Type Safety with JSON Subschema

ANDREW HABIB, TU Darmstadt, Germany

AVRAHAM SHINNAR, IBM Research, USA

MARTIN HIRZEL, IBM Research, USA

MICHAEL PRADEL, University of Stuttgart, Germany

JSON is a popular data format used pervasively in web APIs, cloud computing, NoSQL databases, and increasingly also machine learning. JSON Schema is a language for declaring the structure of valid JSON data. There are validators that can decide whether a specific JSON document is valid with respect to a schema. However, validators cannot show that all documents valid under one schema are also valid under another schema. This paper presents a technique to address this limitation: JSON subschema checking, which can be used for static type checking with JSON Schema. Deciding whether one schema is a subschema of another is non-trivial because of the richness of the JSON Schema specification language. Given a pair of schemas, our approach first canonicalizes and simplifies both schemas, then reasons about the subschema question on the canonical forms, dispatching simpler subschema queries to type-specific checkers. We apply an implementation of our subschema checking algorithm to 8,548 pairs of real-world JSON schemas from different domains, demonstrating that it can decide the subschema question for most schema pairs and is always correct for schema pairs that it can decide. We hope that our work will bring more static guarantees to hard-to-debug domains, such as cloud computing and artificial intelligence.

1 INTRODUCTION

JSON (JavaScript Object Notation) is a data serialization format that is widely adopted to store data on disk or send it over the network. The format supports primitive data types, such as strings, numbers, and Booleans, and two possibly nested data structures: arrays, which represent lists of values, and objects, which represent maps of key-value pairs. JSON is used in numerous applications. It is the most popular data exchange format in web APIs, ahead of XML [21]. Cloud-hosted applications also use JSON pervasively, e.g., in micro-services that communicate via JSON data [17]. On the data storage side, not only do traditional database management systems, such as Oracle, IBM DB2, MySQL, and PostgreSQL, now support JSON, but two of the most widely deployed NoSQL database management systems, MongoDB and Cloudant, are entirely based on JSON. Beyond web, cloud, and database applications, JSON is also gaining adoption in artificial intelligence (AI) [13, 24].

With the broad adoption of JSON as a data serialization format soon emerged the need for a way to describe how a JSON document should look. For example, a web API that consumes JSON data can avoid unexpected behavior if it knows the structure of the data it receives. To describe a JSON document, *JSON Schema* allows users to declaratively define the structure of nested values (documents) via types (schemas) [19]. We adopt the value space definition for *type* as being a set of possible values [18]. A JSON Schema *validator* checks whether a JSON document d conforms to a schema s, denoted $d : s$. There are libraries with JSON Schema validators for many programming languages, and they are widely used to make software more reliable.

Since these $d : s$ validation checks happen at runtime, they can usually only detect problems late during deployment and production. In cloud applications, a wrongly-structured document can cause hard-to-debug failures, it may go unnoticed and simply cause undesired behavior, or even worse, it may exploit a security vulnerability. In AI, a machine learning *pipeline* is a graph of

$s \dashrightarrow^{\text{check if}} t$

canonicalize and simplify | canonicalize and simplify

$s_c \xrightarrow{\text{check if}} t_c$

extract fragments | extract fragments

$s_{1c} \xrightarrow{\text{check if}} t_{1c}$

$\vdots$
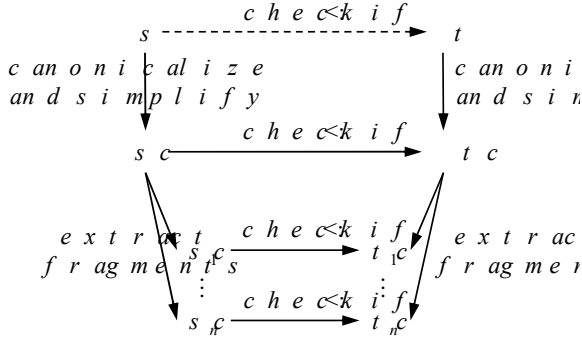
$s_{nc} \xrightarrow{\text{check if}} t_{nc}$

Fig. 1. Overview of JSON subschema checker.

operators for preprocessing and prediction [8]. Any mismatch between training data and a pipeline, between production data and training data, or between data in adjacent steps of a pipeline can cause crashes or poor predictive performance [7]. Sometimes, dynamic data validity checks of the form $d : s$ only trigger after earlier pipeline steps have already performed costly computations.

Given how essential dynamic $d : s$ checks are, we argue that JSON schemas could be even more useful if they were also checked statically. The quintessential static property is whether one schema is a *subschema* of another. We say that a schema s is a subschema of a schema t, denoted $s <: t$, if all documents that validate against s also validate against t. Subschema checks of the form $s <: t$ can find the same mistakes as dynamic schema validation, but with less wasted compute and human time, as they can statically rule out entire classes of errors.

JSON subschema checking has various practical applications, such as statically reasoning about breaking changes of web APIs and type-checking machine learning pipelines before running them. As one example, consider two independently evolving services that communicate over a RESTful API, where both services, as well as the schema that describes the data they exchange, are evolving [10, 16]. JSON subschema checking could detect a number of likely errors before deployment by checking whether the schemas of the services remain compatible. As another example, consider a machine learning pipeline where the data, as well as the input and output of operations, are specified as JSON schemas [13, 24]. JSON subschema checking could type-check the entire pipeline before running it, saving precious developer time caused by bugs that otherwise may be detected only after hours of computation.

JSON schemas have several features that make subschema checking difficult. Schemas for primitive types involve non-trivial features, such as regular expressions for strings and `multipleOf` for numbers. Schemas for compound types involve `uniqueItems` constraints for arrays and regular expressions for object property names. Furthermore, JSON schemas can be composed using schema conjunction, disjunction, and negation, which requires handling of complex, compound types. Finally, the `enum` type is entangled with other types, further exacerbating the above-mentioned difficulties. As a result of these and other features of JSON Schema, a simple lexical check or a structural comparison of JSON schemas is insufficient to address the subschema question.

This paper presents *jsonsubschema*, a formally described and fully implemented subschema checker for JSON Schema. Figure 1 gives a high-level overview of our algorithm. To check whether $s <: t$, the checker first canonicalizes and simplifies s to sc and t to tc. This preprocessing disentangles different cases to consider during subschema checking, while preserving the schema semantics. Next, our algorithm extracts corresponding fragments sc_i and tc_i. These fragments are type-homogeneous,

i.e., they refer to only one basic JSON type each. This homogeneity makes it possible to use separate rules to check whether $sc_i <: tc_i$ for each type.

Our algorithm is sound, i.e., whenever it gives an answer about the subtype relation of two given schemas, then this answer is correct. In principle, the algorithm is incomplete, i.e., for some pairs of subschemas, *jsonsubschema* refuses to answer the question whether one schema is a subtype of the other. However, the approach handles various challenging cases, including almost all that appear in practice. Our evaluation with 8,548 pairs of real-world JSON schemas shows that the algorithm handles the vast majority of all pairs (94%) and that all the given answers are correct.

The most closely related prior work is an open-source project called is-json-schema-subset developed concurrently with our work, but it only handles a small fraction of the features of JSON Schema [12]. The problem of subschema checking has also been explored for XML, where it is called schema containment [25]. However, that approach treats XML schemas as tree automata, which has been shown to be insufficient for JSON schemas because JSON Schema is more expressive than tree automata [19].

In summary, this paper makes the following contributions:

- A canonicalizer and simplifier that converts a schema *s* into a schema *sc* that is simpler to check yet permits the same set of documents (Sections 4.1 and 4.2).
- A subschema checker for canonicalized JSON schemas that uses separate subschema checking rules for each basic JSON type (Section 4.3).
- Empirical evidence on 8,548 pairs of JSON schemas taken from real-world applications in the web, cloud computing, and AI. We show that the implementation of our algorithm successfully answers the subschema question in most cases, and when it does, yields correct answers in reasonable time (Section 6).

Our JSON subschema checker is available as open-source (link elided for double-blind review). We also have plans for an artifact submission.

2 BACKGROUND

This section briefly describes JSON data and JSON Schema validation.

2.1 JSON Data

JSON was conceived as a light-weight, text-based, and programming-language agnostic data interchange format [15].

Definition 2.1 (Jprimitive). The primitive JSON types are:

$$\mathcal{J}primitive = \{\texttt{boolean}, \texttt{null}, \texttt{number}, \texttt{integer}, \texttt{string}\}$$

Besides primitive types, JSON has two structured types: ordered lists of values (arrays) and unordered maps of key-value pairs (objects).

Definition 2.2 (Jstructure). The structured JSON types are: $\mathcal{J}structure = \{\texttt{array}, \texttt{object}\}$

Definition 2.3 (Jtypes). The set of JSON data types is: $\mathcal{J}types = \mathcal{J}primitive \cup \mathcal{J}structure$

A *valid JSON document* (or value) is either a value of one of the basic types, or an array whose elements are valid JSON documents, or an object mapping string keys to valid JSON documents. Some examples of valid JSON documents include `true`, `null`, `5`, `'ab'`, `[]`, `[1, 2, 3]`, `[.5, {}, 'a']`, and `{'foo':1, 'bar':[true, '']}`.

Finally, JSON documents can use JSON references to point to data from the same or other JSON documents. For example, the reference `{'$ref':'#'}` points to the root of the current document and

the reference `{'$ref':'#/p/q}'` assumes that the root is an object whose `'p'` property is a nested object and points to its `'q'` property.

2.2 JSON Schema

JSON Schema is a declarative language for defining the structure and permitted values of a JSON document [27]. JSON Schema itself uses JSON syntax. JSON Schema is an Internet Draft from the Internet Engineering Task Force (IETF). It is continuously evolving, currently at draft-2019-09. This work focuses on draft04 [11], one of the most widely adopted versions of JSON Schema.

To specify which data types are allowed, JSON Schema uses the keyword `'type'` with one type name or a list of type names. For example, schema `{'type':'string'}` accepts strings and schema `{'type':['null','boolean']}` accepts null or Boolean values. Each JSON type has a set of validation *keywords* that restrict the values a schema of this type permits. For example, the schema `{'type':'integer','minimum':0}` uses the keyword `minimum` to restrict integers to be non-negative, while schema `{'type':'string','pattern':'^[A-Za-z0-9]+$'}` uses the keyword `pattern` with a regular expression to restrict strings to be alphanumeric. The upper part of Table 1 lists the set of keywords associated with each JSON type along with their meta-schema. For example, `string` schemas have a `minLength` keyword with a non-negative integer that defaults to 0. The meta-schema acts similarly to a grammar in that it defines the syntax for schemas, while at the same time using and thus illustrating a realistic JSON Schema use-case.

In addition to type-specific keywords, JSON Schema allows enumerating exact values with `enum` and combining different schemas using a set of logic connectives (Table 1, lower part). For example, schema `{'enum':['a', [], 1]}` restricts the set of permitted JSON values to the string literal `'a'`, an empty array, or the integer 1. Logic connectives, such as `anyOf`, `allOf`, and `not`, allow schema writers to express disjunctions, conjunctions, and negations of schemas. The empty schema, `{}`, is the top of the schema hierarchy (all documents are valid for `{}`). The negation of the empty schema, `{'not':{}}`, is the bottom of the schema hierarchy (no documents are valid for `{'not':{}}`). Finally, the keyword `$ref` retrieves schemas using URIs and JSON pointers. JSON validation against a schema with `$ref` has to satisfy the schema retrieved from the specified URI or JSON pointer. We refer the interested reader to the full specification of JSON Schema [11] and its formalization [19].

Given a JSON document d and a JSON schema s, schema validation checks if d conforms to s.

Definition 2.4 (JValid(d, s)). For any JSON document d and any JSON schema s, JValid(d, s) → {*True, False*}, also written $d : s$, decides whether d is valid with respect to s.

This decision problem JValid(d, s) is shown to be PTIME-hard [19] and solvable in linear time when eliminating the `uniqueItems` keyword as it involves sorting. JSON validators have been implemented in most major programming languages and are widely used in several domains.

3 PROBLEM STATEMENT

Given two schemas s and t, our approach tries to decide whether s is a subschema of t. This section defines the subschema relation, gives concrete usage scenarios for it, and describes why deciding the subschema question is non-trivial.

3.1 JSON Subschema

The goal of this work is to tell whether the set of documents that conform to schema s is a subset of the set of documents that conform to schema t. We aim at answering this question without enumerating all valid documents, and we call this question the *JSON Subschema* decision problem.

Definition 3.1 (JSubSchema <:). For any two JSON schemas s and t, the subschema relation, denoted <:, is defined as: $s <: t \iff (\forall d : \text{JValid}(d, s) \implies \text{JValid}(d, t))$

Table 1. Overview of JSON Schema (version draft04).

`null` and `boolean` are basic types without keywords. `$ref` refers to another schema, and `$ref:{'#'}` points to the root of the meta-schema, thus indicating a nested sub-schema. `default:{}` indicates the top schema, which validates for all values.

Description	Keyword	Meta-schema
Schemas for type `string`, e.g.,`{type:'string',minLength:1,pattern:'[a-z]*'}`:		
Restrict length of the string	`minLength`	`{type:'integer',minimum:0, default:0}`
	`maxLength`	`{type:'integer',minimum:0}`
String must match a pattern	`pattern`	`{type:'string',format:'regex'}`
Schemas for types `number` and `integer`, e.g.,`{type:'number',minimum:0.0,multipleOf:0.1}`:		
Restrict range of values	`minimum`	`{type:'number'}`
	`maximum`	`{type:'number'}`
	`exclusiveMinimum`	`{type:'boolean',default:false}`
	`exclusiveMaximum`	`{type:'boolean',default:false}`
Must be a multiple of a given value	`multipleOf`	`{type:'number',minimum:0, exclusiveMinimum:true}`
Schemas for type `array`, e.g.,`{type:'array',items:{type:'string'},minItems:1, uniqueItems:true}`:		
Restrict type of all items .. .. or for each individual item	`items`	`{default:{}, anyOf:[{$ref:'#'}, {type:'array',minItems:1, items:{$ref:'#'}}]}`
Restrict number of items	`minItems`	`{type:'integer',minimum:0, default:0}`
	`maxItems`	`{type:'integer',minimum:0}`
Allow items beyond those specified above	`additionalItems`	`{anyOf:[{type:'boolean'},{$ref:'#'}], default:{}}`
Enforce that items are unique	`uniqueItems`	`{type:'boolean',default:false}`
Schemas for type `object`, e.g.,`{type:'object',properties:{'a':{type:'string'},'b':{enum:[0,1]}}}`:		
Specify key-value pairs allowed in the schema	`properties`	`{type:'object',additionalProperties:{$ref:'#'}, default:{}}`
Restrict number of properties	`minProperties`	`{type:'integer',minimum:0, default:0}`
	`maxProperties`	`{type:'integer',minimum:0}`
Names of properties that must be present	`required`	`{type:'array',items:{type:'string'},minItems:1, uniqueItems:true}`
Allow properties beyond those specified above	`additional-Properties`	`{anyOf:[{type:'boolean'},{$ref:'#'}], default:{}}`
Restrict properties whose name match a regular expression	`pattern-Properties`	`{type:'object',additionalProperties:{$ref:'#'}, default:{}}`
When one property is present, enforce presence or type of another one	`dependencies`	`{type:'object',additionalProperties:{anyOf:[{$ref:'#'}, {type:'array',items:{type:'string'}, minItems:1, uniqueItems:true}]}}`
Specifying types and combining schemas, e.g.,`{'anyOf':[{'type':'string'},{'$ref':'#/some/type'}]}`:		
Restrict value to one type or to an array of types	`type`	`{anyOf:[{$ref:'#/definitions/types'},{type:'array',items:{$ref:'#/definitions/types'}, minItems:1,uniqueItems:true}]}` `{definitions:{types:{enum:['boolean','null','number','integer','string','array','object']}}}`
Enumerate values	`enum`	`{type:'array',minItems:1, uniqueItems:true}`
Negation	`not`	`{$ref:'#'}`
Conjunction	`allOf`	`{type:'array',minItems:1, items:{$ref:'#'}}`
Disjunction	`anyOf`	`{type:'array',minItems:1, items:{$ref:'#'}}`
Exclusive or	`oneOf`	`{type:'array',minItems:1, items:{$ref:'#'}}`

```
{'type': 'object',                    {'anyOf': [
 'required': ['type', 'address'],       {'type': 'object',
 'properties': {                          'required': ['type', 'address'],
   'address': {                           'properties': {
     'description': 'Node address',         'type': { 'enum': ['ExternalIP', 'InternalIP']},
     'type': ['string', 'null']},           'address': {
   'type': {                                 'type': 'string',
     'description': 'Node address             'pattern': '^\d+\.\d+\.\d+\.\d+$'}}},
         type; one of Hostname,         {'type': 'object',
         ExternalIP or InternalIP',      'required': ['type', 'address'],
     'type': ['string', 'null']}}}        'properties': {
                                            'type': { 'enum': ['Hostname']},
            (a)                             'address': {
                                             'type': 'string',
                                             'pattern': '^[A-Za-z0-9.]+$'}}}]}

                                                  (b)
```

Fig. 2. Example for usage scenario of schemas as static types.

The relation <: is a partial order, and equivalence of schemas follows directly by the anti-symmetry of the subtype relation.

Definition 3.2 (JEquivSchema ≡). For any two JSON schemas s and t, the equivalence relation $s \equiv t$ is given by: $s \equiv t \iff (s <: t \land t <: s)$

That is, two schemas are equivalent if they describe the same exact set of JSON documents.

3.2 Usage Scenarios

In many usage scenarios, a membership check with JValid corresponds to a runtime check, whereas a subschema check provided by JSubSchema corresponds to static type checking.

Backward compatibility. One of the most pervasive use cases of JSON schemas is describing requests and responses of web APIs. For example, version 0.6.1 of the Washington Post `ans-schema` contains the following:

```
'category': {
    'type': 'string',
    'enum': ['staff', 'wires', 'freelance', 'other' ]}
```

The continuous development and evolution of these APIs involves regular changes to the corresponding JSON schemas, and developers need to keep a close eye on such changes to avoid breaking backward compatibility. For example, version 0.6.2 of the same schema contains:

```
'category': {
    'type': 'string',
    'enum': ['staff', 'wires', 'freelance', 'stock', 'handout', 'other' ]}
```

Version 0.6.1 is a subschema of version 0.6.2. Assuming the developers intend to retain backward compatibility, this evolution would be fine for an API request argument, but it could break clients when used as an API response.

Schemas as static types. JSON subschema checking can help make code more robust by flagging some mistakes statically. Consider version 0.14.0 of the `NodeAddress` schema from Kubernetes, extracted from its OpenAPI specification, shown in Figure 2a. With this schema, an application can check `NodeAddress` objects at runtime, but runtime errors in distributed, cloud-based systems are difficult to debug. Therefore, client code might define a stricter schema s, as shown in Figure 2b. Schema s uses enums to constrain the values for `'type'` and `patterns` to constrain the values for

'address', with an anyOf to provide two cases. By being stricter, schema *s* can rule out more bugs. The static check *s* <: NodeAddress can validate that *s* is indeed a subschema.

Machine learning pipelines. Machine-learning pipelines are of little use if the data is formatted incorrectly [7]. For example, here is the input schema for the NMF transformer from scikit-learn [8], which performs non-negative matrix factorization.

```
{   'type': 'array',
    'items': {
        'type': 'array',
        'items': {
            'type': 'number', 'minimum': 0.0}}}}}
```

The outer array ranges over rows (samples in machine learning) and the inner array ranges over columns (features in machine learning). NMF requires non-negative numbers as input, captured by the keyword minimum. This transformer is typically used in the middle of a larger machine-learning pipeline, after data cleansing transformers but before a classifier or regressor. This schema can check a specific part of a dataset at a specific point in the pipeline, but this check has to be repeated for production data, which may differ from training data for various reasons.

Instead, several sources argue that the dataset itself should also be associated with a schema [7, 9, 23]. For example, here is a schema for the features of the California Housing dataset.

```
{   'type': 'array',
    'items': {
        'type': 'array', 'minItems': 4, 'maxItems': 4,
        'items': [
        {   'description': 'AveRooms',    'type': 'number', 'minimum': 0.0},
        {   'description': 'Population', 'type': 'number', 'minimum': 0.0},
        {   'description': 'Latitude',    'type': 'number', 'minimum': 0.0},
        {   'description': 'Longitude',   'type': 'number'}]}}
```

The outer array ranges over rows, each one describing a different district in California. The inner array ranges over columns, each one describing a feature that might be helpful for predicting the median house value in that district. Instead of a single schema for all items, there are separate per-item schemas (see the meta-schema for the items keyword in Table 1).

Is the dataset compatible with the input expected by the NMF transformer? California is north of the equator, so the latitude is positive, and west of Greenwich, so the longitude is negative. Due to the negative longitudes (expressed by the absence of the minimum keyword), this schema is not a subschema of the input schema of NMF. Machine learning engineers could add 360 degrees to the longitude and change its schema to {'description':'Longitude','type':'number','minimum':0.0}. Then, the dataset schema would be a subschema of the NMF input schema, thanks to the keywords 'minItems':4, 'maxItems':4 which guarantee the absence of additional columns with unspecified schemas. Later, when this model is used in production, the schema can be used to check the validity of new samples. For instance, if there is a sample that comes not from California but from south of the equator, a schema check could catch the negative latitude early and explain which assumption it violated.

3.3 Challenges

JSON schemas define the nested *structure* and *valid values* permitted in a set of JSON documents. The rich feature set of JSON Schema makes establishing or refuting a subschema relationship between two schemas non-trivial. Even for simple, structurally similar schemas, such as {'enum':[1, 2]} and {'enum':[2, 1]}, equivalence does not hold through textual equality. There are several challenges for algorithmically checking the subtype relation of JSON schemas.

```
                                                                {'allOf': [
                                      {'anyOf': [                    {'anyOf': [
{'type': [                               {'type': 'null'},              {'type': 'null'},
    'null',                              {'type': 'string'}],          {'type': 'string'}],
    'string'],                         'not': {                      'not': {
 'not': {                                'type': 'string',             'type': 'string',
    'enum': ['']}}                       'enum': ['']}}                'enum': ['']}}]}
        (a)                                  (b)                          (c)

{'allOf': [
   {'anyOf': [
       {'type': 'null'},
       {'type': 'string'}],
   {'anyOf': [
       {'type': 'boolean'}, {'type': 'null'},
       {'type': 'number'}, {'type': 'integer'},    {'anyOf': [
       {'type': 'array'}, {'type': 'object'},         {'type': 'null'},
       {'type': 'string', 'pattern': '.+'}]}]}        {'type': 'string', 'pattern': '.+'}]}
                    (d)                                           (e)
```

Fig. 3. Five syntactically different but semantically equivalent schemas for a value that is either a non-empty string or null.

First, the schema language is flexible and the same set of JSON values, i.e., the same type, could be described in several different syntactical forms, i.e., schemas. For example, Figure 3 shows five equivalent schemas describing a JSON value that is either a non-empty string or null.

Second, even for primitive types, such as strings and numbers, nominal subtyping is not applicable. JSON Schema lets users specify various constraints on primitive types, resulting in non-trivial interactions that are not captured by nominal types. For example, one cannot infer that an integer schema is a subtype of a number schema without properly comparing the range and multiplicity constraints of the schemas.

Third, logic connectives combine non-homogeneous types such as 'string' and 'null' in Figure 3b. Moreover, enumerations restrict types to predefined values, which require careful handling, especially when enumerations interact with non-enumerative types, such as in Figures 3a, 3b, and 3c.

Fourth, the schema language allows implicit conjunctions and disjunctions. For example, Figure 3b has an implicit top-level conjunction between the subschemas under anyOf and not. As another example, a schema that lacks a type keyword, such as {'pattern':'.+'}, has an implicit disjunction between all possible types, while still enforcing any type-specific keyword, such as the pattern for strings only. For illustration purposes, Figure 3d makes this implicit disjunction explicit.

Finally, JSON Schema also allows uninhabited types. That is, a schema can be syntactically valid yet semantically self-contradicting, e.g., {'type':'number','minimum':5,'maximum':0}. Such schemas validate no JSON value at all and complicate reasoning about subtyping.

4 ALGORITHM

This section describes how we address the problem of checking whether one JSON schema is a subtype of another. Because JSON schemas are complex, creating a subtype checker for arbitrary schemas directly would necessitate a complex algorithm to handle all of its variability. Instead, we decompose the problem into three steps. The first step canonicalizes a given schema into an equivalent but more standardized schema (Section 4.1). The second step further simplifies a schema by eliminating enumerations, negation, intersection, and union of schemas where possible

Table 2. Canonicalization and simplification guarantees.

Language feature	Use of feature in schemas		
	Full JSON schema	**Canonicalized**	**Simplified**
`null`	Yes	Yes	Yes
`boolean`	Yes	Represented as `enum`	Represented as `enum`
`string`	All keywords	Keyword `pattern` only	Keyword `pattern` only
`number`	All keywords	All keywords	All keywords
`integer`	All keywords	Eliminated	Eliminated
`array`	All keywords	All keywords; `items` is always list	All keywords; `items` is always list
`object`	All keywords	Only `{min,max}Properties`, `patternProperties`, `required` keywords	Only `{min,max}Properties`, `patternProperties`, `required` keywords
`enum`	Heterogeneous, any type	Homogeneous, any type	Only for `boolean`
`not`	Multiple connectives	Single connective	Only for `number`, `array`, `object`
`allOf`	Multiple connectives	Single connective	Only for `not`
`anyOf`	Multiple connectives	Single connective	Only for `not`, `allOf`, `array`, `object`, and disjoint `number`.
`oneOf`	Multiple connectives	Eliminated	Eliminated

(Section 4.2). Table 2 summarizes the first two steps. Finally, the third step checks for two canonicalized and simplified schemas whether one is a subtype of the other by extracting and comparing type-homogeneous schema fragments (Section 4.3).

Notation. We formalize canonicalization and simplification via rewrite rules of the form $s_1 \rightarrow s_2$. The notation $s.k$ indicates access of property k in schema s. For any JSON schema s, helper function $dom(s)$ returns its property names, i.e., the set of keys in the key-value map s. The notation $s[k \mapsto v]$ indicates a substitution, which is a copy of s except that the mapping of key k is added or changed to value v. The notation $[\ldots]$ indicates a JSON array and the notation $\{\ldots\}$ indicates a JSON object. The notation $\{k:v \mid \ldots\}$ indicates a JSON object comprehension. The notation $a \parallel b$ is a default operator that returns a if it is defined and b otherwise.

4.1 JSON Schema Canonicalization

This section introduces a canonicalization procedure that compiles any JSON schema into an equivalent canonical schema. Column "Canonicalized" of Table 2 summarizes the properties that the canonicalizer establishes. Given any JSON schema as input, canonicalization terminates and produces a semantically equivalent, canonical JSON schema as output.

The canonicalization enforces two main simplifications. First, JSON Schema allows schemas to mix specifications of different types. To enable local, domain-specific reasoning in the subtype checker, canonicalization splits up these schemas into smaller, homogeneously typed schemas that are combined using logic connectives. Second, JSON Schema allows many alternative ways to represent the same thing. Additionally, there are keywords that can be omitted and the defaults are then assumed. Canonicalization picks, when possible, one form, and supplies omitted defaults explicitly.

Figure 4 presents non-type-specific canonicalization rules whose purpose it is to enable reasoning about one type or connective at a time. Rule *multiple types* applies to schemas whose type is a list, such as in the example in Figure 3a, making the implicit disjunction explicit using `anyOf`, as shown

$$\text{multiple types} \quad \frac{s.\mathtt{type} = [\tau_1, \ldots, \tau_n]}{s \to \{\mathtt{anyOf:}[s[\mathtt{type} \mapsto \tau_1], \ldots, s[\mathtt{type} \mapsto \tau_n]]\}}$$

$$\text{multiple connectives} \quad \frac{dom(s) \cap \{\mathtt{enum}, \mathtt{anyOf}, \mathtt{allOf}, \mathtt{oneOf}, \mathtt{not}\} \neq \emptyset \quad dom(s) \setminus \{c\} \neq \emptyset}{s \to \{\mathtt{allOf:}[\{c{:}s.c\}, \{k{:}s.k \mid k \in (dom(s) \setminus \{c\})\}]\}}$$

$$\text{missing type} \quad \frac{dom(s) \cap \{\mathtt{type}, \mathtt{enum}, \mathtt{anyOf}, \mathtt{allOf}, \mathtt{oneOf}, \mathtt{not}\} = \emptyset}{s \to s[\mathtt{type} \mapsto \mathit{Jtypes}]}$$

Fig. 4. Non-type-specific canonicalization rules that ensure exactly one type, enum, or logic connective.

in Figure 3b. Rule *multiple connectives* applies to schemas that contain a connective mixed with other connectives, such as in the example in Figure 3b, making the implicit conjunction explicit using `allOf`, as shown in Figure 3c. Rule *missing type* generously assumes all JSON types are possible, yielding an implicit disjunction to be further canonicalized by the multiple-types rule.

Figure 5 presents type-specific canonicalization rules whose purpose it is to reduce the number of cases to handle for later simplification and subschema rules. Rule *missing keyword* adds the default for a keyword if there is a single type and the keyword for that type is missing, using a helper function *default* that returns the default from the meta-schema in Table 1 and maps `minimum` to $-\infty$ and `maximum` to ∞ for convenience. Rule *irrelevant keywords* strips out spurious keywords that do not apply to a given type (or to any type), using a helper function *kw* that returns the relevant keywords in Table 1. Rule *integer* rewrites integer schemas to number schemas with the appropriate `multipleOf`. Rule *heterogeneous enum*, when iterated, ensures that each enumeration contains only values from a single type, using a helper function *typeOf* that maps a concrete JSON value to its type. Rule *oneOf* eliminates the `oneOf` keyword by rewriting the exclusive or into a disjunction of conjunctions.

The two *string* rules eliminate the keywords `minLength` and `maxLength` by compiling them into regular expressions, so that after canonicalization, string schemas only have the keyword `pattern`.

The two *array* rules handle keywords that can be specified in multiple different ways, as indicated by meta-schemas with `anyOf` in Table 1. Rule *array with one schema for all items* changes the keyword `items` from a single schema to a list (empty) of schemas, by moving the schema into `additionalItems`. Rule *array with additionalItems false* changes the keyword `additionalItems` from a boolean to a schema (bottom).

The five *object* rules eliminate the object keywords `properties`, `additionalProperties`, and `dependencies` by rewriting them into the keywords `required` and `patternProperties`, and ensure that `patternProperties` uses non-overlapping regular expressions. The object rules are the most intricate out of the canonicalization rules because in JSON Schema, object schemas have the largest number of special cases. Reducing the cases reduces the complexity of subsequent rules for simplification and subschema checking. Rule *object with additionalProperties false* changes the keyword `additionalProperties` from a boolean to a schema (bottom). Rule *object with properties* eliminates `properties` and `additionalProperties` by rewriting them into `patternProperties`. First, it turns each property key k_i from `properties` into a pattern property with the regular expression `'^k_i$'` that accepts exactly k_i. Second, it subtracts all keys k_i from each of the original pattern properties so they only apply as a fall-back, where the notation $p_1 \setminus p_2$ indicates regular expression subtraction. Third, it creates a tertiary fall-back regular expression that applies when neither the original properties nor the original pattern properties match, using the notation $\neg p$ for the complement of a regular expression, and associates that regular expression with the original `additionalProperties`. Finally, it removes the eliminated keywords `properties` and

$$\text{missing keyword} \quad \frac{s.\mathtt{type} = \tau \quad \tau \in \mathcal{J}types \quad \tau \neq \mathtt{string} \quad k \in kw(\tau) \quad k \notin dom(s)}{s \rightarrow s[k \mapsto default(k)]}$$

$$\text{irrelevant keywords} \quad \frac{s.\mathtt{type} = \tau \quad \tau \in \mathcal{J}types}{s \rightarrow \{k : s.k \mid k \in (dom(s) \cap (kw(\tau) \cup \{\mathtt{type}, \mathtt{enum}\}))\}}$$

$$\text{integer} \quad \frac{s.\mathtt{type} = \mathtt{integer}}{s \rightarrow s[\mathtt{type} \mapsto \mathtt{number}, \mathtt{multipleOf} \mapsto lcm(1, s.\mathtt{multipleOf} \parallel 1)]}$$

$$\text{heterogeneous enum} \quad \frac{s.\mathtt{enum} = [v_1, \ldots, v_n] \quad \exists j, typeOf(v_j) \neq typeOf(v_1)}{s \rightarrow \{\mathtt{anyOf} : [s[\mathtt{enum} \mapsto [v_i \mid typeOf(v_i) = typeOf(v_1)], \mathtt{type} \mapsto typeOf(v_1)], \\ s[\mathtt{enum} \mapsto [v_j \mid typeOf(v_j) \neq typeOf(v_1)]]]\}}$$

$$\text{oneOf} \quad \frac{s.\mathtt{oneOf} = [s_1, \ldots, s_n]}{s \rightarrow \{\mathtt{anyOf} : [\{\mathtt{allOf} : [s_1, \{\mathtt{not} : s_2\}, \ldots, \{\mathtt{not} : s_n\}]\}, \\ \ldots, \\ \{\mathtt{allOf} : [\{\mathtt{not} : s_1\}, \ldots, \{\mathtt{not} : s_{n-1}\}, s_n]\}]\}}$$

$$\text{string without maxlength} \quad \frac{s.\mathtt{type} = \mathtt{string} \quad s.\mathtt{pattern} = p \quad s.\mathtt{minLength} = a \quad \mathtt{maxLength} \notin dom(s)}{s \rightarrow \{\mathtt{type} : \mathtt{string}, \mathtt{pattern} : p \cap \text{'\textasciicircum.}\{a\}\text{'}\}}$$

$$\text{string with maxlength} \quad \frac{s.\mathtt{type} = \mathtt{string} \quad s.\mathtt{pattern} = p \quad s.\mathtt{minLength} = a \quad s.\mathtt{maxLength} = b}{s \rightarrow \{\mathtt{type} : \mathtt{string}, \mathtt{pattern} : p \cap \text{'\textasciicircum.}\{a, b\}\$\text{'}\}}$$

$$\text{array with one schema for all items} \quad \frac{s.\mathtt{type} = \mathtt{array} \quad s.\mathtt{items} = \{\ldots\}}{s \rightarrow s[\mathtt{items} \mapsto [], \mathtt{additionalItems} \mapsto s.\mathtt{items}]}$$

$$\text{array with additionalItems false} \quad \frac{s.\mathtt{type} = \mathtt{array} \quad s.\mathtt{additionalItems} = \mathtt{false}}{s \rightarrow s[\mathtt{additionalItems} \mapsto \{\mathtt{not} : \{\}\}]}$$

$$\text{object with additionalProperties false} \quad \frac{s.\mathtt{type} = \mathtt{object} \quad s.\mathtt{additionalProperties} = \mathtt{false}}{s \rightarrow s[\mathtt{additionalProperties} \mapsto \{\mathtt{not} : \{\}\}]}$$

object with properties
$$\frac{s.\mathtt{type} = \mathtt{object} \qquad s.\mathtt{properties} = \{k_1 : s_{k_1}, \ldots, k_n : s_{k_n}\} \\ s.\mathtt{additionalProperties} = \{\ldots\} \quad s.\mathtt{patternProperties} = \{p_1 : s_{p_1}, \ldots, p_m : s_{p_m}\}}{\begin{aligned} s \rightarrow & s[\mathtt{patternProperties} \mapsto \{\text{'\textasciicircum}k_1\$\text{'} : s_{k_1}, \quad \ldots, \quad \text{'\textasciicircum}k_n\$\text{'} : s_{k_n}, \\ & p_1 \setminus \text{'\textasciicircum}(k_1 | \ldots | k_n)\$\text{'} : s_{p_1}, \quad \ldots, \quad p_m \setminus \text{'\textasciicircum}(k_1 | \ldots | k_n)\$\text{'} : s_{p_m}, \\ & \neg\text{'\textasciicircum}(k_1 | \ldots | k_n)\$|p_1| \ldots |p_m\text{'} : s.\mathtt{additionalProperties}\}] \\ & \setminus\{\mathtt{properties}, \mathtt{additionalProperties}\} \end{aligned}}$$

object with string list dependencies
$$\frac{s.\mathtt{type} = \mathtt{object} \quad s.\mathtt{dependencies} = \{k_i : [k_{i_1}, \ldots, k_{i_n}]\} \cup d_{rest}}{s \rightarrow s[\mathtt{dependencies} \mapsto d_{rest} \cup \{k_i : \{\mathtt{type} : \mathtt{object}, \mathtt{required} : [k_{i_1}, \ldots, k_{i_n}]\}\}]}$$

object with schema dependencies
$$\frac{s.\mathtt{type} = \mathtt{object} \quad s.\mathtt{dependencies} = \{k_i : s_i\} \cup d_{rest}}{s \rightarrow \{\mathtt{allOf} : [s[\mathtt{dependencies} \mapsto d_{rest}], \\ \{\mathtt{anyOf} : [s_i, \{\mathtt{type} : \mathtt{object}, \mathtt{properties} : \{k_i : \mathtt{not} : \{\}\}\}]\}]\}}$$

object with overlapping patternProperties
$$\frac{s.\mathtt{type} = \mathtt{object} \quad s.\mathtt{patternProperties} = \{p_i : s_i\} \cup \{p_j : s_j\} \cup d_{rest}}{s \rightarrow s[\mathtt{patternProperties} \mapsto \\ \{p_i \cap p_j : \{\mathtt{allOf} : [s_i, s_j]\}\} \cup \{p_i \cap \neg p_j : s_i\} \cup \{\neg p_i \cap p_j : s_j\} \cup d_{rest}]}$$

Fig. 5. Type-specific canonicalization rules.

multi-valued enum
$$\frac{s.\texttt{type} = \tau \quad\quad \tau \neq \texttt{boolean} \quad\quad s.\texttt{enum} = [v_1,\ldots,v_n] \quad\quad n > 1}{s \to \{\texttt{anyOf}:[\{\texttt{type}:\tau,\texttt{enum}:[v_1]\},\ldots,\{\texttt{type}:\tau,\texttt{enum}:[v_n]\}]\}}$$

null enum
$$\frac{s.\texttt{type} = \texttt{null} \quad\quad s.\texttt{enum} = [\texttt{null}]}{s \to \{\texttt{type}:\texttt{null}\}}$$

string enum
$$\frac{s.\texttt{type} = \texttt{string} \quad\quad s.\texttt{enum} = [v]}{s \to \{\texttt{type}:\texttt{string},\texttt{pattern}:\texttt{'^}v\texttt{\$'}\}}$$

number enum
$$\frac{s.\texttt{type} = \texttt{number} \quad\quad s.\texttt{enum} = [v]}{s \to \{\texttt{type}:\texttt{number},\texttt{minimum}:v,\texttt{maximum}:v\}}$$

array enum
$$\frac{s.\texttt{type} = \texttt{array} \quad\quad s.\texttt{enum} = [[v_1,\ldots,v_n]]}{s \to \{\texttt{type}:\texttt{array},\texttt{minItems}:n,\texttt{maxItems}:n,\texttt{items}:[\{\texttt{enum}:[v_1]\},\ldots,\{\texttt{enum}:[v_n]\}]\}}$$

object enum
$$\frac{s.\texttt{type} = \texttt{object} \quad\quad s.\texttt{enum} = [\{k_1:v_1,\ldots,k_n:v_n\}]}{s \to \{\texttt{type}:\texttt{object},\texttt{required}:[k_1,\ldots,k_n],\texttt{additionalProperties}:\texttt{false},\\ \texttt{properties}:\{k_1:\{\texttt{enum}:[v_1]\},\ldots,k_n:\{\texttt{enum}:[v_n]\}\}\}}$$

Fig. 6. Simplification rules to eliminate enum except for type boolean.

additionalProperties from the resulting schema. Rule *object with string list dependencies*, when iterated, eliminates dependencies specified as a list of property names by rewriting them into dependencies specified as a schema. Rule *object with schema dependencies*, when iterated, eliminates dependencies of a key k_i on a schema s_i by rewriting them into a conjunction with either s_i or a schema that enforces the absence of k_i. Rule *object with overlapping pattern properties* rewrites a pair of pattern properties with overlapping regular expressions p_i and p_j so their regular expressions match only disjoint keys, by replacing them with different schemas for the cases where (i) both p_i and p_j match, (ii) only p_i and not p_j matches, and (iii) only p_j and not p_i matches. When iterated, this eliminates all overlapping patterns, facilitating local reasoning.

4.2 Simplification of Enumeration, Negation, Intersection, and Union Types

This section describes the second step of our approach: a simplifier that compiles any canonical JSON schema into an equivalent simplified schema. Column "Simplified" of Table 2 summarizes the properties that the simplifier establishes. The simplifier eliminates many cases of enum, not, allOf, and anyOf connectives, thus making subschema checking rules less complicated. Unfortunately, in some cases, JSON schema cannot express schemas without these connectives, so they cannot be completely simplified away.

Figure 6 shows simplification rules for enum, which turn schemas with enums into schemas without enums by using restrictions keywords from their corresponding types instead. Rule *multi-valued enum* puts each non-Boolean enumerated value into an enum of its own. The rules for primitive types (null, string, and number) express a primitive enumerated value via a schema that does not use an enum. For instance, in Figure 3c, the enumerated empty string value is compiled into the regular expression '^\$' before computing its complement '.+' in Figure 3d. The rules for structured types (array and object) push down enums to components; with iteration, the rules eventually reach primitive types and the enums get eliminated. The simplifier does not eliminate Boolean enumerations as the space of values is finite and there is no other way to specify the true or false value.

Figure 7 shows simplification rules for not, which eliminates negation except for numbers, arrays, and objects. Rule *not type* turns a schema of a given type τ into a disjunction of either $\neg s$ (the complement of the values permitted by s in τ) or values of any type other than τ. An example for this rule in action is the rewrite from Figure 3c to Figure 3d, where the complement of a string schema

$$\text{not type} \quad \frac{s.\texttt{type} = \tau \quad \tau \in \mathit{Jtypes}}{\{\texttt{not}:s\} \to \{\texttt{anyOf}:[\neg s, \{\texttt{type}:(\mathit{Jtypes}\setminus\tau)\}]\}} \qquad \text{complement null} \quad \frac{s.\texttt{type} = \texttt{null}}{\neg s \to \{\texttt{not}:\{\}\}}$$

$$\begin{array}{l}\text{complement} \\ \text{boolean}\end{array} \quad \frac{s.\texttt{type} = \texttt{boolean} \quad s.\texttt{enum} = e}{\neg s \to \{\texttt{type}:\texttt{boolean}, \texttt{enum}:\neg e\}} \qquad \begin{array}{l}\text{complement} \\ \text{string}\end{array} \quad \frac{s.\texttt{type} = \texttt{string} \quad s.\texttt{pattern} = p}{\neg s \to \{\texttt{type}:\texttt{string}, \texttt{pattern}:\neg p\}}$$

$$\text{not anyOf} \quad \frac{s = \{\texttt{anyOf}:[s_1, \ldots, s_n]\}}{\{\texttt{not}:s\} \to \{\texttt{allOf}:[\{\texttt{not}:s_1\}, \ldots, \{\texttt{not}:s_n\}]\}}$$

$$\text{not allOf} \quad \frac{s = \{\texttt{allOf}:[s_1, \ldots, s_n]\}}{\{\texttt{not}:s\} \to \{\texttt{anyOf}:[\{\texttt{not}:s_1\}, \ldots, \{\texttt{not}:s_n\}]\}} \qquad \text{not not} \quad \frac{s = \{\texttt{not}:s_1\}}{\{\texttt{not}:s\} \to s_1}$$

Fig. 7. Simplification rules to eliminate negation, except for types number, array, and object.

introduces schemas of all non-string types. The *complement* rules for null, boolean, and string use the bottom type {not:{}}, the complement of the Boolean enumeration, and the complement of the regular expression, respectively. Rules *not anyOf* and *not allOf* use De Morgan's theorem to push negation through disjunction and conjunction, and rule *not not* eliminates double negation. Unfortunately, JSON Schema is not closed under complement for numbers, arrays, and objects. For example, the complement of schema {type:'number',multipleOf:1} is $\mathbb{R} \setminus \mathbb{Z}$, which cannot be expressed in JSON schema without a negation. Similar counter-examples exist for array and object schemas. The case of negated number schemas is handled later during subschema checking.

Figure 8 shows simplification rules for allOf. For example, the intersection type in Figure 3d yields the simplified schema in Figure 3e. Rule *singleton allOf* rewrites a conjunct of just one schema into that schema. Rule *fold allOf* turns an n-ary allOf into a binary one, so the remaining rules need to handle only the binary case. Rule *intersect heterogeneous types* returns the bottom type because intersection of incompatible types is the empty set, so the remaining rules only need to handle homogeneously-typed schemas. Rule *intersect null* rewrites two nulls to one null. Rule *intersect boolean* uses the intersection of enumerations. Rule *intersect string* uses the intersection of regular expressions. Rule *intersect number* uses helper functions *schema2range* and *range2schema* to convert back and forth between number schemas and mathematical ranges, and *lcm* to compute the least common multiple of the multipleOf constraints, where *lcm* handles undefined arguments by returning the other argument if defined, or an undefined value if both arguments are undefined. Rule *intersect array* takes advantage of the canonical form, where items are always specified as lists, to compute an item-wise intersection; undefined per-item schemas default to additionalItems. Rule *intersect object* simply picks up the union of the patternProperties keywords, relying on the rule for objects with overlapping patternProperties to make them disjoint again later. Finally, rule *intersect anyOf* pushes conjunctions through disjunctions by using the distributivity of intersection over union. We choose not to push intersections through negations because we prefer the end result of simplification to resemble distributive normal form to the extent possible.

Figure 9 shows simplification rules for anyOf. In contrast to intersection, union allows incompatible types, e.g., string or null as in Figure 3e. Fortunately, such heterogeneous unions are non-overlapping for inhabited schemas, so they can be handled later during subschema checking. Rule *singleton anyOf* rewrites a disjunct of just one schema into that schema. Rule *fold anyOf* turns an n-ary anyOf into a binary one so the remaining rules only need to handle the binary case. Rule *union null* rewrites two nulls to one. Rule *union boolean* uses the union of enumerations. Rule *union string* uses the union of regular expressions. Rule *union number* turns a binary union with overlap into a ternary non-overlapping union. In other words, while it does not eliminate the union of number schemas, it does simplify subschema checking by at least making the union

$$\text{singleton allOf } \frac{s.\mathtt{allOf} = [s_1]}{s \to s_1} \qquad \text{fold allOf } \frac{s.\mathtt{allOf} = [s_1, s_2, \ldots, s_n] \qquad n \geq 3}{s \to \{\mathtt{allOf} : [s_1, \{\mathtt{allOf} : [s_2, \ldots, s_n]\}]\}}$$

$$\begin{array}{c}\text{intersect hetero-}\\ \text{geneous types}\end{array} \frac{s_1.\mathtt{type} \neq s_2.\mathtt{type}}{\{\mathtt{allOf} : [s_1, s_2]\} \to \{\mathtt{not} : \{\}\}} \qquad \text{intersect null } \frac{s_1.\mathtt{type} = \mathtt{null} \qquad s_2.\mathtt{type} = \mathtt{null}}{\{\mathtt{allOf} : [s_1, s_2]\} \to \{\mathtt{type} : \mathtt{null}\}}$$

$$\text{intersect boolean } \frac{s_1.\mathtt{type} = \mathtt{boolean} \qquad s_2.\mathtt{type} = \mathtt{boolean}}{\{\mathtt{allOf} : [s_1, s_2]\} \to \{\mathtt{type} : \mathtt{boolean}, \mathtt{enum} : s_1.\mathtt{enum} \cap s_2.\mathtt{enum}\}}$$

$$\text{intersect string } \frac{s_1.\mathtt{type} = \mathtt{string} \qquad s_2.\mathtt{type} = \mathtt{string}}{\{\mathtt{allOf} : [s_1, s_2]\} \to \{\mathtt{type} : \mathtt{string}, \mathtt{pattern} : s_1.\mathtt{pattern} \cap s_2.\mathtt{pattern}\}}$$

$$\begin{array}{c}\text{intersect}\\ \text{number}\end{array} \frac{s_1.\mathtt{type} = \mathtt{number} \qquad s_2.\mathtt{type} = \mathtt{number} \qquad r_1 = schema2range(s_1) \qquad r_2 = schema2range(s_2)}{\{\mathtt{allOf} : [s_1, s_2]\} \to range2schema(r_1 \cap r_2) \cup \{\mathtt{multipleOf} : lcm(s_1.\mathtt{multipleOf}, s_2.\mathtt{multipleOf})\}}$$

$$\begin{array}{c}\text{intersect}\\ \text{array}\end{array} \frac{\begin{array}{c}s_1.\mathtt{type} = \mathtt{array} \quad s_2.\mathtt{type} = \mathtt{array} \quad s_1.\mathtt{items} = [s_{1_1}, \ldots, s_{1_k}] \quad s_2.\mathtt{items} = [s_{2_1}, \ldots, s_{2_m}]\\ n = max(k, m)\end{array}}{\begin{array}{l}\{\mathtt{allOf} : [s_1, s_2]\} \to \{\mathtt{type} : \mathtt{array},\\ \qquad \mathtt{minItems} : max(s_1.\mathtt{minItems}, s_2.\mathtt{minItems}),\\ \qquad \mathtt{maxItems} : min(s_1.\mathtt{maxItems}, s_2.\mathtt{maxItems}),\\ \qquad \mathtt{items} : [\{\mathtt{allOf} : [s_{1_1} \parallel s_1.\mathtt{additionalItems}, s_{2_1} \parallel s_2.\mathtt{additionalItems}]\},\\ \qquad \qquad \ldots,\\ \qquad \qquad \{\mathtt{allOf} : [s_{1_n} \parallel s_1.\mathtt{additionalItems}, s_{2_n} \parallel s_2.\mathtt{additionalItems}]\}],\\ \qquad \mathtt{additionalItems} : \{\mathtt{allOf} : [s_1.\mathtt{additionalItems}, s_2.\mathtt{additionalItems}]\},\\ \qquad \mathtt{uniqueItems} : s_1.\mathtt{uniqueItems} \wedge s_2.\mathtt{uniqueItems}\}\end{array}}$$

$$\begin{array}{c}\text{intersect}\\ \text{object}\end{array} \frac{s_1.\mathtt{type} = \mathtt{object} \qquad s_2.\mathtt{type} = \mathtt{object}}{\begin{array}{l}\{\mathtt{allOf} : [s_1, s_2]\} \to \{\mathtt{type} : \mathtt{object},\\ \qquad \mathtt{minProperties} : max(s_1.\mathtt{minProperties}, s_2.\mathtt{minProperties}),\\ \qquad \mathtt{maxProperties} : min(s_1.\mathtt{maxProperties}, s_2.\mathtt{maxProperties}),\\ \qquad \mathtt{required} : s_1.\mathtt{required} \cup s_2.\mathtt{required},\\ \qquad \mathtt{patternProperties} : s_1.\mathtt{patternProperties} \cup s_2.\mathtt{patternProperties}\}\end{array}}$$

$$\text{intersect anyOf } \frac{s_2 = \{\mathtt{anyOf} : [s_{2_1}, \ldots, s_{2_n}]\}}{\{\mathtt{allOf} : [s_1, s_2]\} \to \{\mathtt{anyOf} : [\{\mathtt{allOf} : [s_1, s_{2_1}]\}, \ldots, \{\mathtt{allOf} : [s_1, s_{2_n}]\}]\}}$$

Fig. 8. Simplification rules to eliminate `allOf` except for connective `not`.

disjoint so it can be checked elementwise. Unfortunately, JSON Schema is not closed under union for types number, array, and object. For example, the union of `{type:number,minimum:0}` and `{type: number,multipleOf:1}` is $\mathbb{R}^+ \cup \mathbb{Z}$, which cannot be expressed in JSON schema without `anyOf`. There are similar counter-examples for arrays and objects. The case of unioned number schemas is handled later in subschema checking. As mentioned earlier, we would like simplification to end in schemas that resemble a distributive normal form, so we choose not to push `anyOf` through `allOf` or `not`.

4.3 JSON Subschema Checking

Given two canonicalized and simplified schemas, the third step of our approach checks whether one schema is a subtype of the other. Figure 10 presents inference rules defining the subschema relation on canonical, simplified schemas. All rules are algorithmically checkable, and all rules except for *schema uninhabited* are type-directed. To simplify their presentation, some of the rules use quantifiers, but all quantifiers are bounded and can thus be checked via loops.

$$\text{singleton anyOf} \quad \frac{s.\texttt{anyOf} = [s_1]}{s \rightarrow s_1} \qquad \text{fold anyOf} \quad \frac{s.\texttt{anyOf} = [s_1, s_2, \ldots, s_n] \qquad n \geq 3}{s \rightarrow \{\texttt{anyOf}: [s_1, \{\texttt{anyOf}: [s_2, \ldots, s_n]\}]\}}$$

$$\text{union null} \quad \frac{s_1.\texttt{type} = \texttt{null} \qquad s_2.\texttt{type} = \texttt{null}}{\{\texttt{anyOf}: [s_1, s_2]\} \rightarrow \{\texttt{type}:\texttt{null}\}}$$

$$\text{union boolean} \quad \frac{s_1.\texttt{type} = \texttt{boolean} \qquad s_2.\texttt{type} = \texttt{boolean}}{\{\texttt{anyOf}: [s_1, s_2]\} \rightarrow \{\texttt{type}:\texttt{boolean}, \texttt{enum}:s_1.\texttt{enum} \cup s_2.\texttt{enum}\}}$$

$$\text{union string} \quad \frac{s_1.\texttt{type} = \texttt{string} \qquad s_2.\texttt{type} = \texttt{string}}{\{\texttt{anyOf}: [s_1, s_2]\} \rightarrow \{\texttt{type}:\texttt{string}, \texttt{pattern}:s_1.\texttt{pattern} \cup s_2.\texttt{pattern}\}}$$

$$\text{union number} \quad \frac{\begin{array}{c} s_1.\texttt{type} = \texttt{number} \qquad s_2.\texttt{type} = \texttt{number} \qquad r_1 = schema2range(s_1) \qquad r_2 = schema2range(s_2) \\ r_1 \cap r_2 \neq \emptyset \end{array}}{\begin{array}{l} \{\texttt{anyOf}: [s_1, s_2]\} \rightarrow \{\texttt{anyOf}: [\\ \quad range2schema(r_1 \cap r_2) \cup \{\texttt{multipleOf}: gcd(s_1.\texttt{multipleOf}, s_2.\texttt{multipleOf})\}, \\ \quad range2schema(r_1 \setminus r_2) \cup \{\texttt{multipleOf}: s_1.\texttt{multipleOf}\}, \\ \quad range2schema(r_2 \setminus r_1) \cup \{\texttt{multipleOf}: s_2.\texttt{multipleOf}\}]\} \end{array}}$$

Fig. 9. Eliminating overlapping `anyOf`, except for connectives `not` and `allOf`, and types `array` and `object`.

Rule *schema uninhabited* states that an uninhabited schema is a subtype of any other schema. It uses an auxiliary *inhabited* predicate, which is elided for space but easily computable for primitives (recall that emptiness is decidable for regular languages). For structures, the predicate ensures that the schemas of all required components are inhabited. For logic connectives, the predicate is more involved but decidable. The rule for uninhabited types is the only rule that is not type-directed. Because canonicalization generally separates schemas by type, all other rules check same-typed schemas. We can handle uninhabited schemas independently of their type because there is no actual data of that type that would require type-specific reasoning.

Rule *subschema non-overlapping anyOf* handles `anyOf` schemas for the cases where simplification eliminates overlapping unions. Helper function *nonOverlapping* checks for unions of arrays and objects and conservatively assumes that those might overlap. In the non-overlapping case, it suffices to check the component schemas independently. For each schema on the left, we require a same-typed super schema on the right.

Rule *subschema number* is the most complicated of the subtype rules for primitive types due to `multipleOf` constraints. The simplifier cannot push negation through `multipleOf` constraints, and it cannot combine `allOf` combinations of such negated schemas. As a result, the rule has to handle multiple such constraints on both sides of the relation, with or without negation. We treat simple number schemas as single-element `allOf`s for consistency. This rule verifies that any number allowed by the set of constraints on the left is also allowed by the set of constraints on the right using an auxiliary *subNumber* relation, which is sketched in the following.

The *subNumber* relation first normalizes all schema range bounds by rounding them to the nearest included number that satisfies its `multipleOf` constraint. For each side, it then finds the least and greatest finite bound used. Every unbounded schema is split into two (or three for totally unbounded) schemas: one (or two) that are unbounded on one side, with the least/greatest bound as the other bound. The "middle" part is bounded. All these derived schemas keep the original `multipleOf`. The bounded schemas can all be checked (exhaustively if needed). For the unbounded schemas, we can separately check the positive and negative schemas, since they do not interact in interesting ways over unbounded sets. If *PL* and *PR* are the left and right positive schemas, and *NL*

$$\text{subschema uninhabited} \ \frac{\neg\, inhabited(s_1)}{s_1 <: s_2}$$

$$\text{subschema non-overlapping anyOf} \ \frac{\forall i \in \{1..n\}, \exists j \in \{1..m\}, s_i <: t_j \quad nonOverlapping([t_1, \ldots, t_m])}{\{\texttt{anyOf}:[s_1, \ldots, s_n]\} <: \{\texttt{anyOf}:[t_1, \ldots, t_m]\}}$$

$$\text{subschema null} \ \frac{s_1.\texttt{type} = \texttt{null} \quad s_2.\texttt{type} = \texttt{null}}{s_1 <: s_2}$$

$$\text{subschema boolean} \ \frac{s_1.\texttt{type} = \texttt{boolean} \quad s_2.\texttt{type} = \texttt{boolean} \quad s_1.\texttt{enum} \subseteq s_2.\texttt{enum}}{s_1 <: s_2}$$

$$\text{subschema string} \ \frac{s_1.\texttt{type} = \texttt{string} \quad s_2.\texttt{type} = \texttt{string} \quad s_1.\texttt{pattern} \subseteq s_2.\texttt{pattern}}{s_1 <: s_2}$$

$$\text{subschema number} \ \frac{\begin{array}{c} \forall i \in \{1..k\}, \texttt{not} \notin dom\, s_i \wedge s_i.\texttt{type} = \texttt{number} \\ \forall i \in \{k+1..n\}, \texttt{not} \in dom\, s_i \wedge s_i.\texttt{not.type} = \texttt{number} \\ \forall i \in \{1..l\}, \texttt{not} \notin dom\, t_i \wedge t_i.\texttt{type} = \texttt{number} \\ \forall i \in \{l+1..n\}, \texttt{not} \in dom\, t_i \wedge t_i.\texttt{not.type} = \texttt{number} \\ subNumber(([s_1, \ldots, s_k], [s_{k+1}, \ldots, s_n]), ([t_1, \ldots, t_l], [t_{l+1}, \ldots, t_m])) \end{array}}{\{\texttt{allOf}:[s_1, \ldots, s_k, s_{k+1}, \ldots, s_n]\} <: \{\texttt{allOf}:[t_1, \ldots, t_l, t_{l+1}, \ldots, t_m]\}}$$

$$\text{subschema array} \ \frac{\begin{array}{cc} s_1.\texttt{type} = \texttt{array} & s_2.\texttt{type} = \texttt{array} \\ s_1.\texttt{minItems} \geq s_2.\texttt{minItems} & s_1.\texttt{maxItems} \leq s_2.\texttt{maxItems} \\ s_1.\texttt{items} = [s_{1_1}, \ldots, s_{1_k}] & s_2.\texttt{items} = [s_{2_1}, \ldots, s_{2_m}] \\ \multicolumn{2}{c}{\forall i \in [0, \ldots, max(k,m)+1], s_{1_i} \parallel s_1.\texttt{additionalItems} <: s_{2_i} \parallel s_2.\texttt{additionalItems}} \\ \multicolumn{2}{c}{s_2.\texttt{uniqueItems} \implies (s_1.\texttt{uniqueItems} \vee allDisjointItems(s_1))} \end{array}}{s_1 <: s_2}$$

$$\text{subschema object} \ \frac{\begin{array}{cc} s_1.\texttt{type} = \texttt{object} & s_2.\texttt{type} = \texttt{object} \\ s_1.\texttt{minProperties} \geq s_2.\texttt{minProperties} & s_1.\texttt{maxProperties} \leq s_2.\texttt{maxProperties} \\ \multicolumn{2}{c}{s_1.\texttt{required} \supseteq s_2.\texttt{required}} \\ \multicolumn{2}{c}{\forall p_1 : s_{p_1} \in s_1.\texttt{patternProperties}, p_2 : s_{p_2} \in s_2.\texttt{patternProperties}, p_1 \cap p_2 \neq \emptyset \implies s_{p_1} <: s_{p_2}} \end{array}}{s_1 <: s_2}$$

Fig. 10. JSON Schema subtype inference rules.

and NR are the left and right negative schemas, we verify that the constraints divide each other:

$$\forall_{pl \in PL}, \exists_{pr \in PR}, \ pl.\texttt{multipleOf} \bmod pr.\texttt{multipleOf} = 0$$
$$\forall_{nr \in NR}, \exists_{nl \in NL}, \ nr.\texttt{multipleOf} \bmod nl.\texttt{multipleOf} = 0$$

Rule *subschema array* checks two array schemas. The left array size bounds should be within the size bounds of the right array. Additionally, the schema of every item specified in the former needs to be a subschema of the corresponding specification in the latter. If a schema is not explicitly provided, the schema provided by `additionalItems` is used. Recall that canonicalization adds in a default `additionalItems` schema if it was not specified. Additionally, if the right side specifies that the items must be unique, then the left needs to either specify the same or implicitly enforce this. For example, `{type:array,items:[{enum:[0]},{enum:[1]}]}` is a subschema of `{type:array, uniqueItems:true}`. The *allDisjointItems* predicate checks for this by first obtaining the set of all the effective item schemas: every item schema for an index within the specified min/max bounds, and `additionalItems` if any allowed indices are unspecified. It then verifies that the conjunction of all pairs of effective items schemas are uninhabited.

Rule *subschema object* checks two object schemas. It first verifies that the number of properties of both sides have the appropriate relation, and that the left side requires all the keys that the right side requires. Next, for every regular expression pattern p_1 on the left, if there is an overlapping regular expression pattern p_2 on the right, it checks that the corresponding schemas are subschemas. This check can be done separately for one pattern at a time thanks to the fact that canonicalization eliminates overlapping pattern properties.

5 IMPLEMENTATION

We implemented our subschema checker as a Python tool in around 2,000 lines of code. The implementation builds upon the jsonschema library[1] to validate schemas before running our subtype checking, the greenery library[2] for computing intersections of regular expressions, and the jsonref library[3] for resolving JSON schemas references.

6 EVALUATION

This section evaluates the implementation of our JSON subschema checker, which we refer to as *jsonsubschema*. It answers the following research questions:

RQ₁ How correct is *jsonsubschema* in practice?
RQ₂ How does *jsonsubschema* fare against existing work?
RQ₃ How complete is *jsonsubschema* in practice?
RQ₄ How efficient is *jsonsubschema*?

6.1 Experimental Setup

We evaluate our subschema checker on four datasets of JSON schemas from different domains: JSON schema official test suite, *WP*, *K8s*, and *Lale*.

The official test suite for JSON Schema draft-04 [4] is a widely-used test suite for JSON Schema validators and provides 146 schemas and 531 tests that offer full coverage of the JSON Schema specification. Each test provides a JSON document d to be validated against a specific schema s and the expected correct validation behavior $res \in \{True, False\}$, i.e., tests are of the form $JValid(d, s) = res$. These schemas are designed to cover the JSON schema features, but they are not representative for real-world schemas. Moreover, these schemas were not designed to test the subschema relation we define in this work, so we only use this test suite to evaluate the canonicalization and simplification steps (Sections 4.1 and 4.2) of our approach.

The three other datasets, *WP*, *K8s*, and *Lale*, are used for evaluating the full *jsonsubschema* approach. *WP* is a collection of schemas describing content used by the Washington Post within the *Arc Publishing* content creation and management system.[5] *K8s* is the set of JSON schemas describing the OpenAPI specifications for *Kubernetes*, an open-source system for automating deployment, scaling, and management of containerized applications.[6] Since OpenAPI specifications contain more information beyond the schemas for the REST API endpoints, we use a set of JSON schemas extracted from them.[7] Specifically, we used the standalone flavor of schemas where $ref have been

---

[1] https://github.com/Julian/jsonschema
[2] https://github.com/qntm/greenery
[3] https://github.com/gazpachoking/jsonref
[4] https://github.com/json-schema-org/JSON-Schema-Test-Suite
[5] https://github.com/washingtonpost/ans-schema
[6] https://kubernetes.io/
[7] https://kubernetesjsonschema.dev/

Table 3. Dataset details.

Dataset	Versions	Schemas	Schema pairs	Schema pairs used for evaluation
WP	28	2,604	2,411	2,060
K8s	124	86,461	82,814	6,460
Lale	–	–	28	28
Total				8,548

resolved to local files. The *Lale* dataset is a set of schema pairs from the *Lale* open-source project[8], which is a Python library for type-driven automated machine learning.

The second and third datasets, *WP* and *K8s*, comprise several versions. We apply our subschema checker across each pair of consecutive versions of the same schema that introduces some textual modification, to spot whether a change may impact the compatibility of the corresponding systems. We use the third dataset, *Lale*, to find type errors in AI pipelines where wrong operators could be applied to specific datasets. We consider 4 *Lale* operators and 7 datasets, yielding 28 schema pairs.

Table 3 shows statistics for the three datasets that we use for the full evaluation of *jsonsubschema*. For example, *K8s* has in total 124 versions, with a total of 86,461 schemas. Due to the additions and deletions of schemas between pairs of consecutive versions, the total number of pairs of schemas across all pairs of subsequent versions is 82,814. Finally, since not every new version of an API modifies every schema, we only keep pairs of non-equal files. Overall, the total number of pairs of schemas is 8,548. Many of the schemas are of non-trivial size, with an average of 56KB and a maximum of 1,047KB. The first dataset, the JSON schema test suite, is omitted from Table 3 since it does not have pairs of versions and we only use it to evaluate the canonicalization and simplification steps.

All experiments are performed on an Intel Core i7-4600U CPU (2.10GHz) machine with 16GB of memory running Ubuntu 18.04 (64-bit).

6.2 Correctness in Practice

For RQ$_1$, we evaluate the correctness of the two main steps of our approach: canonicalization and simplification (canonicalization), and subtype checking.

6.2.1 Canonicalization. Canonicalization and simplification aim at producing a valid and simpler canonical schema that is semantically equivalent to the input schema. To test validity, *jsonsubschema* checks the canonicalized schema against the meta-schema of JSON Schema using an off-the-shelf JSON schema validator (Section 5). Across our entire dataset, there is no single case where this check fails.

To check that canonicalization and simplification preserve the semantics of the original input schema, we apply the canonicalization step to all schemas in the JSON Schema official test suite. We tested whether: $\forall s, \forall d, \text{JValid}(d, s) = res \implies \text{JValid}(d, \text{canonical}(s)) = res$. for schema s, JSON document d, and outcome res. In all cases except one where *jsonsubschema* yields a canonicalized schema, this new schema passes all relevant tests in the JSON schema test suite. This single case, we believe, is due to the ambiguity of the specification of JSON schema and hence, a mismatch between our own interpretation and the interpretation of the JSON schema validator of the semantics of the `allOf` connector when combined with the `additionalProperties` object constraint.

These results, of course, show only that canonicalization in most of the cases does not yield an invalid or a more strict schema than the input schema. The following experiments on the

---

Table 4. Effectiveness of *jsonsubschema* and comparison to the existing *issubset* tool.

	Pairs	jsonsubschema					issubset				
		Fail	TP	TN	FP	FN	Fail	TP	TN	FP	FN
<:	35	0	29	6	0	0	10	9	0	6	10
:>	35	0	31	4	0	0	10	21	0	4	0
≡	100	0	100	0	0	0	50	27	0	0	23
≢	100	0	63	37	0	0	0	63	0	37	0
Lale	28	0	12	16	0	0	7	3	10	0	8
Total	298	0	235	63	0	0	77	123	10	47	41

correctness of subtype checking rely on correct canonicalization, and hence, provide additional evidence.

6.2.2 Subtype Checking.

Self-equivalence. As an automated, large-scale correctness check of the subtype checking, we perform a simple sanity check that asks *jsonsubschema* whether a schema is equivalent to itself (Definition 3.2). We randomly sample 1000 schemas from the *K8s* dataset and run our subschema checker using the same schema on both sides of the subtype relation, i.e., checking for a schema s whether $s <: s$. Our subtype checking does not rely on any sort of structural equality. Therefore, in this setup, our implementation is oblivious to the fact that both schemas are the same, so it canonicalizes both schemas and then performs subtype checking. In all 1,000 samples, this test passes correctly.

Comparison against a ground truth. For further validation, we compare the results of *jsonsubschema* against a ground truth. Specifically, we gather pairs of schemas, along with their expected subtype relationship, in three ways. First, we randomly sample pairs that are textually different and manually assess their subtype relationship. Second, we sample consecutive versions of schemas from *WP* and *K8s* and manually assess their subtype relationship. For example, we consider versions 0.7.0 and 0.7.1 of the utils/named_entity schema from the *WP* dataset. This sample of pairs represents the usage scenario where our approach checks whether an evolving API specified through a JSON schema may break an application. Third, for the 28 schema pairs from *Lale*, the ground truth is whether or not the corresponding machine-learning operator throws an exception when training on the corresponding dataset. By checking the schemas statically, our subtype checker can avoid such runtime errors. To focus on the correctness of the subtype checker output, we sampled pairs of schemas from the set of schemas where our approach yields a decision, i.e., ignoring the few cases where *jsonsubschema* cannot decide on the subschema relation.

In total, we gathered 298 pairs with a ground truth subtype relationship, as summarized in Table 4. The <:, :>, ≡, and ≢ symbols represent what test we performed on each pair. For example, for each pair $\langle s, t \rangle$ in the <: row, the ground truth indicates whether $s <: t$ holds (positive, P) or not (negative, N). The *jsonsubschema* part of the table shows the results of applying our subschema checker to each pair. The TP, TN, FP, and FN columns indicate the true positives, true negatives, false positives, and false negatives, respectively. For example, TN means that the tool produces the correct result (T for true) and that the ground truth indicates that the relationship being tested is not expected to hold (N for negative). Our tool produces the correct results for all 298 schema pairs in the ground truth.

6.3 Comparison to Existing Work

As we discuss in Section 7, our work is the first to define the subschema relation on JSON schemas and present an algorithm to perform this check for a large subset of JSON schema features. Therefore, to our knowledge, there is no academic work that we can compare against. However, since our work is motivated by the practical need for subtyping JSON schemas, for RQ$_2$, we compare our implementation to the closest developer tool we could find, is-json-schema-subset (*issubset*) [12]. The *issubset* tool is written in TypeScript and its documentation states the same goal as ours: "Given a schema defining the output of some process A, and a second schema defining the input of some process B, will the output from A be valid input for process B?" We use the most recent version, which is 1.0.6.

We first run *issubset* on the three datasets, *WP*, *K8s*, and *Lale*. It fails to run on some schemas from *WP* and *Lale* due to unsupported schema versions, although the *issubset* documentation does not describe such a limitation. Next, we compare the correctness results of *issubset* against our *jsonsubschema*. The right part of Table 4 shows the correctness results for *issubset* using the methodology described in Section 6.2.2.

The first observation is that *issubset* produces a non-negligible number of true positives, which means it indeed captures some of the semantic of the subtyping relation. However, *issubset* also produces 47 false positives and 41 false negatives, i.e., gives a wrong answer to a subtype query. Overall, the existing tool gives a wrong answer in 40% of the cases where the tool does not fail.

To get a better understanding of the low recall, we inspected the code of *issubset* and tested it on simple hand-crafted schemas. We find that although the tool performs some simple semantic checks, e.g., it correctly reports `{'type':'integer'}` <: `{'type':'number'}`, it lacks the ability to capture the richness of JSON schema in many ways. For instance, it fails to detect `{'type':['string','null']}` ≡ `{'type':['null','string']}`, and is oblivious to uninhabited schemas, such as `{'type':'string','enum':[1]}`.

6.4 Completeness in Practice

Being complete in practice is difficult. To balance completeness and effort, there is a set of features our approach currently cannot deal with. As discussed in Section 4, our approach cannot reason about negation and complement of object, array, and numeric with `multipleOf` schemas. Therefore, for RQ$_3$, we report on two dimensions. The first is the pervasiveness of different validation keywords and which of them *jsonsubschema* supports. The second dimension is failure cases of *jsonsubschema*, i.e, cases where we do not produce a subtype decision, due to a limitation of the approach.

6.4.1 Pervasiveness of Validation Keywords and Supported Features. Figure 11 shows the frequency of validation keywords across all schemas in the *K8s* and *WP* datasets. Validation keywords on the x-axis are sorted by their relevance to each schema type and according to the order of keywords in Table 1. The figure shows that *jsonsubschema* indeed supports the majority of JSON schema features which are used in practice

We observe that JSON schema types `null` and `string` are the two most prevalent schema types present in the dataset. Both types are fully supported in the subtype checking performed by *jsonsubschema* as indicated by the color code in Figure 11.

The keywords `properties`, `additionalProperties`, and `required` for specifying constrains on a JSON `object` show up next on the order of the number of use cases. Of these keywords, *jsonsubschema* fully supports `properties`. The `additionalProperties` and `required` keywords are supported whenever they are not used in union schemas or negated schemas. In general, disjunction of schemas happens rarely (146 occurrence among millions of occurrences of other keywords); while negated schemas are not used at all in our dataset of real-world schemas.
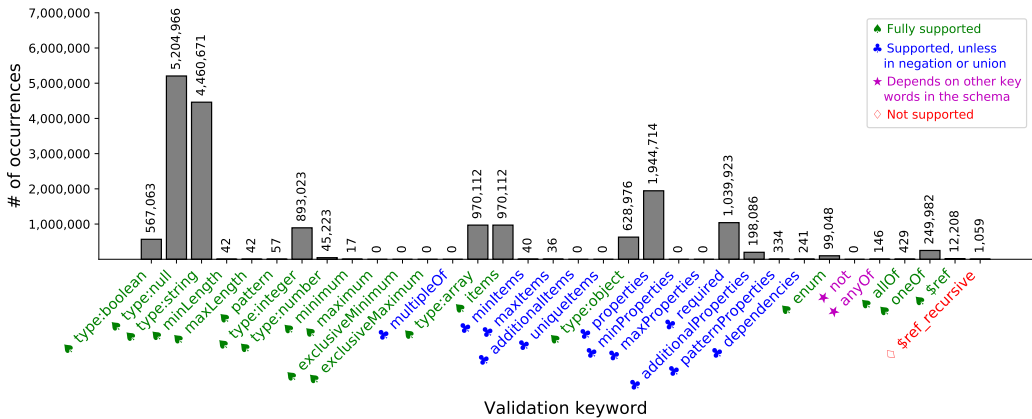
Fig. 11. Prevalence of JSON schema validation keywords in practice and supported features in *jsonsubschema*.

Worth noting here is that the counts in Figure 11 are for schemas that do not use the negated schema keyword `not` at all, which is also evident from its frequency being 0. The reason is that in the dataset of schemas, there is no single use of a negated schema. In fact, the use of negation in JSON schemas is indeed highly discouraged since the purpose of schema validation is to constraint what is allowed rather than filtering out what is disallowed. The NSA security guidelines for using JSON schemas also advises against using negated schemas for the same reason [9]. That said, *jsonsubschema* still supports the negation of all *𝒥primitive* except for the union and negation of numeric schemas with a `multipleOf` constraint. Overall, this shows that the incompleteness of *jsonsubschema* rarely affects its usability in real world, on a large dataset of real-world schemas.

The only feature that is not supported at the moment is recursive references in schemas using `$ref`. Although our approach currently does not handle recursive schemas, we know theoretically that subtyping recursive types is decidable [1].

6.4.2 Unsupported Features in Practice. The *jsonsubschema* tool reports an error and terminates without yielding a decision when an unsupported JSON schema feature is encountered.

Table 5 shows the cases when *jsonsubschema* fails on our dataset. In total, out of 8,548 schemas pairs, the subschema checks fails for 5.69% of the pairs. The table shows three kinds of failures that happen in practice due to limitations of our approach. The first and most dominant failure reason are circular and recursive schemas. As indicated above, this is not an inherent limitation of *jsonsubschema*. The second case is the rare use case of negating objects schemas. As seen in Table 5, only 0.34% of schema pairs fail due to the absence of this feature. It is worth noting here is that the original schemas do not use negated schemas at all as seen in Figure 11. However, they were introduced as part of *jsonsubschema* canonicalization phase (Section 4.1) where `oneOf` constraints are re-written into disjunction of conjunctions with negations (Figure 5). We have discussed in Section 4 that JSON schema is not closed under union or negation of object schemas. The third case is when string schemas or object schemas use non-regular language for specifying textual patterns using the keywords `pattern` amd `patternProperties`, respectively. Inclusion in non-regular languages (e.g., regular expressions with positive and negative look-around) is undecidable and is beyond our scope.

Table 5. Reasons for incompleteness in practice.

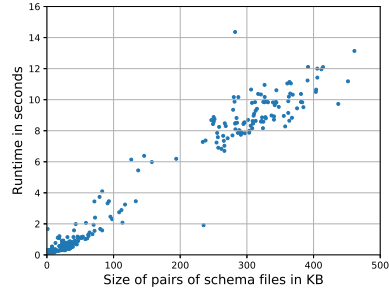| Failure reason	Count	%
Recursive or circular $ref	453	5.30%
Negated object schema	29	0.34%
Non-regular regex pattern	5	0.06%
Total	487	5.69%

Fig. 12. Efficiency of Subschema checking.

For eleven pairs of schemas, at least one of the schemas had an invalid $ref. For two pairs of schemas, at least one of the two files is not a valid JSON document. These thirteen pairs are omitted from Table 5.

For the same incompleteness reasons described here, the canonicalizer failed to canonicalize 22 schemas out of the 146 schemas in the JSON schema test suite. However, the schemas in this test suite indeed cover the entire language features of JSON schema but they do not represent the real-world use cases of JSON schema in practice.

6.5 Efficiency

To evaluate how fast our subschema checker is in practice for RQ_4, we measure the time taken by subschema checks on a sample of 798 pairs of non-equal schemas from Table 3. We took every time measurement ten times and report the average. Figure 12 shows the size of pairs of schema files in KB against the time subschema checking takes in seconds.

In most cases, our subschema checker terminates within a few seconds for moderately sized schemas, with time increasing roughly linearly with the schema file size. However, our subschema approach is lazy and terminates on the first violation of a subtyping rule. On one pair of schemas in our dataset, eliminated from the figure for scaling sake, it took around 2.8 minutes to terminate, which is not optimal for production. We will explore how to improve on this, for instance, by on-demand canonicalization.

7 RELATED WORK

7.1 JSON Schema and Schema Subtyping

Practitioners have significant interest in reasoning about the subtype relation of JSON schemas. Section 6.3 has an experimental comparison against the strongest competitor among the available tools, is-json-schema-subset [12], which was developed concurrently with our work. Another closely related tool[10] relies on simple syntactic checks. For example, that work considers a change as a breaking change whenever a node is removed from the schema. As illustrated by Figure 3, removing nodes (and replacing them by others) may yield not only subtypes but even equivalent schemas. Yet another existing tool[11] checks whether two schemas are equivalent but does not address the subtyping problem.

Pezoa et al. [19] formally define the syntax and semantics of JSON Schema, including the JSON validation problem. An alternative formulation of JSON validation uses a logical formalism [6].

---

[10]https://bitbucket.org/atlassian/json-schema-diff-validator
[11]https://github.com/mokkabonna/json-schema-compare

Baazizi et al. [2] address the problem of inferring schemas for irregular JSON data, but their work does not use the JSON Schema standard we are targeting here. None of the above pieces of work addresses the subschema problem.

There are other schema definition languages for JSON besides JSON Schema. One popular alternative is the Swagger/OpenAPI specification language.[12] While similar to JSON Schema, it is not fully compatible. The swagger-diff tool[13] aims at finding breaking API changes through a set of syntactic checks, but does not provide the detailed checks that we do. Avro is another schema definition language[14], which, however, does not specify a subschema relation. We envision our work to help define subtype relations of these alternative schema definition languages.

7.2 Type Systems for XML, TypeScript, and Python

CDuce is a functional language designed for working with XML, which can reason about the types of XML documents and about subtype relations [4]. [25] address the problem of subschema checking for XML, where it is called schema containment. These approaches treat XML as tree automata, which is impossible for JSON, as JSON Schema is more expressive than tree automata [19].

Both JavaScript and Python have a convenient built-in syntax for JSON documents. Furthermore, there are type systems retrofitted onto both JavaScript [5] and Python [26]. Therefore, a reasonable question to ask is whether JSON schema subtype queries could be decided by expressing JSON documents in those languages and then using the subtype checker of those type systems. Unfortunately, this is not the case, since JSON Schema contains several features that those type systems cannot express. For instance, JSON Schema supports negation, `multipleOf` on numbers, and `pattern` on strings, none of which those type systems support.

7.3 Applications of Subschemas

One application of JSON subschema is for statically reasoning about breaking changes of web APIs. A study of the evolution of such APIs shows that breaking changes are frequent [16]. Another study reports that breaking changes of web APIs cause distress among developers [10]. Since JSON schemas and related specifications are widely used to specify data types, our subschema checker could help identify breaking changes already statically and on the schema-level, instead of relying on testing.

Data validation for industry-deployed machine learning pipelines is of crucial value as such pipelines are usually retrained frequently with new data. In order to validate incoming data, Google TFX [3] synthesizes a custom data schema based on statistics from available data and uses this schema to validate future data instances fed to the TensorFlow pipeline [7]. Amazon production ML pipelines [22] offer a declarative API which allow users to manually define desired constraints or properties of data. Then data quality metrics such as completeness and consistency are measured on real-time data with respect to the pre-defined constraints and anomalies are reported. Both systems are missing an explicit notion of schema subtyping. For instance, TFX uses versioned schemas to track the evolution of inferred data schemas, and reports back to the user whether to update to a more (or less) permissive schema based on the historical and new data instances [3]. LALE uses JSON schemas to specify both correct ML pipelines and their search space of hyperparameters [13]. The ML Bazaar also specifies ML primitives via JSON [24]. Another type-based system for building ML pipelines is described by [20]. These systems could benefit from JSON subschema checking to avoid running and deploying incompatible ML pipelines.

---

[12]https://swagger.io/
[13]https://github.com/civisanalytics/swagger-diff
[14]http://avro.apache.org/

8 CONCLUSION

This paper introduces a subtype checker for JSON Schema. There are several features in JSON Schema that make subtype checking difficult, including a full set of Boolean connectives, enumerations containing values of possibly heterogeneous other types, regular expressions for strings, and multiple-of constraints for numbers. Our checker is the first to effectively handle these cases. The evaluation demonstrates that the tool works well on a large set of examples of high real-world importance, including web APIs, cloud computing, and artificial intelligence.

REFERENCES

[1] Roberto M. Amadio and Luca Cardelli. 1993. Subtyping Recursive Types. *Transactions on Programming Languages and Systems (TOPLAS)* 15, 4 (Sept. 1993), 575–631. http://doi.acm.org/10.1145/155183.155231

[2] Mohamed Amine Baazizi, Dario Colazzo, Giorgio Ghelli, and Carlo Sartiani. 2017. Counting types for massive JSON datasets. In *Symposium on Database Programming Languages (DBPL)*. 9:1–9:12. https://doi.org/10.1145/3122831.3122837

[3] Denis Baylor, Eric Breck, Heng-Tze Cheng, Noah Fiedel, Chuan Yu Foo, Zakaria Haque, Salem Haykal, Mustafa Ispir, Vihan Jain, Levent Koc, Chiu Yuen Koo, Lukasz Lew, Clemens Mewald, Akshay Naresh Modi, Neoklis Polyzotis, Sukriti Ramesh, Sudip Roy, Steven Euijong Whang, Martin Wicke, Jarek Wilkiewicz, Xin Zhang, and Martin Zinkevich. 2017. TFX: A TensorFlow-Based Production-Scale Machine Learning Platform. In *Conference on Knowledge Discovery and Data Mining (KDD)*. 1387–1395. https://doi.org/10.1145/3097983.3098021

[4] Véronique Benzaken, Giuseppe Castagna, and Alain Frisch. 2003. CDuce: an XML-centric general-purpose language. In *International Conference on Functional Programming (ICFP)*. 51–63. https://doi.org/10.1145/944705.944711

[5] Gavin Bierman, Martín Abadi, and Mads Torgersen. 2014. Understanding TypeScript. In *European Conference for Object-Oriented Programming (ECOOP)*. 257–281. https://doi.org/10.1007/978-3-662-44202-9_11

[6] Pierre Bourhis, Juan L. Reutter, Fernando Suárez, and Domagoj Vrgoc. 2017. JSON: Data model, Query languages and Schema specification. In *Symposium on Principles of Database Systems (PODS)*. 123–135. https://doi.org/10.1145/3034786.3056120

[7] Eric Breck, Marty Zinkevich, Neoklis Polyzotis, Steven Whang, and Sudip Roy. 2019. Data Validation for Machine Learning. In *Conference on Systems and Machine Learning (SysML)*. https://www.sysml.cc/doc/2019/167.pdf

[8] Lars Buitinck, Gilles Louppe, Mathieu Blondel, Fabian Pedregosa, Andreas Mueller, Olivier Grisel, Vlad Niculae, Peter Prettenhofer, Alexandre Gramfort, Jaques Grobler, Robert Layton, Jake VanderPlas, Arnaud Joly, Brian Holt, and Gaël Varoquaux. 2013. API Design for Machine Learning Software: Experiences from the scikit-learn Project. https://arxiv.org/abs/1309.0238

[9] James Campbell et al. 2017. Great Expectations: Always know what to expect from your data. https://github.com/great-expectations/great_expectations

[10] Tiago Espinha, Andy Zaidman, and Hans-Gerhard Gross. 2015. Web API growing pains: Loosely coupled yet strongly tied. *Journal of Systems and Software (JSS)* 100 (2015), 27–43. https://doi.org/10.1016/j.jss.2014.10.014

[11] Francis Galiegue and Kris Zyp. 2013. JSON Schema draft 04. http://json-schema.org/draft-04/json-schema-validation.html

[12] Petter Haggholm. 2019. Check if a JSON schema is a subset of another. https://github.com/haggholm/is-json-schema-subset

[13] Martin Hirzel, Kiran Kate, Avraham Shinnar, Subhrajit Roy, and Parikshit Ram. 2019. Type-Driven Automated Learning with Lale. https://arxiv.org/abs/1906.03957

[14] John E. Hopcroft and Jeffrey D. Ullman. 2000. *Introduction to Automata Theory, Languages and Computation, Second Edition*. Addison-Wesley.

[15] ECMA International. 2017. ECMA-404 Standard: The JSON Data Interchange Syntax. http://www.ecma-international.org/publications/files/ECMA-ST/ECMA-404.pdf

[16] Jun Li, Yingfei Xiong, Xuanzhe Liu, and Lu Zhang. 2013. How Does Web Service API Evolution Affect Clients?. In *International Conference on Web Services (ICWS)*. 300–307. https://doi.org/10.1109/ICWS.2013.48

[17] Sam Newman. 2015. *Building Microservices: Designing Fine Grained Systems*. O'Reilly.

[18] David L. Parnas, John E. Shore, and David Weiss. 1976. Abstract types defined as classes of variables. In *Conference on Data: Abstraction, Definition and Structure*. 149–154. https://doi.org/10.1145/800237.807133

[19] Felipe Pezoa, Juan L. Reutter, Fernando Suarez, Martín Ugarte, and Domagoj Vrgoč. 2016. Foundations of JSON Schema. In *International Conference on World Wide Web (WWW)*. 263–273. https://doi.org/10.1145/2872427.2883029

[20] Martin Pilat, Tomas Kren, and Roman Neruda. 2016. Asynchronous Evolution of Data Mining Workflow Schemes by Strongly Typed Genetic Programming. In *International Conference on Tools with Artificial Intelligence (ICTAI)*. 577–584. https://doi.org/10.1109/ICTAI.2016.0094

[21] Carlos Rodríguez, Marcos Baez, Florian Daniel, Fabio Casati, Juan Carlos Trabucco, Luigi Canali, and Gianraffaele Percannella. 2016. REST APIs: A Large-Scale Analysis of Compliance with Principles and Best Practices. In *International Conference on Web Engineering (ICWE)*. 21–39. https://doi.org/10.1007/978-3-319-38791-8_2

[22] Sebastian Schelter, Dustin Lange, Philipp Schmidt, Meltem Celikel, Felix Biessmann, and Andreas Grafberger. 2018. Automating Large-scale Data Quality Verification. In *Conference on Very Large Data Bases (VLDB)*. 1781–1794. https://doi.org/10.14778/3229863.3229867

[23] D. Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips, Dietmar Ebner, Vinay Chaudhary, Michael Young, Jean-François Crespo, and Dan Dennison. 2015. Hidden Technical Debt in Machine Learning Systems. In *Conference on Neural Information Processing Systems (NIPS)*. 2503–2511. http://papers.nips.cc/paper/5656-hidden-technical-debt-in-machine-learning-systems

[24] Micah J. Smith, Carles Sala, James Max Kanter, and Kalyan Veeramachaneni. 2019. The Machine Learning Bazaar: Harnessing the ML Ecosystem for Effective System Development. https://arxiv.org/abs/1905.08942

[25] Akihiko Tozawa and Masami Hagiya. 2003. XML Schema Containment Checking based on Semi-implicit Techniques. In *International Conference on Implementation and Application of Automata (CIAA)*. 213–225. https://doi.org/10.1007/3-540-45089-0_20

[26] Michael M. Vitousek, Andrew M. Kent, Jeremy G. Siek, and Jim Baker. 2014. Design and Evaluation of Gradual Typing for Python. In *Dynamic Languages Symposium (DLS)*. 45–56. http://doi.acm.org/10.1145/2661088.2661101

[27] Kris Zyp. 2009. JSON Schema. http://json-schema.org/