



Do RESTful API design rules have an impact on the understandability of Web APIs?

Justus Bogner¹ · Sebastian Kotstein² · Timo Pfaff³

Accepted: 13 July 2023 / Published online: 18 September 2023
© The Author(s) 2023

Abstract

Context Web APIs are one of the most used ways to expose application functionality on the Web, and their understandability is important for efficiently using the provided resources. While many API design rules exist, empirical evidence for the effectiveness of most rules is lacking.

Objective We therefore wanted to study 1) the impact of RESTful API design rules on understandability, 2) if rule violations are also perceived as more difficult to understand, and 3) if demographic attributes like REST-related experience have an influence on this.

Method We conducted a controlled Web-based experiment with 105 participants, from both industry and academia and with different levels of experience. Based on a hybrid between a *crossover* and a *between-subjects* design, we studied 12 design rules using API snippets in two complementary versions: one that adhered to a *rule* and one that was a *violation* of this rule. Participants answered comprehension questions and rated the perceived difficulty.

Results For 11 of the 12 rules, we found that *violation* performed significantly worse than *rule* for the comprehension tasks. Regarding the subjective ratings, we found significant differences for 9 of the 12 rules, meaning that most violations were subjectively rated as more difficult to understand. Demographics played no role in the comprehension performance for *violation*.

Conclusions Our results provide first empirical evidence for the importance of following design rules to improve the understandability of Web APIs, which is important for researchers, practitioners, and educators.

Keywords Web API design · Understandability · controlled experiment

Communicated by: Sebastian Baltes

✉ Justus Bogner
j.bogner@vu.nl
Sebastian Kotstein
sebastian.kotstein@reutlingen-university.de

¹ University of Stuttgart, Institute of Software Engineering, Stuttgart, Germany

² Reutlingen University, Herman Hollerith Zentrum, Reutlingen, Germany

³ Independent researcher, Germany; research conducted while at University of Stuttgart, Institute of Software Engineering, Stuttgart, Germany

1 Introduction

Technologies for Web Application Programming Interfaces (APIs) like WSDL, SOAP, and HTTP are the technical foundations for realizing modern Web applications (Jacobson et al., 2011). These technologies allow developers to independently implement smaller software components and share their functionality on the Internet. Successfully reusing an existing component, however, depends on the ability to understand its purpose and behavior, especially its API, which hides internal logic and complexity. In cases where neither additional documentation nor developers of the original component can be consulted, the Web API might even be the first and only point of contact with the exposed functionality. Therefore, understandability is an important quality attribute for the design of Web APIs (Palma et al., 2017).

Over the last two decades, HTTP combined with other well-established Web standards like URI has become a popular choice for realizing Web APIs that expose their functionality through Web resources (Schermann et al., 2016; Bogner et al., 2019). In these resource-oriented Web APIs, the role of HTTP has shifted from a transport mechanism for XML-based messages to an application-layer protocol for interacting with the respective API (Pautasso et al., 2008). With Representational State Transfer (REST) (Fielding and Taylor, 2002), there exists an architectural style that formalizes the proper use of Web technologies like HTTP and URIs in Web applications. REST is considered a foundation for high-quality, so-called *RESTful* API design, and it describes a set of constraints for the recommended behavior of Web applications, e.g. HTTP-based Web APIs. However, it does not instruct developers how to implement this behavior (Rodríguez et al., 2016).

Since there exist various interpretations and (mis-)understandings among practitioners how RESTful API design looks like, users and integrators of these services are confronted with a multitude of heterogeneous interface designs, which can make it difficult to understand a given Web API (Palma et al., 2021). Therefore, several works have proposed design rules and best practices to complement the original REST constraints and to guide developers when designing and implementing Web APIs, e.g., Pautasso (2014), Palma et al. (2017), Richardson and Ruby (2007), and Massé (2011). In most cases, however, we do not have sufficient empirical evidence for the effectiveness of these RESTful API design rules, i.e., if they really have a positive impact on the quality of Web APIs.

Additionally, multiple studies analyzed the degree of REST compliance in practice by systematically comparing real-world Web APIs against proposed design rules. Many of these works, e.g., Neumann et al. (2018), Renzel et al. (2012), and Rodríguez et al. (2016), concluded that only a small degree of real-world Web APIs are truly RESTful. This suggests that many practitioners perceive proposed design rules differently in terms of their importance. We have provided confirmation for this in previous work (Kotstein and Bogner, 2021). In a Delphi study, we confronted industry practitioners with 82 RESTful API design rules by Massé (2011) to find out which ones they perceived as important and how they perceived their impact on software quality. Only 45 out of 82 rules were rated with high or medium importance, and maintainability and usability were the most associated quality attributes. Both of these attributes are closely related to understandability.

To confirm these opinion-based results with additional empirical evidence, we conducted a controlled Web-based experiment, in which we presented 12 Web API snippets to 105 participants with at least basic REST-related experience. Each API snippet existed in two versions, one adhering to a design rule and one violating the rule (see, e.g., Fig. 1). The participants' task was to answer comprehension questions about each snippet, while we

GET /groups/{groupId}/members

Parameters

Name	Description
groupId * required	groupId
string	
(path)	

What does this endpoint return?

Choose one of the following answers

- A specific group that a specific member is part of
- An array of groups that a specific member is part of
- A map containing all groups and members
- A specific member that belongs to a specific group
- An array of members that belong to a specific group

GET /groups/{groupId}/member

Parameters

Name	Description
groupId * required	groupId
string	
(path)	

What does this endpoint return?

Choose one of the following answers

- A specific member that belongs to a specific group
- A specific group that a specific member is part of
- A map containing all groups and members
- An array of members that belong to a specific group
- An array of groups that a specific member is part of

Fig. 1 Example of a return value question based on the rule `PluralNoun` (rule on the left, violation on the right, correct answer is checked)

measured the required time. Furthermore, participants also had to rate the perceived difficulty to understand an API snippet. In this paper, we present the design and results of our controlled experiment on the understandability impact of RESTful API design rules.

2 Background and related work

We start with a discussion of terminology around Web APIs and REST and explain how we use these terms in the paper. Furthermore, we mention existing works that propose rules and best practices for RESTful API design, and present existing studies about Web API quality.

2.1 Terminology

In this paper, we focus on resource-oriented HTTP-based Web APIs. In distinction to SOAP-/WSDL-based APIs (Pautasso et al., 2008), we use the term *Web API* for any resource-oriented API that exposes its functionality via HTTP and URIs at the application level. We consider a Web API as *RESTful*, i.e., a so-called *RESTful API*, if the respective Web API satisfies all mandatory REST constraints defined by Fielding and Taylor (2002). As a consequence, a Web API that implements only some RESTful API design rules cannot automatically be considered as RESTful.

2.2 Best practices for REST in practice

To combat the potentially harmful heterogeneity of interface designs among Web APIs, several works tried to translate the REST constraints into more concrete guidelines to instruct developers how to achieve good RESTful design with HTTP. RESTful API design rules and best practices have been proposed in scientific articles, e.g., by Pautasso (2014), Petrillo et al. (2016), Palma et al. (2014), and Palma et al. (2017), but also in textbooks, e.g., by Richardson and Ruby 2007, Massé (2011), and Webber et al. (2010). Moreover, Leonard Richardson developed a maturity model (Martin Fowler, 2010) allowing Web developers to estimate the degree of REST compliance of their Web APIs. Incorporating the principles of REST, the model defines four levels of maturity.

- Level 0: Web APIs offer their functionality over a single URI and use HTTP solely as a transport protocol for tunneling requests through this endpoint by using POST. Examples of level 0 are SOAP- and XML-RPC-based services.
- Level 1: Web APIs use the concept of resources, i.e., expose different URIs for different resources. However, operations are still identified via URIs or specified in the request payload rather than by different HTTP methods.
- Level 2: Web APIs use HTTP mechanisms and semantics, including different HTTP methods for different operations and semantically correct status codes. Level 2 partially aligns with the *Uniform Interface* constraint.
- Level 3: Web APIs additionally conform to the HATEOAS constraint (“Hypermedia As The Engine Of Application State”) by embedding hypermedia controls into responses to advertise semantic relations between resources and to offer navigational support to clients.

Despite the existence of these design rules and guidelines, multiple studies, e.g., by Neumann et al. (2018), Renzel et al. (2012), and Rodríguez et al. (2016), revealed that only a small number of existing Web APIs are indeed RESTful, although many Web APIs claimed to be RESTful (Neumann et al., 2018). This suggests that there is still no common understanding of how a RESTful API should look like in industrial practice. Providing empirical evidence may therefore help to identify the effective and important rules from the large collection of existing guidelines.

2.3 Related work

Several works investigated the quality of real-world Web APIs, mainly by analyzing interfaces, their descriptions, and exchanged HTTP messages.

Rodríguez et al. (2016) analyzed more than 78 GB of HTTP traffic to gain insights into the use of best practices in Web API design. They applied API-specific heuristics to extract API-related messages from the whole data set and validated these extracted API requests and responses against 18 heuristics aligned with REST design principles and best practices. For a few heuristics, they described their negative effect on maintainability and evolvability when violating associated design principles and best practices. Moreover, they mapped the heuristics to the levels of the Richardson maturity model to estimate the level of REST compliance of investigated Web APIs. The paper concluded that only a few APIs reach level 3 of the maturity model, but the majority of investigated APIs complied with level 2.

The governance of RESTful APIs was the focus for Haupt et al. (2018): using a framework developed by Haupt et al. (2017), they conducted a structural analysis of 286 real-world Web

APIs. In detail, the framework takes an interface description, converts it into a canonical metamodel, and calculates several metrics to support API governance. As a usage example, they demonstrated how calculated metrics can be used to estimate the user-perceived complexity of an API, which is related to API understandability. For this, they randomly selected 10 of their 286 APIs and let 9 software developers rank them based on their perceived complexity. Based on their knowledge and experience, the authors then defined several metrics for user-perceived complexity and calculated these metrics for the 10 APIs. As a result, some calculated metrics coincided with the developers' judgments and were proposed for automatic complexity estimation. A follow-up confirmatory study with a larger sample size to substantiate this exploratory approach is missing so far.

An approach similar to Haupt et al. (2018) was used in a study by Bogner et al. (2019). They proposed a modular framework, the *RESTful API Metric Analyzer* (RAMA), which calculates maintainability metrics from interface descriptions and enables the automatic evaluation of Web APIs. More precisely, RAMA converts an interface description into a hierarchical model and calculates 10 service-based maintainability metrics. In a benchmark run, the authors applied RAMA to a set of 1,737 real-world APIs, and calculated quartile-based thresholds for the metrics. However, relating the metrics to other software quality correlates is missing to evaluate their effectiveness.

The impact of good and poor Web API design on understandability and reusability has been investigated in a series of publications by Palma et al.: in (Palma et al., 2017), the authors defined 12 linguistic patterns and antipatterns focusing on URI design in Web APIs, which may impact understandability and reusability of such APIs. Moreover, they proposed algorithms for their detection and implemented them as part of the *Service-Oriented Framework for Antipatterns* (SOFA). They used SOFA to detect linguistic (anti-)patterns in 18 real-world APIs, with the result that most of the investigated APIs used appropriate resource names and did not use verbs within URI paths. However, URI paths often did not convey hierarchical structures.

In another study, Palma et al. (2021) tried to answer whether a well-designed RESTful API also has good linguistic quality and, vice versa, whether poorly designed Web APIs have poor linguistic quality. For this, they used SOFA to analyze 8 Google APIs and to detect 9 design patterns and antipatterns, as well as 12 linguistic patterns and antipatterns. However, their statistical tests revealed only negligible relationships between RESTful design and linguistic design qualities.

Subsequently, they extended SOFA with further linguistic (anti-)patterns, improved approaches for their detection, and applied the linguistic quality analysis on Web APIs from the IoT domain (Palma et al., 2022b) or compared the quality between public, partner, and private APIs (Palma et al., 2022a).

In summary, existing studies assessed the quality of real-world Web APIs by collecting metrics and detecting (anti-)patterns. The latter are somewhat related to RESTful API design rules and best practices that should, in theory, improve several quality aspects of an API. However, in many cases, there is no empirical evidence for the effectiveness of the impact of these design rules and best practices on software quality, especially on understandability. Haupt et al. (2018) and Kotstein and Bogner (2021) used subjective ratings, but no studies in which human participants solve comprehension or maintenance tasks have been conducted. To the best of our knowledge, our experiment is the first study that investigated the understandability impact of violating RESTful API design rules from the perspective of human API consumers.

Table 1 Experiment overview

Goal	Study the impact of design rules on the understandability of Web APIs
Study objects	12 design rules compiled from Massé (2011) and Richardson and Ruby (2007), two functionally equivalent API snippets per rule (one follows the rule, one violates it)
Participants	105 people with at least basic REST-related experience (both students and professionals, both from academia and industry)
Setting	Online experiment via LimeSurvey
Tasks	Answering comprehension question about API snippets for RQ1 (12 per participant), rating the difficulty to understand an API snippet for RQ2 (12 per participant)
Dependent variables	Timed actual understandability (TAU) for RQ1, perceived difficulty rating for RQ2
Treatments	API snippet follows a <i>rule</i> (version 1) or API snippet contains a <i>violation</i> of a rule (version 2)
Other independent variables	Demographic attributes like REST-related experience or current role (RQ3)
Design	Hybrid between crossover and between-subjects design (2 counterbalanced sequences, each with 6 tasks for <i>rule</i> and 6 for <i>violation</i>)

3 Research design

In this section, we describe the details of our methodology. We roughly follow the reporting structure for software engineering experiments proposed by Jedlitschka et al. (2008). Inspired by the experiment characteristics discussed by Wyrich et al. (2022), Table 1 provides a quick overview of the most important characteristics of the study. For transparency and reproducibility, we publish our experiment artifacts on Zenodo¹.

3.1 Research questions

We investigated three different research questions in this study.

RQ1: Which design rules have a significant impact on the understandability of Web APIs?

Our hypothesis for this central, confirmatory RQ was that each selected rule should improve understandability, i.e., the effectiveness and efficiency of grasping the functionality and intended purpose of a Web API endpoint.

RQ2: Which design rules have a significant impact on software professionals' perceived difficulty while understanding Web APIs?

For the confirmatory part of this RQ, we hypothesized that API snippets with rule violations are rated as more difficult to understand. Additionally, we analyzed the correlation between actual and perceived understandability in an exploratory part to identify potential differences.

RQ3: How do participant demographics influence the effectiveness and perception of design rules for understanding Web APIs?

We did not have strong hypotheses for this exploratory RQ. Nonetheless, we had some intuitions about attributes that may be interesting to analyze. For example, it could be possible that adhering to the design rules mostly has an influence on experienced professionals but not on students (or vice versa). Furthermore, some rules might require the participant to know about the Richardson maturity model or some rule violations might be perceived as more

¹ <https://doi.org/10.5281/zenodo.7381500>

critical by participants from academia or from industry. During the study design phase, we selected some general demographic attributes plus several specific ones for the experiment context.

3.2 Participants and sampling

The only requirements for participation were basic knowledge of REST and HTTP, as well as the ability to understand English. Our goal was to attract participants from diverse backgrounds and experience levels, e.g., both students and professionals, both participants from industry and academia, etc. We used *convenience sampling* mixed with *referral-chain sampling*, i.e., we distributed the call for participation within our personal networks via email, and kindly asked for forwarding to relevant circles (Baltes and Ralph, 2022). A similar message was displayed after the experiment to encourage sharing. Students were recruited via internal mailing lists of several universities. Moreover, we advertised the study via social media, such as Twitter², LinkedIn³, XING⁴, and in several technology-related subreddits⁵.

3.3 Experiment objects

In this experiment, the objects under study were 12 design rules for RESTful APIs that have been proposed in the literature. They are summarized in Table 2 together with short identifiers that we use throughout the rest of the paper. Rule selection was guided by the results of our previous Delphi study (Kotstein and Bogner, 2021), i.e., we focused on rules from Massé (2011) that were perceived as very important by industry experts, with an influence on maintainability or usability. Additionally, we included three instances of the *PathHierarchy* rule proposed by Richardson and Ruby (2007). While this rule was not part of our previous study, it has strong relationships to Massé's rules *Variable path segments may be substituted with identity-based values* and *Forward slash separator (/) must be used to indicate a hierarchical relationship*. Both of these rules fulfill the above criteria, i.e., high importance plus influence on maintainability or usability. In the following, the rule descriptions are taken from Massé (2011) and Richardson and Ruby (2007) respectively. For each rule, we also present the concrete endpoint pair that was used in the experiment, one version for following the rule and one for violating it. These concrete Web API examples and the *violation* versions are based on our industry experience, but also on existing public APIs that we identified via the APIs Guru repository⁶. For our experiment, the chosen real-world examples were adapted to simplify them and to avoid that participants are already familiar with the presented HTTP endpoints. During the pilot, we discussed the created pairs of *rule* and *violation* with external experts to validate if the *violation* snippets were not strongly exaggerated. Several snippets were adapted based on this feedback, and one task was dropped entirely.

² <https://www.twitter.com>

³ <https://www.linkedin.com>

⁴ <https://www.xing.com>

⁵ <https://www.reddit.com>

⁶ <https://apis.guru>

Table 2 Selected design rules for RESTful APIs

ID	Rule	Source	Category
PluralNoun	A plural noun should be used for collection and store names	Massé (2011)	URI Design
VerbController	A verb or verb phrase should be used for controller names	Massé (2011)	URI Design
CRUDNames	CRUD function names should not be used in URIs	Massé (2011)	URI Design
PathHierarchy (3 versions)	Use path variables to encode hierarchy	Richardson and Ruby (2007)	Hierarchy Design
NoTunnel	GET and POST must not be used to tunnel other request methods	Massé (2011)	Request Methods
GETRetrieve	GET must be used to retrieve a representation of a resource	Massé (2011)	Request Methods
POSTCreate	POST must be used to create a new resource in a collection	Massé (2011)	Request Methods
NoRC200Error	200 (OK) must not be used to communicate errors in the response body	Massé (2011)	HTTP Status Codes
RC401	401 (Unauthorized) must be used when there is a problem with the client's credentials	Massé (2011)	HTTP Status Codes
RC415	415 (Unsupported Media Type) must be used when the media type of a request's payload cannot be processed	Massé (2011)	HTTP Status Codes

3.3.1 URI design

The three rules in this category are concerned with the concrete design of URI paths in an API.

PluralNoun Massé (2011) defines this rule for both *collections* and *stores*, which we merge into a single rule for simplicity. In both cases, it prescribes to use a plural noun as the name in the URI. A violation of this rule would be to use a singular noun for a *collection* or *store* name instead.

Rule:	GET /groups/{groupId}/members
Violation:	GET /groups/{groupId}/member

VerbController A controller provides an action that cannot be easily mapped to a typical CRUD operation on a resource. In relation to function names in source code, Massé (2011) proposes to always use a verb or verb phrase for controller resources. Using a noun instead would be a violation.

Rule:	POST /servers/{serverId}/backups/{backupId}/restore
Violation:	POST /servers/{serverId}/backups/{backupId}/restoration

CRUDNames Based on the invoked HTTP method, a RESTful API selects the semantically equivalent CRUD operation to perform. Therefore, Massé (2011) prescribes not to use CRUD function names like “create” or “update” in URIs, especially not with incorrect HTTP verbs. Adhering to this rule means solely relying on the HTTP verb to indicate the wanted CRUD operation. Our chosen example also includes the rule “DELETE must be used to remove a resource from its parent” (Massé, 2011) for the violation.

Rule:	DELETE /messaging-topics/{topicId}/queues/{queueId}
Violation:	GET /messaging-topics/{topicId}/delete-queue/{queueId}

3.3.2 Hierarchy design

A prominent rule from Richardson and Ruby (2007) prescribes the use of path parameters to encode the hierarchy of resources in a URI. Since there are several possibilities to apply and interpret this, we created three specific rules based on this idea (PathHierarchy1 to PathHierarchy3).

PathHierarchy1 (path params vs. query params) Version 1 explores the difference between using path parameters (*rule*) and query parameters (*violation*) for retrieving a hierarchically structured resource.

Rule:	GET /shops/{shopId}/products/{productId}
Violation:	GET /shops/products?shopId={shopId}&productId={productId}

PathHierarchy2 (top-down vs. bottom-up) In version 2, the difference between structuring the hierarchy top-down / from left to right (*rule*) and bottom-up / from right to left (*violation*) is tested.

Rule:	GET /companies/{companyId}/employees
Violation:	GET /employees/companies/{companyId}

PathHierarchy3 (hierarchical path vs. short path) Lastly, version 3 analyzes differences during the creation of a resource when either using a long hierarchical path with parameters (*rule*) or a short path with parameters in the request body (*violation*).

Rule:	POST /customers/{customerId}/environments/{environmentId}/servers
Violation:	POST /servers

3.3.3 Request methods

In accordance with level 2 of the Richardson maturity model, Massé (2011) proposes several rules that prescribe that each HTTP method should exclusively be used for its semantically equivalent operation. We selected three of these rules for our experiment.

NoTunnel One of these rules states that GET and POST must not be used to tunnel other request methods, which might seem tempting for the sake of simplicity. In our example, an API correctly uses PUT to update a resource (*rule*), whereas the *violation* always uses POST and tunnels the update operation via an additional query parameter.

Rule:	PUT /trainings/{trainingId}/organizers/{organizerId}
Violation:	POST /trainings/{trainingId}/organizers/{organizerId}?operation=update

GETRetrieve In similar fashion, another rule from Massé (2011) states that the HTTP method GET must be used to retrieve a representation of a resource. Since GET requests have no request body, it may seem tempting to use POST in some cases to be able to use a JSON object instead of overly complex query parameters. A typical example of this is a search resource. Adhering to the *rule* requires using GET plus query parameters, while reverting to POST plus a request body with the search options is a *violation*.

Rule:	GET /events?date=2022-10-03&category=music
Violation:	POST /events/search

POSTCreate Lastly, we tested the complementary rule for POST, namely that this method must be used to create resources in a collection (Massé, 2011). A typical violation of this rule is the use of PUT to create a resource, as seen in our chosen example.

Rule:	POST /customers/{customerId}/orders
Violation:	PUT /customers/{customerId}/orders

3.3.4 HTTP status codes

Another important theme for RESTful API design is the correct usage of HTTP status codes with response messages. Massé (2011) provides a number of rules in this area, from which we selected three in total, namely the rules for 200 (OK), 401 (Unauthorized), and 415 (Unsupported Media Type). Contrary to the previous three categories, the two versions of our used examples (*rule* and *violation*) do not differ in the displayed endpoint, but only for the displayed response including the response code.

NoRC200Error The first rule in this category states that the response code 200 (OK) must not be used in case of error. Instead, client-side (4XX) or server-side (5XX) error codes must be used accordingly, based on the nature of the error.

Rule:	A required parameter is missing in the request body. The server indicates there was a problem and correctly responds with 400 (Bad Request).
Violation:	A required parameter is missing in the request body. The server indicates there was a problem, but incorrectly responds with 200 (OK).

RC401 A similar rule prescribes the use of the status code 401 (Unauthorized) in case of issues with client credentials, e.g., during a login attempt. Using a different client-side response code like 400 (Bad Request) or 403 (Forbidden) should be avoided for such errors. Since the nuances between 401 (not logged in or login failed) and 403 (logged-in user does not have the required privileges) might not be fully clear to all participants, we opted for the more generic code 400 (Bad Request) in the violation example.

Rule:	A secured resource is accessed with an empty Bearer token in the Authorization header. The server indicates there was a problem and correctly responds with 401 (Unauthorized).
Violation:	A secured resource is accessed with an empty Bearer token in the Authorization header. The server indicates there was a problem, but incorrectly responds with 400 (Bad Request).

RC415 The final rule in this category focuses on the correct usage of the status code 415 (Unsupported Media Type), which must be returned if the client uses a media type for the request body that cannot be processed by the server. A typical example is a request body in XML when the server only supports JSON. Returning a different client-side error code, e.g., 400 (Bad Request), should be avoided in this case.

Rule:	A request body contains XML, even though the server only accepts JSON. The server indicates there was a problem and correctly responds with 415 (Unsupported Media Type).
Violation:	A request body contains XML, even though the server only accepts JSON. The server indicates there was a problem, but incorrectly responds with 400 (Bad Request).

3.4 Material

To incorporate the created API rule examples into our experiment, we relied on the OpenAPI specification format⁷, one of the most popular ways to document Web APIs many practitioners and researchers should be familiar with. Using the Swagger editor⁸, we created two OpenAPI documents, one with the examples following the rules and one with the violation examples. Each document contained 12 endpoints, one per rule. We then created screenshots of the graphical representation of each resource, with the purpose of showing them to participants with each task (see, e.g., Fig. 1).

To reach a larger and more diverse audience, we decided to conduct an online experiment via a web-based tool. We selected the open-source survey tool LimeSurvey⁹ for this purpose, as it provides all the features we need. Additionally, we had access to an existing LimeSurvey

⁷ <https://www.openapis.org>

⁸ <https://editor.swagger.io>

⁹ <https://www.limesurvey.org>

instance via one of our universities. It supports several types of questions, is highly customizable, and also allows measuring the duration per task, i.e., survey question, which we needed for our experiment. Lastly, random assignment of participants to sequences is also possible. This setup meant that participants exclusively used LimeSurvey for the experiment via a computing device and web browser of their choice. All necessary information was provided this way.

3.5 Tasks

The participants' main task was to inspect and understand several Web API snippets that were presented in the graphical representation of the Swagger editor. To evaluate understanding, participants had to answer one comprehension question per snippet. Each of these questions was a single choice question with five different options presented in a random order, with exactly one of them being correct. The options were the same regardless of whether the version adhering to the rule or the one violating the rule was displayed. Opting for single choice comprehension questions had several advantages. While free-text answers might reflect a more in-depth understanding of participants, they are much harder to correct. Additionally, writing free-text answers takes more time and effort for participants, which may increase the drop-out rate. Lastly, participants may specify answers with different level of details. This not only complicates the grading, but it also influences the time to answer, which we include in our comprehension measures. Depending on the shown snippet, exactly one of three different types of comprehension questions was asked.

Return value: Participants had to determine the return value of an endpoint, i.e., the type of entity and if a single object or a collection was returned. This question type was used for snippets with GET requests, namely `PluralNoun`, `PathHierarchy1`, and `PathHierarchy2`. An example is shown in Fig. 1.

Endpoint purpose: Participants had to determine the purpose of the shown endpoint, i.e., what operation or functionality was executed on invocation. This question type was used for snippets which do not always return entities, e.g., `POST`, `PUT`, and `DELETE` requests, namely `VerbController`, `CRUDNames`, `PathHierarchy3`, `NoTunnel`, `GETRetrieve`, and `POSTCreate`. An example is presented in Fig. 2.

Response code reason: Participants had to determine why a certain request failed or what the outcome of a request was. This question type was exclusively used for snippets of the category *HTTP Status Codes*, namely `NoRC200Error`, `RC401`, and `RC415`. An example is presented in Fig. 3.

After each comprehension question, participants also had to rate the perceived difficulty of understanding the API snippet on a 5-point ordinal scale, ranging from *very easy* (1) to *very hard* (5). For this purpose, the last snippet was shown again, so participants did not have to rely solely on memory.

3.6 Variables and hypotheses

Two **dependent variables** were used in this experiment. For RQ1, we needed a metric for understandability. To operationalize this quality attribute, we collected both the *correctness* and *duration* for each task per participant. Since every task had a single correct answer, *correctness* was a binary variable, with 0 for false and 1 for correct. The required *duration* per task was documented in seconds. To combine these two measures into a single variable, we adapted an aggregation procedure from Scalabrino et al. (2021), namely *Timed Actual*

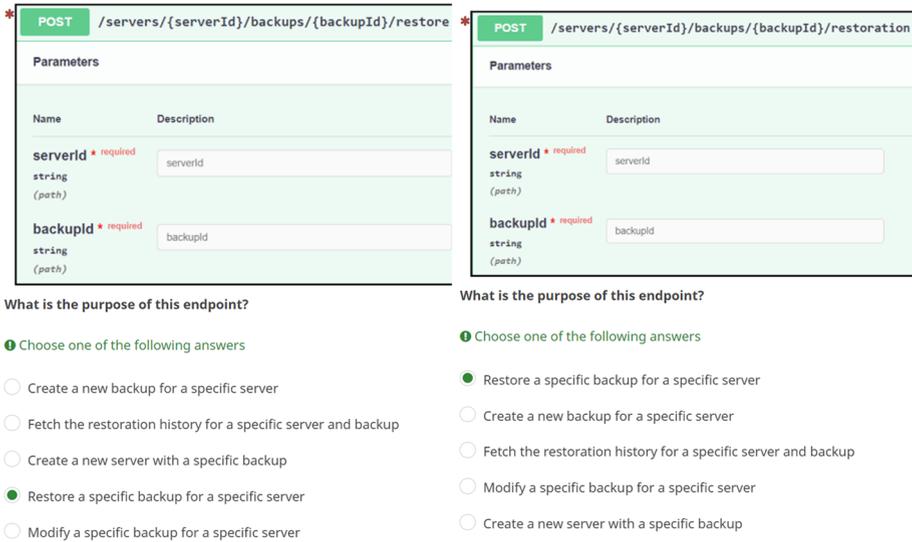


Fig. 2 Example of an endpoint purpose question based on the rule *VerbController* (*rule* on the left, *violation* on the right, correct answer is checked)

Understandability (TAU). In our experiment, TAU for a participant *p* and task *t* was calculated as follows:

$$TAU_{p,t} = correctness_{p,t} \times (1 - \frac{duration_{p,t}}{max(durations_t)}) \tag{1}$$

TAU produces values between 0 and 1, with values closer to 1 indicating a higher degree of understandability. For an incorrect answer, TAU is always 0. For a correct answer, the task duration is set in relation to the maximum duration that was recorded for this task. This is then inverted by subtracting it from 1, meaning the faster the correct answer was found, the greater is TAU. As such, TAU represents a pragmatic aggregation of *correctness* and *duration* that respects differences between participants in the sample and leads to easily interpretable values. Even though it leads to unusual distributions (see, e.g., Fig. 8), we therefore chose TAU as the dependent variable for RQ1. For RQ2, the dependent variable was the *perceived difficulty*, i.e., the rating that participants had to give after each task using a 5-point ordinal scale, ranging from *very easy* (1) to *very hard* (5).

As the directly controlled **independent variable**, we used the version number of the respective snippet. This was either version 1 (*rule*) that followed the rule or version 2 (*violation*) that violated the rule. Additionally, we had uncontrollable independent variables that we collected for RQ3, namely various demographic attributes of participants such as their current role, years of experience with REST, or knowledge of the Richardson maturity model.

Based on these variables, we formulated hypotheses for the confirmatory questions RQ1 (difference in actual understandability between *rule* and *violation*) and RQ2 (difference in perceived understandability). They are displayed in Table 3. In both cases, we expected that the version following the design rule would lead to significantly better results in comparison to the version violating the rule. While we only list two hypotheses, each of the 12 rules was tested individually to clearly identify which rules have an impact and which do not. Finally, the exploratory RQ3 did not have a clear hypothesis.

*You send a request to the following API:

```
GET /live-streams/{liveStreamId}
```

Parameters

Name	Description
liveStreamId * required	liveStreamId
string	
(path)	

The header of your request looks like this:

```
GET /live-streams/123456 HTTP/1.1
Host: ws.api.video
Accept: application/json
Accept-Language: en-us,en;q=0.5
Accept-Charset: ISO-8859-1,utf-8
Authorization: Bearer ""
Keep-Alive: 300
Connection: keep-alive
Cache-Control: no-cache
```

The server responds with the following message:

```
401 - Unauthorized
```

What was the reason for the failed request?

Choose one of the following answers

- The request body is malformed
- The bearer token is empty
- An error on the server side
- The location of the resource has changed
- No permission for accessing the specified live stream

*You send a request to the following API:

```
GET /live-streams/{liveStreamId}
```

Parameters

Name	Description
liveStreamId * required	liveStreamId
string	
(path)	

The header of your request looks like this:

```
GET /live-streams/123456 HTTP/1.1
Host: ws.api.video
Accept: application/json
Accept-Language: en-us,en;q=0.5
Accept-Charset: ISO-8859-1,utf-8
Authorization: Bearer ""
Keep-Alive: 300
Connection: keep-alive
Cache-Control: no-cache
```

The server responds with the following message:

```
400 - Bad Request: There was a problem with your request
```

What was the reason for the failed request?

Choose one of the following answers

- The bearer token is empty
- The request body is malformed
- The location of the resource has changed
- No permission for accessing the specified live stream
- An error on the server side

Fig. 3 Example of a response code reason question based on the rule RC401 (rule on the left, violation on the right, correct answer is checked)

Table 3 Null hypotheses with their alternatives for the two confirmatory RQs

RQ	Metric	Null Hypothesis	Alternative Hypothesis
RQ1	Timed actual understand-ability (TAU)	H_0^1 : Web API snippets adhering to design rules are equally or less understandable than snippets violating rules.	H_1^1 : Web API snippets adhering to design rules are more understandable than snippets violating rules.
RQ2	Perceived difficulty	H_0^2 : Web API snippets adhering to design rules are rated as equally or more difficult to understand than snippets violating rules.	H_1^2 : Web API snippets adhering to design rules are rated as less difficult to understand than snippets violating rules.

3.7 Experiment design

We opted for a mixed experiment design that can be described as a hybrid between a *crossover design* (Vegas et al., 2016) and a *between-subjects design* (Wohlin et al., 2012), thereby exploiting specific benefits of both. In a crossover design (a special form of a within-subjects design), each participant receives each treatment at least once, but the order in which participants receive the treatments differs. Our participants worked on six tasks with the *rule* version and six tasks with *violation*. However, in typical crossover experiments in software engineering, each participant applies all treatments to a given experiment object (Wohlin et al., 2012), which was not the case for our experiment. Only half of our participants worked on each version of a specific task, which is the between-subject characteristic of our experiment. This mixed design is still fairly robust regarding inter-participant differences because participants are not assigned to a single treatment. Letting each participant work on both treatments for the same rule would be even more robust, but would also create problems with familiarization effects. It also would increase the experiment duration considerably, leading to fewer participants and higher mortality or fewer rules that could be tested. Lastly, this design still provides us with many observations to overall compare *rule* vs. *violation*.

To avoid suboptimal task orderings regarding treatments or task categories that could cause carryover or order effects, we did not randomize the order of tasks. Instead, we consciously designed two sequences with counterbalanced task orders (see Table 4), which is common for crossover designs (Vegas et al., 2016). Our goals were to spread out the tasks of the same category, to ensure the same number of *rule* and *violation* tasks, to avoid having the same treatment too often in a row, and to ensure that *rule* and *violation* for the same snippet appear at the same position per sequence. Participants were randomly assigned to one of these sequences at the start of the experiment.

This experiment design was the result of several iterations with internal discussion, followed by a pilot with external reviewers. Based on the pilot results, we refined the tasks and

Table 4 Task orders for the two experiment sequences

#	Rule	Category	Sequence 1	Sequence 2
1	VerbController	URI Design	violation	rule
2	RC415	HTTP Status Codes	violation	rule
3	PathHierarchy2	Hierarchy Design	rule	violation
4	GETRetrieve	Request Methods	violation	rule
5	PathHierarchy3	Hierarchy Design	rule	violation
6	NoRC200Error	HTTP Status Codes	violation	rule
7	NoTunnel	Request Methods	rule	violation
8	PluralNoun	URI Design	rule	violation
9	POSTCreate	Request Methods	violation	rule
10	PathHierarchy1	Hierarchy Design	violation	rule
11	CRUDNames	URI Design	rule	violation
12	RC401	HTTP Status Codes	rule	violation

1 GET **/cars/{carId}/extras/{extraId} 2**

Parameters

Name	Description
carId * required string (path)	carId
extrald * required string (path)	extrald

3

Fig. 4 Example API snippet from the introduction page, explanations were provided for the numbered elements

survey text. Furthermore, we removed a few previously planned tasks to keep the required participation time below 15 minutes, leading to the total of 12 tasks.

3.8 Experiment execution

Participation in the online experiment via LimeSurvey was open for a period of approximately three weeks. At the start of and during this period, we actively promoted the experiment URL within in our network. Participants starting the experiment were first presented with a welcome page containing some general information about the study. We explained that our goal with this survey was to investigate the understandability of different Web API designs, and that basic REST knowledge was the only requirement for participation. We mentioned that the experiment should take between 10-15 minutes and should ideally be finished in one sitting. For each of the first 100 participants, we pledged to donate 1 € to UNICEF¹⁰ as additional motivation.

Afterward, we described the tasks, namely that we will present 12 different API snippets. For each snippet, they have to answer a multiple-choice question with five different options, e.g., “What is the purpose of this endpoint?”. We would measure the time it takes to answer, but answering correctly would be more important. After each comprehension question, they would rate how difficult it was to understand the API snippet without time measurements. Finally, there would be a few demographic questions. To familiarize participants with the used API snippet visualization, we also presented an example snippet with additional explanations about the different elements (see Fig. 4).

The last part of the welcome page contained our privacy policy. We explained what data would be requested and that data would be used for statistical analysis and possibly be included in scientific publications. Additionally, we emphasized that participation was strictly

¹⁰ <https://www.unicef.org>

voluntary and anonymous, and that participants could choose to abandon the experiment at any time, resulting in the deletion of the previously entered data. Participants had to consent to these terms before being able to start the experiment.

Upon accepting our privacy policy, participants started working on the tasks of their randomly assigned sequence (see Table 4). For each of the 12 tasks, participants first analyzed the presented API snippet and answered the comprehension question. Afterward, participants had to rate the difficulty of understanding the studied API snippet. The snippet was displayed again for this purpose. For the experiment part, participants answered a total of 24 questions, one comprehension and one rating question for each task.

Finally, participants answered the demographic questions, namely their country of origin, current role, technical API perspective (either API user / client developer, API developer / designer, or both), years of professional experience with REST, knowledge of the Richardson maturity model, and (if they knew about the model) their opinion about the minimal required maturity level a Web API should possess to be considered RESTful. We also provided an optional free-text field for any final remarks or feedback participants wanted to give. Finally, an outro page was presented, where we thanked participants and kindly asked them to forward the experiment URL to suitable colleagues.

3.9 Experiment analysis

To analyze the experiment results, we first exported all responses as a CSV file. We then performed the following data cleaning and transformation steps:

- Removing 48 incomplete responses, i.e., participants that aborted the experiment at some point (rationale: avoid including responses without demographics, participants not fully committed, or participants who realized themselves that their background was not suitable)
- Resolving and harmonizing free-text answers for current role and technical API perspective (“Other:”)
- Harmonizing country names
- Adding binary variables (1 or 0) for `is_Student`, `is_Academia` (an academic professional), and `is_Germany`; if both `is_Student` and `is_Academia` were 0, the participant was an industry practitioner

No substantial ambiguity was identified during the harmonization, and no final participant comments had to be considered for adapting the assigned correctness score. However, since TAU is sensitive to outliers of the measured duration, we also analyzed the durations for all comprehension questions. If answering the question took less than 5 seconds or more than 3 minutes, we judged this as invalid and removed the individual answer for all three dependent variables (correctness, time, and subjective rating). The rationale for this was that we did not want to consider any responses (not even the subjective rating) where it was very likely that the participant had not concentrated fully on the task, e.g., by answering before reading everything or leaving the experiment for more than a few seconds to do something else. Since questions of the type *HTTP Status Codes* were more verbose, the allocated thresholds for these questions were 10 seconds and 4 minutes. If these duration thresholds were triggered three or more times for the same participant, we removed their complete response instead of only the individual answers. Due to this filtering, 2 complete responses (initially, we had 107 complete responses) and 12 individual answers were removed. The number of valid responses considered in the analysis is displayed for each individual task and treatment in Table 5.

Table 5 Number of valid responses per individual task and treatment

#	Rule	# of responses for <i>rule</i>	# of responses for <i>violation</i>
1	VerbController	53	52
2	RC415	53	52
3	PathHierarchy2	52	53
4	GETRetrieve	53	52
5	PathHierarchy3	52	52
6	NoRC200Error	53	50
7	NoTunnel	50	52
8	PluralNoun	52	52
9	POSTCreate	53	51
10	PathHierarchy1	53	50
11	CRUDNames	51	53
12	RC401	51	53

The cleaned CSV file was then imported by an analysis script written in R¹¹, which had also been tested and refined during the pilot study. The script performs basic data transformation, calculates TAU, and provides general descriptive statistics as well as diagrams to visualize the data, such as box plots. To select a suitable hypothesis test for RQ1 and RQ2, we first analyzed the data distributions with the Shapiro-Wilk test (Shapiro and Wilk, 1965). For both dependent variables, the test resulted in a p-value $\ll 0.05$, i.e., the data did not follow a normal distribution, which called for a test without the assumption of normality. We therefore opted for the non-parametric Wilcoxon-Mann-Whitney test (Neuhäuser, 2011), which has a mature R implementation in the `stats` package¹². To combat the multiple comparison problem (in our case, the testing of 12 rules), we applied the Holm-Bonferroni correction (Shaffer, 1995). We used the `p.adjust()` function from the `stats` package¹³ to adjust the computed p-values. When an adjusted p-value was less than our targeted significance level of 0.05, we rejected the null hypothesis and accepted the alternative. To judge the effect size of accepted hypotheses, we additionally calculated Cohen's *d* (Cohen, 1988) using the `effectsize` package¹⁴. Following Sawilowsky (2009), the values can be interpreted as follows:

- $d < 0.2$: very small effect
- $0.2 \leq d < 0.5$: small effect
- $0.5 \leq d < 0.8$: medium effect
- $0.8 \leq d < 1.2$: large effect
- $1.2 \leq d < 2.0$: very large effect
- $d \geq 2.0$: huge effect

For the exploratory RQ3, we first used a correlation matrix for visual exploration. Correlations of identified variable pairs were then further analyzed with Kendall's Tau (Kendall, 1938), as it is more robust and permissive regarding assumptions about the data than other

¹¹ <https://www.r-project.org>

¹² <https://www.rdocumentation.org/packages/stats/versions/3.6.2/topics/wilcox.test>

¹³ <https://www.rdocumentation.org/packages/stats/versions/3.6.2/topics/p.adjust>

¹⁴ <https://www.rdocumentation.org/packages/effectsize/versions/0.8.2>

Table 6 Participant demographics by sequence (RMM: Richardson Maturity Model)

Sequence	#	REST exp. in years (median)	# of industry participants	# of academic participants	# of students	# with RMM knowledge
1	52	3.5	31 (60%)	9 (17%)	12 (23%)	9 (17%)
2	53	4	28 (53%)	9 (17%)	16 (30%)	19 (36%)
Total	105	4	59 (56%)	18 (17%)	28 (27%)	28 (27%)

methods. Lastly, we used linear regression¹⁵ to further analyze combined effects of demographic attributes and to explore the potential for predictive modelling.

4 Results and discussion

In this section, we first present some general statistics about our participants. Afterward, we provide the results for each RQ, starting with descriptive statistics and then presenting the hypothesis testing or correlation results.

4.1 Participant demographics

After the cleaning procedure, we were left with 105 valid responses, which is more than most software engineering experiments have.¹⁶ 96 of these were complete, while we had to exclude one or two answers for 9 of them. 52 participants were randomly assigned to sequence 1 and 53 to sequence 2. Table 6 compares the attributes for the two sequences, which are pretty similar in most cases. Overall, our 105 participants had between 1 and 15 years of experience with REST, with a median of 4 years. 59 participants were from industry (56%), 18 were professionals from academia (17%), and 28 were students (27%).

Most of our participants, namely 70 of 105 (67%), were located in Germany, followed by Portugal (12), the US (6), and Switzerland (5). The remaining 12 responses were distributed across 7 countries with between 1 and 3 participants. Regarding the roles of the 59 industry participants, the majority of them were software engineers (41). Eight were consultants, two were software architects, and two were team managers. The remaining six each had a role that was only mentioned once, e.g., test engineer. Concerning the technical API perspective, most participants reported to be both active in API development and usage (76), while 16 exclusively were API users / client developers, 8 exclusively were API developers, and 5 did not provide an answer to this optional question. Interestingly, only 28 participants (27%) reported knowing the Richardson maturity model. When these 28 were then asked their opinion about the minimal maturity level that a Web API should have to be RESTful, 22 chose level 2 (79%) and 6 level 3 (21%). No one selected level 0 and 1. This may be an indication that the Richardson maturity model is not very well known and that HATEOAS is not perceived as an important requirement for RESTfulness by most professionals. Especially the latter is in line with previous findings (Kotstein and Bogner, 2021).

¹⁵ <https://www.rdocumentation.org/packages/stats/versions/3.6.2/topics/lm>

¹⁶ See, e.g., the mapping study on code comprehension experiments by Wyrich et al. (2022): median number of participants was 34 (for journals alone, it was 61)

4.2 Impact on understandability (RQ1)

For our central research question, we wanted to analyze which design rules have a significant impact on understandability, as measured via comprehension tasks and TAU. We visualize the results for the percentage of correct answers (Fig. 5), the required time to answer (Fig. 6), and TAU (Fig. 7) per individual rule and treatment. For a more detailed comparison, Tables 8, 9, and 10 in the appendix list the descriptive statistics for this per task.

For 11 of the 12 tasks, participants with the *rule* version performed better than participants with the *violation* version, i.e., mean TAU was higher for *rule*. Correctness was often the deciding factor, e.g., for the tasks `CRUDNames` (96% vs. 38%), `GETRetrieve` (100% vs. 60%), or `NoRC200Error` (98% vs. 58%). In other cases, correctness was much closer between the two treatments, but participants with *violation* required more time. This is, e.g., visible for the tasks `PluralNoun` (100% vs. 94%, but 20.82 s vs. 32.23 s) or `RC401` (75% vs. 70%, but 43.42 s vs. 67.21 s). The only surprising exception was the task `PathHierarchy3`, where participants with *violation* performed notably better (mean TAU of 0.6141 vs. 0.7505).

To further visualize the experiment performance, we created strip plots of TAU for all rules, which makes it easier to understand its unusual distribution and to compare the two treatments. Figure 8 shows the TAU distributions for the rules in the categories *URI Design* and *Hierarchy Design*. Incorrect answers are displayed as dots at the bottom (TAU = 0), with the median value being displayed as an orange diamond. For all three *URI Design* tasks (1-3), it immediately becomes apparent that *violation* (red) performed worse, i.e., the median of *violation* is below *rule* in all cases. This difference is especially large for the task `CRUDNames`, which constitutes the worst performance for *violation*. Additionally, the values are more spread out for *violation* in all three tasks.

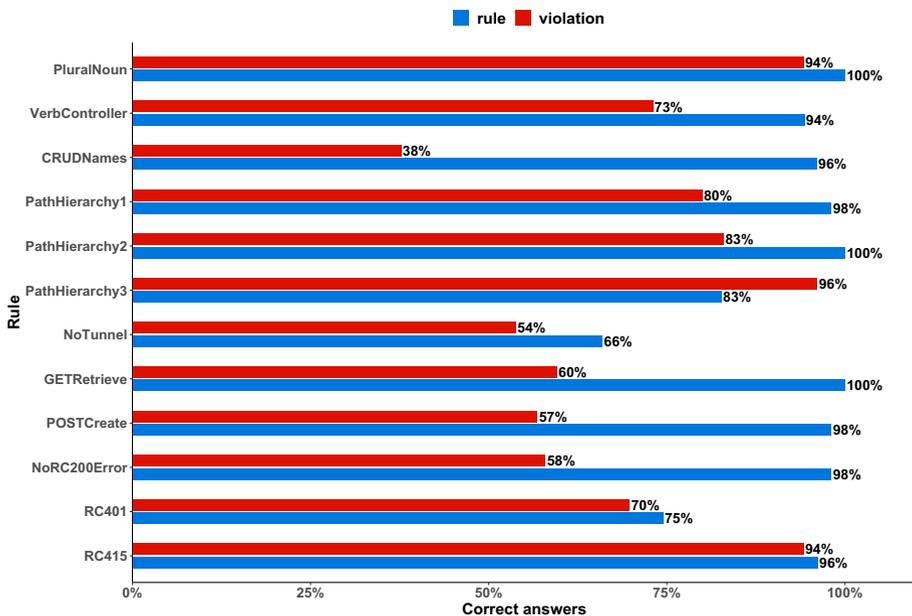


Fig. 5 Comparison of correctness between treatments *rule* and *violation* (higher is better)

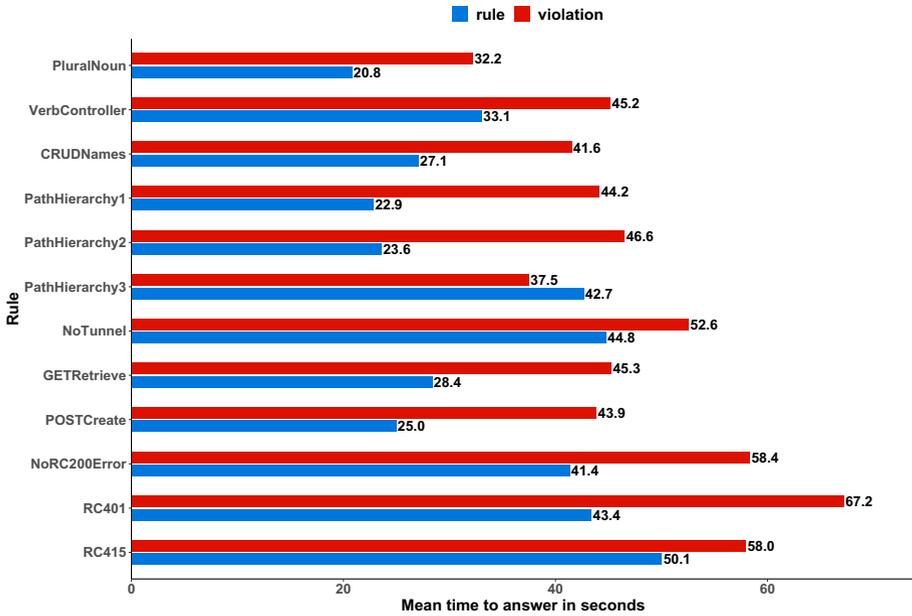


Fig. 6 Comparison of duration between treatments *rule* and *violation* (lower is better)

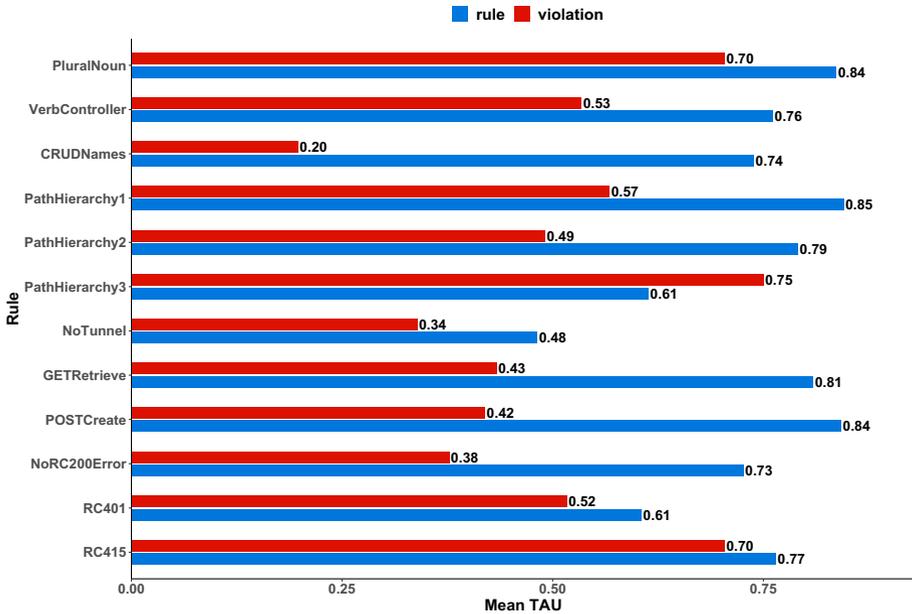


Fig. 7 Comparison of TAU between treatments *rule* and *violation* (higher is better)

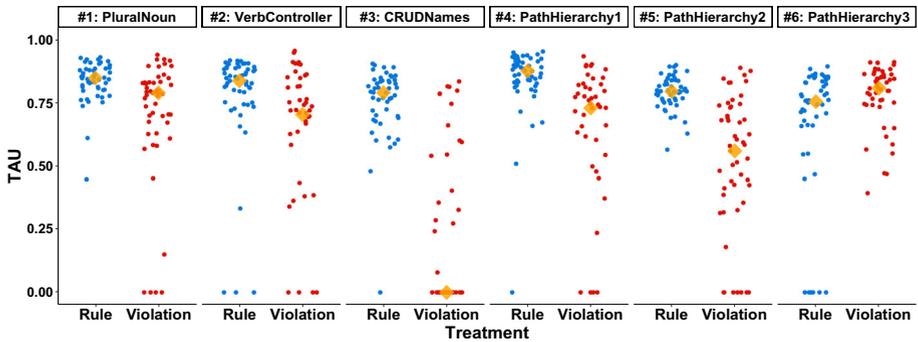


Fig. 8 TAU distributions for the categories *URI Design* (1-3) and *Hierarchy Design* (4-6), median values are displayed as orange diamonds

For the rules of the *Hierarchy Design* category (4-6), we see a similarly improved performance for *rule* in the case of *PathHierarchy1* and *PathHierarchy2*. However, the exception is *PathHierarchy3*, which shows a visibly better performance and less spread for *violation*. A potential explanation could be that the general difference in path length might have been the deciding factor, i.e., `POST /customers/{customerId}/environments/{environmentId}/servers` vs. the much simpler *violation* `POST /servers`. Even though following the rule provides richer details about the resource hierarchies (“there are customers, who have IT environments like *Production*, in which servers are placed”), it seems to make the endpoint purpose more difficult to understand, in this case that a new server is created. Additionally, from the five response options, three were about creating a new resource, which fits to the `POST` of the shown endpoint. However, it was probably clearer to exclude the options “Create multiple new servers...” and “Create an environment...” for *violation*, as “environment” did not occur in the path. The needed time was not that different (42.69 s vs. 37.51 s), but only 83% answered correctly for *rule*, while 96% did the same for *violation*. Extrapolating from the other *PathHierarchy* rules (1 and 2), we might state that using a path hierarchy correctly, i.e., from left to right with sequential pairs of `{collectionName} / {id}`, is significantly better than using it *incorrectly*, but not automatically better than *not using it at all*. However, more research is needed to confirm this, e.g., by analyzing other variants of *PathHierarchy3*, e.g., with `GET` instead of `POST`, and identifying potential thresholds below which the path length might become irrelevant.

Figure 9 presents the TAU distributions for the categories *Request Methods* and *HTTP Status Codes*. For *Request Methods* (7-9), the *rule* version shows again better performance in all three tasks, with substantial distance between treatment medians. However, in the `NoTunnel` task, the spread of TAU values is more similar between *rule* and *treatment* than for other tasks. Overall, `NoTunnel` was the worst performance for *rule*, with only 66% answering it correctly. One explanation may be that several practitioners may not see it as a violation to use `PUT` to create a new resource.

In the final *HTTP Status Codes* category (10-12), the performance of *rule* is also better in all three cases, however, to varying degrees. For `NoRC200Error`, the differences are substantial, but for `RC401` and especially `RC415`, the median values are not as far apart. In the case of `RC401`, the spread in TAU values is also much more similar between the two treatments. This was the second-worst performance for *rule*, with only 75% providing the correct answer.

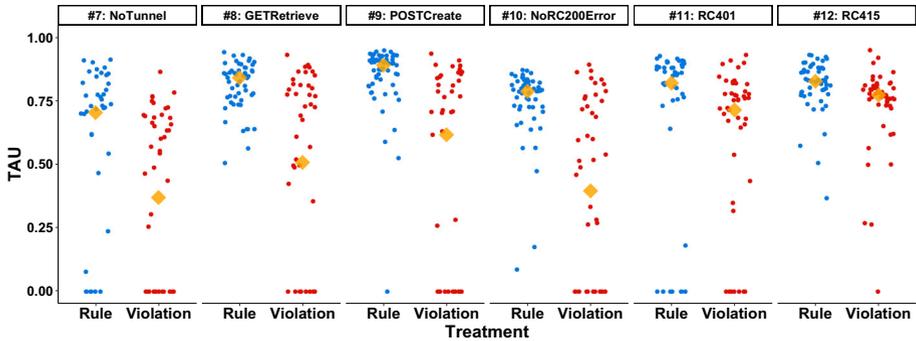


Fig. 9 TAU distributions for the categories *Request Methods* and *HTTP Status Codes*, median values are displayed as orange diamonds

To verify if the visually identified differences between *rule* and *violation* were statistically significant, we continued with hypothesis testing. Despite the applied Holm-Bonferroni correction, **11 of our 12 hypothesis tests produced significant results**, i.e., 11 of 12 rules had significant impact on understandability. The only outlier was *PathHierarchy3*, where *violation* had performed better. Additionally, we calculated the effect size for the significant tests via Cohen's *d*. We visualize the *d* values in Fig. 10. For a more detailed comparison of the test results, please refer to Table 11 in the appendix.

The rules *NoTunnel*, *RC401*, and *RC415* produced a *small* effect ($0.2 < d < 0.5$) and *PluralNoun* and *VerbController* a *medium* one ($0.5 < d < 0.8$). The remaining six rules showed even stronger effects, namely a *large* one for *PathHierarchy1* ($0.8 < d < 1.2$), a *very large* one for *NoRC200Error*, *GETRetrieve*, *POSTCreate*, and *PathHierarchy2* ($1.2 < d < 2.0$), and even a *huge* one for *CRUDNames* ($d > 2.0$). It is difficult to determine the most impactful category of rules, as the six rules with a Cohen's *d* > 1.0 cover all four categories. For *URI Design*, violating *CRUDNames* had by far the largest effect ($d = 2.17$), but *VerbController* ($d = 0.75$) and *PluralNoun* ($d = 0.72$) were in the bottom half. The categories *Request Methods* and *Hierarchy Design* also had two impactful rules, but *NoTunnel* produced only a small effect and *PathHierarchy3*

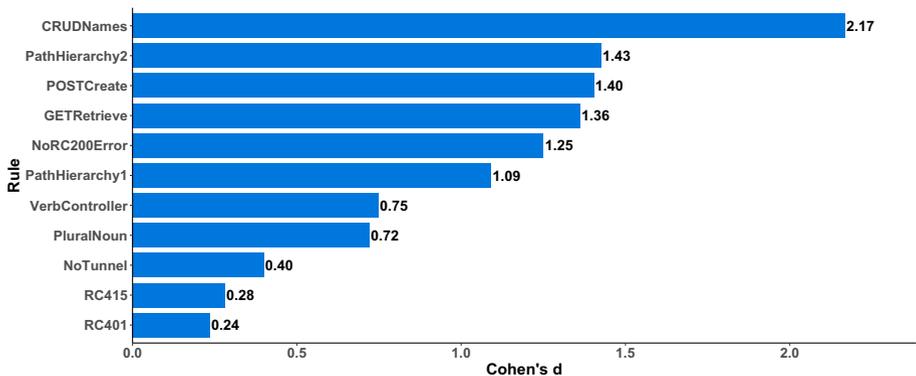


Fig. 10 Effect sizes for the 11 significant differences between TAU for treatments *rule* and *violation*, ordered by Cohen's *d*, insignificant *PathHierarchy3* not included

none at all. Lastly, *HTTP Status Codes* was the least impactful category. Even though *NoRC200Error* had a Cohen's $d = 1.25$, both RC401 and RC415 resulted in small effects.

Results for RQ1: For 11 of the 12 tasks (*PathHierarchy3* being the exception), adhering to the design rule resulted in significantly better comprehension performance. Effect sizes were between 0.24 and 2.17, with 6 rules resulting in a Cohen's $d > 1.0$ (*large* up to *huge* effects). Violating the rules *CRUDNames*, *PathHierarchy2*, and *POSTCreate* had the strongest negative impact on understandability.

4.3 Impact on perceived difficulty (RQ2)

After showing that the majority of design rules led to significantly better comprehension performance, we analyzed if the results were similar for the subjective *perceived* understandability. Table 7 summarizes the results for the perceived difficulty ratings and sets them in relation to the corresponding TAU values. The mean difficulty rating is lower for *rule* in all 12 tasks, while the median difficulty rating is only lower for 10, the exceptions being RC415 and again *PathHierarchy3* (median rating of 2 for both treatments). In four tasks, namely *VerbController*, *CRUDNames*, *GETRetrieve*, and *NoRC200Error*, the difference in median rating is 1 point, while it is even higher for the remaining six tasks (between 1.5 and 3 points). To further analyze differences, we visualized the results with a comparative bar plot of the ratings 1 (*very easy*) and 2 (*easy*) in Fig. 11. For a full Likert plot, please refer to Fig. 14 in the appendix.

The difference between *rule* and *violation* becomes apparent for many of the tasks here, e.g., for *CRUDNames*, *PathHierarchy1*, *PathHierarchy2*, and RC401. However, in some tasks, the treatment ratings also appear to be decently close to each other, despite

Table 7 Descriptive statistics for RQ2, perceived difficulty rating ranging from 1 (*very easy*) to 5 (*very difficult*), mean TAU provided for comparison

Task	Median difficulty		Mean difficulty		Mean TAU	
	<i>rule</i>	<i>violation</i>	<i>rule</i>	<i>violation</i>	<i>rule</i>	<i>violation</i>
PluralNoun	1	2.5	1.25	2.40	0.8369	0.7046
VerbController	2	3	2.13	2.50	0.7617	0.5345
CRUDNames	2	3	1.73	3.36	0.7392	0.1982
PathHierarchy1	1	3	1.34	2.94	0.8457	0.5678
PathHierarchy2	1	4	1.23	3.36	0.7916	0.4913
PathHierarchy3	2	2	2.31	2.42	0.6141	0.7505
NoTunnel	2	4	2.34	3.27	0.4820	0.3396
GETRetrieve	2	3	1.83	2.77	0.8095	0.4337
POSTCreate	1	3	1.45	2.75	0.8427	0.4201
NoRC200Error	2	3	1.75	2.80	0.7268	0.3776
RC401	2	4	1.80	3.45	0.6058	0.5170
RC415	2	2	2.00	2.35	0.7654	0.7044

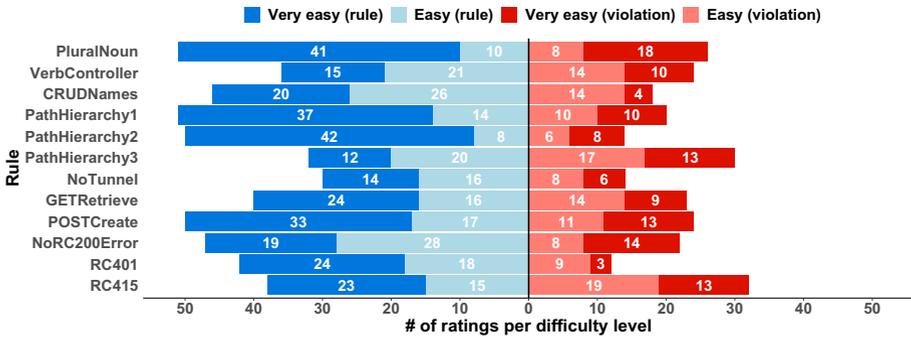


Fig. 11 Bar plots of perceived difficulty (RQ2) for the ratings 1 (very easy) and 2 (easy), rule ratings on the left, violation ratings on the right

the median rating being different, e.g., for VerbController or GETRetrieve. In general, the violations perceived as the least difficult to understand were PathHierarchy3, RC415, and PluralNoun.

To confirm if the differences between rule and violation are significant, we again applied hypothesis testing. This time, we only found significant differences for 9 of the 12 tasks in the perception of rule and violation. In addition to PathHierarchy3 and RC415 with equal medians, the difference between treatments for VerbController was also not significant. We visualize the values for Cohen’s *d* in Fig. 12. For a more detailed comparison, please refer to Table 12 in the appendix. The effect sizes for the significant tasks ranged from 0.79 (medium) to 2.06 (huge), with 8 rules producing a Cohen’s *d* > 0.8 (large and higher). Violating the rules PathHierarchy2, RC401, PathHierarchy1, and CRUDNames had the strongest impact on difficulty perception.

Especially for PathHierarchy2 with Cohen’s *d* = 2.06 (huge), we see that the constructed URI for violation (GET /employees/companies/{companyId}) was a very extreme case leading to much confusion. Mixing the URI path segments in this way seems to have made it much more difficult for participants to identify the cardinalities of the original domain model (all employees belonging to a specific company). For PathHierarchy3, the missing effect was to be expected and complementary to the insignif-

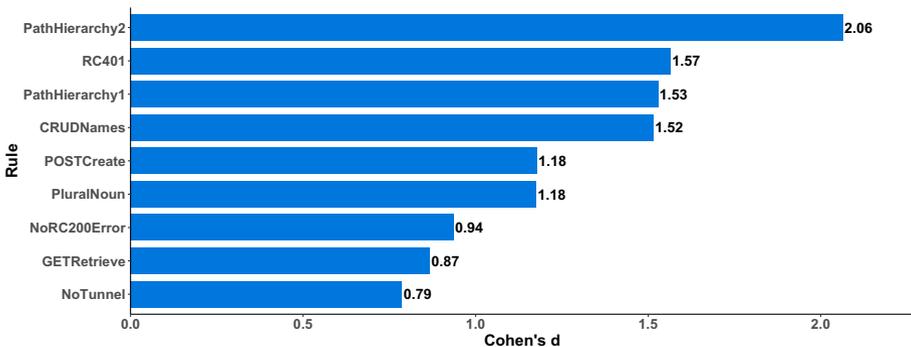


Fig. 12 Effect sizes for the 9 significant differences between perceived difficulty for treatments rule and violation, ordered by Cohen’s *d*, insignificant PathHierarchy3, RC415, and VerbController not included

ificant results for TAU. However, `VerbController` and `RC415` are not as easy to explain, as both of them had a significant impact on the actual understandability. While *violation* was perceived as slightly more difficult than *rule* in both tasks, this difference was not statistically significant. For `RC415`, this might be explained by the fairly small effect size for the TAU difference ($d = 0.28$) and the similar levels of correctness per treatment. If participants only needed a bit more time but were still fairly certain to have found the correct answer, they might not have directly associated this task with a high difficulty. For `VerbController`, however, this explanation is not applicable, as both correctness and time differed per treatment, and it produced a medium effect ($d = 0.75$). This makes it an especially dangerous rule to violate because it has considerable impact, but many people might not notice that something is ambiguous or unclear with the endpoint.

Other rule violations that may be especially critical can be identified by analyzing the correlations between TAU and perceived difficulty for the *violation* treatment. Ideally, we would like to have significant negative correlations between the two dependent variables for all *violation* API snippets, i.e., the worse the comprehension performance of a rule violation, the higher the perceived difficulty ratings should be. Rule violations where this is not the case may indicate that participants felt confident in their performance despite answering incorrectly or slowly because they unknowingly misunderstood the API snippet. For nine tasks, there was *no* significant negative correlation between TAU and perceived difficulty, the exceptions being `PathHierarchy3`, `VerbController`, and `GETRetrieve`. For the detailed correlation results, please refer to Table 13 in the appendix. As a measure of explanatory power, we visualize the adjusted R^2 values for the regression between TAU and the perceived difficulty ratings for all *violation* API snippets in Fig. 13. For many of them, the values are close to zero, meaning that the perceived difficulty ratings cannot explain any variation in TAU. Among them are also rule violations that have a substantial impact on understandability, like `CRUDNames`, `PathHierarchy2`, or `POSTCreate`, which makes these rule violations especially problematic.

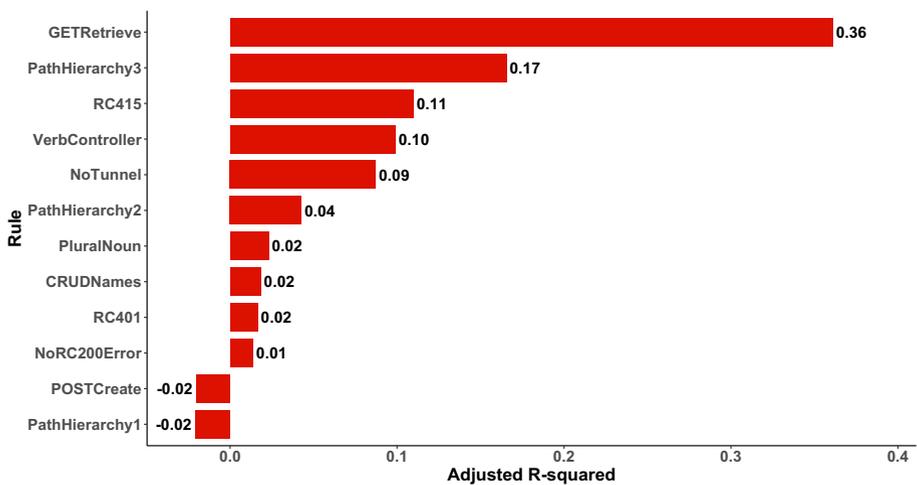


Fig. 13 Adjusted R^2 values for the regression between TAU and perceived difficulty ratings for *violation*, ordered by adjusted R^2 , higher is better

Results for RQ2: For 9 of the 12 tasks (except for PathHierarchy3, RC415, and VerbController), adhering to the API design rule resulted in significantly smaller perceived difficulty ratings than violating it. Effect sizes were between 0.79 and 2.06, with 6 rules resulting in a Cohen's $d > 1.0$ (*large* up to *huge* effects). Additionally, there was no significant relationship between perceived difficulty and TAU for many *violation* tasks, making these rule violations especially dangerous.

4.4 Relationships with demographic attributes (RQ3)

For RQ3, we analyzed if there were any relationships between the dependent variables (TAU and the perceived difficulty ratings) and demographic attributes of our participants. We explored this question separately for each treatment (*rule* and *violation*), and compared the results. Studied predictors were being from Germany, being from academia (vs. industry), being a student, being an API developer / designer, years of professional experience with REST, having knowledge of the Richardson maturity model, and the preferred minimal maturity level. Even though this was an exploratory RQ, we used Holm-Bonferroni adjusted p-values and a significance level of $\alpha = 0.05$.

Overall, we did not find many significant relationships, and no strong ones at all. Years of experience with REST had a small positive correlation with TAU for *rule* (Kendall's $\tau = 0.1956$, $p = 0.0311$), i.e., for the API snippets adhering to the rules, participants with more experience had a slight tendency to perform better. However, this was not the case for API snippets violating the rules. Similarly, participants knowing the Richardson maturity model also tended to perform slightly better for *rule*, with a Kendall's $\tau = 0.1941$. However, after adjusting the original p-value ($p = 0.0158$), this correlation was no longer significant ($p = 0.0791$). The deciding factor here should be years of experience, though, as it was also positively correlated with knowing the Richardson maturity model in our sample (Kendall's $\tau = 0.3760$, $p < 0.001$). Conversely, knowledge of the Richardson maturity model was also positively correlated with the perceived difficulty ratings for *violation*, i.e., if participants knew about this model, they tended to rate the API snippets violating the rules as slightly more difficult to understand (Kendall's $\tau = 0.3035$, $p = 0.0014$). This correlation was absent for *rule*. All other demographic attributes did not produce any significant relationships.

While the identified correlations were small, they still seem to highlight differences between the two treatments that may provide an explanation. Experience and knowledge about REST is only linked to better experiment performance if no design rules are violated. If rules are violated, it does not matter much if people are more experienced: their performance still suffers. However, people who know advanced REST-related concepts are at least more likely to *notice* that something is wrong with the rule-violating API snippets, even though this does not help to understand them better. This theory also seems to be supported by our results of trying to build linear regression models to predict TAU per treatment based on the demographic attributes. While both models (*rule* and *violation*) were unable to provide reliable predictions for the majority of our sample, the model for *violation* performed considerably worse. The model for *rule* was able to explain a decent percentage of variability (adjusted $R^2 = 0.2171$, $p = 0.0618$), while the one for *violation* had no explanatory power at all (adjusted $R^2 = -0.1792$, $p = 0.9674$).

Results for RQ3: Participants with more REST-related experience and knowledge of the Richardson maturity model performed slightly better within *rule*, but not within *violation*. Conversely, people knowing the Richardson maturity model perceived the *violation* API snippets as slightly more difficult to comprehend. This seems to indicate that violating design rules decreases understandability regardless of demographic attributes like experience.

5 Threats to validity

This section describes how we tried to mitigate potential threats to validity, and which threats and limitations to our results remain. We discuss these issues mainly through the perspectives provided by Wohlin et al. (2012).

Construct validity is concerned with relating the experiment and especially its collected measures to the studied concepts. This includes whether our dependent variables were adequate representations of the constructs. Our measure for understandability, namely TAU as a combination of correctness and time (Scalabrino et al., 2021), has been used in several studies before. It provides the advantage of increasing the information density in a single statistical test. However, TAU's trade-off between time and correctness is not ideal when correctness is binary. TAU therefore definitely has limitations, but looking at time and correctness separately to interpret the results allows reducing these threats. Regarding the perceived understandability, it is an accepted practice to use ordinal scales that are balanced and symmetrical for subjective ratings.

To avoid *hypothesis guessing* (Wohlin et al., 2012), participants were told that the goal of the experiment was to analyze the understandability of different Web API designs. We see it as very unlikely that participants figured out the true goal of the experiment and then deliberately gave worse responses for *violation* tasks. Lastly, it is possible that our experiment may be slightly impacted by a *mono-operation bias* (Wohlin et al., 2012), i.e., the underrepresentation of the construct by focusing on a single treatment. While we used different categories with several rules each in the experiment, we still only compared one exemplary rule implementation to its complementary violation in isolation. We did not combine several rules or involved other RESTful API concepts, which could have led to a richer theory.

Internal validity can suffer from threats that may impact the dependent variables without the researchers' knowledge, i.e., *confounders* (Wohlin et al., 2012). In general, a crossover design is fairly robust against many confounders by reducing the impact of inter-participant differences, e.g., large expertise or experience differences between participants. The relatively short experiment duration of 10-15 minutes caused by our between-subjects characteristic also made most *history* (Wohlin et al., 2012) or *maturation* (Wohlin et al., 2012) effects unlikely, e.g., noticeable changes in performance due to tiredness or boredom. However, due to the similar task structures, learning effects are very likely. Randomization of the task order could have mitigated this, but it could also have led to suboptimal sequences of

treatments and task categories, which could have introduced harmful carryover effects. Our fixed sequences at least guaranteed that both treatments for the same task appeared at the same position, thereby ensuring a similar level of maturation per experiment object. Additionally, randomization was used to assign participants to the two sequences and to display the different answers per comprehension question.

This study was conducted as an online experiment, with considerably less control over the experiment environment. While *random irrelevancies in the experimental setting* (Wohlin et al., 2012) might have occurred in some cases, such as reduced concentration due to loud noise or an interruption by colleagues or family members, this potential increase in variance still did not impact our hypothesis tests. Moreover, we also had no means to prevent people from participating several times, as responses were anonymous. Due to no reasonable incentive for this, we deem this threat as very unlikely.

External validity is the extent to which the results are generalizable to other settings or parts of the population. Our sample was fairly diverse, with participants with different levels of experience from both industry and academia. With 105 participants, our sample size was also decently large when compared to many other software engineering experiments, even though we need to remember that only half of our participants worked on each treatment per individual rule due to our hybrid design. While we reached statistical significance for most hypotheses, a larger sample would have improved the generalizability of the results even more. Nonetheless, we see it as unlikely that the *interaction of selection and treatment* threat (Wohlin et al., 2012) could impact our results. Even though most of our sample was located in Germany (67%), we do not believe that country-specific differences might have a noticeable influence on the results. To combat the threat *interaction of setting and treatment* (Wohlin et al., 2012), we ensured using realistic API concepts and violations inspired by real-world examples. The Swagger editor is also a very popular tool in the area of Web APIs.

Finally, we need to emphasize that the understandability of Web APIs was tested in a consciously constructed setting, without a hosted implementation for manual exploration or additional API documentation. Participants could only rely on the provided API snippets in the visual notation from the Swagger editor. This is obviously different from a real-world environment, e.g., in industry, where a software engineer trying to use a Web API may have access to the API documentation or even a running API instance for manual testing. However, in many cases, access to more documentation or the API itself is also not available in the real world. Additionally, using these artifacts also requires additional time, i.e., being able to understand the purpose of an endpoint without having to consult other materials is still preferable. Moreover, it is also plausible that some rules might even have a stronger effect in a real-world setting, e.g., RC401 and RC415: getting access to HTTP header information requires a lot more effort there, while we conveniently presented the headers in the experiment. All in all, we believe that our results are decently transferrable to less controlled real-world environments. To analyze the full degree of this generalization, follow-up research is necessary, e.g., based on repository mining or industrial case studies.

6 Conclusion and future work

In this paper, we presented the design and results of our controlled Web-based experiment on the understandability impact of 12 RESTful API design rules from Massé (2011) and Richardson and Ruby (2007). In detail, we presented 12 Web API snippets to 105 participants, asked them comprehension questions about each snippet, and let them rate the

perceived understandability. For 11 of 12 rules, we identified a significant negative impact on understandability for the *violation* treatment. Effect sizes ranged from *small* to *huge*, with Cohen's d between $d = 0.24$ and $d = 2.17$. Furthermore, our participants also rated 9 of 12 rule violations as significantly more difficult to understand.

All in all, our results indicate that violating commonly accepted design rules for RESTful APIs has a negative impact on understandability, regardless of REST-related experience or other demographic factors. For several rule violations, we could also show that they are prone to misinterpretation and misunderstandings, making them especially dangerous. Practitioners should therefore respect these rules during the design of Web APIs, as understandability is linked with important quality attributes like maintainability and usability. This becomes especially important for publicly available APIs that are meant to be used by many external people. Providing comprehensive API documentation may partly mitigate understandability problems caused by rule violations, but is still an insufficient solution.

In the future, additional experiments should try to replicate these findings with other samples of the population or other violations of the same rules, and potentially extend the evidence to other rules not tested in our study. To enable such studies, we publicly share our experiment materials¹⁷. Additionally, tool-supported approaches to automatically identify these rule violations will be helpful for practitioners trying to ensure the quality of their Web APIs. Such tool support would also pave the way for large-scale studies of these rule violations in less controlled environments, such as software repository mining or industrial case studies.

Appendix

¹⁷ <https://doi.org/10.5281/zenodo.7381500>

Table 8 Descriptive statistics for RQ1

Task	Mean TAU		Mean duration (s)		# of correct answers	
	<i>rule</i>	<i>violation</i>	<i>rule</i>	<i>violation</i>	<i>rule</i>	<i>violation</i>
PluralNoun	0.8369	0.7046	20.82	32.23	52 (100%)	49 (94%)
VerbController	0.7617	0.5345	33.07	45.22	50 (94%)	38 (73%)
CRUDNames	0.7392	0.1982	27.10	41.55	49 (96%)	20 (38%)
PathHierarchy1	0.8457	0.5678	22.85	44.19	52 (98%)	40 (80%)
PathHierarchy2	0.7916	0.4913	23.62	46.55	52 (100%)	44 (83%)
PathHierarchy3	0.6141	0.7505	42.69	37.51	43 (83%)	50 (96%)
NoTunnel	0.4820	0.3396	44.80	52.60	33 (66%)	28 (54%)
GETRetrieve	0.8095	0.4337	28.40	45.29	53 (100%)	31 (60%)
POSTCreate	0.8427	0.4201	24.99	43.88	52 (98%)	29 (57%)
NoRC200Error	0.7268	0.3776	41.37	58.35	52 (98%)	29 (58%)
RC401	0.6058	0.5170	43.42	67.21	38 (75%)	37 (70%)
RC415	0.7654	0.7044	50.05	57.97	51 (96%)	49 (94%)

Table 9 Extended descriptive statistics for TAU (RQ1)

Task	Min TAU		Max TAU		Variance of TAU	
	<i>rule</i>	<i>violation</i>	<i>rule</i>	<i>violation</i>	<i>rule</i>	<i>violation</i>
PluralNoun	0.4479	0	0.9322	0.9418	0.0067	0.0606
VerbController	0	0	0.9197	0.9572	0.0577	0.1276
CRUDNames	0	0	0.9065	0.8362	0.0320	0.0913
PathHierarchy1	0	0	0.9550	0.9359	0.0206	0.1118
PathHierarchy2	0.5652	0	0.9013	0.8902	0.0043	0.0836
PathHierarchy3	0	0	0.8959	0.9129	0.1008	0.0378
NoTunnel	0	0	0.9139	0.8657	0.1436	0.1112
GETRetrieve	0.5064	0	0.9433	0.9326	0.0095	0.1441
POSTCreate	0	0	0.9502	0.9374	0.0219	0.1621
NoRC200Error	0	0	0.8730	0.8941	0.0323	0.1263
RC401	0	0	0.9280	0.9323	0.1533	0.1312
RC415	0	0	0.9324	0.9511	0.0468	0.0477

Table 10 Extended descriptive statistics for response time (RQ1)

Task	Min duration (s)		Max duration (s)		Variance of duration	
	<i>rule</i>	<i>violation</i>	<i>rule</i>	<i>violation</i>	<i>rule</i>	<i>violation</i>
PluralNoun	8.65	7.43	70.48	127.66	109.92	510.73
VerbController	11.55	7.38	172.27	113.62	671.21	660.98
CRUDNames	11.05	12.59	61.47	118.14	130.32	660.62
PathHierarchy1	7.45	6.53	81.19	165.46	181.03	934.83
PathHierarchy2	11.19	12.45	49.28	113.35	54.90	495.81
PathHierarchy3	16.05	14.20	162.99	140.22	638.31	607.81
NoTunnel	13.40	17.02	155.65	134.77	971.09	874.44
GETRetrieve	8.45	10.05	73.58	149.08	211.75	851.61
POSTCreate	8.89	9.50	84.65	178.55	258.03	1146.90
NoRC200Error	20.10	16.75	144.48	158.21	556.36	1083.60
RC401	11.54	16.08	237.67	192.83	1588.84	1476.35
RC415	15.38	11.13	227.46	167.35	1550.69	1001.35

Table 11 Hypothesis testing results for TAU (RQ1), Holm-Bonferroni adjusted p-values, $\alpha = 0.05$, sorted by effect size

Task	Test statistic (U)	p-value	Cohen's <i>d</i>	Accepted
CRUDNames	2434.0	< 0.001	2.17	yes
PathHierarchy2	2359.0	< 0.001	1.43	yes
POSTCreate	2281.5	< 0.001	1.40	yes
GETRetrieve	2195.0	< 0.001	1.36	yes
NoRC200Error	2077.5	< 0.001	1.25	yes
PathHierarchy1	2225.5	< 0.001	1.09	yes
VerbController	1978.0	< 0.001	0.75	yes

Table 11 continued

Task	Test statistic (U)	p-value	Cohen's <i>d</i>	Accepted
PluralNoun	1928.5	< 0.001	0.72	yes
NoTunnel	1696.0	0.0093	0.40	yes
RC415	1884.5	0.0024	0.28	yes
RC401	1756.0	0.0093	0.24	yes
PathHierarchy3	944.0	0.9961	-	no

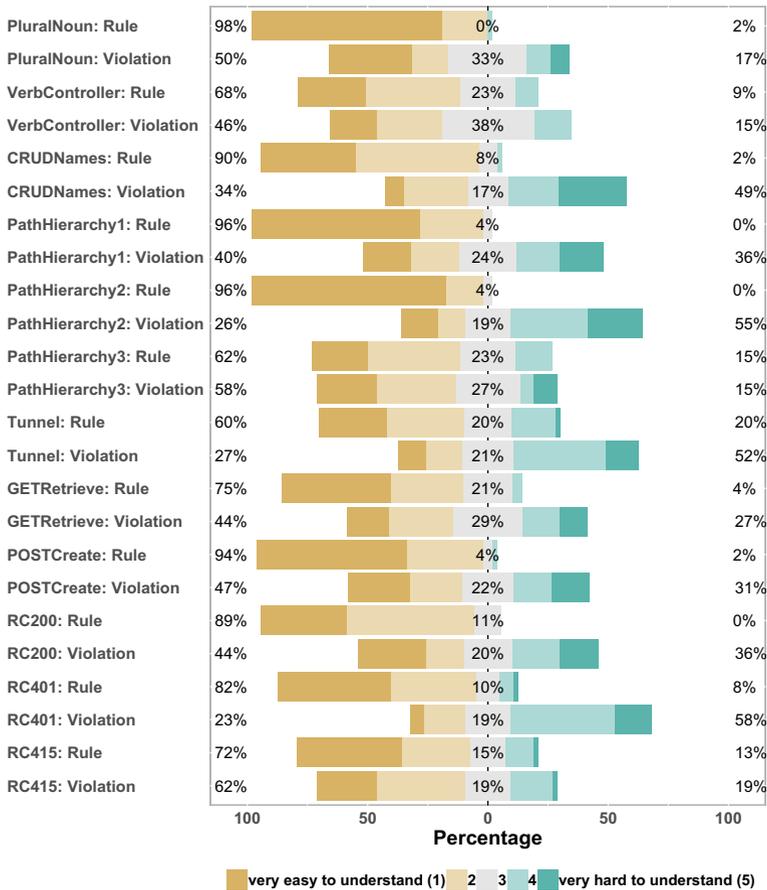


Fig. 14 Likert plots of perceived difficulty ratings (RQ2)

Table 12 Hypothesis testing results for perceived difficulty ratings (RQ2), Holm-Bonferroni adjusted p-values, $\alpha = 0.05$, sorted by effect size

Task	Test statistic (U)	p-value	Cohen's <i>d</i>	Accepted
PathHierarchy2	2462.0	< 0.001	2.06	yes
RC401	2297.5	< 0.001	1.57	yes

Table 12 continued

Task	Test statistic (U)	p-value	Cohen's <i>d</i>	Accepted
PathHierarchy1	2203.0	< 0.001	1.53	yes
CRUDNames	2254.5	< 0.001	1.52	yes
POSTCreate	2076.0	< 0.001	1.18	yes
PluralNoun	2069.5	< 0.001	1.18	yes
NoRC200Error	1851.0	< 0.001	0.94	yes
GETRetrieve	1972.5	< 0.001	0.87	yes
NoTunnel	1839.5	< 0.001	0.79	yes
VerbController	1676.0	0.0701	-	no
RC415	1642.5	0.0774	-	no
PathHierarchy3	1388.0	0.4051	-	no

Table 13 Correlation between TAU and perceived difficulty for the treatment *violation*, sorted by correlation strength, Holm-Bonferroni adjusted p-values, $\alpha = 0.05$, insignificant correlations are marked with †, difficulty ratings and TAU provided for comparison

Task	Kendall's τ	p-value	Median difficulty <i>rule</i>	<i>violation</i>	Mean TAU <i>rule</i>	<i>violation</i>
NoTunnel	0.2264	1.0000 †	2	4	0.4820	0.3396
CRUDNames	0.2124	1.0000 †	2	3	0.7392	0.1982
PathHierarchy1	-0.0574	1.0000 †	1	3	0.8457	0.5678
POSTCreate	-0.0634	1.0000 †	1	3	0.8427	0.4201
RC415	-0.0950	0.9238 †	2	2	0.7654	0.7044
NoRC200Error	-0.1548	0.4996 †	2	3	0.7268	0.3776
RC401	-0.2516	0.0692 †	2	4	0.6058	0.5170
PathHierarchy2	-0.2582	0.0609 †	1	4	0.7916	0.4913
PluralNoun	-0.2595	0.0609 †	1	2.5	0.8369	0.7046
PathHierarchy3	-0.3435	0.0062	2	2	0.6141	0.7505
VerbController	-0.3543	0.0062	2	3	0.7617	0.5345
GETRetrieve	-0.5078	< 0.001	2	3	0.8095	0.4337

Acknowledgements We kindly thank all our experiment participants for their valuable time! We also thank the experts participating in our pilot for their detailed feedback! Lastly, we thank Dr. Daniel Graziotin (University of Stuttgart) and Dr. Sira Vegas (Universidad Politécnica de Madrid) for discussing the experiment design and its terminology with us.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Baltes S, Ralph P (2022) Sampling in software engineering research: A critical review and guidelines. *Empirical Software Engineering* 27(4):94. <https://doi.org/10.1007/s10664-021-10072-8>
- Bogner J, Fritzsche J, Wagner S, Zimmermann A (2019) Microservices in Industry: Insights into Technologies, Characteristics, and Software Quality. In: 2019 IEEE International Conference on Software Architecture Companion (ICSA-C), IEEE, Hamburg, Germany, pp 187–195. <https://doi.org/10.1109/ICSA-C.2019.00041>, <https://ieeexplore.ieee.org/document/8712375/>
- Cohen J (1988) *Statistical Power Analysis for the Behavioral Sciences*, zeroth edn. Routledge. <https://doi.org/10.4324/9780203771587>, <https://www.taylorfrancis.com/books/9781134742707>
- Fielding RT, Taylor RN (2002) Principled Design of the Modern Web Architecture. *ACM Trans Internet Technol* 2(2):115–150
- Haupt F, Leymann F, Scherer A, Vukojevic-Haupt K (2017) A Framework for the Structural Analysis of REST APIs. 2017 IEEE International Conference on Software Architecture (ICSA). IEEE, Gothenburg, Sweden, pp 55–58
- Haupt F, Leymann F, Vukojevic-Haupt K (2018) API governance support through the structural analysis of REST APIs. *Computer Science - Research and Development* 33(3–4):291–303
- Jacobson D, Brail G, Woods D (2011) *APIs: A Strategy Guide*. O'Reilly Media, Inc
- Jedlitschka A, Ciolkowski M, Pfahl D (2008) Reporting Experiments in Software Engineering. In: Guide to Advanced Empirical Software Engineering, Springer London, London, pp 201–228. https://doi.org/10.1007/978-1-84800-044-5_8, http://link.springer.com/10.1007/978-1-84800-044-5_8
- Kendall MG (1938) A New Measure of Rank Correlation. *Biometrika* 30(1–2):81–93. <https://doi.org/10.1093/biomet/30.1-2.81>
- Kotstein S, Bogner J (2021) Which RESTful API Design Rules Are Important and How Do They Improve Software Quality? A Delphi Study with Industry Experts. In: Service-Oriented Computing, SummerSOC 2021. Communications in Computer and Information Science, Vol 1429, Springer International Publishing, pp 154–173. https://doi.org/10.1007/978-3-030-87568-8_10, http://dx.doi.org/10.1007/978-3-030-87568-8_10 https://link.springer.com/10.1007/978-3-030-87568-8_10
- Martin Fowler (2010) Richardson Maturity Model. <https://martinfowler.com/articles/richardsonMaturityModel.html>. Last accessed 26 March 2021
- Massé M (2011) *REST API Design Rulebook*. O'Reilly Media, Inc., Sebastopol, CA, USA, <https://www.oreilly.com/library/view/rest-api-design/9781449317904/>
- Neuhäuser M (2011) Wilcoxon-Mann-Whitney Test. In: Lovric M (ed) *International Encyclopedia of Statistical Science*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp 1656–1658. https://doi.org/10.1007/978-3-642-04898-2_615, http://link.springer.com/10.1007/978-3-642-04898-2_615
- Neumann A, Laranjeiro N, Bernardino J (2018) An Analysis of Public REST Web Service APIs. *IEEE Transactions on Services Computing* PP(c):1–1
- Palma F, Dubois J, Moha N, Guéhéneuc YG (2014) Detection of rest patterns and antipatterns: A heuristics-based approach. *Service-Oriented Computing*. Springer, Berlin Heidelberg, Berlin, Heidelberg, pp 230–244
- Palma F, Gonzalez-Huerta J, Founi M, Moha N, Tremblay G, Guéhéneuc YG (2017) Semantic Analysis of RESTful APIs for the Detection of Linguistic Patterns and Antipatterns. *International Journal of Cooperative Information Systems* 26(02):1742001
- Palma F, Zarraa O, Sadia A (2021) Are developers equally concerned about making their apis restful and the linguistic quality? a study on google apis. In: *Computing Service-Oriented*, International Springer (eds) Hacid H, Kao O, Mecella M, Moha N, Paik Hy. Publishing, Cham, pp 171–187
- Palma F, Olsson T, Wingkvist A, Ahlgren F, Toll D (2022a) Investigating the linguistic design quality of public, partner, and private rest apis. In: 2022 IEEE International Conference on Services Computing (SCC), pp 20–30. <https://doi.org/10.1109/SCC55611.2022.00017>
- Palma F, Olsson T, Wingkvist A, Gonzalez-Huerta J (2022b) Assessing the linguistic quality of rest apis for iot applications. *J Syst Softw* 191(C), <https://doi.org/10.1016/j.jss.2022.111369>,
- Pautasso C (2014) *RESTful web services: Principles, patterns, emerging technologies*. Web Services Foundations, vol 9781461475. Springer, New York, New York, NY, pp 31–51
- Pautasso C, Zimmermann O, Leymann F (2008) Restful web services vs. “big” web services: Making the right architectural decision. In: *Proceedings of the 17th International Conference on World Wide Web*, Association for Computing Machinery, New York, NY, USA, WWW '08, p 805–814. <https://doi.org/10.1145/1367497.1367606>,
- Petrillo F, Merle P, Moha N, Guéhéneuc YG (2016) Are REST APIs for Cloud Computing Well-Designed? An Exploratory Study. *Service-Oriented Computing*. Springer International Publishing, Cham, pp 157–170

- Renzel D, Schlebusch P, Klamma R (2012) Today's Top "RESTful" Services and Why They Are Not RESTful. *Web Information Systems Engineering - WISE 2012*. Springer, Berlin Heidelberg, Berlin, Heidelberg, pp 354–367
- Richardson L, Ruby S (2007) *RESTful Web Services*. O'Reilly Media, Sebastopol, CA, USA
- Rodríguez C, Baez M, Daniel F, Casati F, Trabucco J, Canali L, Percannella G (2016) REST APIs: A large-scale analysis of compliance with principles and best practices. In: *Lecture Notes in Computer Science*, Springer, vol 9671
- Sawilowsky SS (2009) New Effect Size Rules of Thumb. *Journal of Modern Applied Statistical Methods* 8(2):597–599, <https://doi.org/10.22237/jmasm/1257035100>, <http://digitalcommons.wayne.edu/jmasm/vol8/iss2/26>
- Scalabrino S, Bavota G, Vendome C, Linares-Vásquez M, Poshyvanyk D, Oliveto R (2021) Automatically assessing code understandability. *IEEE Trans Software Eng* 47(3):595–613. <https://doi.org/10.1109/TSE.2019.2901468>
- Schermann G, Cito J, Leitner P (2016) All the Services Large and Micro: Revisiting Industrial Practice in Services Computing. In: Norta A, Gaaloul W, Gangadharan GR, Dam HK (eds) *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol 9586, Springer Berlin Heidelberg, Berlin, Heidelberg, pp 36–47, https://doi.org/10.1007/978-3-662-50539-7_4, <http://link.springer.com/10.1007/978-3-662-50539-7>
- Shaffer JP (1995) Multiple Hypothesis Testing. *Annual Review of Psychology* 46(1):561–584. <https://doi.org/10.1146/annurev.ps.46.020195.003021>
- Shapiro SS, Wilk MB (1965) An Analysis of Variance Test for Normality (Complete Samples). *Biometrika* 52(3/4):591. <https://doi.org/10.2307/2333709>, <https://www.jstor.org/stable/2333709?origin=crossref>
- Vegas S, Apa C, Juristo N (2016) Crossover Designs in Software Engineering Experiments: Benefits and Perils. *IEEE Transactions on Software Engineering* 42(2):120–135. <https://doi.org/10.1109/TSE.2015.2467378>, <http://ieeexplore.ieee.org/document/7192651/>
- Webber J, Parastatidis S, Robinson I (2010) *REST in Practice: Hypermedia and Systems Architecture*, 1st edn. O'Reilly Media, Inc., Sebastopol, USA
- Wohlin C, Runeson P, Höst M, Ohlsson MC, Regnell B, Wesslén A (2012) Planning. In: *Experimentation in Software Engineering*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp 89–116, https://doi.org/10.1007/978-3-642-29044-2_8
- Wyrich M, Bogner J, Wagner S (2022) 40 years of designing code comprehension experiments: A systematic mapping study. <https://doi.org/10.48550/ARXIV.2206.11102>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.