statements. The graph contains an additional link (`owl:sameAs`) between the representation of the movie in DBpedia and Freebase

2.3 Delivering useful information

The third Linked Data principle demands to provide useful information when a client accesses a URI. This can happen in different manners: users can dereference a URI with their web browser, or an application can send a request to a URI in order to retrieve structured metadata for the resource. In the first case, the server that handles the request should return HTML to be rendered by the Web browser and presented to the user, while in the second case it should return RDF to be processed by the client application. It is common practice to implement this requirement by using *HTTP Content Negotiation*. The server can then decide which representation to return depending on the value of the `Accept` header field sent by the client as part of the HTTP request.

We have mentioned before that real-world entities (which cannot be conveyed by a digital message, like persons, movies, or places) and documents describing these entities (like Web pages or RDF documents) should be assigned distinct URIs (cf. Section 2.2.1). Based on this idea, the server can also issue an HTTP 303 "See also" response depending on the value of the HTTP `Accept` header field, thus redirecting the client to a descriptive document when a request to a resource without digital representation is issued. For instance, a client request to http://dbpedia.org/resource/The_Shining_(film) (the movie) is redirected either to http://dbpedia.org/page/The_Shining_(film) or http://dbpedia.org/resource/The_Shining_(film), depending on the value of the HTTP `Accept` header.

As an alternative to providing different representations of a resource via different URIs, *RDFa* [1] has been proposed as a means to include structured information within the content of an XHTML web page. RDFa significantly reduces the effort needed to publish Linked Data because it avoids the need for additional server configuration and the provisioning of separate RDF documents. The simplicity of RDFa has also accelerated the inclusion of Semantic Web technologies into content management systems, as demonstrated e.g., with Drupal [23].

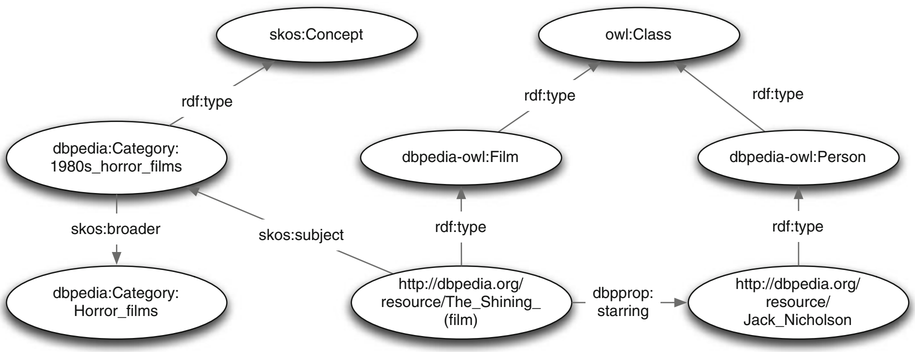

Fig. 3 An RDF graph showing how OWL and SKOS are used in DBpedia for classification and categorization of the movie *The Shining*. The classes `dbpedia-owl:Film` and `dbpedia-owl:Person` are defined as OWL classes. The property `skos:subject` links the movie to the category `dbpedia:Category:1980s_horror_films`, which in turn is `skos:Concept` and has a broader concept `dbpedia:Category:Horror_films`

RDFa may not only be included in XHTML but in other XML-based formats. *Scalable Vector Graphics* (SVG) is one example, which allows to include structured metadata about parts of the graphics. In particular, the recent *SVG Tiny 1.2 Specification* [4] provides two mechanisms for this purpose: first, metadata may be attached to graphics elements using extensible *metadata attributes*, which may be used to incorporate RDFa; second, a designated `metadata` element has been defined that provides a container for arbitrary metadata (e.g., RDF graphs serialized as XML) to be included in the SVG document.

3 Producing Linked Data

Following the Linked Data Principles formulated in the previous section, everyone is able to contribute information to the Web of Data simply by publishing RDF data and linking entities to things described in external RDF datasets. In the context of multimedia these data are typically descriptive metadata about media resources.

If metadata are already available in RDF and stored in a SPARQL-compliant data store, they can easily be published on the Web by applying generic and light-weight tools on top of the SPARQL endpoint. Single RDF files can also be served by a web server supporting content negotiation and correctly handling RDF MIME types as described in [14]. If metadata are available in other formats (e.g., as part of file metadata such as EXIF, ID3, etc. or in a relational database) they need to be converted to RDF first. Technically this can be achieved by wrapper components. In the following we will cover these two cases and present existing software tools. Then we also present existing solutions for generating links among published RDF sources.

3.1 Publishing RDF data on the Web

Generic tools for publishing Linked Data from existing RDF sources are, for instance, *Pubby*[7] and *Paget*.[8] The former is a Java-based Linked Data front-end for SPARQL endpoints, which means that it is able to serve Linked Data from an existing dataset which is accessible via SPARQL. This way, information from any RDF store or wrapper that supports SPARQL can be easily published as Linked Data. Figure 4 shows how Pubby can be used: at the backend it maintains a connection via SPARQL to an RDF store or a wrapper, while on the front-end it serves HTML and Linked Data.

Paget can also be used to publish RDF files and collections of RDF files or datasets stored in the *Talis Platform*.[9] Using additional wrappers as explained in the next section, both tools can be used to publish Linked Data from multimedia sources.

[7] Pubby: http://www4.wiwiss.fu-berlin.de/pubby/

[8] Paget: http://code.google.com/p/paget/

[9] Talis Platform: http://www.talis.com/platform/

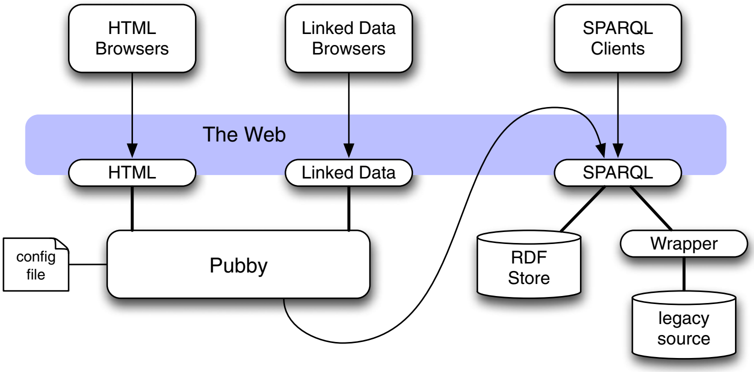

Fig. 4 Exposing SPARQL endpoints as Linked Data via Pubby

3.2 Wrapping and publishing non-RDF data on the Web

In most situations, however, data are not primarily available in RDF. If structured metadata are available for multimedia content, generic *wrappers* may be used to transform them to RDF; otherwise, features and metadata need to be extracted first. While some wrappers support a formal approach and virtual translation based on mappings,[10] others are used to transform all source data into RDF at once (dumping). It usually depends on the application, the size of the dataset, and its update-frequency which approach to follow. For very specific applications, when there is no wrapper available, it may also be easier to hard-code RDF generation.

In order to wrap relational databases to RDF many solutions exist today. One of the most prominent RDB-to-RDF wrappers is *D2RQ* [13], which also contains a SPARQL endpoint. Other popular tools are *Triplify* [6], *R2O* [54], and *OpenLink Virtuoso RDF Views* [26]. Endpoints supporting the *Open Archives Initiative Protocol for Metadata Harvesting* (OAI-PMH) can be exposed using *OAI2LOD* [33], and the contents of spreadsheets can be exposed with tools like XLWrap [44] and *RDF123* [30]. Further, RDF wrappers exist for metadata extracted from files (e.g., EXIF, IPCT, Adobe XMP), web services and web APIs (e.g., Flickr, Google Base, Amazon), calendar and email applications, LDAP sources, XML, and so forth.[11] The *Aperture Framework*[12] supports the extraction of RDF metadata from a broad range of different file formats. File type-specific extractors can be integrated into *TripFS* [59], which is a framework for publishing file systems and file metadata as Linked Data.

[10] In this case it is not required to materialize the complete source dataset into RDF because parts of it can be transformed upon request.

[11] Lists are maintained at http://esw.w3.org/topic/ConverterToRdf and http://simile.mit.edu/wiki/RDFizers.

[12] Aperture: http://aperture.sourceforge.net/

Interactive tools for creating RDF-based annotations (e.g., image regions) are also available.[13] Although some of these tools already generate RDF output, most of them require additional RDF wrappers because the primary output format is either proprietary or some non-RDF standard.

3.3 Linking data

Setting links between different data sources is one of the Linked Data principles and allows to discover more information. Linked Data browsers, crawlers, and applications can automatically follow these links and retrieve information from various sources. The property that is used to link data depends on the application domain and the intended semantics. For instance, the property `owl:sameAs` denotes that two URIs refer to the same entity, whereas `foaf:knows` expresses social relationships between persons. Any property can be used to set links between URIs and there exists a wide range of vocabularies that provides commonly used properties (cf. Section 4.1). It is advisable to utilize properties from popular vocabularies as this increases reusability.

Linking data basically involves three steps:

1. Identify local resources that could be enriched with external information
2. Look up URIs for further information in external datasets
3. Choose a property to set the link

In very small datasets like personal FOAF profiles the linking can be done manually. In order to find a suitable URI to link to in external datasets, lookup services (cf. Section 5.1) can be used. Such a manual approach is however not feasible for larger datasets. The challenge of linking related Web resources is closely related to the record-linkage problem in database research, as URIs in different datasets need to be found that describe the same resource. If there exists an identifying feature—such as an ISBN for books or ISO-3166 country codes—in both datasets, the link generation is almost straightforward. With the aid of tools it is also possible to automatically link data from different datasets in more sophisticated cases. An overview of such tools can be found in [60].

Linking tools usually take as input two or more datasets and a linkage specification, and return a set of links between the datasets. The *SILK Link Discovery Framework* [66] is a popular aid in link generation that uses a declarative language to specify which RDF links between datasets should be discovered under which conditions. It supports different string comparison techniques and similarity measures (e.g., Jaro distance, string similarity based on q-grams, etc.). The tool's documentation also offers an example for setting links between DBpedia movies and directors in LinkedMDB (cf. Section 4.2). In this example movies in DBpedia are linked via a `dbpedia:director` property to their respective directors in LinkedMDB by a simple label match. To achieve this, only movies from DBpedia and only directors from LinkedMDB are selected, and the similarity of the DBpedia movie's

[13] Listed at http://www.w3.org/2005/Incubator/mmsem/wiki/Tools_and_Resources.

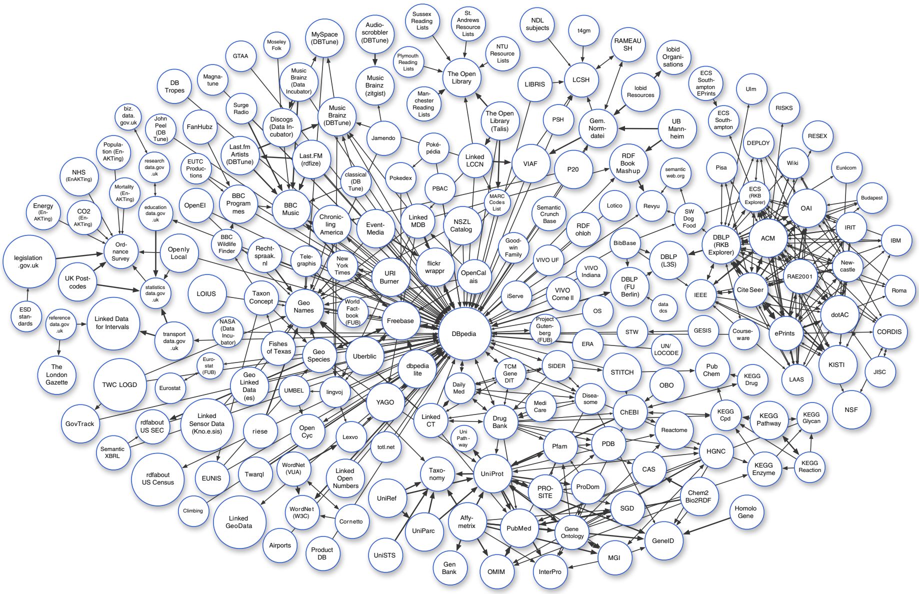

Fig. 5 Linked Data Cloud (most recent version as of September 2010) [24]

director `rdfs:label` to both the `rdfs:label` and the `movie:director_name` of directors in LinkedMDB are evaluated. If the similarity is beyond a user-defined threshold the link is created.

4 Existing vocabularies and data sets

The Web strives its enormous power from the amounts of information that it can provide, which are distributed across millions of servers, that is ready to be used by any client user or application. Recently the Semantic Web has experienced a significant increase of available data as part of the Linked Data initiative, which can be similarly used within applications, as described in the previous section. Currently the so-called *Linked Data Cloud* (cf. Fig. 5), which is a graphical depiction of the most popular data sources and the links between them, is estimated to contain more than 13 billions RDF triples,[14] and its size is continuously increasing. In this section we explore some of the most important data sets that are relevant for a multimedia context and discuss their origins, the data they contain, and the vocabularies they use.

4.1 Vocabularies

A vocabulary in the context of the Web of Data is a *collection of identifiers* (URIs) with well-defined meanings, which are defined within a common *namespace*.

[14]Cf. http://esw.w3.org/topic/TaskForces/CommunityProjects/LinkingOpenData/DataSets/Statistics

Normally a vocabulary consists of identifiers for *resource types*, and identifiers for *properties*. Often these identifiers are related through formal constraints (e.g., domain/range and sub-/superclass restrictions), which are usually expressed using the vocabulary description languages discussed in Section 2.2.3.

During the last years the Semantic Web and Linked Data communities have developed several vocabularies that are of relevance in the context of multimedia (meta)data. It is considered good practice to reuse them or, if they do not provide appropriate modeling primitives, extend them [14]. A detailed discussion of these vocabularies is out of the scope of this paper; however all presented vocabularies provide descriptive information that can retrieved by dereferencing their namespace URIs.

- *Dublin Core* (DC, http://purl.org/dc/elements/1.1/) and *Dublin Core Terms* (DCT, http://purl.org/dc/terms/) define generic classes and properties for the annotation of human-created artifacts. Their roots lie in the domain of bibliographic metadata, and therefore the vocabularies mainly consist of properties to describe the provenance, format, language, and rights of digital or physical items.
- *Friend of a Friend* (FOAF, http://xmlns.com/foaf/0.1/) is a vocabulary that can be used to describe persons, organizations, and the relationships between them. The underlying idea of FOAF is to model a global social network, where personal information is published in a decentralized style, and related persons link themselves through the `foaf:knows` property. A prominent property in the context of multimedia is `foaf:depicts`, which can be used to indicate things depicted in an image.
- *Basic Geo Vocabulary* (WGS84, http://www.w3.org/2003/01/geo/wgs84_pos#) defines a small set of properties for the representation of geographical coordinates (latitude, longitude, and altitude).
- *Music Ontology* (MO, http://purl.org/ontology/mo/) defines terms for representing a wide range of music-related information, ranging from the representation of musical works over physical media (CD, tape) to information about performances and artists.
- *Creative Commons* (CC, http://creativecommons.org/ns#) provides terms and classes for representing legal information about works, their associated licenses, and usage and distribution permissions.
- *Review Vocabulary* (REV, http://purl.org/stuf\/f/rev#) consists of terms that represent reviews, ratings, and comments for arbitrary objects; e.g., multimedia content.

There exist many more vocabularies of relevance for multimedia data (cf. [19]), like the *Multimedia Metadata Ontology* (M3O) [55], the Core Ontology for Multimedia (COMM) [5], the *Video Vocabulary*,[15] or the *W3C Exif Vocabulary*;[16] however they are not (yet) widely used in the context of Linked Data. With the increasing adoption of multimedia-related Linked Data, we except the importance of carefully designed ontologies to increase. Since ontologies differ widely in their modeling focus and usage, Semantic Web search engines as described in Section 5.1 can help the user or developer to find relevant vocabularies to be reused and extended.

[15]Video Vocabulary: http://digitalbazaar.com/media/video

[16]W3C Exif Vocabulary: http://www.w3.org/2003/12/exif/

Table 1 Sizes[a] and vocabularies of selected Linked Data sets

Data set	Size (triples)	Vocabularies in use					
		DC	FOAF	WGS84	MO	CC	REV
ACM RKB	>12,000,000	•					
AudioScrobbler	>60,000,000	•	•	•	•		
BBC	>20,000,000	•	•		•		•
Bio2RDF	>2,000,000,000	•	•				
data.gov	>5,000,000,000	•	•				
DBpedia	>470,000,000		•	•			
Freebase	>100,000,000					•	
FOAF profiles	>60,000,000		•				
GeoNames	>90,000,000		•	•			
LinkedGeoData	>3,000,000,000	•		•		•	
LinkedMDB	>3,500,000	•	•				
MusicBrainz	>60,000,000	•	•		•		
RDF Book Mashup	>100,000,000	•	•				•
US Census	>1,000,000,000	•		•			
Wordnet	>2,700,000	•					

[a]Size data taken from http://esw.w3.org/topic/TaskForces/CommunityProjects/LinkingOpenData/DataSets/Statistics

4.2 Data sets

The number of data sets that are part of the Web of Data is too big to discuss them all in detail here. Therefore we focus on a number of selected data sets that are either of general interest for metadata-centric applications, or are especially relevant because they contain multimedia-related information. Table 1 gives an overview on these relevant data sets, together with their approximate sizes in RDF triples and the vocabularies that are used therein.

DBpedia DBpedia [16], often denoted as "*nucleus*", "*hub*", or "*crystallization point*" of the Web of Data, is an RDF representation of data extracted from Wikipedia pages, and contains especially semi-structured information from the wiki data structure and from info-boxes found on many pages. Therefore DBpedia provides information about an immensely broad range of topics and, even more importantly, provides identifiers for nearly every concept one can imagine. Because of this, DBpedia is of high importance for connecting different data sets through shared references.

To get an impression on data that is available from DBpedia, consider the record about the movie "*The Shining*", of which a snippet is depicted in Fig. 6.[17] In addition to data taken from the different language editions from Wikipedia (labels in different languages, line 2) and data extracted from info-boxes (lines 3–6), we find different categorizations based on Wikipedia's category system (line 7) and a `owl:sameAs` reference to another dataset; in this example, Freebase, which contains more information about this movie.

[17]We use Turtle syntax [9] for representing RDF examples.

```
<http://dbpedia.org/resource/The_Shining_%28film%29>
    rdfs:label "The Shining"@en , "Shining"@de , "Osvícení"@cs ;
    dbpedia-owl:director <http://dbpedia.org/resource/Stanley_Kubrick> ;
    dbpedia-owl:starring <http://dbpedia.org/resource/Shelley_Duvall> ,
        <http://dbpedia.org/resource/Jack_Nicholson> ;
    dbpedia-owl:releaseDate "1980-05-23"@en ;
    rdf:type dbpedia-owl:Work , dbpedia-owl:Film ;
    owl:sameAs
        <http://rdf.freebase.com/ns/guid.9202a8c04000641f800000000046c3da> .
```

Fig. 6 DBpedia data about "The Shining"

All DBpedia URIs are dereference-able, and data are served in compliance with Linked Data principles, either as XHTML+RDFa or as RDF/XML. In addition, several search and retrieval services have been established around DBpedia, including the DBpedia lookup service[18] and a faceted search application.[19]

BBC programmes and music The British Broadcasting Corporation (BBC) is publishing significant amounts of internal databases in Linked Data form. Most notably this encompasses data about broadcast programs and their content, artists, and events. BBC uses the Music Ontology [52] to model their data, and includes references to DBpedia entries. Most notably, BBC employs Linked Data concepts not only to expose their data to the public, but also to internally connect heterogeneous data sources [43]. Thus, not only organization's internal data management becomes more efficient because entities that are present in multiple databases are connected; but it additionally brings better user experience for customers since these links are also included in the public BBC web appearance, allowing users to seamlessly navigate across different services.

DBtune DBtune is a collection of music-related data sets, which are exposed as part of the Web of Data. Data published under this label includes, amongst others, MusicBrainz (a community-maintained collection of metadata about published media), AudioScrobbler (a database containing listening habits from customers of the Last.fm online music service), and data extracted from the MySpace social network. Similarly to the BBC Music services, DBtune represents data using the Music Ontology [52]; these data sets are therefore fully compatible. DBtune also provides *Henry*,[20] an online agent that performs signal analysis on media resources found on the Web. This process can be initiated and controlled via SPARQL queries, and the processing results are in turn published as Linked Data. The different parts of the DBtune service are interlinked, and are connected to external data sets such as DBpedia and GeoNames.

Linked Movie Data Base The Internet Movie Data Base (IMDB) is one of the most popular web services around motion pictures. The Linked Movie Data Base [36] is a Linked Data representation of a fraction of this database. It contains information

[18] DBpedia Lookup: http://lookup.dbpedia.org/

[19] Faceted Wikipedia Search: http://dbpedia.neofonie.de/browse/

[20] Henry: http://dbtune.org/henry/

about around 38,000 movies and 29,000 actors. Parts of the dataset are links to external data sources, including DBpedia, GeoNames, and MusicBrainz.

GeoNames and Linked Geo Data The GeoNames service provides Linked Data-based descriptions for over 6.2 millions geographical entities, such as countries, regions, and cities. GeoNames resources are interlinked using appropriate relationship types (e.g., `geo:inCountry` or `geo:nearbyFeature`) and contains geographical coordinates for entities according to the W3C WGS84 vocabulary.[21] GeoNames is used as link target for a number of external data sources (including DBpedia) and therefore can be considered as a Linked Data hub for geography-related data. LinkedGeoData [7], on the other hand, is an RDF representation of geographical information from the *Open Street Map* service, which contains community-maintained map data from all over the world. It allows clients to retrieve spatial data based on a REST-style approach, where coordinates are encoded in the request URI. The LinkedGeoData set is connected to DBpedia and provides, in addition to the Linked Data interface, a graphical faceted spatial browser.

In addition to these multimedia-related data sets, there exist a large number of datasets on general topics (e.g., RDF representations of the *Freebase* service), statistical and governmental data (e.g., *data.gov* or *US Census*), data from the life sciences (e.g., *Bio2RDF*), and scientific publications (e.g., *ACM RKB*). Furthermore, one can consider the large amount of published FOAF profiles[22] as part of the Web of Data, since usually they are interlinked with `foaf:knows` properties and commonly refer to other data sets (e.g., DBpedia) using properties such as `foaf:interest`. In summary, we can observe that there exists a significantly large set of multimedia-related data sets, which can be considered as an important step to solve the bootstrapping problem for multimedia-related Linked Data. It is now up to application developers to utilize and exploit these data sets, which will in turn lead to more feedback and increased quality of data.

5 Consuming Linked Data

Linked Data can be consumed in a variety of ways by humans and machines. Several tools exist that aid in the use of Linked Data. Basically there are three distinguishable tasks when interacting with Linked Data that will be discussed in the following.

5.1 URI discovery

The first thing to start with when consuming Linked Data is the URI of a resource's description. According to the Linked Data principles this allows to retrieve useful information when the URI is looked up. Usually the URI of a resource is not known beforehand but there exist several ways for discovering URIs that identify the thing one is interested in.

[21] W3C WGS84 vocabulary: http://www.w3.org/2003/01/geo/

[22] FOAF: http://www.foaf-project.org

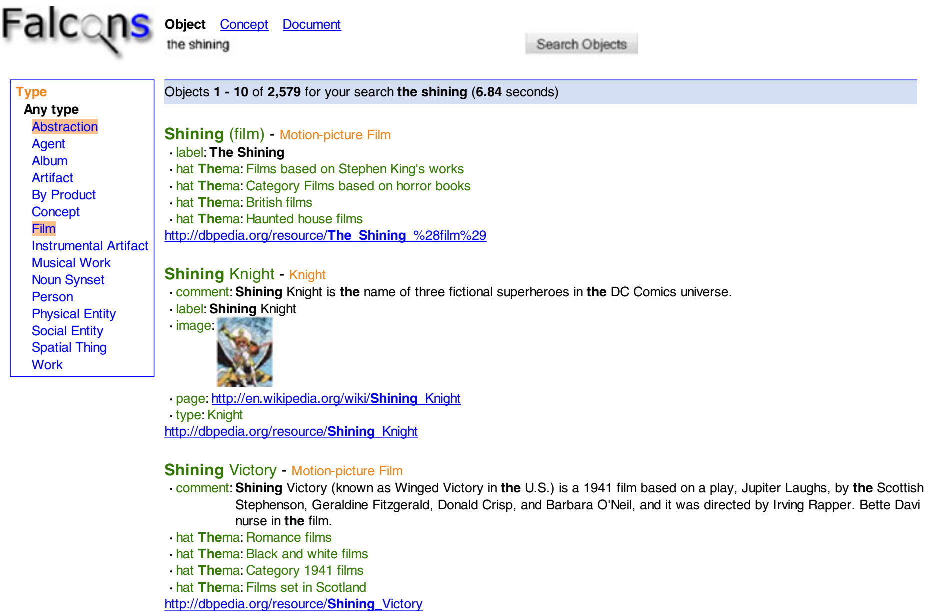

Fig. 7 Results of a search for *The Shining* on Falcons search engine

Search engines for the Web of Linked Data provide a keyword-based search across different datasets. The most widely used search engines are *Falcons* [22], *Sindice* [65], *SWSE* [31], and *Watson* [25]. The results of a query for "The Shining" on the Falcons search engine are depicted in Fig. 7. In the search box the keywords for objects to be found have to be entered. For multimedia objects this will most likely be the label of the object if information about the entire object should be retrieved. For multimedia objects it might also be relevant to find more information about fragments of the object (such as for instance an image region) where the search also needs a keyword describing the desired information. The search engine further allows to refine the search by types (box on the left-hand side) where the category "Film" is already suggested appropriately in the depicted search case. In addition to a graphical user interfaces, most search engines also provide APIs to programmatically access search results.

Many datasets offer dedicated keyword-based search interfaces. The already mentioned lookup feature provided by DBpedia, shown in Fig. 8, returns DBpedia URIs for the search terms. Another method to discover URIs are SPARQL queries on endpoints that are either provided by the dataset providers themselves (e.g., DBpedia) or by data consolidators that allow the access to many different datasets. One example is the *Open Link LOD Cloud Cache*,[23] which aggregates data from several sources into a single query endpoint.

[23]Open Link LOD Cloud Cache: http://lod.openlinksw.com/sparql

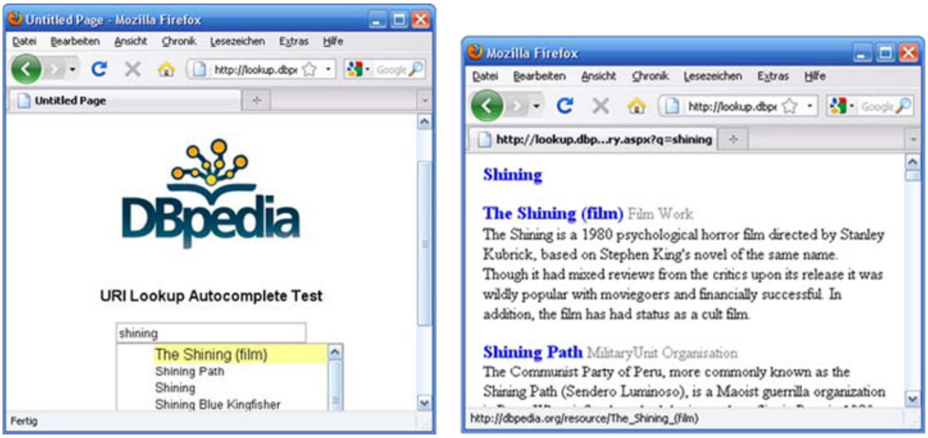

Fig. 8 DBpedia lookup feature based on keywords for the exemplary query "shining"

5.2 Data discovery

Given the URI it is easily possible to retrieve all information that is available by looking it up. The following approaches allow to discover more data that is available in the Web of Data:

- *Link traversal*: The "*follow-your-nose*" approach suggests to follow links such as `owl:sameAs` or `rdfs:seeAlso` and is one of the Linked Data principles.
- *Co-reference services*: Dedicated services such as http://sameas.org allow to find different URIs that refer to the same thing.
- *Search engines*: Many Semantic Web search engines also allow to query by URI and thus act like a co-reference service.

All three approaches can also be used programmatically inside an application. For link traversal the information from links has to be retrieved. The co-reference services and search engines provide APIs for easy integration.

5.3 Dataset discovery

The discovery of datasets is primarily interesting to find programmatically accessible endpoints where data can be retrieved from. A list of publicly known SPARQL endpoints is available from http://esw.w3.org/topic/SparqlEndpoints. Apart from manually selecting the datasets there is also the possibility of (semi-)automatically discovering datasets that are described using the Vocabulary of Interlinked Datasets (voiD) [3]. Using the voiD vocabulary it is possible to describe the topics, terms of usage, interlinkage, and discovery mechanisms for data sets.

These dataset descriptions are in turn published according to Linked Data principles, which enables them to be linked and to be discovered by search engines. Additionally, specialized *description indices* have been established (e.g., the voiD collection on RKBExplorer[24]), which provide descriptions of more than 40 datasets that are mainly focused on people, publications, research areas, and projects. From such collections of dataset descriptions it is possible to identify further relevant sources of information.

5.4 Accessing and consuming Linked Data

There exist several tools for interacting with Linked Data as a user client and there are also libraries available that enable software developers to directly interact with Linked Data. Some of the available solutions will be presented below. However, an exhaustive list of all available solutions cannot be provided here given the plethora of available tools.

If Linked Data are to be consumed by humans, Linked Data browsers can be used, which are usually implemented as Web applications running inside a browser. They display data returned when looking up one or more URIs in a tabular form. Linked Data browsers such as *Tabulator* [11] and *OpenLink Data Explorer Extension*[25] are available as extensions for the Firefox web browser and provide views of the data associated with an URI. *ZitGist DataViewer*[26] and *Marbles*[27] are examples of Linked Data browsers that can be used without installation with any traditional web browser as an HTML view of the RDF data is rendered.

However, Linked Data are primarily designed to be consumed by machines. One of the core ideas of Linked Data is to reuse Web technologies for this purpose as far as possible. This means that for programmatic access to Linked Data, only an HTTP client implementation is needed, which is part of most modern programming language. If such a client is available, dereferencing a resource's URI (together with content negotiation, cf. Section 2.3) yields direct access to RDF descriptions, which can be parsed using appropriate tools like the *Jena Semantic Web Framework* [21] for Java, *RDFLib*[28] for Python, or SemWeb[29] for C#/.NET environments. These libraries usually provide convenient, object-oriented representations of RDF graphs, resources, and statements, and allow to load, manipulate, and write RDF graphs in various formats.

Because SPARQL is based on HTTP, programmatic access to SPARQL endpoints is possible with standard HTTP clients. Several convenient client APIs for executing remote SPARQL queries and processing result sets exist.[30] It is also possible to process federated queries across multiple SPARQL endpoints where various implementations exist such as *SemWIQ* (Semantic Web Integrator and Query Engine) [45] or *Semantic Web Client library* (SWClib) [32], which are all based on Jena's query processor *ARQ*. While SemWIQ is based on a mediator architecture

[24] RKBExplorer: http://void.rkbexplorer.com/sparql/

[25] OpenLink Data Explorer: http://ode.openlinksw.com/

[26] ZitGist DataViewer: http://dataviewer.zitgist.com/

[27] Marbles: http://www5.wiwiss.fu-berlin.de/marbles/

[28] RDFLib: http://www.rdflib.net

[29] SemWeb: http://razor.occams.info/code/semweb/

[30] List of SPARQL implementations: http://esw.w3.org/SparqlImplementations

and algebra-based query federation, SWClib is executing queries based on dynamic link traversal, i.e., following links of distributed RDF graphs.

6 Selected applications

The previous sections introduced the underlying technologies and recent research prototypes in the field of Linked Data. This section shows that the discussed technologies are actually of use in two concrete multimedia-related applications: first, we show how structured semantic multimedia annotations can be integrated with the Web of Data by reusing concepts thereof and re-publishing annotations according to Linked Data principles. Second, we demonstrate the concrete application of Linked Data technologies in the desktop context, in particular for the problem of managing personal media collections.

6.1 Multimedia annotations as part of the Web of Data

Annotations have a long research history and can be traced back to Vannevar Bush's 1945 vision of the *Memex* [20], where users can build associative trails of interest by inserting personal comments into microfilm frames. In the literature we can find various kinds of annotations, different workflows and environments where they are created, and therefore also various interpretations of the term *annotation* (cf. [2, 28, 48, 49, 51]). In recent years, a new form of annotations has emerged: *social tagging*. According to Hunter [39] we can conceive these systems as a sub-class, hence specialized form, of annotation systems. For a recent and detailed analysis of existing annotations research and solutions we refer to [34] and [39].

Here we conceive an annotation simply as being information that is attached to another piece of information. Consequently, a *multimedia annotation* is (multimedia) information attached to a multimedia object. Since this attached information typically describes the multimedia object, we can regard its data representation as descriptive metadata. Figure 9 shows an example annotation on a specific region of an image of Jack Nicholson.

In the Web of Data context, the target of an annotation (e.g., an image of Jack Nicholson) is a multimedia resource identified by a dereference-able HTTP URI. The annotation itself is also identified by an HTTP URI, and when dereferencing this URI, one retrieves the annotation data in a suitable format. The principle of linking resources is inherent to annotations: since an annotation is always attached to another resource, it must refer to that resource and thereby creates a link between itself and the annotated resource.

The questions that arise when using annotations in the Web of Data are: how to represent annotation data, and how to identify fragments in media objects, both in an interoperable way. In the following we discuss existing Web-enabled annotation models and also discuss how they deal with the fragment identification problem.

6.1.1 The Annotea model

The *Annotea* project [42] was the first adopter of Semantic Web technologies for annotations. Annotea defines a client-server based architecture that allows users

Fig. 9 A sample annotation on an image showing Jack Nicholson. A user has marked a certain region (*fragment*) within the image and added her comment on that region

to enhance Web sites with notes and bookmarks. The specification comprises a protocol for client-server communication and a model for representing annotation data in RDF. Over the years it has become a de-facto standard for semantic annotation models and has been implemented in several client-server environments (e.g., *Annozilla*[31]).

Figure 10 depicts an example of a Web site annotation. It clearly shows that Annotea already fulfills most of the Linked data requirements: the annotation itself as well as the annotated Web site are identified by dereference-able HTTP URIs (lines 4 and 7), and the annotation data are represented in RDF and are therefore machine-readable. The content of the annotation is represented in the annotation body (line 13) and links to a Web site containing a user comment. An XPointer [67] expression identifies the second paragraph in the Web site's `Main` element as annotated region (lines 8 and 9).

Annotea was originally designed for the annotation of Web sites and therefore offers only limited capabilities for annotating multimedia objects. Others (e.g., [61]) have extended the Annotea model to furthermore enable the annotation of segments in multimedia objects on the Web.

[31] Annozilla: http://annozilla.mozdev.org

```
<r:RDF xmlns:r="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
       xmlns:a="http://www.w3.org/2000/10/annotation-ns#"
       xmlns:d="http://purl.org/dc/elements/1.1/">
    <r:Description about="http://www.example.org/Annotation/3ACF6D754">
        <r:type resource="http://www.w3.org/2000/10/annotation-ns#Annotation"/>
        <r:type resource="http://www.w3.org/2000/10/annotationType#Comment"/>
        <a:annotates r:resource="http://example.com/some/page.html"/>
        <a:context r:resource="http://example.com/some/page.html#
            xpointer(id('Main')/p[2])"/>
        <d:creator>John Die</d:creator>
        <a:created>2010-01-15T12:10Z</a:created>
        <d:date>2010-01-17-T12:10Z</d:date>
        <a:body r:resource="http://www.example.com/mycomment.html"/>
    </r:Description>
</r:RDF>
```

Fig. 10 Sample Web page annotation represented using the Annotea model

6.1.2 The LEMO annotation framework

LEMO [34] is a multimedia annotation framework that implements the aforementioned Linked Data principles. It provides a *core* model which serves as common denominator for all types of annotations created within the framework. Since this model is based on the Annotea model, LEMO annotations are also fully compatible with systems built on top of Annotea. LEMO adopters also have the possibility to define so-called *annotation profiles*, which are extensions of the core model that serve application-, context-, and content-type specific requirements. For instance, additional properties required by an annotation (e.g., `dc:subject`) can be defined in a context-specific application profile.

Fig. 11 The LEMO map annotation client

```
<r:RDF xmlns:r="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
       xmlns:a="http://lemo.mminf.univie.ac.at/annotation-core#"
       xmlns:x="http://lemo.mminf.univie.ac.at/annotation-video#">
    <r:Description about="http://www.example.org/Annotation/1">
        <a:annotates>http://www.univie.ac.at/test.mpg</a:annotates>
        <a:fragment>
            http://www.univie.ac.at/test.mpg#mp(~time('npt','30','40'))
        </a:fragment>
        <x:time_fragment xmlns:mpeg21="http://lemo.mminf.univie.ac.at/annotation_mpeg21#">
            <mpeg21:uri_fid>
                http://www.univie.ac.at/test.mpg#mp(~time('npt','30','40'))
            </mpeg21:uri_fid>
            <mpeg21:time_scheme>npt</mpeg21:time_scheme>
            <mpeg21:start_time>30</mpeg21:start_time>
            <mpeg21:end_time>40</mpeg21:end_time>
        </x:time_fragment>
    </r:Description>
</r:RDF>
```

Fig. 12 LEMO's dual fragment identification approach

The goal of the LEMO framework is to serve as a common annotation data store for various types of clients. Figure 11 shows a screenshot of an annotation client that allows the annotation of historic maps. Similar to the video annotation client it proposes semantically relevant tags for given regions. These tags are in fact links to possibly relevant DBpedia entries. For further LEMO use cases, we refer to [35, 63].

LEMO relies on a dual fragment identification approach: first, it proposes to encode the identified regions within a multimedia object using the MPEG-21 *media pointer scheme* [40], which is an extension of the XPointer language and provides a rich vocabulary to address parts of MPEG resources and encode this information in a URI fragment. This approach has the drawback that MPEG-21 fragment information has only limited means to express shapes and that MPEG-21 fragments are encoded in a URI string, which complicates structured access to fragment information. Therefore, as an alternative fragment identification approach, LEMO also supports custom fragment definitions in annotation profiles. The semantics of the custom and MPEG-21 fragment definitions should, however, be the same. Figure 12 gives an example of this approach, where the addressed region within a video is addressed by means of an MPEG-21 compatible URI (line 7) as well as in a structured way using a custom fragment element (lines 9–16).

6.1.3 W3C Media Fragments Working Group

The W3C Media Fragments Working Group[32] is currently working on a recommendation for a "*media-format independent, standard means of addressing media fragments on the Web using Uniform Resource Identifiers (URI)*" [73]. This includes a media fragments syntax for addressing the temporal, spatial, and track dimensions of media objects on the Web. Additionally it supports the so-called *named* dimensions, which allows named-based selection of fragments that are denoted by names such as chapters in certain container formats.

Another aspect of that recommendation is the processing of media fragments by servers and clients. This should allow clients (e.g., Web browsers) to fetch only parts of online media objects when dereferencing their URIs, which is important especially for large media resources (e.g., videos).

[32]W3C Media Fragments Group: http://www.w3.org/2008/WebVideo/Fragments/

6.1.4 Other related work in the area of annotations

Complementary approaches to the ones introduced so far have been defined in the multimedia semantics community to semantically describe media resources in an interoperable way. While many of these approaches provide means to describe and locate fragments, most of them do not define the concept of annotation. A brief overview of these endeavors is given in [19].

Only recently the W3C Media Annotations working group has been launched with the goal to improve interoperability between multimedia metadata schemes.[33] In order to do so, it defined the so-called Media Ontology which consists of a set of core properties [64]. This set is based on a list of the most commonly used annotation properties from media metadata schemas currently in use. The ontology defines 20 descriptive and 8 technical metadata properties, an API to access media descriptions and a set of mappings between popular standards. The intention is more to define a vocabulary to describe media resources than an annotation model in the sense of the models introduced in the previous sections. In order to apply descriptions to a fragment of a resource, the ontology provides properties identify fragments using specifications from the W3C Media Fragments group. Another example that defines an annotation model based on so-called ontology design patterns is the *Multimedia Metadata Ontology* (M3O) [55]. One pattern defined in the M3O is the *Annotation Pattern* that formally expresses annotations that can be assigned to arbitrary information entities.

Another example for annotations in the area of online-videos is tele-TASK [56]. The descriptive video metadata are stored in an underlying RDF store and included as RDFa in the web pages containing the videos. This allows external applications to link their own resources with tele-TASK data, while tele-Task data can be augmented with external semantic data.

DBpedia Mobile [8] is a noteworthy service enabling image annotations. It relates photos taken with a mobile device to DBpedia locations and allows links between photos and locations to be published as Linked Data.

LODr [47] is a system that allows users to semantically enrich existing tagged data with meaning. Thereby it establishes relationships between existing tags and linked data resources from arbitrary sources.

6.2 Personal semantic multimedia

In the previous sections, the fundamental Linked Data concepts have been described, which form the basis for an envisioned global Web of Data. The motivation behind the development of these concepts was mostly that information on the Web exists only in human-interpretable, weakly structured form (HTML pages), which makes it difficult to build more elaborate services that utilize these data.

Let us now consider the so-called *Personal Space of Information*, i.e., the set of information that is relevant to a single person's activities. This space consists not only of resources on the web, but even importantly on resources stored on the devices that the user is using. Predominantly this includes computers, notebooks,

[33] W3C Media Annotations Group: http://www.w3.org/2008/WebVideo/Annotations/

and mobile phones, but also non-public company file servers or databases. With the increasing multimedia capabilities of personal devices, the question of efficient media management becomes apparent.

As already mentioned, metadata are key for the management of personal media collections. However, current systems expose significant drawbacks in this field. First, personal devices mostly rely on hierarchical file systems, which are known to be inefficient for the management of large numbers of items. Moreover, hierarchical file systems are usually not capable of metadata-centric search and retrieval, and are therefore of limited use in the context of multimedia management. This has lead to the tendency to implement metadata-centric applications (into which category multimedia-related tools usually fall) not based on files, but on application-specific databases. In consequence, because of the lack of a commonly accepted standard for the representation and interpretation of media metadata in the personal sphere of information, personal media management solutions use proprietary and closed mechanisms to store media descriptions. This leads to a lack of cross-media interoperability, and to the establishment of redundant representation of structures [17].

As an example, consider the suite of multimedia-related applications on the Apple Mac platform. The most important multimedia tools are iTunes (for music and video) and iPhoto (for photos). Both of them provide rich facilities for managing metadata, but store them using different, application-specific mechanisms: iTunes uses one single XML file, while iPhoto uses a set of file-backed SQLite databases. Both variants cannot be manipulated using the file browser, therefore forcing the user to open the corresponding application if they want to access or manipulate a certain item.

To a certain extent, these media tools are integrated with other platform tools; for instance, photos in iPhoto can be annotated with names taken from the address book. However, both the Address Book and iPhoto do not offer ways to follow this link and view photos depicting a certain person, or view the contact details of a depicted person. Other features of the Apple media suite are not linked at all: for instance, iPhotos offers the concept of *events* to group photos, but these events cannot be linked to entries in the Calendar application. Similarly, the integration with external (Web-based) services is practically non-existing: for instance, the Safari web browser allows the user to directly import pictures from web pages into iPhoto; however, the photo is not annotated with its source URL, therefore the connection between the image in the media library and the web page (i.e., its original context) is lost.

6.2.1 Unified personal metadata management

These functional limitations are grounded in two fundamental weaknesses that the desktop exposes: first, there exists no standardized way of *identifying* items, and second, there exists no way of *linking* related items. The (Semantic) Web technology stack and, in particular, the Linked Data philosophy provide solutions for both problems. Thus it is obvious to also apply their principles to the desktop context. This idea has lead to the development of the *(Social) Semantic Desktop*, an environment where each information item can be identified, described, annotated, and manipulated according to Semantic Web standards [58]. In other words, the Semantic Desktop can be seen as a miniature web of data, encompassing all resources a single user is interested in, as well as semantic descriptions thereof. As such the

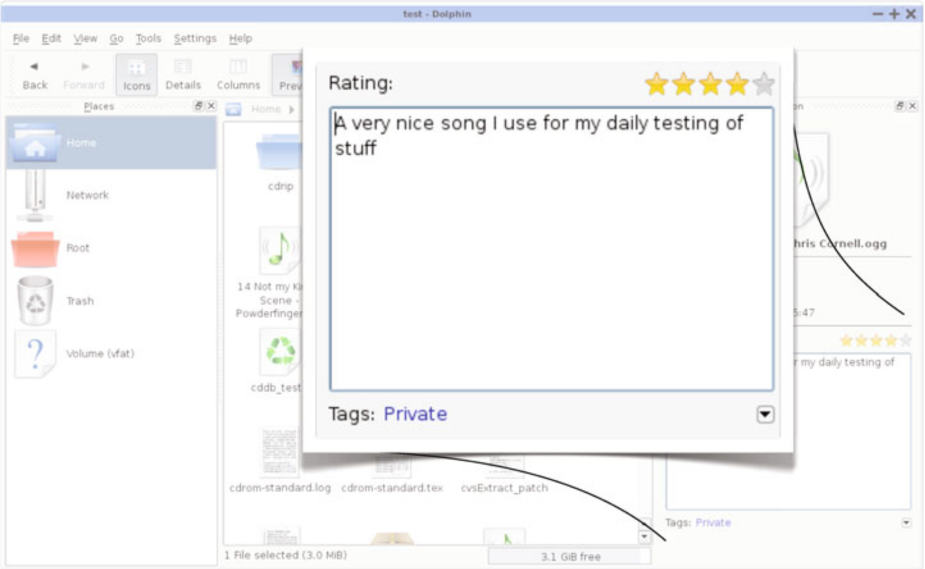

Fig. 13 Semantic Desktop features in KDE Dolphin: rating, comments, tags (screenshot taken from http://liquidat.wordpress.com)

Semantic Desktop is an ideal foundation for the management of personal media collections.

6.2.2 Personal semantic multimedia in practice

As an outcome of the *NEPOMUK* research project,[34] initial building blocks of a Semantic Desktop infrastructure have been integrated into KDE[35] since version 4, equipping each user with an RDF store and a SPARQL query engine on their devices, and providing a comprehensive Semantic Desktop API to KDE application developers. Moreover, a number of ontologies for core data have been specified and are now maintained by a public organization.[36] Using this API and the corresponding ontologies, metadata for media objects can be read, written, and queried in an application-independent manner.

Several KDE applications already make use of this technology; the most prominent example is the KDE file explorer (called *Dolphin*). It allows the user to annotate files with a rating, comments, and freely-chosen keywords, and stores these annotations in the local RDF store (cf. Fig. 13). From there they can be accessed by any application via the Nepomuk-KDE API. An example of such an application is *Gwenview*,[37] which allows the user to search, filter, and browse images based on their ratings or associated keywords.

[34]NEPOMUK Project: http://nepomuk.semanticdesktop.org

[35]Nepomuk-KDE: http://nepomuk.kde.org

[36]OSCAF Ontologies: http://sourceforge.net/apps/trac/oscaf/wiki/Ontologies

[37]Gwenview: http://gwenview.sourceforge.net

In the domain of multimedia, the *Bangarang* media player[38] uses the local RDF store to save metadata about media items. This includes not only metadata extracted from files in the user's media library (like ID3 tags from MP3 files), but also user-generated data like ratings (using the same vocabulary as Dolphin, therefore ratings appear equally in both applications), or automatically collected data like the play count for single tracks. Moreover it stores semantic relationships between media files: for instance, a movie (represented as a video file) can be related to the audio tracks that constitute its soundtrack, if they are available in the media library.

Although they are not strictly "semantic", the features of the applications described before show that Semantic Web technology can be successfully integrated into end-user applications. To an increasing amount the data that forms the basis for these applications can be seen as a Web of Data on end user's personal devices. Because of the technical homogeneity of the personal data network and Linked Data (usage of URIs to identify entities, and usage of the RDF model to express resource descriptions), it is straightforward to connect these two worlds and therefore forming a unified Personal Space of Information that actually encompasses all relevant information for an individual and therefore increases their knowledge work efficiency. An example that stresses the need for such an integrated approach is the recent *Chrome OS* which provides a single metaphor for interacting with both web and local resources.

For end users dealing with multimedia data, the Semantic Desktop provides a common platform for metadata management, search, and retrieval, which has the potential to unify the nowadays heterogeneous data silos of different applications. It allows users to connect their media objects and relate them to other entities in their information space, therefore bringing their data closer together and allowing for more efficient workflows.

7 Summary and outlook

This paper summarizes the current state of research in the field of multimedia-related Linked Data. The analysis of best practices to publish and consume data on a Web scale shows that the technological foundations for a global Web of Data are profound, and a significant amount of data are already published that can be used by applications at any time. Tools and libraries as well as discovery and linking services are available, and it is now up to application-oriented research and practitioners to adopt these technologies and data sets in order to provide new applications and services that bring benefit to end users.

However, there are still a number of unsolved issues [15], and we can observe that there is much space left for more multimedia-related content. Until now, only a limited number of multimedia-related data sets and applications have been developed. We suspect that one reason for this deficit is a lack of cross-community awareness and understanding. A number of ongoing initiatives, however, aim to overcome this weakness and aim to provide a better integration of the two research fields. The *Open Annotation Collaboration* (OAC), for instance, is a recent activity

[38] Bangarang: http://bangarangkde.wordpress.com

in the area of annotation models in the context of Linked Data. It will define an annotation model[39] that enables the sharing of annotations across annotation clients, collections, and applications. One of its goal is to make use of the current developments in the W3C Media Fragments Working Group (cf. Section 6.1.3). At the time of writing, the OAC model is still under development.

Besides that, both the *W3C Ontology for Media Resource 1.0* [46] and the *Media Fragments URI 1.0* [73] are undergoing the W3C recommendation process and are expected to be final in the middle of 2010. We expect both specifications to play a significant role in weaving multimedia into the Web of Data.

Regarding the increasing adoption of Linked Data for the purposes of multimedia metadata management, a number of prerequisites have to be met in order to weave multimedia seamlessly into the Web of Data [38]. Amongst others this demands for means to *address* and *describe* multimedia fragments, as has been highlighted in Section 6.1. Both aspects bear technical challenges due to a vast majority of specifications which can be used for both purposes. The challenges that are prevalent with respect to the application of the Linked Data to multimedia are especially the third and fourth Linked Data principle. No commonly agreed solution has been identified e.g., for the purpose of serving media fragments in a Linked Data context; while a number of options are discussed e.g., in [38].

When it comes to tools supporting the value chain of Linked Data in general, we can observe discrepancies between the need of multimedia data and what is currently offered for publication and consumption of Linked Data. Foremost, regarding the provisioning of links to other URIs to enable discovery (i.e., the fourth Linked Data principle), we can recognize a lack of methods to generate links for and between media. The fourth principle demands for methods to interlink resources on the semantic level. In [18, 38] a number of options are discussed to generate links (semi-)automatically or with human intervention. While potential methods are discussed at the conceptual level, actually working implementations are currently lacking.

In conclusion, we can observe that Linked Data publishing frameworks serve well the demands of textual content; for lifting media descriptions to the Web of Data, however, specific solutions need to be available respecting the characteristics of rich media, taking for instance MPEG-7 descriptions and making them available in a meaningful way on the Web of Data. A conceptual framework for publishing multimedia metadata on the Semantic Web has already been presented [37], but more instantiations that also adhere to Linked Data principles are needed in practice. Likewise, there is currently a lack of Linked Data publication frameworks for media management solutions, which is amongst others due to the need for access control, licensing, and IPR mechanisms for Linked Data which are currently not existing [15].

Many applications in the field will furthermore require not only reading, but also updating of Linked Data, and more precisely to synchronize links and datasets and update information in them. This also requires support for provenance and trust mechanisms to assess who provided which information at a certain level of

[39] An initial draft of the OAC model is available at http://www.openannotation.org/documents/OAC-Model_UseCases-alpha.pdf.

granularity. Due to the subjectivity of media interpretation, this is more urgently required for Linked Media rather than for the Linked Data cloud as such. A lack of methods and tools can not only be observed on the publishing side, but also at the consumption side. Here we see a demand for novel interfaces inspired by works in hypermedia, and better tool support for integrated programmatic handling of Web-based multimedia resources, fragments, and descriptions, in order to experience rich interlinked media collections. On the interface level, given the growth of the Linked Data cloud, there will be an increased need for supporting end users in selecting resources in the Linked Data cloud, for navigating the cloud and for consolidated presentation of information about resources. Linked Data browsers and viewers have to endorse multimedia features in order to allow navigation in and between multimedia resources and datasets.

Acknowledgements Parts of this work have been funded by FIT-IT grant 815133 from the Austrian Federal Ministry of Transport, Innovation, and Technology, and the EU eContentPlus project EuropeanaConnect.

References

1. Adida B, Birbeck M, McCarron S, Pemberton S (2008) RDFa in XHTML: syntax and processing. World Wide Web Consortium. Available at http://www.w3.org/TR/rdfa-syntax/
2. Agosti M, Ferro N, Frommholz I, Thiel U (2004) Annotations in digital libraries and collaboratories—facets, models and usage. In: Heery R, Lyon L (eds) ECDL. Lecture notes in computer science, vol 3232. Springer, pp 244–255
3. Alexander K, Cyganiak R, Hausenblas M, Zhao J (2009) Describing linked datasets—on the design and usage of void, the 'vocabulary of interlinked datasets'. In: Proceedings of the 2nd international workshop on Linked Data on the Web (LDOW). Madrid, Spain
4. Anderson O, Berjon R, Dahlström E, Emmons A, Ferraiolo J, Grasso A, Hardy V, Hayman S, Jackson D, Lilley C, McCormack C, Neumann A, Northway C, Quint A, Ramani N, Schepers D, Shellshear A (2008) Scalable Vector Graphics (SVG) tiny 1.2 specification (W3C recommendation 22 December 2008). World Wide Web Consortium. Available at http://www.w3.org/TR/SVGTiny12/
5. Arndt R, Troncy R, Staab S, Hardman L, Vacura M (2008) COMM: designing a well-founded multimedia ontology for the Web. In: Aberer K, Choi KS, Noy N, Allemang D, Lee KI, Nixon L, Golbeck J, Mika P, Maynard D, Mizoguchi R, Schreiber G, Cudré-Mauroux P (eds) The Semantic Web: ISWC 2007 + ASWC 2007. Lecture notes in computer science, vol 4825. Springer, Berlin, pp 30–43. doi:10.1007/978-3-540-76298-0_3
6. Auer S, Dietzold S, Lehmann J, Hellmann S, Aumueller D (2009) Triplify—light-weight Linked Data publication from relational databases. In: Proc. of 18th international World Wide Web conference
7. Auer S, Lehmann J, Hellmann S (2009) LinkedGeoData—adding a spatial dimension to the Web of Data. In: Proc. of 7th international Semantic Web conference (ISWC)
8. Becker C, Bizer C (2008) DBpedia Mobile: a location-enabled Linked Data browser. In: Linked Data on the Web workshop 2008, CEUR-WS, vol 369. Beijing, China
9. Beckett D, Berners-Lee T (2008) Turtle—terse RDF triple language (W3C team submission 14 January 2008). World Wide Web Consortium, http://www.w3.org/TeamSubmission/turtle/, retrieved 8 August 2008
10. Berners-Lee T (2006) Linked Data. World Wide Web Consortium. Available at http://www.w3.org/DesignIssues/LinkedData.html, retrieved 8 August 2008
11. Berners-Lee T, Chen Y, Chilton L, Connolly D, Dhanaraj R, Hollenbach J, Lerer A, Sheets D (2006) Tabulator: exploring and analyzing Linked Data on the Semantic Web. In: Proceedings of the 3rd international Semantic Web user interaction workshop (SWUI'06). Athens, GA, USA

12. Berruta D, Phipps J (2008) Best practice recipes for publishing RDF vocabularies. W3C Semantic Web deployment working group. Available at: http://www.w3.org/TR/swbp-vocab-pub/
13. Bizer C, Cyganiak R (2006) D2R server—publishing relational databases on the Semantic Web. In: 5th international Semantic Web conference
14. Bizer C, Cyganiak R, Heath T (2007) How to publish Linked Data on the Web. URL: http://www4.wiwiss.fu-berlin.de/bizer/pub/LinkedDataTutorial
15. Bizer C, Heath T, Berners-Lee T (2009) Linked Data—the story so far. Int J Semantic Web Inf Syst 5(3):1–22. doi:10.4018/jswis.2009081901
16. Bizer C, Lehmann J, Kobilarov G, Auer S, Becker C, Cyganiak R, Hellmann S (2009) DBpedia—a crystallization point for the Web of Data. J Web Sem 7(3):154–165
17. Boardman R, Spence R, Sasse MA (2003) Too many hierarchies? The daily struggle for control of the workspace. In: Proceedings of HCI international 2003, vol 1, pp 616–620
18. Bürger T, Hausenblas M (2008) Interlinking multimedia—principles and requirements. In: First international workshop on interacting with multimedia content on the social Semantic Web, co-located with SAMT 2008, 3–5 December 2008
19. Bürger T, Hausenblas M (2010) Metadata standards and ontologies for multimedia content. In: Handbook of metadata, semantics and ontologies (to be published). World Scientific, Singapore
20. Bush V (1945) As we may think. Atl Mon 176(1):101–108. URL http://www.ps.uni-sb.de/~duchier/pub/vbush/
21. Carroll J, Dickinson I, Dollin C, Reynolds D, Seaborne A, Wilkinson K (2004) Jena: implementing the Semantic Web recommendations. In: Proceedings of the 13th international World Wide Web conference on alternate track papers and posters (WWW'04). New York, NY, USA
22. Cheng G, Ge W, Qu Y (2008) Falcons: searching and browsing entities on the Semantic Web. In: Proceedings of the 17th international conference on World Wide Web (WWW'09). Beijing, China
23. Corlosquet S, Delbru R, Clark T, Polleres A, Decker S (2009) Produce and consume Linked Data with Drupal! In: Proceedings of the 8th international Semantic Web conference (ISWC 2009)
24. Cyganiak R, Jentzsch A (2009) The linking open data dataset cloud. URL http://richard.cyganiak.de/2007/10/lod/
25. d'Aquin M, Baldassarre C, Gridinoc L, Angeletou S, Sabou M, Motta E (2007) Characterizing knowledge on the Semantic Web with Watson. In: Proceedings of 5th international EON workshop (EON2007) at the 6th international Semantic Web conference. Busan, Korea
26. Erling O, Mikhailov I (2007) RDF support in the virtuoso DBMS. In: CSSW, pp 59–68
27. Fielding RT (2000) Architectural styles and the design of network-based software architectures. PhD thesis, University of California, Irvine. URL http://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm
28. Frisse ME (1987) Searching for information in a hypertext medical handbook. In: Hypertext'87 proceedings, 13–15 November 1987. Chapel Hill, North Carolina, USA. ACM, pp 57–66
29. Gruber T (1993) A translation approach to portable ontology specifications. Knowl Acquis 5(2):199–220
30. Han L, Finin T, Parr C, Sachs J, Joshi A (2008) RDF123: from spreadsheets to RDF. In: 7th international Semantic Web conference (ISWC2008)
31. Harth A, Hogan A, Delbru R, Umbrich J, Decker S (2007) SWSE: answers before links. In: Semantic Web challenge at 6th international Semantic Web conference. Busan, Korea
32. Hartig O, Bizer C, Freytag JC (2009) Executing SPARQL queries over the Web of Linked Data. In: Proceedings of the 8th international Semantic Web conference (ISWC'09). Washington, DC, USA
33. Haslhofer B, Schandl B (2008) The OAI2LOD server: exposing OAI-PMH metadata as Linked Data. In: International workshop on Linked Data on the Web (LDOW2008). Beijing, China
34. Haslhofer B, Jochum W, King R, Sadilek C, Schellner K (2009) The LEMO annotation framework: weaving multimedia annotations with the Web. Int J Digit Libr 10(1):15–32
35. Haslhofer B, Momeni Roochi E, Gay M, Simon R (2010) Augmenting Europeana content with Linked Data resources. In: Linked Data triplification challenge, co-located with I-Semantics 2010. Graz, Austria
36. Hassanzadeh O, Consens M (2009) Linked movie data base. In: Proceedings of the 2nd international workshop on Linked Data on the Web (LDOW). Madrid, Spain
37. Hausenblas M, Bailer W, Bürger T, Troncy R (2007) ramm.x: deploying multimedia metadata on the Semantic Web. In: Proceedings of SAMT 2007, December 4–7. Genova, Italy

38. Hausenblas M, Troncy R, Raimond Y, Bürger T (2009) Interlinking multimedia: how to apply Linked Data principles to multimedia fragments. In: Proceedings of the 2nd international workshop on Linked Data on the Web (LDOW). Madrid, Spain
39. Hunter J (2009) Collaborative semantic tagging and annotation systems. In: Cronin B (ed) Annual review of information science and technology, vol 43, chap 2. American Society for Information Science & Technology
40. ISO/IEC (2006) Multimedia framework (MPEG-21)—part 17: fragment identification of MPEG resources. International Organization for Standardization, Geneva, Switzerland
41. Jacobs I, Walsh N (2004) Architecture of the World Wide Web, vol 1. URL http://www.w3.org/TR/webarch/
42. Kahan J, Koivunen MR (2001) Annotea: an open RDF infrastructure for shared Web annotations. In: WWW '01: proceedings of the 10th international conference on World Wide Web. ACM Press, New York, pp 623–632
43. Kobilarov G, Scott T, Raimond Y, Oliver S, Sizemore C, Smethurst M, Bizer C, Lee R (2009) Media meets Semantic Web—how the BBC uses DBpedia and Linked Data to make connections. In: Proceedings of the 6th European Semantic Web conference. Springer, Berlin, pp 723–737
44. Langegger A, Wöß W (2009) XLWrap—querying and integrating arbitrary spreadsheets with SPARQL. In: Proceedings of the 8th international Semantic Web conference (ISWC 2009). Springer
45. Langegger A, WößW, Blöchl M (2008) A Semantic Web middleware for virtual data integration on the Web. In: Proceedings of the European Semantic Web conference 2008. Tenerife
46. Lee W, Bürger T, Sasaki F, Malaise V, Stegmaier F, Söderberg J (2009) Ontology for media resource 1.0. W3C working draft, 18 June 2009
47. LODr—a linking open data tagging system. In: Workshop on social data on the Web (SDoW2008), vol 405. CEUR-WS.org
48. Marshall CC (2000) The future of annotation in a digital (paper) world. In: Harum S, Twidale M (eds) Successes and failures of digital libraries. University of Illinois, Urbana-Champaign, pp 97–117
49. Marshall CC, Brush BAJ (2004) Exploring the relationship between personal and public annotations. In: JCDL '04: proceedings of the 4th ACM/IEEE-CS joint conference on digital libraries. ACM Press, New York, pp 349–357
50. Network Working Group (NWG) (2005) RFC3986—uniform resource identifier (URI): generic syntax. URL http://www.gbiv.com/protocols/uri/rfc/rfc3986.html
51. Ovsiannikov IA, Arbib MA, McNeill TH (1999) Annotation technology. Int J Hum-Comput Stud 50(4):329–362
52. Raimond Y, Sandler MB (2008) A Web of musical information. In: Bello JP, Chew E, Turnbull D (eds) ISMIR, pp 263–268
53. Richardson L, Ruby S (2007) RESTful Web services. O'Reilly Media Inc
54. Rodríguez JB, Corcho O, Gómez-Pérez A (2004) R2O, an extensible and semantically based database-to-ontology mapping language. In: Proceedings of the 2nd workshop on Semantic Web and databases (SWDB2004)
55. Saathoff C, Scherp A (2009) M3O: the multimedia metadata ontology. In: Proceedings of the workshop on semantic multimedia database technologies (SeMuDaTe 2009)
56. Sack H, Baumann B, Groß A, Meinel C (2009) Linking tele-TASK video protal to the Semantic Web. In: Proceeding of the 9th international conference on innovative community systems (IICS), June 15–17. Jena, Germany
57. Sauermann L, Cyganiak R (2008) Cool URIs for the Semantic Web. W3C Semantic Web activity. Available at: http://www.w3.org/TR/cooluris/
58. Sauermann L, Bernardi A, Dengel A (2005) Overview and outlook on the semantic desktop. In: Proceedings of the 1st semantic desktop workshop, CEUR workshop proceedings, Galway, Ireland, vol 175
59. Schandl B, Popitsch N (2010) Lifting file systems into the Linked Data cloud with TripFS. In: Proceedings of the 3rd international workshop on Linked Data on the Web (LDOW2010). Raleigh, North Carolina, USA
60. Scharffe F, Euzenat J (2009) Alignments for data interlinking: analysed systems. http://melinda.inrialpes.fr/systems.html
61. Schroeter R, Hunter J (2007) Annotating relationships between multiple mixed-media digital objects by extending annotea. In: Proceedings of the 4th European Semantic Web conference, ESWC 2007

62. Sheth A, Klas W (1998) Multimedia data management: using metadata to integrate and apply digital media. Mcgraw-Hill Education, New York
63. Simon R, Sadilek C, Korb J, Baldauf M, Haslhofer B (2010) Tag clouds and old maps: annotations as linked spatiotemporal data in the cultural heritage domain. In: Workshop on linked spatiotemporal data 2010, held in conjunction with the 6th international conference on geographic information systems (GIScience 2010). Zurich, Switzerland
64. Stegmaier F, Bailer W, Bürger T, Döller M, Höffernig M, Lee W, Malaise V, Poppe C, Troncy R, Kosch H, de Walle RV (2009) How to align media metadata schemas? Design and implementation of the media ontology. In: Proceedings of the workshop on semantic multimedia database technologies (SeMuDaTe 2009)
65. Tummarello G, Delbru R, Oren E (2007) Sindice.com: weaving the open Linked Data. In: Proceedings of the 6th international the Semantic Web and 2nd Asian conference on Asian Semantic Web conference. Busan, Korea
66. Volz J, Bizer C, Gaedke M, Kobilarov G (2009) Silk—A link discovery framework for the Web of Data. In: Proceedings of the 2nd international workshop on Linked Data on the Web (LDOW). Madrid, Spain
67. W3C (2003) XPointer framework. Available at: http://www.w3.org/TR/xptr-framework/
68. W3C (2004) RDF vocabulary description language 1.0: RDF schema. W3C Semantic Web activity—RDF core working group. Available at: http://www.w3.org/TR/rdf-schema/
69. W3C (2004) RDF/XML syntax specification (revised). W3C Semantic Web activity—RDF core working group. Available at: http://www.w3.org/TR/2004/REC-rdf-syntax-grammar-20040210/
70. W3C (2004) Resource Description Framework (RDF). W3C Semantic Web activity—RDF core working group, URL http://www.w3.org/RDF/
71. W3C (2008) SPARQL query language for RDF. W3C Semantic Web activity—RDF Data Access working group, URL http://www.w3.org/TR/rdf-sparql-query/
72. W3C (2009) OWL 2Web ontology language. W3C Semantic Web activity—OWL working group, URL http://www.w3.org/TR/owl2-overview/
73. W3C Media Fragments Working Group (2009) Media fragments URI 1.0. W3C working draft. Available at: http://www.w3.org/TR/media-frags/
74. W3C Semantic Web Deployment Group (2009) SKOS simple knowledge organization system reference. URL http://www.w3.org/TR/2009/REC-skos-reference-20090818/

Bernhard Schandl is a researcher and entrepreneur in the field of semantic technologies. He received his doctoral degree in computer science from University in Vienna in 2009, where he has carried out research on the Semantic Desktop, Semantic Web, Linked Data, and Multimedia, and has been actively involved in national and international research projects. He is co-founder of Gnowsis.com, designing software that helps people to remember.

Bernhard Haslhofer is a postdoctoral researcher at the University of Vienna since 2008. His research interests are in the area of Linked Data, with a special focus on data quality, data linking, and the exploitation of user-contributed data (annotations). He has contributed to the W3C Linking Open Data Initiative from the very beginning (2007) and served as chair, PC member, and reviewer for several conferences and journals in this field.

Tobias Bürger is currently a senior research at Salzburg Research in the Knowledge and Media Technologies Group. He holds a diploma in computer science from the University of Passau and a PhD (Dr. rer. nat.) from the University of Innsbruck. His research interests are in the area of multimedia semantics, Linked Data, and Social Media. He currently contributes to the W3C Media Annotation Working Group and is involved in several initiatives around Linked Data.

Andreas Langegger was working as a scientist at the Johannes Kepler University Linz, Austria from 2005 to 2010. He holds a doctoral degree in computer science and his current research interests include novel data management, query and analysis approaches, ontologies, information extraction, and semantic technologies. He has given many lectures, tutorials, invited talks, and he has published several research papers in popular journals and major international conferences. After his academic career he has started working for a major IT consulting company.

Wolfgang Halb is key researcher for semantic technologies at JOANNEUM RESEARCH in the research area of intelligent information systems. He holds a master degree in software development-economy from the Graz University of Technology and is working on a PhD thesis on applied Semantic Web technologies. He is actively contributing to a range of projects related to the application of semantic technologies and is currently member of the W3C Library Linked Data Incubator Group and RDB2RDF Working Group.