

Semantic Description of REST APIs

Ruben Verborgh¹, Andreas Harth², Maria Maleshkova²,
Steffen Stadtmüller² Thomas Steiner³, Mohsen Taheriyani⁴, and
Rik Van de Walle²

¹ Ghent University – iMinds, ELIS – Multimedia Lab, Belgium

² Institute AIFB, Karlsruhe Institute of Technology, Germany

³ LSI Department, Universitat Politècnica de Catalunya, Spain

⁴ Information Science Institute, University of Southern California, USA

1 Introduction

The REST architectural style assumes that client and server form a contract with content negotiation, not only on the data format but implicitly also on the semantics of the communicated data, *i.e.*, an agreement on how the data have to be interpreted [55]. In different application scenarios such an agreement requires vendor-specific content types for the individual services to convey the meaning of the communicated data. The idea behind vendor-specific content types is that service providers can reuse content types and service consumers can make use of specific processors for the individual content types. In practice however, we see that many RESTful APIs on the Web simply make use of standard non-specific content types, *e.g.*, *text/xml* or *application/json* [33]. Since the agreement on the semantics is only implicit, programmers developing client applications have to manually gain a deep understanding of several APIs from multiple providers.

Common Web APIs are typically either exclusively described textually⁵, or far less frequently—and usually based on third-party contributions—a machine-readable WADL [18] API description exists⁶. However, neither human-focused textual, nor machine-focused WADL API descriptions carry any *machine-processable* semantics, *i.e.*, do not describe *what* a certain API does. Instead, they limit themselves to a description of machine-readable in- and output parameters in the case of WADL, or a non-machine-readable prose- and/or example-driven description of the API in the case of textual descriptions. While this may suffice the requirements of developers in practice, the lack of semantic descriptions hinders many more advanced use cases such as API discovery or API composition.

Machine-interpretable descriptions can serve several purposes when developing client applications:

- One can generate textual documentation from the standardised machine-interpretable descriptions, which leads to a more coherent presentation of the APIs, similar to what JavaDoc has achieved in the Java world (see also Knuth’s idea of literal programming [22]).

⁵ For example, the Twitter REST API: <https://dev.twitter.com/docs/api>

⁶ A large archive of WADL descriptions is available on GitHub: <https://github.com/apigee/wadl-library>

- A standardised way to access REST APIs introduce a higher degree of automation for high level tasks such as composition.
- Machine-interpretable descriptions facilitate a more structured approach to developing APIs, which means that automated tools can check coherence and the RESTful-ness of an API (e.g., according to the Richardson Maturity Model⁷).

In a RESTful interaction with Web resources, only the constraint set of HTTP methods can be applied to the resources. The semantics of the HTTP methods itself is defined by the IETF [12] and do not need to be explicitly described for individual resources. We can distinguish between safe and non-safe methods, where safe methods guarantee not to affect the current state of resources. Additionally, some of the methods require additional input data to be provided for their invocation. The communicated input data can be subject to requirements that need to be described to allow an automated interaction. Furthermore, the effect on the resources state of an application of a non-safe method has needs to be assessable before the actual invocation to allow clients to decide how to interact with the resources. The effected change of resources after applying an HTTP method can also depend on the communicated input data. This dependency between communicated input and the resulting state of resources has also to be subject of a description.

In the chapter we classify the various approaches for providing machine-interpretable descriptions of Web APIs. Chapter 2 surveys lightweight semantic descriptions. Chapter 3 introduces descriptions based on graph patterns (a subset of the SPARQL query language for RDF, a graph-structured data format). Chapter 4 covers logic-based descriptions. Chapter 5 covers JSON-based descriptions. Chapter 6 contains a description of two tools for annotating existing APIs, and Chapter 8 concludes.

2 Lightweight semantic descriptions

2.1 Syntactic REST API descriptions

Web services enable the publishing and consuming of functionalities of existing applications, facilitating the development of systems based on decoupled and distributed components. **WSDL** [58], [7] is an XML-based language for describing the interface of a Web service. A WSDL service description specifies: 1) the supported operations for consuming the Web service; 2) its transport protocol bindings; 3) the message exchange format; and 4) its physical location. In this way, the WSDL description contains all information necessary for invoking a service and since it is XML-based, conforming to the WSDL XML schema, it is also machine-processable.

WSDL Similarly, to Web services, which provide access to the functionality of existing components, Web APIs and Web applications conforming to the REST

⁷ <http://martinfowler.com/articles/richardsonMaturityModel.html>

paradigm, provide access to resources by using the WWW as an infrastructure platform. Based on this parallel between the two types of services, WSDL was extended to Version 2.0 [58], which can also be used for formally describing RESTful services. As a result WSDL is a machine-processable, platform- and language-independent form of describing Web services and RESTful services alike.

The difficulty of using WSDL for RESTful services, is that it was not especially designed for resource-oriented services and as a result, everything has to be described in an operation-based manner. In addition, WSDL introduces some difficulties with specifying the message exchange format and limits HTTP authentication methods. Moreover, the most important drawback is that it lacks support for simple links. There is no mechanism in WSDL 2.0 to describe new resources that are identified by links in other documents. However, one of the most important characteristic of RESTful services is that they consist of collections of interlinked resources. Finally, the adoption of WSDL as means for describing Web APIs would require that all providers update or completely change their websites with documentation, moving away for using only text in HTML form. Similarly developers would need to learn to deal with WSDL instead of simply reading a natural language description.

WADL In contrast to WSDL, the Web Application Description Language (WADL, [18]) was especially designed for describing RESTful services in a machine-processable way. It is also XML-based and is platform and language independent. As opposed to WSDL, WADL models the resources provided by a service, and the relationships between them in the form of links. A service is described using a set of resource elements. Each resource contains descriptions of the inputs and method elements, including the request and response for the resource. In particular, the request element specifies how to represent the input, what types are required and any specific HTTP headers. The response describes the representation of the service's response, as well as any fault information, to deal with errors.

Currently, neither WADL nor WSDL are widely accepted and used for Web APIs and RESTful services. A Google search for WADL files returns only 49 unique results ⁸, while from the popular Web 2.0 applications only delicious ⁹ and YahooSearch ¹⁰ have WADL descriptions. The main difficulty of using WADL descriptions is that they are complex, in comparison to text-based documentation, and require that developers have a certain level of training and tool support that enables the application development on top of WADL. This complexity contradicts with the current proliferation of Web APIs, which can be greatly attributed to simplicity and direct used of the infrastructure provided by the Web, which enable the easy retrieving of resources only thought an HTTP request, directly in the Web browser.

Web APIs evolve rather autonomously without conforming to a shared set of guidelines or standards, especially evident by the fact the documentation is usually given in natural language as part of a webpage. The developer has

⁸ Search done on October 5th, 2009

⁹ <http://delicious.com/>

¹⁰ <http://search.yahoo.com/>

to decide what structure to use and what information to provide. As a result, everyone who is able to create a Web page is also able to create a Web API description. However, plain text/HTML descriptions, in contrast to WSDL and WADL descriptions, are not meant for automated machine-processing, which means that if developers want to use a particular service, they have to go to an existing description Web page, study it and write the application manually. Therefore, current research proposes the creation of machine-interpretable descriptions on top of existing HTML descriptions by using **microformats** [19]. Microformats offer means for annotating human-oriented Web pages in order to make key information automatically recognisable and processable, without modifying the visualization or the content.

hRESTS One particular approach for creating machine-processable descriptions for RESTful services by using microformats is hRESTS (HTML for RESTful Services) [24]. hRESTS enables the marking of service properties including operations, inputs and outputs, HTTP methods and labels, by inserting HTML tags within the HTML. In this way, the user does not see any changes in the Web page, however, based on the tags, the service can be automatically recognized by crawlers and the service properties can directly be extracted by applying a simple XSL transformation. The result is an HTML page that also contains the syntactical information of the described Web API and therefore, no longer relies solely on human interpretation.

RDFa An alternative to using hRESTS is offered by RDFa [1] that enables the embedding of RDF data in HTML. RDFa is similar to using microformats, but is somewhat more complex and offers more HTML marking options, as opposed to hRESTS. Approaches, based on making existing RESTful service descriptions machine-processable by using HTML tags are simpler and more lightweight as opposed to WSDL and WADL. In addition, as already mentioned, they can be applied directly on already available descriptions, rather than creating new service descriptions from scratch. The adoption by developers is also easier, since the creation of a machine-processable RESTful service description is equivalent to Web content creation or modification.

2.2 MicroWSMO / SA-REST

In contrast to research in the area of Semantic Web Services, which has been quite prolific, the number of semantic frameworks targeted at capturing Web API characteristics is relatively limited. Web APIs have only recently achieved greater popularity and wider use, thus raising the interest of the research community. In this section we discuss the two main approaches – MicroWSMO and SA-REST, aiming to support a greater level of automation of common service tasks through employing semantics. We also consider further description languages and ontologies, including ReLL and ROSM.

MicroWSMO MicroWSMO [23] is a formalism for the semantic description of Web APIs, which is based on adapting the SAWSDL [11] approach. MicroWSMO uses microformats for adding semantic information on top of HTML service documentation, by relying on hRESTS for marking service properties and making

the descriptions machine-processable. It uses three main types of link relations: 1) **modelReference**, which can be used on any service property to point to appropriate semantic concepts identified by URIs; 2) **liftingSchemaMapping** and 3) **loweringSchemaMapping**, which associate messages with appropriate transformations (also identified by URIs) between the serialization format such as XML and a semantic knowledge representation format such as RDF. Therefore, MicroWSMO, based on hRESTS, enables the semantic annotation of Web APIs in the same way in which SAWSDL, based on WSDL, supports the annotation of Web services.

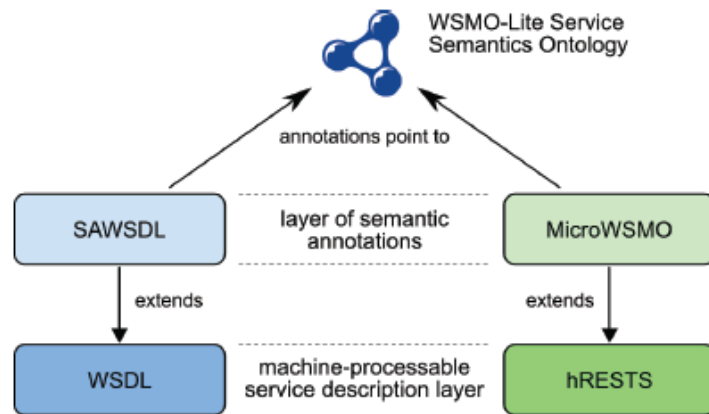


Fig. 1. Unifying SAWSDL and MicroWSMO through WSMO-Lite

In addition, MicroWSMO can be complemented by the WSMO-Lite service ontology specifying the content of the semantic annotations (see Figure 1¹¹). Since both Web APIs and WSDL-based services can have WSMO-Lite annotations, this provides a basis for integrating the two types of services. Therefore, WSMO-Lite enables unified search over both Web APIs and WSDL-based services, and tasks such as discovery, composition and mediation can be performed based on WSMO-Lite, completely independently from the underlying Web service technology (WSDL/SOAP or REST/HTTP).

SA-REST Another formalism for the semantic description of RESTful services is SA-REST [42], which also applies the grounding principles of SAWSDL but instead of using hRESTS relies on RDFa [1] for marking service properties. Similarly to MicroWSMO, SA-REST enables the annotation of existing HTML service descriptions by defining the following service elements: input, output, operation, lifting, lowering, or fault and linking these to semantic entities. The main differences

¹¹ <http://www.wst.univie.ac.at/workgroups/sem-nessi/index.php?t=semanticweb>

between the two approaches are not the underlying principles but rather the implementation techniques. In addition, MicroWSMO aims to add semantics as means for Web service automation, while SA-REST is more oriented towards enabling tool support in the context of service discovery and composition by mashup or smashup [41].

2.3 Minimal service model

The Minimal Service Model (MSM) represents an operation-based approach towards describing Web APIs. It is a simple RDF(S) ontology that supports the annotation of common Web API descriptions. It also aims to enable the reusability of existing Semantic Web service approaches by capturing the maximum common denominator between existing conceptual models for services. Additionally, as opposed to most semantic Web services research up to date, MSM targets to support both traditional Web services, as well as Web APIs with procedural view on resources, so that they can be handled in a unified way.

Originally the MSM was introduced together with hRESTS [24], aiming to cover for the fact that Web API descriptions do not typically have any structure in terms of the resources handled nor the operations exposed. Although its initial purpose was to provide structure to hRESTS, it has been subsequently adjusted and updated to its current form, in order to provide means for integrating heterogeneous services (*i.e.*, WSDLs and Web APIs), and together with WSMO-Lite it has been used as a means to facilitate a common framework covering the largest common denominator of the most used Semantic Web service formalisms on the Web. On the basis of MSM and WSMO-Lite, generic publication and discovery machinery has been developed that supports SAWSDL, WSMO-Lite, hRESTS/MicroWSMO, and OWL-S services [38]. The Semantic Web Services exposed by this infrastructure [37] put the emphasis on reducing the complexity of conceptual models and integrating services with existing Linked Data [6]. This integration serves both as a means to simplify the creation and management of Semantic Web Services through reuse, as well as it provides a new view over semantic Web services understood as a means to support the generating and processing Linked Data on the Web. Subsequent work around the Minimal Service Model is focused on supporting the invocation and authentication of Web APIs [31,32].

The original MSM [24] defined a Web API in terms of a *Services* that has a number of *Operations*, which have an *Input*, an *Output*, and *Faults*. However, the original MSM, which was used as a basis for SA-REST [42] and MicroWSMO [23], fails to capture some significant parts of the descriptions, such as optional, default and mandatory parameters, which can have a crucial effect on discovery and invocation. In addition, it does not enable the description of the inputs, parts of the input, and parts of the parts of the input. As a result, the MSM has been extended [38] to support further service characteristics.

The MSM, in its current version [38], is visualized in Figure 2. The MSM, denoted by the *msm* namespace, defines *Services* which have a number of *Operations*. *Operations* in turn have input, output and fault *MessageContent* descriptions. A *MessageContent* may be composed of *MessageParts*, which can be mandatory

their associated aspects (*e.g.*, parameters or response codes). In summary, ROSM represents a simple ontology for describing resource-centered services, in terms of resources, collection of resources, addresses and HTTP methods. The specific semantics of the ROSM annotations can be described by linking to the WSMO-Lite ontology.

ServONT is an ontology-based hybrid approach where different kinds of match-making strategies are combined together to provide an adaptive, flexible and efficient service discovery environment [9]. Semantic-enriched frameworks are considered a key issue to enforce timely discovery and dynamic composition of services. In contrast, the ServFace project aims at creating a model-driven service engineering methodology for an integrated development process for service-based applications¹³. Finally, the Simple Semantic Web Architecture and Protocol (SSWAP) is the driving technology for the iPlant Semantic Web Program¹⁴. It combines Web service functionality with an extensible semantic framework to satisfy the conditions for high throughput integration [15]. SSWAP originates from the Semantic MOBY project, which is a branch of BiOMOBY project [57]. Using SSWAP, users can create scientific workflows based on the discovery and execution of Web services and RESTful services.

3 SPARQL-based descriptions

Following the motivation to look beyond the exposure of fixed datasets, an extension of Linked Data with REST technologies has been proposed and explored for some time [3,56,25,44,53,36]. These approaches extend the traditional use of HTTP in Linked Data, consistent with REST, by allowing all HTTP operations to be applied to Linked Data resources. REST services in this context are often perceived RDF prosumers, *i.e.*, the state of resources is made available encoded in RDF, at least as an alternative via Content Negotiation. Clients can interact with Linked Data resources by submitting RDF input data with HTTP methods (*e.g.*, POST, PUT) resulting in the manipulation or creation of resources. The output data a client receives after the successful submission of data describes the effected state change of resources (*i.e.*, their new content after creation or manipulation) and is serialized in RDF as well.

Such REST services contribute to the Web of Data by interlinking output data with existing Linked Data sets.

RESTful Linked Data resources consider how the data that results from computation over input can carry explicit semantics and base their service descriptions on the notion that Linked Data provides a description for resources' input and output requirements: the graph patterns provided by the SPARQL query language or the N3 notation. Graph pattern provide the advantage of a more thorough description of what should be communicated, familiarity to Linked Data producers and consumers, and the possibility for increased tool support.

¹³ <http://www.servface.eu/>

¹⁴ <http://sswap.info/>

The rationale behind the descriptions for RESTful Linked Data resources is that the current state of a resource can be retrieved with an `HTTPGET`, while the data exchange that constitutes a manipulating interaction with a resource is described with two graph patterns:

- A graph pattern that describes the RDF input data that is submitted to a resource (*e.g.*, with `HTTPPOST`), which is necessary to invoke the manipulation.
- A graph pattern that describes the RDF output data the client receives after a successful call.

The description implies that the input pattern has to match the input data to invoke the service and the output pattern will match the output data returned by the service.

To illustrate the use of graph pattern-based IO descriptions we introduce an example based on an API of a social network platform. A common feature of social networks is to post a message to the timeline of a user. The approach in the context of Linked Data would be to wrap this call to provide the information in RDF, reusing existing vocabularies. Here the natural choice is the `SIOC` vocabulary¹⁵.

The description of a resource representing the timeline of a user, would make explicit as required input for a `HTTPPOST` call the content of the message and its creator identified in that social network. Furthermore the output would be described as the created message with enriched with additional information (*e.g.*, the creation date). A client would receive data matching this pattern together with the HTTP status code indicating the success of the call. The resulting input and output descriptions for the Linked API are represented in Table 1.

```
Input:  { ?post a sioc:Post.
           ?post sioc:content ?content.
           ?post sioc:has-creator ?user. }
```

```
Output: { ?post sioc:content ?content.
            ?post dcterms:created ?date. }
```

Table 1. Example description for a timeline resource in a social network

Note that the reuse of the variables `?post` and `?content` across input and output implies that this variable will have the same binding in the output as provided in the input. In SPARQL terms these would be ‘safe variables’ if we considered the service to be a `SPARQLCONSTRUCT` query from the input to the output patterns. Variables in the output pattern that do not appear in the input pattern are bound by the functionality of the resource, *i.e.*, the binding is a result of the manipulation or creation of a resource.

Graph pattern based descriptions allow for a thorough descriptions of what a client has to communicate to successfully interact with a RESTful Linked

¹⁵ <http://sioc-project.org/ontology/>

Data resource. The output descriptions that share variables with the input descriptions allow clients to anticipate the result of their interaction with respect to the input they intended to provide. Such an anticipation is due to the circumstance that the actual output message of an interaction is intended to convey the effected state change of an interaction. The predictability of effects of manipulating actions is essential to enable (semi-)automated clients that use Linked Data REST services.

4 Logic-based descriptions

Since its inception, the Semantic Web has always had a strong link with logic [5]. The impact of logic is visible in essential building stones such as RDF [20], whose *open world* assumption prohibits conclusions based on the absence of certain triples. Two logic families are predominant on the Semantic Web: *description logic*, the underlying model of OWL [35], is the most widespread, followed by *first-order logic*, which is typically expressed in extensions of RDF with rules. Description logic is in essence a decidable fragment of first-order logic, at the cost of a loss in expressivity.

This reduced expressivity is a motivator for LOS and LIDS to adopt SPARQL, and the reason for a few other methods use rule-based expression languages. For example, in OWL-S [34], one of the early semantic description formalisms for traditional Web services, rule languages such as KIF or SWIRL have to be used to express various aspects that extend beyond the expressivity of RDF. These languages capture expressions for pre- and postconditions, results, and effects of a Web service invocation. In general, rules are a straightforward mechanism to express dynamic relationships, such as those that occur with Web APIs.

4.1 RESTdesc

RESTdesc [53,54] is a logic-based description method that focuses on exposing the functional aspect of Web APIs in machine-processable form. Concretely, RESTdesc descriptions are rules expressed in the Notation3 language (\mathbb{N}_3 , [4]), which adds support for variables and quantification to RDF. The latter is a prerequisite to natively describe statements such as “all requests to y result in x ”, which is not directly supported in RDF (*i.e.*, only through the use of modeling vocabularies). In fact, RESTdesc uses quantification to express functionality as follows:

$$\forall x \text{ preconditions}(x) \Rightarrow \exists r(\text{request}(r, x) \wedge \text{postconditions}(x))$$

Herein, the predicates are expressed as RDF graphs, called *formulae* in \mathbb{N}_3 . Reasoners interpret the above as “for every situation x where certain preconditions are met, there exists a specified request r that makes certain postconditions true.” Moreover, for specific situations x_n , reasoners can instantiate the above rule, thereby creating an RDF representation of a concrete HTTP request r_n , which can be executed by any RDF-compatible HTTP client.

Descriptions can be used to express the effects that occur as a result of a `POST` request, including the description of that request itself. Representations are *not* described, as `RESTdesc` aims to be representation-independent, but it can describe properties of these representations. This is not unlike the graph patterns used in SPARQL-based methods (see Section 3), but represented in a syntax that integrates the request description. `RESTdesc` descriptions can also be useful to describe requests with safe methods such as `GET`, in order to explain what properties the response will satisfy, again without constraining the representation. This can be useful in two cases: *a*) if the client does not want to retrieve the resource (for instance, if there are too many) or *b*) if the resource does not exist yet, but can be created as the result of another action. In both cases, clients can then predict certain properties of the resource, without actually accessing it.

Since `RESTdesc` descriptions are expressed as rules, they inherit the logical functionality, in particular the *chaining* property:

$$P \Rightarrow Q, Q \Rightarrow R \quad \vdash \quad P \Rightarrow R$$

Therefore, existing Semantic Web reasoners with `N3` support can create compositions of `RESTdesc`-described Web APIs by chaining `RESTdesc` descriptions [51], either in a forward or backward (goal-driven) manner. This enables agents to respectively discover possible actions and to find a sequence of actions to satisfy a predetermined goal.

Further work on `RESTdesc` includes the incorporation of quality parameters, possibly subjective, in Web API descriptions [52].

5 JSON-based descriptions

So far we have explored solutions that are directly based on Semantic Web technologies. However, many Web developers are reluctant to integrate `RDF`, `SPARQL` or `N3` in their applications. Lanthaler and Gütl provide three reasons for what they call *semaphobia* [27]. Firstly, they observe that the perceived complexity of the Semantic Web, in combination with its background in artificial intelligence, makes developers assume integration will be difficult. Furthermore, some of them are unsure whether the integration cost will be worthwhile. After all, the Semantic Web is sometimes regarded as a solution in search for a problem, suffering from the chicken-and-egg syndrome and thus still waiting for a *killer application* (although the `W3C` maintains a list of case studies and use cases [2]). Finally, the Semantic Web is often incorrectly considered a disruptive technology that is hard to implement in an evolving ecosystem [43].

Whatever the reasons for semaphobia, adoption is often a decisive factor for technologies. Therefore, description formats anchored on a technology Web developers are already familiar with have a head start. As of 2012, more than 44% of all APIs on ProgrammableWeb, the largest API index [10], communicate in `JSON`. `JSON` is the JavaScript Object Notation language, which quickly became popular on the Web as it was natively parsed by JavaScript, the only scripting language supported by the majority of browsers. Its simplicity and extensibility make it also an interesting target for many other environments [40].

5.1 SEREDASj

As the letter *j* in its name suggests, SEREDASj (Semantic RESTful Data Services, [27]) is a semantic description language format expressed in JSON. The motivation behind SEREDASj is to provide simpler descriptions (in contrast to OWL-S [34] or WSMO [30]) that contain necessary semantics (such as SA-REST and MicrowSMO, see Section 2.2) in a widely used machine-targeted media type (*application/json*). SEREDASj defines a syntactic structure for JSON documents and a corresponding interpretation. It enables the representation of hypermedia links, which are not natively present in JSON. It is, however, not a strictly validated approach such as JSON schema [60], which enforces the presence of certain elements and properties in a JSON document.

The authors of SEREDASj put forward three use cases [27]. Firstly, they envision it as a means of creating documentation, both in machine-processable and human-readable form. The human-readable counterpart is generated from labels of predicates defined in ontologies, which are referenced by URI in the SEREDASj description. The benefit here is reuse—on the one hand by having a single description for humans and machines, and on the other hand by referencing to existing ontological definitions. Secondly, the goal is to enable more flexibility in Web API clients. By adding support for hyperlinks to JSON, SEREDASj descriptions become a hypermedia format through which clients can navigate a Web API in accordance with the hypermedia constraint [13]. Thirdly, SEREDASj aims to facilitate data integration, as its annotations enable the transformation of JSON into RDF. However, as we will argue below, other means to this end exist, so employing SEREDASj specifically for this last use case might prove suboptimal.

Generally speaking, one SEREDASj description is created per resource type. As SEREDASj currently only supports JSON, it is assumed that resources of this type have at least a JSON representation. The accompanying SEREDASj description documents the elements in these JSON documents by mapping them to predicates of ontologies in RDF format. This principle can be applied to the whole hierarchy of the document, including arrays, which represent multi-valued properties. Every subtree can additionally be associated with an RDF type. Next to this, SEREDASj describes the controls to navigate through the Web API in the form of URI templates [17], as well as the RDF predicates they correspond to, capturing their meaning. SEREDASj can furthermore detail the format of entities for use in `PUT` or `POST` requests.

Part of the functionality of SEREDASj is currently offered by JSON-LD [28], whose specification is currently a W3C editor's draft [45]. JSON-LD similarly provides predicate and type annotations that allow JSON data to be translated into RDF, but these annotations are included in the JSON document itself as opposed to a separate SEREDASj description document. However, the question arises whether it would not be more beneficial for servers to provide separate JSON and RDF representations of a resource, allowing clients to indicate through content negotiation with which representation they would like to proceed.

Further work on SEREDASj includes an architecture to integrate Web APIs into the Web of Data, which makes use of SEREDASj descriptions [29].

6 Tools

In previous sections, different solutions for describing semantics of REST APIs have been investigated. However, there are some obstacles preventing them from being widely adopted. First, writing semantic service descriptions by hand is a time consuming and tedious task. Furthermore, to model APIs, most of these approaches require some degree of expertise in Semantic Web languages such as RDF, SPARQL, and N3 in addition to the domain knowledge. Tools can play a significant role by providing a user interface to rapidly build semantic descriptions, making the complexity of formal specification transparent to the user.

The here introduced formalisms, including MicroWSMO, SA-REST and the MSM, which make HTML service descriptions machine-processable and enable the adding of semantic information, provide the means for creating semantic descriptions of Web APIs. However, without supporting tools or guidelines, developers would have to modify and enhance the descriptions manually by using a simple text/HTML editor. In addition, the complete annotation process would have to be completed manually, if there are no tools, which enable the search for suitable domain ontologies or the reuse of annotations of previously semantically described services.

6.1 Karma

Karma¹⁶ [48] is a Web-based framework for integrating structured data from a variety of sources. Users can load data from relational databases, spreadsheets, delimited text files, KML (Keyhole Markup Language) files, and semi-structured Web pages. Users can clean and normalize data with a programming by example interface [59]. Then, Karma semi-automatically builds a semantic model of the source by mapping it a domain ontology chosen by the user [21]. Karma models each column in terms of the classes and data properties defined in the ontology, and models the relationships among columns using object properties. Once data is modeled, Karma can translate the data into a variety of formats including RDF. The semantic models also enable Karma to integrate information from multiple sources and to invoke services to compute new information.

The main goal in Karma is to make it easy and efficient for users to perform all information integration tasks. Karma enables users to perform operations on a small set of input instances, and then learns from these examples a general procedure that it can apply to all inputs. Compared to other data integration tools, Karma significantly reduces the time and effort needed to perform the data integration tasks. Recently, Karma has been extended with the capability to build semantic descriptions of Web APIs as a foundation to compose data sources and Web services [46]. In this section, we explain how it enables users to rapidly generate semantic service descriptions.

To model services, Karma first asks the user to provide samples of the API invocations URLs. This conforms to the main idea of Karma that examples are the basis to carry out the tasks. These sample URLs can also be automatically

¹⁶ <http://www.isi.edu/integration/karma>

extracted from the documentation pages of the APIs. Next, the user interactively builds a semantic model of the API by mapping the service inputs and outputs to the domain ontology. Building semantic models is the central part of the approach and it is very similar to how Karma models the data from other sources. Once the semantic model is built, Karma formalizes it using a new expressive RDF vocabulary that represents both the syntactic part and the functionality of the API. Karma stores the service specifications in a triple store, enabling users to query the service models using SPARQL. Finally, Karma deploys a Linked API that consumes and produces Linked Data. The Linked API provides REST interfaces enabling users to send RDF data in the body of a `POST` request and get back RDF output linked to the input. In the following paragraphs, we explain each step in more detail.

The input to the system are examples of the API requests. Karma parses the URLs and extracts the individual input parameters along with their values. For each invocation example, Karma calls the API and extracts the output attributes and their values from the XML or JSON response. Then, Karma joins the input and the output values into one table and shows that in a worksheet. Karma treats this table as a regular data source and applies its source modeling technique to build a semantic model of the API.

The goal of semantic modeling is to express the API functionality in terms of classes and properties defined in a domain ontology. The modeling process consists of two steps. The first step is to identify the type of data by assigning a *semantic type* to each column. A semantic type can be either an ontology class or a pair consisting of a data property and its domain. Karma uses a conditional random field (CRF) [26] model to learn the assignment of semantic types to columns of data [16]. Karma uses this classifier to automatically suggest semantic types for new data columns. If the correct label is not among the suggested labels, users can browse the ontology through a user friendly interface to find the appropriate type. The system re-trains the CRF model after these manual assignments.

The second part of the modeling process is to extract the relationships between the inferred semantic types. Given the domain ontology and the assigned semantic types, Karma creates a graph that defines the space of all possible mappings between the source and the ontology [21]. The nodes in this graph represent classes in the ontology connected by direct and inferred properties. Once Karma constructs the graph, it computes the API model as the minimal tree that connects all the semantic types. It is possible that multiple minimal trees exist, or that the correct model of the data is captured by a non-minimal tree. In these cases, Karma allows the user to interactively impose constraints on the algorithm to build the correct model. Karma provides an easy-to-use GUI where the user can adjust the relationships between the columns.

The models that Karma builds are themselves represented in RDF according to an ontology¹⁷ reusing existing ontologies such as SWRL¹⁸ and hRESTS¹⁹ [46]. This ontology is semantically richer than WSMO-Lite²⁰ and Minimal Service Model (MSM) [39] because in addition to annotating each input and output with semantic types, it also explicitly represents the relationships among inputs and outputs. Another advantage of the Karma models is that they are represented in RDF, making it possible for clients to query and discover models using SPARQL. Other approaches use graph patterns to represent the service models, so it is not possible to use SPARQL to query the models. The Karma API models are expressive enough that it would be possible to export them to other formal specifications such as LOS and MSM. This is a direction for future work [47].

Karma also has a Web server where the modeled API will be deployed as a Linked API. The Linked API implements a REST interface allowing clients to send RDF data in a POST request. One benefit of service descriptions in Karma is that they contain all the information needed to automatically execute APIs and do the required lowering and lifting, obviating the need to manually write separate instructions using formalisms such as XSLT and SPARQL. Once the Web server receives the user POST request, it uses the service description in the Linked API repository to lower the RDF data and build the invocation URL. Then, it invokes the Web API and again uses the service description to automatically lift the XML or JSON response to generate linked data. Karma is available as an open source²¹ software and users can use it to model the APIs based on their own needs.

6.2 SWEET

To facilitate the easier adoption of semantic description of Web APIs by supporting users in their creation, KMi has developed SWEET: Semantic Web sErVICES Editing Tool ²².

SWEET is developed as a Web application, that can be launched in a common Web browser and does not require any installation or additional configuration. It provides key functionalities for modifying the HTML Web API descriptions in order to include markup that identifies the different parts of the API, such as operations, inputs and outputs, and also supports the adding of semantic annotations by linking the different service parts to semantic entities. As a result, SWEET enables the creation of complete semantic Web API descriptions, based on the previously introduced models, given only the existing HTML documentation. More importantly, the tool hides formalism and annotation complexities from the user by simply visualizing and highlighting the parts of the API that are already annotated and produces an HTML description that is visually equivalent

¹⁷ <http://isi.edu/integration/karma/ontologies/model/current#>

¹⁸ Semantic Web Rule Language: <http://www.w3.org/Submission/SWRL>

¹⁹ <http://purl.org/hrests/current#>

²⁰ <http://www.w3.org/Submission/WSMO-Lite>

²¹ <https://github.com/InformationIntegrationGroup/Web-Karma-Public>

²² <http://sweet.kmi.open.ac.uk>

to the original one but is enhanced with metadata that captures the syntactical and semantic details of the APIs. The resulting HTML description also serves as the basis for extracting an RDF-based semantic Web API description, which can be published and shared in a service repository, such as iServe, enabling service browsing and search.

SWEET²³ is a Web application developed using JavaScript and ExtGWT²⁴, which is started in a Web browser by calling the host URL. It is part of a fully-fledged framework, developed within the scope of the SOA4All European project²⁵, for supporting the lifecycle of services, particularly targeted at enabling the creation of semantic Web API descriptions. SWEET takes as input an HTML Web page documenting a Web API and offers functionalities for annotating service properties and for associating semantic information with them. A current version of the tool can be found online²⁶.

SWEET is designed as a classical three-layered Web application. The architecture of SWEET consists of three main components, including the visualization component, the data preprocessing component and the annotations recommender. The visualization component is based on a model-view-controller architecture design pattern, where the model implements an internal representation of the annotated Web API, in accordance with the elements foreseen by the semantic formalisms detailed in the previous sections. In this way, every time the user adds a new annotation via the interface, the model representation of the Web API description is automatically updated. Similarly, when parts of the model representation are altered or deleted, the highlighting and visualization in the user interface is also adjusted. When the annotation process is complete, the resulting HTML and RDF Web API descriptions are generated based on the produced internal model.

The GUI of the visualization component is shown in Figure 3 and it has three main panels. The HTML of the Web API is loaded in the *Navigator* panel, which implements a reverse proxy that enables the communication between the annotation functions and the HTML by rerouting all sources and connections from the original HTML through the Web application. Based on this, the HTML DOM of the RESTful service can freely be manipulated by using functionalities of the *Annotation Editor* panel. The current status of the annotation is visualized in the form of a tree structure in the *Semantic Description* panel, which is implemented by automatically synchronizing the visualization of the service annotation with an internal model representation, every time the user manipulates it.

In addition to these three main panels, SWEET offers a number of supplementary useful functionalities. It guides the user thorough the process of marking service properties with hRESTS tags, by limiting the available tags depending on the current state of the annotation. This implements measures for reducing possible mistakes during the creation of annotations. In addition, based on the

²³ <http://sweet.kmi.open.ac.uk>

²⁴ <http://extjs.com/products/gxt/>

²⁵ SOA4AllEU project FP7 - 215219, <http://soa4all.eu/>

²⁶ <http://sweetdemo.kmi.open.ac.uk/soa4all/MicroWSMOeditor.html>

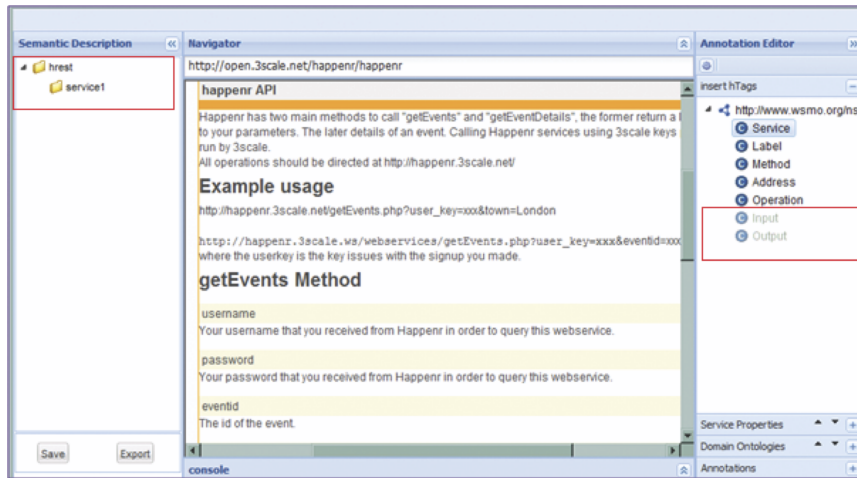


Fig. 3. SWEET: Inserting hRESTS Tags

hRESTS tagged HTML, which provides the structure of the Web API, the user can link service properties to semantic content. This is done by selecting service properties, searching for suitable domain ontologies by accessing Watson [8] in an integrated way, and by browsing ontology information. Based on this details the user can decide to associate a service property with particular semantic information by inserting a SAWSDL model reference tag.

In summary, SWEET takes as input the HTML Website description of the Web API, which is loaded in the central panel, and returns a semantically annotated version of the HTML or a RDF semantic description. In order to do this the user needs to complete the following four main steps:

1. Identify service properties (service, operation, address, HTTP method, input, output and label) by inserting hRESTS tags in the HTML service description.
2. Search for domain ontologies suitable for annotating the service properties.
3. Annotate service properties with semantic information.
4. Save or export the annotated Web API.

The first step can easily be completed by simply selecting the part of the HTML, which describes a particular service property, and clicking on the corresponding tag in the *insert hTags* pane. In the beginning, only the *Service* node of the hRESTS tree is enabled. After the user marks the body of the service, additional tags, such as the *Operation* and *Method*, are enabled. In this way, the user is guided through the process of structuring the RESTful service description and is prevented from making annotation mistakes. After the user structures the HTML description and identifies all service properties, the adding of semantic information can begin. The new version of SWEET, just like the bookmarklet, supports users in searching for suitable domain ontologies by providing an integrated search with Watson [8]. The search is done by selecting a service property and sending it as

a search request to Watson. The result is a set of ontology entities, matching the service property search. Once the user has decided, which ontology to use for the service property annotation, he/she can do an annotation by selecting a part of the service HTML description and clicking on *Semantic Annotation* in the *Service Properties* context menu. This results in inserting a `model` attribute and a reference pointing to the URI, of the linked semantic concept.

The resulting descriptions can be directly posted to iServe [38] or can be reposted on the Web. The use of the microformat tags enables the automated search and crawling for APIs, since it serves as a basis for distinguishing simple HTML websites from Web API descriptions. Furthermore, when posted to iServe, the Web API descriptions can be browsed and searched alongside with WSDL-based services. Since the semantic Web API descriptions use SAWSDL-like annotations in combination with the WSMO-Lite service ontology, both WSDL-based services and APIs can be retrieved by using the same queries. For example, a search query for music services would return the Last.fm description as well as all other APIs or services related to music. As a result, all type of services, can be retrieved in a unified way.

7 Open problems / future work

7.1 Cross-Origin Resource Sharing (CORS)

The XMLHttpRequest specification [50] defines an API that provides scripted client functionality for transferring data between a client and a server. In today's common Web applications like online spreadsheets, word processors, presentation tools *etc.*—and even more in so-called mash-up applications²⁷—the majority of data transfers between server and client happen based on XMLHttpRequest. An important security aspect of XMLHttpRequest, however, is the so-called Same Origin Policy (SOP). This policy permits scripts running on pages originating from the *same* site to access each other's methods and properties with *no* specific restrictions, however, *prevents* access to most methods and properties across pages on *different* sites. While providing at least some protection from rogue Web pages accessing private data, SOP also has severe implications for cases where cross-origin data transfers are actually legit. Past attempts to legally circumvent SOP include using proxy servers, Adobe Flash, and JSON-P,²⁸ however, more recently, the tendency goes in the direction of properly handling cross-origin resource sharing (CORS) through a mechanism documented in a w3C Working Draft [49]. The CORS standard works by adding new HTTP headers that allow servers to serve resources to permitted origin domains. Browsers support these headers and enforce the restrictions they establish. While not all APIs support

²⁷ Web applications that combine data from multiple sources to create new services, many of them listed in [10].

²⁸ First documented appearance of JSON-P in Bob Ippolito's December 2005 blog post: <http://bob.ippoli.to/archives/2005/12/05/remote-json-jsonp/>.

CORS yet, there is a remarkable momentum of Web API and data providers in general to open up their data and becoming CORS-enabled²⁹.

7.2 Authentication

The IEEE defines³⁰ authentication as *the act of confirming the truth of an attribute of a datum or entity. This might involve confirming the identity of a person [...], or assuring that a computer program is a trusted one.* In the world of APIS, simple authentication paradigms include (but are not limited to) API keys (codes passed in by computer programs calling an API to identify the calling program, its developer, or its user to the API provider), HTTP Basic authentication, HTTP headers, or HTTP cookies. In recent times, different versions of the *authorization* protocol OAuth³¹ gain traction as the *de facto* default standard for authorization. *Authentication* is the mechanism whereby systems may securely identify their users. *Authorization*, by contrast, is the mechanism by which a system determines what level of access a particular authenticated user should have to secured resources controlled by the system³². In the case of OAuth, if the user grants access to a resource, the application can retrieve the unique identifier for establishing the identity in turn by using the particular API calls, and thus effectively enabling pseudo-authentication using OAuth.

[32] provides an extensive overview of currently used authentication approaches, the required credentials, ways of transmitting the credentials and used authentication mechanisms. The provided solution is based on defining authentication extensions to the MSM defined in Section 2.

7.3 CORS and authentication in API descriptions

To the best of our knowledge, neither authentication nor CORS are covered by the before-mentioned API description formats, the honorable exception being the Web Application Description Language (WADL, [18]), where (authentication) HTTP headers can be described for API query parameters. While the implementation status of CORS can be determined at runtime by examining a sample API request and checking for the existence of the particular HTTP header, there is no general way to discover the authentication requirements of an API at runtime. In consequence, both CORS and authentication and ways to semantically describe them remain a field for future research.

8 Conclusion

We have surveyed the current state-of-the-art in descriptions of Web APIs and classified the various approaches. The main strength of RESTful APIS, the

²⁹ <http://enable-cors.org/>

³⁰ <http://technav.ieee.org/tag/2585/authentication>

³¹ <http://oauth.net/>

³² <http://www.duke.edu/~rob/kerberos/authvauth.html>

flexibility which with the APIs can be designed and deployed, at the same time burdens client application developers with the manual work of understanding, interpreting, and reconciling the various approaches to API design. Almost all of today's Web APIs come with a textual description, lacking coherence. A little structure in architecting and documenting the APIs could greatly benefit application developers and reduce the amount of manual effort required when integrating multiple APIs.

Acknowledgments. R. Verborgh and R. Van de Walle are funded by Ghent University, the Interdisciplinary Institute for Broadband Technology (iMinds), the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT), the Fund for Scientific Research Flanders (FWO Flanders), and the European Union. A. Harth and S. Speiser acknowledge the support of the European Commission's Seventh Framework Programme FP7/2007-2013 (PlanetData, Grant 257641). S. Stadtmüller has been supported by a Software Campus grant.

References

1. Adida, B., Birbeck, M., McCarron, S., Herman, I.: RDFa core 1.1. W3C Recommendation (Jun 2012), <http://www.w3.org/TR/2012/REC-rdfa-core-20120607/>
2. Baker, T., Noy, N., Swick, R., Herman, I.: Semantic Web case studies and use cases (2007–2012), <http://www.w3.org/2001/sw/sweo/public/UseCases/>
3. Berners-Lee, T.: Read-write linked data (Aug 2009), <http://www.w3.org/DesignIssues/ReadWriteLinkedData.html>
4. Berners-Lee, T., Connolly, D.: Notation3 (N3): A readable RDF syntax. W3C Team Submission (Mar 2011), <http://www.w3.org/TeamSubmission/n3/>
5. Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web. *Scientific American* 284(5), 34–43 (2001)
6. Bizer, C., Heath, T., Berners-Lee, T.: Linked Data – The Story So Far. *International Journal On Semantic Web and Information Systems* 5(3), 1–22 (2009)
7. Christensen, E., Curbera, F., Meredith, G., Weerawarana, S.: Web Services Description Language (WSDL) 1.1. W3C Note (Mar 2001), <http://www.w3.org/TR/wsdl>
8. d'Aquin, M., Sabou, M., Motta, E., Angeletou, S., Gridinoc, L., Lopez, V., Zablith, F.: What can be done with the Semantic Web? an overview of Watson-based applications. In: 5th Workshop on Semantic Web Applications and Perspectives (2008)
9. Devis Bianchini, Valeria Antonellis, M.M.: Flexible semantic-based service match-making and discovery. In: Proceedings international conference on World Wide Web. vol. 11, pp. 227–251 (2008)
10. DuVander, A.: 7,000 APIs: Twice as many as this time last year. *ProgrammableWeb* (Aug 2012), <http://blog.programmableweb.com/2012/08/23/7000-apis-twice-as-many-as-this-time-last-year/>
11. Farrell, J., Lausen, H.: Semantic Annotations for WSDL and XML Schema (SAWSDL). W3C recommendation (Aug 2007), <http://www.w3.org/TR/2007/REC-sawsdl-20070828/>

12. Fielding, R.T., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., Berners-Lee, T.: Hypertext Transfer Protocol – HTTP/1.1. IETF Request for Comments (Jun 1999), <http://www.ietf.org/rfc/rfc2616>
13. Fielding, R.T.: REST APIs must be hypertext-driven. Untangled – Musings of Roy T. Fielding (Oct 2008), <http://roy.gbiv.com/untangled/2008/rest-apis-must-be-hypertext-driven>
14. Fischer, F.F., Norton, B.: D3.4.6 MicrowSMO v2 – defining the second version of MicrowSMO as a systematic approach for rich tagging. Soa4all project deliverable (2010)
15. Gessler, D., Schiltz, G., May, G., Avraham, S., Town, C., Grant, D., Nelson, R.: SSWAP: A simple semantic Web architecture and protocol for semantic Web services. BMC Bioinformatics 10, 309 (2009)
16. Goel, A., Knoblock, C.A., Lerman, K.: Exploiting Structure within Data for Accurate Labeling Using Conditional Random Fields. In: Proceedings of the 14th International Conference on Artificial Intelligence (ICAI) (2012)
17. Gregorio, J., Fielding, R., Hadley, M., Nottingham, M., Orchard, D.: URI template. IETF Request for Comments (Mar 2012), <http://tools.ietf.org/html/rfc6570>
18. Hadley, M.: Web Application Description Language. w3c Member Submission (Aug 2009), <http://www.w3.org/Submission/wadl/>
19. Khare, R., Celik, T.: Microformats: a pragmatic path to the semantic web (poster). In: Proceedings of the 15th international conference on World Wide Web (2006)
20. Klyne, G., Carrol, J.J.: Resource Description Framework (RDF): Concepts and Abstract Syntax. w3c Recommendation (Feb 2004), <http://www.w3.org/TR/2004/REC-rdf-concepts-20040210/>
21. Knoblock, C., Szekely, P., Ambite, J.L., Goel, A., Gupta, S., Lerman, K., Muslea, M., Taheriyani, M., Mallick, P.: Semi-Automatically Mapping Structured Sources into the Semantic Web. In: Proceedings of the 9th Extended Semantic Web Conference (ESWC) (2012)
22. Knuth, D.E.: Literate programming. Comput. J. 27(2), 97–111 (1984)
23. Kopecký, J., Vitvar, T.: MicrowSMO. CMS Working Draft (Feb 2008), <http://www.wsmo.org/TR/d38/v0.1/20080219/>
24. Kopecký, J., Gomadam, K., Vitvar, T.: hRESTS: an HTMLMicroformat for Describing RESTful Web Services. In Proceedings of the 2008 IEEE/WIC/ACM International Conference on Web Intelligence (WI-08) (2008)
25. Krummenacher, R., Norton, B., Marte, A.: Towards Linked Open Services. In: FIS (2010)
26. Lafferty, J., McCallum, A., Pereira, F.: Conditional Random Fields: Probabilistic Models for Segmenting and Labeling Sequence Data. In: Proceedings of the 18th International Conference on Machine Learning. pp. 282–289 (2001)
27. Lanthaler, M., Gütl, C.: A semantic description language for RESTful data services to combat semaphobia. In: Proceedings of the 5th IEEE International Conference on Digital Ecosystems and Technologies Conference. pp. 47–53 (Jun 2011), http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5936597
28. Lanthaler, M., Gütl, C.: On using JSON-LD to create evolvable RESTful services. In: Proceedings of the Third International Workshop on RESTful Design. pp. 25–32. ACM, New York, NY, USA (2012), <http://doi.acm.org/10.1145/2307819.2307827>
29. Lanthaler, M., Gütl, C.: Seamless integration of RESTful services into the Web of Data. Advances in Multimedia 2012, 1:1–1:14 (Jan 2012), <http://dx.doi.org/10.1155/2012/586542>
30. Lausen, H., Polleres, A., Roman, D.: Web Service Modeling Ontology (WSMO). w3c Member Submission (Jun 2005), <http://www.w3.org/Submission/WSMO/>

31. Li, N., Pedrinaci, C., Maleshkova, M., Kopecký, J., Domingue, J.: Omnivoke: A framework for automating the invocation of Web APIs. In: Proceedings of Fifth IEEE International Conference on Semantic Computing. pp. 380–387 (2011)
32. Maleshkova, M., Pedrinaci, C., Domingue, J., Alvaro, G., Martinez, I.: Using semantics for automating the authentication of Web APIs. International Semantic Web Conference (ISWC) (2010), shanghai, China
33. Maleshkova, M., Pedrinaci, C., Domingue, J.: Investigating Web APIs on the World Wide Web. In: The 8th IEEE European Conference on Web Services (2010), <http://oro.open.ac.uk/24320/>
34. Martin, D., Burstein, M., Hobbs, J., Lassila, O.: OWL-S: Semantic Markup for Web Services. w3C Member Submission (Nov 2004), <http://www.w3.org/Submission/OWL-S/>
35. McGuinness, D.L., van Harmelen, F.: OWL web ontology language. w3C Recommendation (Feb 2004), <http://www.w3.org/TR/owl-features/>
36. Norton, B., Stadtmüller, S.: Scalable discovery of linked services. In: RED (2011)
37. Pedrinaci, C., Domingue, J.: Toward the next wave of services: Linked services for the web of data. Journal of Universal Computer Science pp. 1694–1719 (2010)
38. Pedrinaci, C., Liu, D., Maleshkova, M., Lambert, D., Kopecký, J., Domingue, J.: iserve: a linked services publishing platform. Workshop: Ontology Repositories and Editors for the Semantic Web at 7th Extended Semantic Web Conference (2010)
39. Pedrinaci, C., Domingue, J.: Toward the Next Wave of Services: Linked Services for the Web of Data. Journal of Universal Computer Science 16(13) (2010)
40. Severance, C.: Discovering JavaScript Object Notation. Computer 45(4), 6–8 (Apr 2012), http://ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=6178118
41. Sheth, A., Verma, K., Gomadam, K.: Semantics to energize the full services spectrum. Comm. ACM 49, 55–61 (2006)
42. Sheth, A.P., Gomadam, K., Lathem, J.: SA-REST: Semantically interoperable and easier-to-use services and mashups. IEEE Internet Computing 11(6):91-94 (2007)
43. Soto, C.A.: 5 technologies that will change the market (Aug 2010), <http://washingtontechnology.com/articles/2010/08/02/cover-5-disruptive-technologies.aspx>
44. Speiser, S., Harth, A.: Integrating Linked Data and Services with Linked Data Services. In: The Semantic Web: Research and Applications, Lecture Notes in Computer Science, vol. 6643, pp. 170–184. Springer (2011)
45. Sporny, M., Longley, D., Kellogg, G., Lanthaler, M., Birbeck, M.: JSON-LD syntax 1.0. w3C Editor’s Draft (Oct 2012), <http://json-ld.org/spec/latest/json-ld-syntax/>
46. Taheriyani, M., Knoblock, C.A., Szekely, P., Ambite, J.L.: Rapidly Integrating Services into the Linked Data Cloud. In: Proceedings of the 11th International Semantic Web Conference (ISWC). Boston, USA (2012)
47. Taheriyani, M., Knoblock, C.A., Szekely, P., Ambite, J.L.: Semi-Automatically Modeling Web APIs to Create Linked APIs. In: Proceedings of the Linked APIs for the Semantic Web Workshop (LAPIS). Heraklion, Crete, Greece (2012)
48. Tuchinda, R., Knoblock, C.A., Szekely, P.: Building Mashups by Demonstration. ACM Transactions on the Web (TWEB) 5(3) (2011)
49. van Kesteren, A.: Cross-Origin Resource Sharing. w3C Working Draft (Apr 2012), <http://www.w3.org/TR/cors/>
50. van Kesteren, A.: XMLHttpRequest Level 2. w3C Working Draft (Jan 2012), <http://www.w3.org/TR/XMLHttpRequest/>

51. Verborgh, R., Haerinck, V., Steiner, T., Van Deursen, D., Van Hoecke, S., De Roo, J., Van de Walle, R., Gabarró Vallés, J.: Functional composition of sensor Web APIs. In: Proceedings of the 5th International Workshop on Semantic Sensor Networks (Nov 2012)
52. Verborgh, R., Steiner, T., Gabarró Vallés, J., Mannens, E., Van de Walle, R.: A social description revolution—describing Web APIs’ social parameters with RESTdesc. In: Proceedings of the AAAI 2012 Spring Symposia (Mar 2012), <http://www.aaai.org/ocs/index.php/SSS/SSS12/paper/viewFile/4283/4665>
53. Verborgh, R., Steiner, T., Van Deursen, D., Coppens, S., Gabarró Vallés, J., Van de Walle, R.: Functional descriptions as the bridge between hypermedia APIs and the Semantic Web. In: Proceedings of the Third International Workshop on RESTful Design. ACM (Apr 2012)
54. Verborgh, R., Steiner, T., Van Deursen, D., De Roo, J., Van de Walle, R., Gabarró Vallés, J.: Capturing the functionality of Web services with functional descriptions. Multimedia Tools and Applications (2013), <http://www.springerlink.com/index/d041t268487gx850.pdf>
55. Webber, J.: REST in Practice: Hypermedia and Systems Architecture. O’Reilly (2010)
56. Wilde, E.: REST and RDF granularity (May 2009), <http://dret.typepad.com/dretblog/2009/05/rest-and-rdf-granularity.html>
57. Wilkinson, M.D., Links, M.: BiOMoby: An open source biological web services proposal. Briefings in Bioinformatics 3, 331–341 (2002)
58. w3C WS Description Working Group: Web service description language (WSDL) Version 2.0, w3C proposed recommendation (May 2007), <http://www.w3.org/TR/wsd120-primer>
59. Wu, B., Szekely, P., Knoblock, C.A.: Learning Data Transformation Rules through Examples: Preliminary Results. In: 9th International Workshop on Information Integration on the Web (IIWeb) (2012)
60. Zyp, K.: A JSON media type for describing the structure and meaning of JSON documents. IETF Request for Comments (Nov 2010), <http://tools.ietf.org/html/draft-zyp-json-schema-03>