

select * where { :I :trust :you }

How to Trust Interlinked Multimedia Data

Wolfgang Halb and Michael Hausenblas

Institute of Information Systems & Information Management,
JOANNEUM RESEARCH, Steyrergasse 17, 8010 Graz, Austria
firstname.lastname@joanneum.at

Abstract. Finding, accessing and consuming multimedia content on the Web is still a challenge. In this position statement we discuss three still widely neglected issues arising when one is interacting with multimedia content in social media environments: provenance, trust, and privacy. We will introduce a generic model allowing us to identify potential risks and problems, further discuss this model regarding multimedia content and finally outline how Semantic Web technologies can help.

1 Motivation

It is a trivial truth that, in order to use any kind of content or service on the Web, one must know how to access it (that is, to know the URI). What is true for the Web of Documents is equally true for the Web of Data. While with the rise of linked data [1, 2] the situation has changed—publishing and consuming data is now possible straight forward—there are still a number of issues in the discovery process. For example, with our multimedia interlinking demonstrator CaMiCatzee [3] we have identified issues around trust and believe of information regarding linked data in general and multimedia content in special. CaMiCatzee allows the FOAF-profile-based search for person depictions on flickr. However, when looking at Fig. 1, how can we find out if one of the depicted persons actually is Sir Tim Berners-Lee?




Fig. 1. Exemplary result from CaMiCatzee.

*select * where { :I :trust :you }*

Motivated by this observation we will address—from a practical point of view—widely neglected issues arising when one is interacting with multimedia content in social platforms:

- **Provenance.** Where does content stem from? Who provided annotations?
- **Trust.** Is a person that provided an annotation trustworthy? Is the interlinking eligible?
- **Privacy.** When interacting with content—what are the consequences?

In the following, we will first introduce a generic model addressing the three above listed issues. Based on our experience with interlinked multimedia data we will have a detailed look at the consequences when this model is applied to multimedia content.

2 The Abstract Provenance-Trust-Privacy Model

In order to identify issues with the usage of content on the Web, we have developed the “Provenance-Trust-Privacy” (PTP) model (Fig. 2). Basically, two aspects are covered by this model: the real life and the online world, the Web. In the PTP model we deal with three orthogonal, nevertheless interdependent dimensions, (i) the **social** dimension, (ii) the **interaction** dimension, and (iii) the **content** dimension.

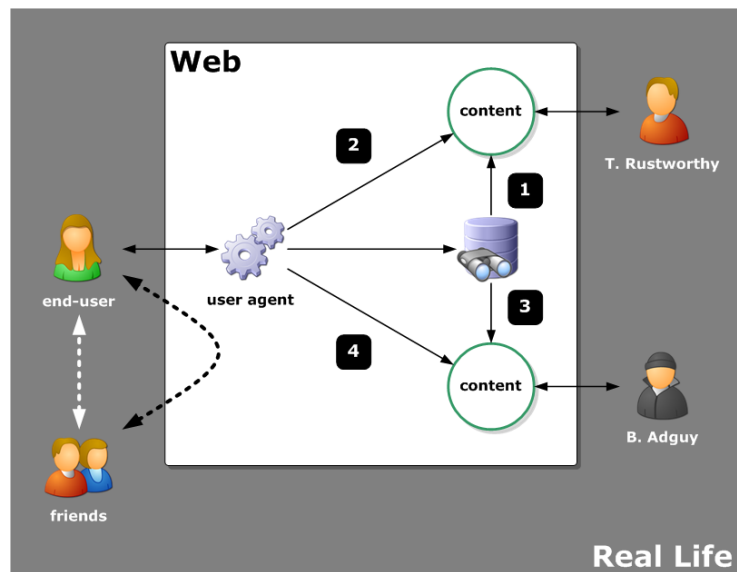


Fig. 2. The abstract PTP model.

`select * where { :I :trust :you }`

The Social Dimension. The dotted arrows in Fig. 2 represent social relations between humans, either in the real life or online. While it is straight-forward to deal with the case when people know each other in real life (and maybe continue this relation in the online world), the other way round can cause substantial problems. For example, only because I have added someone to my “buddy list” on a social platform such as LinkedIn does not mean that I really know the person and that this person also (wants to) know me.

The Interaction Dimension. In the realm of the PTP model we understand the interaction dimension of taking place online-only. Again, referring to Fig. 2, everything that happens between the user-agent (being instructed by a human) and other participants (server, etc.) on the Web. In general, two interaction patterns for the discovery and access of resources on the Web can be observed:

- A **direct access** of the content. In this case the URI of the content is known in advance by the end-user who instructs her user-agent to access the content. The URI may originally stem from a newspaper advertisement, a friend may have pointed it out in an e-mail, etc.
- An **indirect access** of the content by means of consulting an intermediate such as a search engine, a recommendation system, etc.

Based on these two interaction patterns, four possible paths can be identified:

1. The user-agent, equipped with the end-user’s profile and her desire consults an intermediate. For example, a user enters a search string into Google and is presented a list of URIs. The end-user happens to select a trustworthy source.
2. The user-agent, equipped with an URI from the end-user, accesses a trustworthy source.
3. The user-agent, equipped with the end-user’s profile and her desire consults an intermediate. This time, the end-user happens to select a troublesome source.
4. The user-agent, equipped with an URI from the end-user, accesses a troublesome source.

Obviously, the first two situations are desirable. The end-user has—for example based on previous experiences or trust in the search results—found some content that she can use and which is not causing her troubles (a virus, Trojan horse, etc.). However, equally, we are after avoiding the latter two cases where the end-user actually finds herself using content and/or services that are harmful and/or violate her privacy.

The Content Dimension. Regarding the content dimension one generally can state that the more is known about the content, the easier it is to assess its usefulness and its capabilities regarding a potential damage. Wherever possible, we are after self-descriptive resources, that is we require a minimum level of metadata being available. In our case, we focus on multimedia content. In the next section we will therefore initially discuss this kind of content and along the metadata in greater detail.

3 Multimedia Content in the Provenance-Trust-Privacy Model

In the position paper at hand we focus on multimedia content. We will in the following discuss characteristics of spatio-temporal multimedia content in the context of the emerging interlinking multimedia effort¹. Further, we apply the above introduced PTP model to multimedia content and try to derive requirements for it.

Characteristics of Multimedia Content and Multimedia Metadata. Multimedia content has some specific characteristics that allow and/or request special treatment. We have reported on this in great detail elsewhere [4]. A basic observation, however, with impact on many parts of the interaction process is that with multimedia content we are dealing almost always with spatio-temporal dimensions.

From the prosumers point-of-view, multimedia content is cheap to produce and available in high volumes (mobile phones, etc.). Further, most of the current content in that regard is publicly and freely available (Flickr, youtube). Business models remain vague. On the other hand, for professionally created content for very specific domains such as broadcaster’s archives, adult content, etc. the fees are considerably high.

Multimedia content is in general good for consumption in mobile environments (as opposed to reading longish text on a mobile).

In general it is hard and expensive to create good and detailed multimedia content descriptions (for example in MPEG-7, etc.). This leads to a problem regarding the fine-grained search and automated summaries.

In Table 1 the above discussed characteristics are summarised, and weighted regarding the content itself on the one hand and the metadata on the other hand.

Issue	Content	Metadata	Remark
production (prosumers)	++	-	easy to produce high-volume content (e.g., mobile phone)
production (professionals)	++	-	esp. high-level semantic content descriptions expensive
consumption	++	-	easy to consume (also in mobile environments)
search	-	--	practically, only global descriptions are available
summaries	--	--	little automation possible

Table 1. Overview on Multimedia Content Characteristics.

Applying the PTP Model to Multimedia Content. With the above listed characteristics of multimedia content in mind we claim the following regarding the application of the PTP model:

- Any solution addressing PTP issues must at least avoid accessing “bad” content and should support the discovery and consumption of “good” content.

¹ <http://www.interlinkingmultimedia.info/>

*select * where { :I :trust :you }*

- Existing and deployed multimedia metadata formats (such as Exif, ID3, etc.) have to be taken into account.
- The solution at hand needs to scale to the size of the Web.
- It must be practically relevant in terms of availability in widely used platforms such as Drupal, MediaWiki, etc. (for example as a plug-in, etc.; however it needs to be integrated).
- Provider must be able to easily offer and administer it (e.g., “enable” it with little configuration effort).
- Consumer must be able to use it in a non-disruptive way, for example as a part of their everyday tools.

In the next section we will report on how Semantic Web technologies can be used in combination with other, deployed technologies (such as for identification and authentication) in order to address the above listed requirements.

4 How Semantic Web Technology Can Help

We strongly believe that Semantic Web technologies can help to realise a PTP model for multimedia content. Lots of research is already available², however with little practical impact to this end. Apart from avoiding unreliable or even malicious content, the main aim of applying PTP to multimedia content is to help the user in finding trustworthy information sources. In a first step, we consider all content created by a trusted person or authority to be trustworthy. Solving this issue implies that there need to be techniques that can ensure a content’s provenance and the content producer’s identity respectively. Consequently means have to be made available that can decide which person to trust or not.

In the case of multimedia content it also has to be taken into account that information associated with a single content item can potentially have a multitude of contributors. A photo along with some metadata (title, description) on Flickr for instance might be uploaded by the fictitious trusted user “T. Rustworthy”. Another user, “B. Adguy”, could add a fake note about who is depicted in the photo. When accessing the photo and the associated metadata it is thus not sufficient to only consider who contributed the image but we would also need to be able to figure out who contributed the metadata about it. Taking this further to video content it might also be relevant to take into account who contributed which parts of a video (consider, for example, advertisements inserted into a video stream).

In the following we will discuss already available technologies that may be able to address the PTP issues along the three identified dimensions. However, to date only isolated solutions exist and there is still a lack of systems that incorporate all available technologies for the user’s benefit. We envision a framework that would allow to combine the below listed technologies and develop plug-ins for widely used platforms (Flickr, Youtube, etc.) and systems (Drupal, etc.).

² <http://www4.wiwiwiss.fu-berlin.de/bizer/SWTSGuide/>

`select * where { :I :trust :you }`

Social Dimension. For the identification as well as for the authentication several technologies are available. A user can for example provide her OpenID³ to identify against a service. Further, OAuth⁴ can be used for publishing and interacting with protected data. Big players such as Google are already offering support for the above mentioned technologies⁵. It seems advisable to build on this and contemplate what might be missing to align it with the Web of Data, being based on RDF [5].

While FOAF-based white listing and other related approaches [6, 7] are available, there is still a need for an up-front agreed way to deploy it in widely used systems. The same issue can be observed with privacy: there are proposals on the table (for example P3P⁶) but no measurable uptake in terms of users, systems that offer it, etc. can be stated.

Interaction Dimension. Especially for data provenance it seems to us that named graphs [8] offer a solid and scalable solution. With the rise of RDFa⁷ one can think of new provenance mechanisms, as the hosting document can actually be understood as the “name” of the graph. Just imagine Flickr (already offering licensing information in RDFa) to include provenance information on both the content and the metadata, based on vocabularies such as the “Semantic Web Publishing Vocabulary” [9, Chapter 6]. Finally, we note that regarding the discovery and usage of linked (multimedia) data we are currently working on VoiD, the “Vocabulary of Interlinked Datasets”⁸—again, provenance and trust issues are in scope, here.

Content dimension. In our understanding, the content dimension of the PTP model requires the most attention. Basic mechanisms proposed to represent the type of multimedia content in RDF⁹ are available. Still, practical ways for creating and consuming rich multimedia content descriptions are missing. Recently, we have proposed ramm.x [10] which allows to use existing multimedia metadata formats such as MPEG-7, Exif, ID3, etc. in the Web of Data. However, we expect a fair amount of further research being required to address provenance and trust issues properly and make tools and applications available in a real-world setup.

References

1. C. Bizer, T. Heath, D. Ayers, and Y. Raimond. Interlinking Open Data on the Web (Poster). In *4th European Semantic Web Conference (ESWC2007)*, pages 802–815, 2007.
2. C. Bizer, R. Cyganiak, and T. Heath. How to Publish Linked Data on the Web. <http://sites.wiwiss.fu-berlin.de/suhl/bizer/pub/LinkedDataTutorial/>, 2007.

³ <http://openid.net/>

⁴ <http://oauth.net/>

⁵ <http://googledataapis.blogspot.com/2008/06/oauth-for-google-data-apis.html>

⁶ <http://www.w3.org/P3P/>

⁷ <http://www.w3.org/TR/xhtml-rdfa-primer/>

⁸ <http://community.linkeddata.org/MediaWiki/index.php?VoID>

⁹ <http://www.w3.org/TR/Content-in-RDF/>

`select * where { :I :trust :you }`

3. M. Hausenblas and W. Halb. Interlinking Multimedia Data. In *Linking Open Data Triplification Challenge at the International Conference on Semantic Systems (I-Semantics08)*, 2008.
4. T. Bürger and M. Hausenblas. Why Real-World Multimedia Assets Fail to Enter the Semantic Web. In *International Workshop on Semantic Authoring, Annotation and Knowledge Markup (SAAKM07)*, Whistler, Canada, 2007.
5. G. Klyne, J. J. Carroll, and B. McBride. RDF/XML Syntax Specification (Revised). W3C Recommendation, RDF Core Working Group, 2004.
6. J. Golbeck. Combining Provenance with Trust in Social Networks for Semantic Web Content Filtering. In *Provenance and Annotation of Data, International Provenance and Annotation Workshop, IPAW 2006, Chicago, IL, USA, May 3-5, 2006, Revised Selected Papers*, volume 4145 of *Lecture Notes in Computer Science*, pages 101–108. Springer, 2006.
7. A. Harth, A. Polleres, and S. Decker. Towards a social provenance model for the Web. In *Workshop on Principles of Provenance (PrOPr)*, Edinburgh, Scotland, 2007.
8. J. J. Carroll, C. Bizer, P. Hayes, and P. Stickler. Named Graphs, Provenance and Trust. In *WWW 05: Proceedings of the 14th international conference on World Wide Web*, pages 613–622. ACM Press, 2005.
9. C. Bizer. *Quality-Driven Information Filtering in the Context of Web-Based Information Systems*. PhD thesis, Freie Universität Berlin, 2007.
10. M. Hausenblas, W. Bailer, T. Bürger, and R. Troncy. Deploying Multimedia Metadata on the Semantic Web (Poster). In *2nd International Conference on Semantics And digital Media Technologies (SAMT 07)*, 2007.