

# Hyperdata: Update APIs for RDF Data Sources (Vision Paper)\*

Jacek Kopecký

Knowledge Media Institute, The Open University, UK  
j.kopecky@open.ac.uk

**Abstract.** The Linked Data effort has been focusing on how to publish open data sets on the Web, and it has had great results. However, mechanisms for updating linked data sources have been neglected in research. We propose a structure for Linked Data resources into named graphs, connected through hyperlinks and self-described with light metadata, that is a natural match for using standard HTTP methods to implement application-specific (high-level) public update APIs.

## 1 Vision

A major function of Web APIs is to give users a way to contribute to data sources (whether they be called social networks, photo sharing sites, or anything else) through rich scripted web sites, rather than through simple web forms, and also through external (even 3rd-party) tools. Facebook API, Flickr API and so on, support interactive Web interfaces as well as mobile apps or desktop tools.

Some of the data in these apps then gets published as Linked Data, a machine-friendly representation suitable for combining with other data. Commonly, there is a technologies disconnect, though, between the Linked Data read-only view on the data source (which employs RDF and URIs), and the update APIs (with JSON or XML, and non-URI identifiers).

In this paper, we describe a vision of *hyperdata*<sup>1</sup> — data that is not only hyperlinked and self-describing in terms of its schema, but also self-describing on how it can be updated.

As we've discussed in [1], update access cannot practically be provided through protocols such as SPARQL Update. Indeed, public update access should be through a data-source-specific application layer that enforces consistency and security. There are several reasons for this: 1) data dependencies, where an update needs to propagate into dependent data, 2) security, where low-level access policies for RDF stores are harder to manage than if they were policies on the level of application-specific resources, 3) data constraints and validation, for example to guide the users in the structure of the accepted data (for example preventing well-meaning users from using the wrong ontology by mistake), and

---

\* This work has been funded by the European projects SOA4All ([www.soa4all.eu](http://www.soa4all.eu)) and PlanetData ([www.planet-data.eu](http://www.planet-data.eu)).

<sup>1</sup> The term “hyperdata”, which predates the Web, has been used in connection to the Web of Data, for example in <http://www.novaspivack.com/technology/the-semantic-web-collective-intelligence-and-hyperdata>.

4) creation of identifiers, because SPARQL Update does not provide the equivalent of an `AUTO_INCREMENT` field in an SQL database, and leaving the creation of identifiers to clients is undesirable due to the potential for conflicts.

With self-describing read-write hyperdata, applications that consume Linked Data can easily add update functionalities, currently generally missing from mash-ups and other Linked-Data-based apps. Further, data browsers such as Tabulator, which currently supports SPARQL Update and WebDAV [2], will be able to provide edit/update capabilities over a wider range of resources.

We may compare Linked Data to Web 1.0: the latter was mostly read-only documents, and the former is mostly read-only RDF views on some databases. The Web of Data should be more like Web 2.0, with many sites allowing (and even relying on) contributions from their users. With hyperdata, that is possible, because hyperdata is not only linked to other data, but also to its update APIs.

In our vision, the optimal update API should fit well with the structure of Linked Data (including the principle of *following your nose* to discover update capabilities), it should rely as much as possible on the methods of HTTP (adhering to REST's *uniform interface* constraint), and it should easily accommodate application-specific update authorization, validation and propagation logic.

## 2 Use Case Description

Our hyperdata approach was developed within a use case of the SOA4All project, an application called “Offers4All”, which allows diverse companies to advertise offers to subscribers of the service (more detail in [1]). These offers might be “last-minute” travel deals, predefined campaign offers of restaurants, and so on.


The application is backed by an RDF database that stores information about offer providers, their offers, and the users registered to receive the offers. Users can specify what offer categories they are interested in, and they can also choose to “like” an offer which allows social-networking-style recommendations to be used to increase the uptake of offers.

For read and update access, the database is façaded by a custom API, whose functionalities can be seen as the following types of operations:

- `listUsers()` returns a list of the known users, `getUser(id)` returns user data
- `addUser(data)` creates a new user record
- `getUserInterests(id)` returns the offer categories of interest to the user
- `addUserInterest(id, uri)`, `deleteUserInterest(interest-id)`
- `deleteAllUserInterests(id)` clears the list of interests
- *and so on for the various properties of the various objects in the database.*

The granularity of these operations corresponds to the intended uses of the system: these are the types of operations that clients of such a database want to perform, and they are a good input for analyzing access control.

Following the principles of Linked Data and REST (useful even if the data is not published openly), the database is split into a number of resources: a single container *users resource*, multiple *user resources* (one per known user), container *user interests resources* (one per known user), and concrete *interest value resources* (one per a stated interest of a user), etc.



**Fig. 1.** Hyperdata structure of the API

In a read-only data source, this fine level of granularity could be seen as too much, as retrieving all the data about a user does not present much overhead even if the client is only interested in the user’s interests. However, with all these resources in place, update operations naturally map to HTTP methods.

Figure 1 shows the RDF graph of a user who likes two specific offers and has interest in one category. For brevity, the figure doesn’t show the container resource for users. Along with the actual data triples, the figure also displays the self-description aspects, discussed in the next section.

### 3 Hyperdata Approach

The API in our use case consists of the following generic four types of resources: 1) containers of instances (users, offers etc.), 2) the instances themselves, 3) containers of property values, 4) concrete property values. Listing 1 illustrates the self-description metadata and hyperlinks, also shown in Figure 1; it starts on line 1 with (a subset of) the actual data about the particular user.

Line 3 indicates the graph that is the description of the user instance, making it possible for a client to infer that an HTTP DELETE request can remove the instance. Line 4 links the instance graph with one of the property graphs (`/users/1345/likes`), and with the high-level class graph.

```

1 </users/1345#this> a uc:User ;          uc:likes </offers/439>, </offers/637> .
2
3 </users/1345> a g:Graph ;              g:defines </users/1345#this> ;
4   g:contains </users/1345/likes> ;    g:isContainedIn </users> .
5
6 </users/1345/likes> a g:Graph ;
7   g:contains </users/1345/likes/43905>, </users/1345/likes/43906> ;
8   g:defines [ a rdf:Statement ;
9     rdf:subject </users/1345#this> ; rdf:predicate uc:likes ; rdf:object [ ]
10 ] .
11 </users/1345/likes/43905> a g:Graph ;
12   g:defines [ a rdf:Statement ;
13     rdf:subject </users/1345#this> ; rdf:predicate uc:likes ; rdf:object </offers/439>
14 ] .

```

**Listing 1.** Example graph description triples (truncated)

Lines 6–10 describe the property graph: it contains concrete value graphs, and a reified triple pattern (lines 8–10) that indicates that the graph includes statements of the form `/users/1345#this uc:likes something` (note the blank node as object). The triple pattern is meant to indicate what kind of data can be POSTed to the property resource to add a value, and what subset of the data about the user can be expected when GETting the property resource.

Finally, lines 11–14 describe a concrete value graph. The client can use PUT or DELETE here to update or remove a particular statement.

The metadata uses a few very simple concepts to communicate much information: a **Graph** is a resource that besides GET may also accept update and delete requests (actually available methods can be discovered with HTTP OPTIONS).

The meaning of updates depends on the contents of the graph, described through reified statements. The reified statement may indicate a concrete triple like on line 13 (meaning that it represents a specific value, to be updated with PUT or removed with DELETE), or it may use blank nodes to indicate a collection (accepting POST with new items). The listing indicates a collection of property values for `uc:likes` on line 9. A reified statement of the form `something rdf:type uc:User` would be shown on the user container graph to indicate that it contains instances of the given ontology class, and that’s what can be POSTed.

## 4 Conclusion

The Web included update capabilities from the start—the first browser<sup>2</sup> was also an editor—but still Web 1.0 was mostly read-only. A significant boom came with the advent of the Web 2.0, which embodies the attitude that anybody on the Web can—and should be allowed to—contribute.

The Web of Data so far remains on the Web 1.0 level where contributions to it must happen outside it. Hyperdata linked APIs can bring update capabilities to the Web of Data, and make it more like Web 2.0. That’s what gave us Wikipedia, after all, which is now the center point of Linked Data.

We have developed a proof-of-concept triple-store wrapper (also described in [1]) that uses very simple configuration to realize the hyperdata API. Such a wrapper can easily provide hooks for access control, validation and other data-source-specific processing of updates. As future work, we would like to develop a client access library for hyperdata, and to extend Tabulator to support it.

An interesting open question: can the links and metadata help clients adapt to changes in evolving hyperdata APIs?

## References

1. Kopecký, J., Pedrinaci, C., Duke, A.: RESTful Write-oriented API for Hyperdata in Custom RDF Knowledge Bases. In: Proceedings of the International Conference on Next Generation Web Service Practices (NWeSP), Salamanca, Spain (2011)
2. Berners-Lee, T., Hollenbach, J., Kanghao Lu, Presbrey, J., Prud’hommeaux, E., mc schraefel: Tabulator Redux: Browsing and Writing Linked Data. In: Proceedings of the WWW 2008 Workshop on Linked Data on the Web, Beijing, China (2008)

---

<sup>2</sup> <http://www.w3.org/People/Berners-Lee/WorldWideWeb.html>