

A Web API Ecosystem through Feature Based Reuse

Citation for published version (APA):

Verborgh, R., & Dumontier, M. (2018). A Web API Ecosystem through Feature Based Reuse. *Ieee Internet Computing*, 22(3), 29-37. <https://doi.org/10.1109/MIC.2018.032501515>

Document status and date:

Published: 01/01/2018

DOI:

[10.1109/MIC.2018.032501515](https://doi.org/10.1109/MIC.2018.032501515)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

A Web API Ecosystem through Feature-Based Reuse

Ruben Verborgh
Ghent University – imec

Michel Dumontier
Maastricht University

The fast-growing web API landscape brings clients more options than ever before—in theory. In practice, they cannot easily switch between different providers offering similar functionality. We discuss a vision for developing web APIs based on reuse of interface

parts called features. Through the introduction of five design principles, we investigate the impact of feature-based reuse on web APIs. Applying these principles enables a granular reuse of client and server code, documentation, and tools. Together, they can foster a measurable ecosystem with cross-API compatibility, opening the door to a more flexible generation of web clients.

Why can people easily navigate websites they've never encountered before? The answer is simple: interactions patterns are reused across websites. Usability expert Jakob Nielsen observed that “users spend most of their time on other websites,” reminding interaction designers and information architects to compose websites from shared patterns rather than custom ones. For example, the user interfaces to update a status message on Facebook, Twitter, and Instagram are all highly similar. The strength of reuse is confirmed by the survival of visually less-obvious patterns, such as a site's logo doubling as a homepage link, or three parallel lines forming a menu button on mobile devices. From a technical perspective, we could say people are “loosely coupled” to well-designed websites, because they bind to generic interaction patterns rather than specific interfaces. The human web maintains usability by tapping into an evolving ecosystem of interaction patterns, only inventing new interfaces as a last resort.

In contrast, the machine-based web is characterized by a near-total lack of interface-level reuse, as even web APIs with highly similar functionality often expose very different machine interfaces. This results in a lack of substitutability¹ with non-native services;² clients programmed for a specific API task (such as posting a photo on Facebook) cannot perform that same task with another web API (posting that same photo on Twitter or Flickr). Regardless of whether we classify the coupling between a client and an API as “loose” or “tight,”³ switching API providers

proves difficult⁴ as clients are forced to bind to a provider-specific interface rather than a provider-independent, abstract interface. Case in point: Whereas Facebook and Twitter show a near-identical user interface for updating a status (a light-colored textbox with an encouraging question and a camera icon), the corresponding interactions with their web APIs require a different number of HTTP requests with entirely different JSON bodies. If the human web were designed in such a way, information consumption would slow down significantly, as accessing any new website would involve studying its documentation first—as is the case with web APIs.⁵

The absence of an ecosystem of web API interaction patterns means that every client needs custom manual programming for each provider—even if they had already been equipped with support for functionally similar or identical APIs. Because new integrations cannot build upon earlier ones, developers need to selectively pick the APIs they can support. As a result, as the number of web APIs continues to grow at a rapid pace,⁶ each application only gives users access to an increasingly smaller fraction of APIs with relevant functionality.

If we want our applications to access similar web APIs with the same flexibility as people browse similar webpages, we will need to fundamentally rethink the way in which the interfaces to those web APIs are designed. Given the increasing importance and scale of the web API landscape,⁴ it is time to evolve web API design from a craft into a measurable discipline, focused on establishing and repeating interaction designs. The ultimate goal is a new generation of clients that are compatible with web APIs beyond those for which they were explicitly programmed.

In this article, we study the potential of adopting a pattern-based approach to web API design. Through the introduction of five principles, we port the lessons learned in human interface design to machine interface design, evaluate their impact, and discuss possible concrete technologies for implementing them. Although the description and argumentation of such a vision does not imply its immediate realization, this work aims to inspire community-wide discussions on cross-API interoperability through granular reuse.

ISSUES WITH THE CURRENT WEB API LANDSCAPE

When discussing reuse, we should distinguish an interface from its implementation, as the term “web API” often refers to both simultaneously. We will discuss (partial) interface reuse across different implementations, because we aim for service substitutability. Reuse of both human and machine interfaces can be considered from the interface user and the API implementer perspectives. Interestingly, interface-level reuse can facilitate implementation-level reuse; for instance, if different websites need the same widget, developers can implement it with the same library.

It might seem contradictory to discuss a lack of reuse in the context of web APIs; after all, APIs are designed to enable reuse.⁷ Indeed, a web API enables the reuse of an implementation offered by a third-party provider,⁸ and we target reuse of interfaces across providers.¹ However, web APIs are different in that the question is not whether to reuse or re-implement (we have already decided we need third-party functionality), and we cannot provide the user with a viable alternative of our own. For example, when providing share functionality, users will not be satisfied if their photos appear on our website; instead, they want to post the photos on their existing social networks. Another difference is the scale of reuse—with the ever-growing number of web APIs,⁶ new providers for existing functionality appear every day. While integrating a single API poses non-trivial but manageable problems,^{4,5} integrating the same functionality from multiple APIs—or switching between them—proves far more difficult.^{1,9}

Integration concerns of individual web APIs include initial implementation effort, coupling, and evolution. The integration of a web API involves manual labor to (implicitly or explicitly) construct the appropriate HTTP requests programmatically. In absence of a standard,⁴ these requests are different for all APIs, even for those with related functionality.¹ Often, insufficient documentation is available.⁵ Web APIs are frequently cited as providing “loose coupling,”²⁴ yet research revealed that coupling is a multi-faceted concept, with different architectural styles ranking divergently in multiple dimensions of coupling.³ However, the predicted rankings are only obtained in case of strict adherence to an architectural style, which is seldom the case for web APIs.

Furthermore, some dimensions, such as evolution and granularity, always depend on the implementation.³ Coupling is not only a problem for cross-API compatibility, but also for cross-version compatibility within the same API.^{4,9} Because clients necessarily depend on the web API provider's pace of evolution, upgrading between versions can come at a considerable cost, which is forced upon client developers.⁴

Yet the biggest problems occur when trying to integrate the same functionality from different web APIs. While the interface might loosely couple clients to the server's underlying implementation of functionality, that same interface, in practice, tightly couples clients to one specific provider of that functionality. No coupling is loose enough to enable switching providers without changing client code (except for the few cases where entire APIs are standardized). These issues arise because, while hiding the implementation, the interface is the result of a unilateral decision¹ that insufficiently abstracts its functionality. Nonetheless, substitutability of different APIs is important to enable competition on factors such as price and quality,² let end users decide where to read and/or write data, and reduce the costs of client development. At the moment, web APIs still act as silos within walled gardens, partially due to economic motives, but also in a substantial way due to technical obstacles and missing guidance on how to achieve sustainable interoperability.

Several strategies for circumventing the above intra- and cross-API problems providing symptomatic relief instead of addressing the root causes, thereby unwittingly enabling those problems to persist and grow worse. Documentation can simplify integration of a client to one API,⁵ but the necessity of studying documentation beforehand becomes a significant burden when setting up multiple integrations. Recall in this context that the corresponding webpages—which achieve the same functionality as their API—never require a manual. Some APIs offer software development kits (SDKs) for different programming languages, hiding the details of HTTP requests behind language-specific abstractions. While convenient to speed up development and hide minor API evolutions, the creation and maintenance of one or multiple SDKs is expensive, and their provider-specific abstraction still does not achieve substitutability. Web API descriptions can automate integration,⁴ but while technologies such as SOAP and WSDL offer the technological possibility of server-side interface reuse with different implementations, in practice, they mostly facilitate the conversion of Java/.NET/C++ methods to RPC APIs. Such descriptions can also allow client-side code generation, which more recently was provided for JSON-based HTTP APIs by the OpenAPI specification (formerly Swagger). While speeding up development, code generation tightens coupling³ and thereby hampers API substitutability. Migration tools⁹ and adapters¹⁰ cannot bridge the gap between highly different yet functionally similar interfaces such as Facebook's and Twitter's.

From the above, we conclude that today's clients still cannot request conceptually identical functionality from a different web API provider without changing client-side code. The coupling between a client and the interface of a web API is insufficiently loose for substitutability between providers. To overcome this, we propose designing APIs at an abstraction level that enables clients to automatically determine semantic compatibility and invocation mechanics¹ at runtime.

BOTTOM-UP INSTEAD OF TOP-DOWN

The coupling between clients and provider-specific interfaces follows from the fact that current web APIs are constructed and consumed like monoliths. Clients approach APIs in a top-down manner (see Figure 1): They consider the interface a single custom entity with its own principles and conventions. While several APIs incorporate similar or identical functionalities, these are exposed through totally different interfaces. As such, clients cannot easily verify semantic compatibility with others, or discover invocation mechanics beyond parameter names and types. This sharply contrasts with interfaces on the human web in which people recognize smaller interface components across websites (search bars, share buttons, address forms, and so on), and these parts guide us through the entire functionality of the interface.

Therefore, we will explore the impact of similarly approaching web interfaces for machines in a bottom-up manner, by composing them of reusable features. A feature is a part of an interface that identifies, describes, and affords a certain kind of interaction across web APIs. Examples of

such features include interfaces to perform full-text search, auto-completion, file uploads, and status updates. Clients couple to one or more features instead of the entire API and use the feature’s interface to determine semantic compatibility and invocation mechanics. Multiple APIs can reuse features without requiring entire APIs to be identical. As an additional benefit, clients and servers can reuse feature implementations.




Figure 1. Instead of the current top-down web APIs, which couple a client to a single provider, we propose a bottom-up interface structure, which couples a client to reusable, provider-independent features.

This feature-based approach improves upon the granularity and evolution facets of coupling³ within web APIs in general, by shifting the anchor point of coupling from the entire API to one or more features. For the same reason, it improves the binding, generated code, and model facets³ within RPC-based APIs. Within REST APIs, it realizes the self-descriptive and hypermedia constraints of the REST architectural style¹¹ in a modular way by defining them per feature.

Feature-Based Interface Principles

Now that we have established the “feature” concept, we introduce five incremental principles for the design of feature-based web APIs. We discuss each principle in a technology-agnostic way and describe the situation before and after its application, followed by current and/or future examples. To aid concrete implementations, we suggest possible technologies for realizing each principle; however, the choice for specific technologies and an assessment of their adequacy for the ecosystem should be discussed within the web API community.

Principle 1: Web APIs should consist of features that implement a common interface.

This principle breaks up a monolithic custom interface into clearly delineated interface parts that can be reused across different APIs. The input and output of features should reuse specific subsets/profiles¹² of common media types, such as JSON, to achieve interoperability. Features can be compared to visual components that coexist on the same page but are interacted with separately, such as a search form, upload form, or order button. Additionally, features can be switched on or off, depending on factors such as authentication, selected plan or payment options, and the version of the API.

Before: Clients bind to a provider-specific interface to access certain functionality. Neither client- nor server-side code related to the interface can be reused, because each server has a different interface.

After: Clients bind to individual features to access functionality, so they are unaffected by changes in other features. Server-side code related to the feature's interface and implementation can be reused in a granular way, as reusers are not required to copy the entire interface.

Possible technologies: reusable software components, declarative software configuration (to enable or disable features), microservices...

Examples: accessing and updating a list of items, obtaining a sorted view of items, pagination, updating a status message, uploading a photograph, the OpenSearch specification for search and auto-completion, the Atom standard for viewing and managing collections...

This principle is based on the observation that many clients use only a fraction of a service's functionality.² It is thereby related to the interface-segregation principle, which states that clients should not depend on methods they do not use. It differs from the concept of a microservice architecture in that features are interface parts of a larger web API's interface, whereas a microservice is an independently deployable service that can run in its own process.¹³ Features might be realized as microservices internally, but clients do not need to know whether this is the case.

Principle 2: Web APIs should partition their interface to maximize feature reuse.

While each web API is different, similar functionality appears in many of them, at a scale that grows faster than is the case with operating system or software framework APIs. This second principle therefore encourages architects to first check whether a part of the API is already available as a feature elsewhere, before implementing their own feature.

Before: Each web API offers its own specific kinds of features.

After: Because features are reused across APIs, clients that are compatible with a set of features can perform their task with any API that offers these features, regardless of provider. Client- and server-side documentation, tooling, and libraries for certain features can be reused across implementations.

Possible technologies: repositories for features (instead of full APIs), feature-specific SDKs (instead of provider-specific SDKs), shared schemas for GraphQL APIs...

Examples: using Atom for blog posts across providers, using Atom for a collection of tweets (instead of the current Twitter-specific interface), a generic status update feature (instead of a provider-specific one)...

Assuming that every web API can be cleanly separated into existing features would, of course, be unrealistic, so this principle does not demand that. Instead, architects should prioritize feature reuse where applicable and package (only) the remaining provider-specific functionality as separate features. For features to be widely applicable, they should possess a certain extent of flexibility, which is discussed in the next two principles. Features can have a dependency on each other; for instance, sorting a list can depend on browsing a list. However, they do not directly interact with each other—the client decides whether to use multiple features in combination.

Principle 3: Web API responses should advertise the presence of each relevant feature.

When servers include or link to the list of features they support, clients can automatically verify their semantic compatibility with the web API. Support should be explicitly indicated in-band inside of the HTTP response, either through headers or inside of the response body.

Before: Clients are hard-coded with the knowledge of what an API supports and cannot determine their compatibility with a new API.

After: Clients can determine whether a given API offers the features they require. If feature support varies dynamically (based on authentication or other factors), clients can find out at runtime. In case of incompatibility, clients can explain what features are missing. Clients can ignore the presence of non-supported features.

Possible technologies: Outside the response body: the “Link” HTTP header with the “profile” relationship. Inside the body: hypermedia-based media types (such as HAL and SIREN) or hypermedia vocabularies (such as Hydra)...

Examples: W3C Activity Streams 2.0 (through “profile”), GitHub API (through hypermedia)...

This principle realizes the self-describing messages constraint of REST,¹¹ relying on message extensibility³ rather than provider-specific media types¹⁴ (for instance, through the use of profiles¹² to constrain more generic media types). In addition to signaling the presence of the features themselves, the server might indicate to what extent they are supported (for instance, which optional parts are implemented).

Principle 4: Each feature should describe its own invocation mechanics and functionality.

Describing invocation mechanics and functionality might seem redundant, given that Principle 3 already mandates feature identification. However, this principle decouples features from specific URL or request body templates, because these details are now communicated at runtime. Descriptions can be provided in-band in the response body or, if the identifier is a URL, by dereferencing the identifier. Through descriptions of functionality—the possibilities and limitations of which vary by formalization mechanism—new kinds of features can be discovered at runtime.

Before: Clients only receive a feature’s identifier, so they are coupled to out-of-band details that have to be known beforehand. All APIs have to implement a feature with the exact same URL structure and parameter names.

After: Clients are coupled only to an abstract interface of a feature, the details of which are communicated at runtime. Servers are free to choose their own URLs and names, as they explicitly indicate their mapping to the feature.

Possible technologies: hypermedia controls (such as HAL and Hydra), functional description languages or vocabularies (such as Schema.org actions), semantic web technologies (such as JSON-LD and ontologies)...

Examples: hypermedia-driven clients and servers (such as GitHub, Amazon AppStream, and PayPal APIs), the Hydra console (which generates interactive forms for unknown features)...

This principle concerns the hypermedia constraint of REST¹¹ and provides an explicit way for customizing and extending features by describing optional or additional parameters. Cross-API integration of features is facilitated by allowing each API to describe its own URL and form templates, authentication, content types and profiles, and so on. Beyond the substitutability this brings current clients, new possibilities for future clients arise. For instance, interactive clients can automatically generate forms for features they have not encountered before. By additionally describing functionality, clients can gain an understanding of what a feature does and, depending on the formalism and capabilities of the client, access features for which they have not been explicitly designed.

Principle 5: The impact of a feature on a web API should be measured across implementations.

This final principle ensures that the properties of features become quantified, such that architects can understand the impact of enabling certain features on their web API. On the human web, the effectiveness of interface components is routinely measured through in-browser analytics, and interfaces are tweaked to optimize user efficiency. By testing conformance with a certain feature and measuring clients, servers, and caches, we can determine how a certain application task is affected by a feature or combination of multiple features. While server load, bandwidth, and cache efficiency are typically important aspects, the relevance of dimensions depends on the feature. Importantly, we focus on measuring the interface across implementations to understand its inherent tradeoffs. Note that this principle differs from monitoring, which assesses the current state of a web API; instead, we aim to analyze the estimated impact of a feature beforehand.

Before: API architects cannot *a priori* predict the impact of enabling a certain feature on the API; they can only monitor the effects afterwards.

After: Measurements of the impact of a feature are documented, such that *a priori* estimations of their impact are possible.

Possible technologies: conformance tests, benchmarks, benchmarking software...

Example: The expressivity of a query-based interface feature has a strong impact on how a client completes a certain information task, as it influences the required number of requests, bandwidth, cache effectiveness, maximum per-request cost, and server load.¹⁵

Depending on what features are available, a client might be able to complete a certain task with different quality attributes. For example, if the server only supports a pagination feature, the client needs to download all pages of a list to find the five cheapest products. If the client additionally detects a sorting feature, it can obtain the same answer with a single request. However, in the second case, the API exposes its data in more variations, which impacts cache efficiency and requires additional computations, which impacts server load. This indicates that the behavior of a web API under different circumstances is determined by the presence of certain features. Crucially, this effect is not only implementation-dependent, but also interface-dependent, because similar trends can occur per interface regardless of underlying implementations.¹⁵ Providers could use this information to decide whether to enable a feature in the free or paid version of an API, and clients might opt to pay for more expressive features if this allows them to complete tasks more efficiently.

Potential Adoption Obstacles

This article examines the potential impact of a feature-based ecosystem for web APIs. However, bridging the gap from the theory to an actual realization is non-trivial, as the initiative needs to be carried by a larger community. It is crucial that we agree on the base technologies for each principle, especially for the description of features, for which several competing technologies exist. Furthermore, we need to specify or standardize data structures and profiles of media types to achieve message-level interoperability. Otherwise, we risk fragmentation of the community based on differences in the technological stacks, for the mechanisms surrounding a feature's description and interaction. Good initial use cases are those that demand sustainability over a longer period of time, such as the long-term publication of data by archives and other institutions, who currently employ a mix of standardized and non-standardized APIs. Feature-based interfaces recently gained interest for data integration purposes, as they countered problems with existing APIs in a methodical way.¹⁵ To make this approach spread to other communities, potential obstacles need to be considered, some of which we discuss below.

The cost of reuse is well-documented in literature.¹⁶ However, calculations of that cost typically consider whether a use case benefits from reuse. In the case of web APIs, we have already decided we need to reuse server-side functionality; so, the associated expenses will occur in any case. The question is rather how we can reduce those expenses while offering an equal or higher number of API integrations. Our bottom-up approach addresses this by limiting the dependency of clients to a number of reusable features instead of a provider-specific interface, allowing reuse of client and server code, documentation, and tooling across implementations. However, developing an ecosystem of features also comes at a cost. As mentioned above, a community will need to reach consensus on their definition and design, and standardization can evolve slowly.

It is crucial that we agree on the base technologies for each principle.

Incentives for client and server developers are an important issue. Convincing major players to reuse API features might be difficult when they have the power to push any API they desire,¹ especially if they benefit commercially from non-compatibility with competitors' APIs. However, improving the developer experience is also in their best interest. Nonetheless, we might want to approach adoption from a bottom-up perspective, as well, and focus on small to medium

players. For them, reusing API features means making their API compatible with existing clients. For example, consider a local restaurant that aims to provide a reservation web API: When other restaurants reuse a table reservation feature, it is more likely that an adequate client will already exist. In absence of reuse, and despite hosting their own website for humans, restaurants now resort to paid centralized APIs, which come with provider-specific machine clients.

CONCLUSIONS AND FUTURE DIRECTIONS

Automated clients on the web do not enjoy the same flexibility as people—switching between providers of a similar service is labor-intensive. Web APIs lack recognizable interface patterns, even though machines need these more than humans. Hence, we proposed a feature-based method to construct the interface of web APIs, favoring reuse over reinvention, analogous to component-driven interaction design on the human web. We emphasized a quantifiable approach to arrive at an ecosystem of reusable features with a well-understood impact on web APIs.

Such an ecosystem changes web server and client development in three major ways. First, clients gain the ability to interact with multiple similar web APIs rather than only one. This reduces development costs because client-side code and frameworks can be reused across different web APIs and over time. Server interface design starts with selecting appropriate patterns based on functional and non-functional characteristics. Second, an ecosystem broadens a user's choice from specific client/provider combinations to a variety of provider-independent clients for a given task. For example, instead of choosing between a certain Facebook or Twitter client, users could open their preferred status update application to interface with any social network. This sets the stage for even more generic clients, which are not explicitly preprogrammed for certain interactions but interact with self-describing features. Third, an ecosystem of reusable features contributes to measurable improvements of machine-to-machine interfaces. This elevates web API design from a best practice-based craft to a measurable discipline, where interface decisions are guided by measurable evidence.

Furthermore, a feature-based design opens up new possibilities to scale individual APIs by enabling and disabling features at runtime. For example, during peak hours, server load can be reduced by selectively switching off features, and clients can automatically detect this and adjust. Alternatively, certain features could be activated depending on clients' subscription plans. More complex functionality (with higher server load) can be reserved for paying customers; less complex operations that achieve the same result (with higher client investment) can be available freely or at a lower price.

While this article outlines the principles for an ecosystem and suggests potential technologies to demonstrate its feasibility, we purposely leave the choice for a specific technological stack open. As experience with standards such as SOAP and WSDL has shown, the main obstacle is not technological, but rather a question of how technologies are applied to achieve reuse and, in the case of web APIs and features, at what granularity.

What an ecosystem needs foremost is a community of adopters to foster it. This involves creating the right incentives and encouraging an appropriate mindset. Especially if economic motives against substitutability exist, it is tempting to create customized web APIs. Reuse, at its core, is about opening up development, dissolving borders, and realizing cross-fertilization between parties. This challenges current web API practices and the business models created around them. It forces us to think at a longer temporal scale, perhaps even further ahead in the future than the typical lifetime of many web APIs. This will require an active community that maintains a repository of features in the long term.

Our plea for reuse, however, is only a means to an end; the ultimate goal is empowering automated clients on the web. After all, the major innovation of the web is its uniform interface—the web browser—to different providers of information and services. Despite a uniform protocol, machine clients remain confined to provider-specific interaction mechanisms, much like we were before the web. The logical next step is the realization of provider independence for machine clients, so they can also freely interact with the open web.

REFERENCES

1. S. R. Ponnekanti and A. Fox, "Application-service interoperability without standardized service interfaces," *Proceedings of the First IEEE International Conference on Pervasive Computing and Communications*, 2003, pp. 30–37.
2. S. R. Ponnekanti and A. Fox, "Interoperability Among Independently Evolving Web Services," *International Conference on Distributed Systems Platforms and Open Distributed Processing*, 2004, pp. 331–351.
3. C. Pautasso and E. Wilde, "Why is the web loosely coupled? – A multi-faceted metric for service design," *Proceedings of the 18th international conference on World Wide Web*, 2009, pp. 911–920.
4. T. Espinha, A. Zaidman, and H.-G. Gross, "Web API growing pains: Loosely coupled yet strongly tied," *Journal of Systems and Software*, vol. 100, 2015, pp. 27–43.
5. M. Robillard and R. DeLine, "A field study of API learning obstacles," *Empirical Software Engineering*, vol. 16, no. 6, 2011, pp. 703–732.
6. W. Tan et al., "From the service-oriented architecture to the Web API economy," *IEEE Internet Computing*, vol. 20, no. 4, 2016, pp. 64–68.
7. J. Sameting, *Software engineering with reusable components*, Springer Science & Business Media, 1997.
8. A. Sillitti, T. Vernazza, and G. Succi, "Service Oriented Programming: A New Paradigm of Software Reuse," *Software Reuse: Methods, Techniques, and Tools*, 2002, pp. 269–280.
9. M.-E. Fokaefs and E. Stroulia, "WSDarwin: automatic web service client adaptation," *Conference of the Center for Advanced Studies on Collaborative Research*, 2011, pp. 703–732.
10. B. Benatallah et al., "Developing Adapters for Web Services Integration," *Advanced Information Systems Engineering*, 2005, pp. 415–429; Developing Adapters for Web Services Integration.
11. R. T. Fielding, "Architectural styles and the design of network-based software architectures," 2000.
12. *The 'profile' link relation type*, E. Wilde, IETF RFC RFC6906, IETF, 2007.
13. A. Balalaie, A. Heydarnoori, and P. Jamshidi, "Microservices Architecture Enables DevOps: Migration to a Cloud-Native Architecture," *IEEE Software*, vol. 33, no. 3, 2016, pp. 42–52.
14. R. Verborgh, "Your JSON is not my JSON – a case for more fine-grained content negotiation," *Workshop on Smart Descriptions & Smarter Vocabularies*, 2016.
15. R. Verborgh et al., "Triple Pattern Fragments: a low-cost knowledge graph interface for the Web," *Journal of Web Semantics*, vol. 37–38, 2016, pp. 184–206.
16. M. Morisio, M. Ezran, and C. Tully, "Success and failure factors in software reuse," *IEEE Transactions on Software Engineering*, vol. 28, no. 4, 2002, pp. 340–357.

ABOUT THE AUTHORS

Ruben Verborgh is a professor of semantic web technology at Ghent University – imec, as well as a postdoctoral fellow of the Research Foundation – Flanders. His research interests include semantic web technologies and web architecture to build more intelligent clients. Contact him at ruben.verborgh@ugent.be.

Michel Dumontier is a Distinguished Professor of Data Science at Maastricht University. His research focuses on computational methods of organizing data and services for biomedical knowledge discovery. Contact him at michel.dumontier@gmail.com.