

# Search-Based Fuzzing For RESTful APIs That Use MongoDB

Hernan Ghianni  
University of Buenos Aires  
Buenos Aires, Argentina

Juan P. Galeotti  
University of Buenos Aires; CONICET  
Buenos Aires, Argentina  
Kristiania University of Applied Sciences  
Oslo, Norway

Man Zhang\*  
Beihang University  
Beijing, China

Andrea Arcuri  
Kristiania University of Applied Sciences  
Oslo, Norway  
Oslo Metropolitan University  
Oslo, Norway

## Abstract

In RESTful APIs, interactions with a database are a common and crucial aspect. When generating white-box tests, it is essential to consider the database's state (i.e., the data contained in the database) to achieve higher code coverage and uncover more hidden faults. This article presents novel techniques to enhance search-based software test generation for RESTful APIs interacting with NoSQL databases. Specifically, we target the popular MongoDB database by dynamically analyzing its state—via automated code instrumentation—during test generation, with the goal of producing non-empty results for the queries executed by the generated tests. Additionally, to achieve better results, our novel approach also allows inserting automatically generated NoSQL data directly from test cases. This is particularly beneficial when testing read-only microservices, or when it is challenging or time-consuming to generate the correct sequence of HTTP calls to populate the NoSQL database. Our novel techniques are implemented as an extension of EvoMaster, the only open-source tool for white-box fuzzing RESTful APIs. Experiments conducted on six RESTful APIs demonstrated significant improvements in code coverage, with increases of up to 18% compared to existing white-box approaches. To better highlight the improvements of our novel techniques, comparisons are also carried out with four state-of-the-art black-box fuzzers.

## CCS Concepts

• **Software and its engineering** → **Software verification and validation**; Search-based software engineering.

## Keywords

REST API, Fuzzing, Search-based Software Testing, NoSQL Database, Test Case Generation

\*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

Conference'17, Washington, DC, USA

© 2025 ACM.

ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## ACM Reference Format:

Hernan Ghianni, Man Zhang, Juan P. Galeotti, and Andrea Arcuri. 2025. Search-Based Fuzzing For RESTful APIs That Use MongoDB. In . ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 1 Introduction

Nowadays, microservices play a crucial role in our everyday life. Many critical domains (e.g., banking, insurance, and healthcare) rely on microservices to provide their day-to-day operations [42, 52]. A microservice [52] is a specific software design pattern where a large application is divided into smaller, independent *services* that are orchestrated to achieve an overall behaviour.

A popular style to implement such microservices is the Representational State Transfer API (RESTful API or REST API, for short) [33]. Here, microservices achieve interoperability by handling interactions through HTTP(S) calls for seamless communication. Data, objects or any other meaningful entities are represented as *resources*. Each resource is unequivocally identified with a Uniform Resource Identifier (i.e., URIs), and can be manipulated using the semantics of an HTTP call, e.g., GET requests to fetch data, POST to create new data, PUT/PATCH to modify existing data and DELETE to remove data. Request inputs may be provided via path elements, query parameters, HTTP headers, or body payloads. Although there is no universal format for data transfer, one of the most common is JavaScript Object Notation (JSON).

In order to persist data changes to the application, REST APIs usually depend on database systems [37]. Structured Query Language (SQL)-based databases offer a relational model to store uniformized data into a table-based structure (rows and columns) [13]. In contrast, NoSQL databases [36] (originally, “Not Only SQL”) allow storing data into an unstructured, loose, format. This proved to allow applications both scalability and flexibility.

Software testing [53] is the preferable technique for assessing that a given RESTful API complies with its expected behaviour. However, manually writing test cases to detect mismatches between the intended and the expected output (i.e., defects or *bugs*) is a task both difficult and expensive. Automatic test case generators [31, 53] are tools that are specifically tailored for creating test cases on a given program. The most common automatic test case generators for RESTful APIs are black-box ones [38]. More specifically, the inner behaviour of the System Under Test (SUT) is unknown to the automatic test generator during the generation process. It can only rely on the outputs that are produced on each request by the RESTful API. In contrast, a white-box RESTful

test case generator might profit from runtime information from the SUT, aiming at generating test cases that execute previously non-covered regions of the program. At the time of the writing of this article, EvoMASTER [18, 25] is the only white-box test case generator for RESTful APIs, supporting languages that can run on the Java Virtual Machine (JVM). EvoMASTER is a search-based testing tool that profits from runtime information. To guide the test case generation, EvoMASTER instruments the SUT to retrieve both the number of covered lines and a computed *branch distance* [20] heuristic score for each reached branch, as well as computing other advanced white-box heuristics [24, 26].

Existing work of handling SQL databases [23] has presented effectiveness to improve coverage and fault detection for fuzzing RESTful APIs. In a nutshell, it introduces a *secondary* goal for each SQL query that returned no rows. The intuition behind this is that, whenever such SQL queries return a non empty result, it might trigger code paths that were previously unexecuted, thus enhancing code coverage and increasing the likelihood of covering undetected failures. Therefore, the SUT instrumentation computes a new heuristic score (similar to the branch distance) that approximates how close the database was from satisfying the criteria to return a non empty result when executing the SQL query. Additionally, to handle scenarios where data insertion in the database is not feasible (i.e., the RESTful API is *read-only*), or the sequence of interactions with the RESTful API is too complex to craft the appropriate data, EvoMASTER is enabled to synthesize and insert data *directly* into the relational database.

Inspired by the aforementioned approach, we present in this article a new technique to specifically handle NoSQL databases. This approach does not merely repeats the same solution presented in [23] in the context of NoSQL database, since these specific databases pose new research challenges, like their lack of a rigid schema throughout the execution of the SUT. In particular, we focus on MongoDB, which is the most popular NoSQL database, with more than 40 million downloads.<sup>1</sup>

In particular, this article provides the following research and engineering novel contributions:

- We provide novel NoSQL heuristics that can be integrated into search-based test generation tools for system testing.
- We provide novel techniques that automatically infer the expected format of mongoDB documents, enabling their direct generation and integration with the regular search for the SUT's inputs.
- To enable the replicability of our experiments, our extensions to EvoMASTER and a replication package are released as open-source software, available at

<https://github.com/WebFuzzing/fuzzing-nosql>.

## 2 Motivating Example

In this section, we will briefly show by means of an example how a RESTful API can interact with MongoDB, a widely used NoSQL database. The example is written in Java, using the Spring Boot framework [11], with the Spring Data MongoDB [12] extension for handling persistence. Figure 1 presents an excerpt of the class

```
@RestController
@RequestMapping(path = "/api/nosqlrest")
public class NoSQLRest {
    @Autowired private NoSQLRepository r;
    @RequestMapping(method = POST)
    public void post() {
        NoSQLEntity e = new NoSQLEntity();
        e.setX(42);
        e.setY('b');
        e.setZ(1.0);
        r.save(e);
    }
    @RequestMapping(path =("/{x}/{y}/{z}", method=GET,
        produces=APPLICATION_JSON_VALUE)
    public ResponseEntity get(
        @PathVariable("x") int x,
        @PathVariable("y") char y,
        @PathVariable("z") double z) {
        List<NoSQLEntity> l = r.findByXAndYAndZ(x+42,
            y+1,z/3.0);
        if (!l.isEmpty()) {
            // isNotEmpty branch
            return ResponseEntity.status(OK).build();
        } else {
            // isEmpty branch
            return ResponseEntity.status(NOT_FOUND).build();
        }
    }
}

public interface NoSQLRepository
    extends MongoRepository<NoSQLEntity, String> {
    List<NoSQLEntity> findByXAndYAndZ(int x,
        char y,
        double z);
}
```

**Figure 1: Example of RESTful API in Java SpringBoot, with two endpoints (a GET and a POST) accessing a database via Spring Data MongoDB.**

```
{
  "_id": {"$oid": "e3...ac"},
  "x": 42,
  "y": "b",
  "z": 1.0
}
```

**Figure 2: A JSON representation of the object that is stored in the database. The `_id` field stands for the object identifier, a 24-character hexadecimal string (12 bytes), which MongoDB uses as the unique identifier for documents within a single collection.**

NoSQLRest and interface NoSQLRepository. We do not show all the classes and needed configuration files to run this API, but just those we consider necessary to understand the presented example.

In this example, we have only two operations POST and GET on two different endpoints: `/api/nosqlrest/` (for the POST operation) and `/api/nosqlrest/{x}/{y}/{z}` (for the GET operation).

The POST operation saves a NoSQLEntity into the MongoDB repository. In order to do so, the Spring MongoDB library translates the NoSQLEntity instance into the Binary JSON (BSON) format, which is the native data format of MongoDB. This representation will contain the values of the translated instance for fields x, y and z (namely, 42, "b" and 1.0). Figure 2 presents a human readable version (in JSON format) of the BSON that is stored in the database.

<sup>1</sup><https://www.mongodb.com/resources/basics/databases/nosql-explained/most-popular-nosql-database>, accessed 11th September 2024.

Observe that a field `_id` is added with an automatically generated identifier to identify unequivocally the element in the MongoDB database. By default, the Spring Data MongoDB will create and store all the JSON objects that are saved in the repository into a *collection* called “*nosqlentities*”. We could modify the expected collection name that is created and managed by MongoDB using the `@Document` annotation in the `NoSQLEntity` class declaration.

The GET operation retrieves a subset of instances of `NoSQLEntity` that are stored in the underlying “*nosqlentities*” collection in the MongoDB database. Given the `x`, `y` and `z` path parameters in the endpoint, the Spring MongoDB library will select all the JSON objects that satisfy that field `x` is equal to path parameter `x` plus 42, `y` is equal to path parameter `y` plus one, and path parameter `z` is equal to `z` divided by 3.0, constructing a *filter* (i.e., a MongoDB query) to match documents where all three fields are equal to those computed values. By means of reflection, new `NoSQLEntity` instances will be built and fields will be modified accordingly to the data fetched from the MongoDB underlying collection.

Notice that the code that actually performs the search on the MongoDB collection is not shown. The `@Autowired` annotation signals the Spring framework that it must instantiate a *singleton* repository in field repository in order to save and load `NoSQLEntity` instances from the MongoDB database. Given the naming convention that complies the method declared in the `NoSQLRepository`, Spring discovers that the particular method is intended to fetch data with given values for fields `x`, `y` and `z`. Internally, Spring will produce the filter  $F = \{“x” : 82, “y” : “b”, “z” : -1.0\}$  by using the values passed as arguments to that method. Filters could be directly written as JSON objects, or they could be constructed programmatically using the `Filter` API provided by MongoDB.

In this example, handling the scenario where the returned list is non-empty (i.e., the `isNotEmpty` branch) presents a recurring challenge for search-based test generators. These techniques rely on heuristic methods to determine whether a test case is “*closer*” to reaching a specific execution path. A common heuristic such as branch distance assigns a fixed non-zero value (representing covering the `isEmpty` branch) until the `isNotEmpty` branch is covered, at which point a zero value is returned. Essentially, the search remains largely random until it happens upon a test case that triggers the desired execution path. This recurrent challenge is known as the *flag problem* [30, 39]. In order to overcome this limitation, the search-based test generator requires satisfying the filter that is being created and applied on the *nosqlentities* collection in the MongoDB database. In other words, it is required to find values for the path parameters such that filter produced when invoking `findByXAndYAndZ()` is satisfied.

In this article, we introduce a novel technique to provide gradient when retrieving MongoDB data. This is done by computing an heuristic score approximating how close are the documents stored in the collection to satisfying a desired filter. The technique monitors all filters submitted to each collection in the underlying database, recording those that upon execution did not return any element. After test execution, for each such filter, the heuristic score is computed. By extending `EvoMASTER` with the presented technique, we are able to generate test cases as the one shown in Figure 3. As Java integers have  $2^{32}$  possible values, Java `char`(acter) values have a domain of  $2^{16}$  values and Java floating-point double

```
@Test
public void test1 () throws Exception {
    given().accept ( "**/*" )
        .post(baseUrlOfSut + "/api/nosqlrest")
        .then()
        .statusCode(200);

    given().accept ( "**/*" )
        .get(baseUrlOfSut + "/api/nosqlrest/0/a/3.0")
        .then()
        .statusCode(200);
}
```

**Figure 3: A test case generated by EvoMaster that satisfies the true branch in Figure 1.**

```
@Test
public void test1 () throws Exception {
    List<MongoInsertionDto> insertions = mongo()
        .insertInto("mymongodb", "nosqlentities")
        .d("{\"x\": 53, \"y\": \"f\", \"z\": 11.0 }")
        .dtos();

    MongoInsertionResultsDto insertionsresult = controller
        .execInsertionsIntoMongoDatabase(insertions);

    given().accept ( "**/*" )
        .get(baseUrlOfSut + "/api/nosqlrest/11/e/33.0")
        .then()
        .statusCode(200);
}
```

**Figure 4: A test case generated by `EvoMASTER` that satisfies the true branch in Figure 1 without using the POST operation.**

precision values are  $2^{64}$ , finding a specific  $\langle x, y, z \rangle$  has a uniform probability of  $\frac{1}{2^{32} \times 2^{16} \times 2^{64}} = \frac{1}{2^{112}}$ . Note that in this example, a technique like seeding [56], which could potentially enhance the search space, proves ineffective. This is because the path parameters are modified before being used as MongoDB filter values.

Additionally, we introduce a technique to handle scenarios where certain operations, such as the POST operation, are unavailable (e.g., when the database is read-only). Unlike relational databases, NoSQL databases like MongoDB allow the storage of objects with varying attributes in the same collection. For instance, if our technique inserts a document into the *nosqlentities* collection that does not conform to the format in Figure 2, all subsequent GET operations will throw an exception signaled in the Spring MongoDB library. This issue makes dealing with NoSQL databases much harder than dealing with SQL ones, as those have precisely defined schemas. Due to the lack of a fixed schema, our `EvoMASTER` extension dynamically infers an expected type for the elements in a MongoDB collection at runtime. Figure 4 shows a test case generated by `EvoMASTER` using our extended Domain-Specific Language (DSL) for inserting documents into MongoDB collections. This test first adds a JSON document with fields `x=53`, `y="f"`, and `z=11.0` to the *nosqlentities* collection. Subsequently, the `findByXAndYAndZ()` method is invoked with the same values (since  $11 + 42 = 53$ ,  $"f" + 1 = "e"$ , and  $33.0 / 3.0 = 11.0$ ), leading to covering the true branch in the shown example.

### 3 Background

#### 3.1 NoSQL Databases

NoSQL databases [36] refer to a collection of non-relational databases designed to handle large volumes of diverse data, with high flexibility and scalability, especially for distributed systems. Unlike traditional relational databases (i.e., SQL databases), NoSQL databases do not require a fixed schema and allow data to be stored in a variety of formats, such as key-value pairs, document-based structures, wide-column stores, or graph-based representations. This flexibility is advantageous for systems where the data structure is dynamic or where schema evolution is frequent.

In NoSQL databases, the traditional concepts of tables, rows, and columns do not apply. Instead, data can be stored in more flexible structures. MongoDB [6] is a document-based database where data is stored as documents in collections. A document  $D_i$  is encoded in Binary JSON (BSON) and can contain a varied number of key-value pairs, allowing for schema flexibility. These documents may contain nested structures and arrays, enabling complex data relationships to be represented within a single document. As an example, a document  $d$  can contain:

```
d = {"name": "Foo", "year": 2025,
     "location": {"city": "Oslo", "country": "Norway"}}
```

NoSQL databases support a range of query and data manipulation mechanisms depending on the data model. As an example, in MongoDB, the query language allows for complex filtering with commands like `find()` and aggregation pipelines. These queries can include conditions, projections, and sorting, and aggregation functions for performing computations on collections of documents, such as summing values or averaging fields.

#### 3.2 Search-Based Software Testing

Over the years, search-based software engineering has emerged as a highly effective approach to tackle a broad spectrum of software engineering challenges [40], particularly in the realm of software testing [14]. Notable tools such as EvoSuite [34], for unit testing, and Sapienz [48], for Android testing, exemplify the success of these techniques.

Software testing can be casted as an optimization problem, maximizing both code coverage and fault detection in the generated test suites. Once a fitness function is established for the testing problem, a search algorithm can explore the solution space, which in this case represents all possible test cases.

Several types of search algorithms exist, with Genetic Algorithms being among the most popular. Specialized search algorithms (such as Whole Test Suite [35], MOSA [54] and MIO [19]) have been developed for the specific problem of generating test suites (in contrast to solely generating test cases).

One common heuristic for guiding the search process is the branch distance, which helps find input data that satisfies the constraints in logical predicates of branches [45]. The branch distance for any execution of a predicate is computed using a set of recursively defined rules, as shown in Table 1. For instance, for the predicate  $x \geq 100$  where  $x = 42$ , the branch distance to the true branch is  $100 - 42 + K$ , with  $K > 0$ . In practice, calculating the branch distance requires instrumenting each predicate to track the

**Table 1: Branch distance for relational operands.  $a$  and  $b$  are numeric values, whereas  $x$  and  $y$  are boolean predicates/expressions.**

Relational operation	Branch distance $\rho$
$a > b$	$\rho$ : if $a > b$ then 0 else $(b - a)$
$a \geq b$	$\rho$ : if $a \geq b$ then 0 else $(b - a) + K$
$a < b$	$\rho$ : if $a < b$ then 0 else $(a - b) + K$
$a \leq b$	$\rho$ : if $a \leq b$ then 0 else $(a - b)$
$a = b$	$\rho$ : if $a = b$ then 0 else $abs(a - b)$
$x    y$	$\rho$ : $min(\rho(x), \rho(y))$
$x \& \& y$	$\rho$ : $\rho(x) + \rho(y)$

distance during execution. For more details on branch distance computation, readers are referred to [45].

#### 3.3 EvoMaster

EvoMASTER [18, 25] is an open-source tool designed to generate system-level test cases for RESTful APIs [20]. It leverages evolutionary algorithms, such as MIO [19] and MOSA [54]. EvoMASTER automatically instruments the SUT at startup to collect runtime heuristics, such as the aforementioned branch distance. It generates fully self-contained test cases in several output formats, such as Java, Kotlin, Javascript or Python. Each generated test consists of a series of HTTP calls to the SUT, leveraging the widely used RestAssured library [10]. EvoMASTER is open-source [1].

### 4 Related Work

In the last few years, there has been an increasing interest in the research community on automated testing of REST APIs [38]. In the literature, several tools have been proposed, like for example (in alphabetic order): ARAT-RL [43], bBoxrt [46], EvoMASTER [21], Mor-est [47], RestCT [58], Restest [51], RESTLER [29], RestTestGen [57], and SCHEMATHESIS [41]. Different studies have been carried out to compare these tools, like for example as done in two recent empirical studies for existing fuzzers for REST APIs [44, 59]. Apart from EvoMASTER that supports white-box testing, all other tools only support black-box testing.

As far as we know, what presented in this paper is the first approach in the literature aimed at analyzing SUT's interactions with MongoDB databases, and use such information to improve how the fuzzing process is carried out. The closest work in the literature to what we present in this paper is perhaps what presented in [23], which is aimed at SQL databases. However, there are major differences and research challenges between dealing with SQL and NoSQL databases. In particular, the lack of formally defined schemas in MongoDB significantly complicates the runtime analyses.

### 5 NoSQL Heuristics

When testing a RESTful API, multiple criteria may need to be optimized. For instance, we might aim to maximize statement coverage, branch coverage, mutation testing scores, fault detection, or even a combination of these criteria [55]. These metrics can be evaluated by executing test cases in an integrated development environment (IDE) or on a Continuous Integration server. To achieve these objectives, a common approach in search-based software testing is to use the branch distance [45], which provides a gradient for solving

```

Query<NoSQLEntity> q = datastore
.createQuery(NoSQLEntity.class)
.filter("x =", x+42)
.filter("y =", y+1)
.filter("z =", z/3.0);
List<NoSQLEntity> l = q.find().toList();

```

**Figure 5: Snippet of the RESTful API using the Object-Document Mapper Morphia instead of Spring Data MongoDB**

constraints within the SUT's branch statements. This requires instrumenting the SUT's code to collect branch distances. However, such heuristics become ineffective when constraints are translated into filters that are submitted to the NoSQL database within a third-party library, such as the Spring Data MongoDB library. Given the example shown in Figure 1, one might automatically construct such constraints by analyzing both method signatures and names of each found class extending the `MongoRepository` class. By inspecting the SUT, we could conclude that, in order to return a non-empty result when calling `findByXAndYAndZ(int, char, double)` we need to match the values of the arguments to the fields `x`, `y` and `z` in the MongoDB collection. Nevertheless, this solution is highly dependant to how (a given version of) Spring Data MongoDB behaves. Not surprisingly, other Java/Kotlin Object-Document Mappers offer different mechanisms to query stored data. For example, Morphia [8] requires to build a special `Query` object with the filters being passed as pairs of string-values. Figure 5 shows the equivalent implementation of our motivating example but using the Morphia API instead of Spring Data MongoDB.

In this article, we propose an approach that is independent of the Object-Document Mappers (ODMs) or third-party library hiding the querying to the MongoDB database. Our approach builds on [23], extending and adapting it to tackle the unique challenges associated with handling NoSQL databases.

- (1) We monitor all commands issued to the NoSQL database at the MongoDB driver level, ensuring visibility into every database interaction.
- (2) Whenever the SUT performs on a MongoDB collection a `find()` operation that returns no documents, we calculate a heuristic to estimate how close it was to retrieving non-empty data. This heuristic is based on the evaluation of the boolean predicates in the *filter* passed to the `find()` operation.
- (3) These heuristic values, generated for each call to a `find()` in a MongoDB collection, are added as secondary objectives for optimization.

The rest of this section presents a detailed discussion of these three points.

## 5.1 Monitoring NoSQL commands

Object-Document Mappers (ODMs) are frameworks that simplify the interaction between applications and NoSQL databases, such as MongoDB. Popular ODMs for MongoDB include Spring Data MongoDB [12], Morphia [8], KMongo [4] and Micronaut Data MongoDB [5]. These libraries abstract the complexities of MongoDB's query language and operations, providing a high-level interface for developers. They automatically convert user-defined objects into MongoDB documents for storage and retrieval, facilitating

interaction between application objects and the database. This abstraction significantly simplifies common tasks, such as saving and querying data. All these libraries rely on the MongoDB Java drivers [7], which ultimately provides a low-level API to manage database connections and execute queries. Specifically, when these frameworks retrieve documents from a MongoDB collection, they invoke one of the many overloaded variants of the `MongoCollection.find()` method. Therefore, by monitoring calls to this family of methods, we can effectively capture all interactions between the application and the MongoDB database. Following the technique introduced in [26], we monitor each call to `MongoCollection.find()` (whether in the SUT or any third-party library) recording the applied filter.

## 5.2 NoSQL Distance

Given a document  $d$  and a filter  $F$ , we compute the heuristic score  $H_d(F)$  by recursively applying the definition presented in Table 2. Let us explain this definition by means of an example. Suppose we have the following filter  $F = \{ "x" : \{ "$eq" : 17 \} \}$ . This filter states that field `x` equals 17. For presentation purposes, let us copy the document  $d$  taken from Figure 2:

$$d = \{ \_id : \text{<auto-generated-Id>}, "x" : 42, "y" : "b", "z" : 1.0 \}$$

Based on the definition of the heuristic score  $H_d(F)$  for document  $d$  and a filter  $F$ , we need to compute the score  $H_d("x", \{ "$eq" : 17 \})$ . Table 3 provides the definition of the heuristic score for a condition (in this case,  $C = \{ "$eq" : 17 \}$ ) applied to a specific field (here,  $f = "x"$ ). Following the definition of  $H_d("f", C)$ , we deduce that the branch distance  $\rho(d["x"] = 17)$  must be calculated. Therefore, according to the branch distance definitions in Table 1, we find that  $H_d(F) = \rho(42 = 17) = \text{abs}(42 - 17) = 25$ .

Now, let us consider a more complex filter  $F'$ :

$$F' = \{ \text{"$and"} : [ \{ "x" : \{ "$eq" : 17 \} \}, \{ "y" : \{ "$eq" : "c" \} \} ] \}$$

This filter requires not only that the field `x` equals 17, but also that the field `y` equals "c". To compute  $H_d(F')$ , we recursively calculate  $H_d(\{ "x" : \{ "$eq" : 17 \} \})$  and  $H_d(\{ "y" : \{ "$eq" : "c" \} \})$ . Instead of simply summing these scores, we apply the normalization function  $v(x) = \frac{x}{x+1}$  as introduced in [16]. This normalization function ensures that each value is mapped to the range  $[0, 1]$ , preventing any one score from disproportionately dominating the total. Without this normalization, large scores could skew the relative contributions of each value.

As shown in Tables 2 and 3, our prototype already supports most filter operators defined by MongoDB [9], including all comparison, logical, element, and array operators. For string comparisons, we use the Levenshtein distance [15].

Each time we detect a call to a `MongoCollection.find()` method, our instrumentation saves the filter  $F$  and the name of the collection  $c$  on which the filter  $F$  was applied. Once test execution has finished, we execute a new `find` command for each collected filter on the corresponding target collection. If the target collection is not empty, and the filtering retrieves no documents, we compute a heuristic score  $H_d(F)$  on each document  $d$  on collection  $c$ . Finally, the NoSQL distance that approximates how close are the existing  $d \in c$  documents to returning any document on filter  $F$  is computed as  $H_c(F) = \min_{d \in c} \{ H_d(F) \}$ .



**Table 2: NoSQL heuristic scores for filter  $F$  on document  $d$** 

Filter $F$	Heuristic score $H_d(F)$
$\{“f” : cond\}$	$H_d(“f”, cond)$
$\{“$and” : [cond_1, cond_2, \dots]\}$	$\sum_{cond_i} \{v(H_d(cond_i))\}$
$\{“$or” : [cond_1, cond_2, \dots]\}$	$\min_{cond_i} \{H_d(cond_i)\}$
$\{“$nor” : [cond_1, cond_2, \dots]\}$	$\sum_{cond_i} \{v(H_d(\neg cond_i))\}$

**Table 3: NoSQL heuristic scores for field  $“f”$  and condition  $C$  on document  $d$ .**

Condition $C$	Heuristic score $H_d(“f”, C)$
$\{“$eq” : v\}$	$\rho(d[“f”] = v)$
$\{“$ne” : v\}$	<b>if</b> $d[“f”] \neq v$ <b>then</b> 0 <b>else</b> $K$
$\{“$gt” : n\}$	$\rho(d[“f”] > v)$
$\{“$gte” : n\}$	$\rho(d[“f”] \geq v)$
$\{“$lt” : n\}$	$\rho(d[“f”] < v)$
$\{“$lte” : n\}$	$\rho(d[“f”] \leq v)$
$\{“$in” : [v_1, v_2, \dots]\}$	$\min_{v_i} \{\rho(d[“f”] = v_i)\}$
$\{“$nin” : [v_1, v_2, \dots]\}$	<b>if</b> $\nexists v_i d[“f”] = v_i$ <b>then</b> 0 <b>else</b> $K$
$\{“$mod” : [div, rem]\}$	$\rho(d[“f”] \bmod div = rem)$
$\{“$exists” : true\}$	$\min_{f_i \in fields(d)} \{\rho(“f” = “f_i”)\}$
$\{“$exists” : false\}$	<b>if</b> $“f” \notin fields(d)$ <b>then</b> 0 <b>else</b> $K$
$\{“$size” : n\}$	$\rho(array\_length(d[“f”]) = n)$
$\{“$type” : type\}$	$\rho(“type(f)” = “type”)$
$\{“$all” : [v_1, v_2, \dots]\}$	$\sum_{v_i} \{\min_{w_j \in d[“f”]} \{v(\rho(w_j = v_i))\}\}$
$\{“$not” : cond\}$	$H_d(“f”, \neg cond)$

It is worth mentioning that, between the moment in which the filter was observed and the instant in which the NoSQL distance is computed, the database state might have changed (i.e., documents being added, deleted, or updated). Despite a potentially less accurate gradient, this approach balances simplicity and performance.

### 5.3 NoSQL Fitness Function

Once we have computed a heuristic NoSQL distance for each executed `find()` command that returns no documents, the next step is to use these distances to enhance the search-based test generation. In EvoMASTER, for each target goal  $g$  (e.g., a branch, a statement or a HTTP status code per endpoint) there is a heuristic value  $h(t, g) \in [0, 1]$ , where  $h(t, g) = 0$  means the target goal  $g$  is not reached when executing test case  $t$ , and  $h(t, g) = 1$  means the target goal  $g$  is covered when executing test case  $t$ . Intermediate values guide the search toward covering these targets.

Besides the primary heuristic  $h(t, g)$ , EvoMASTER also allows having a secondary heuristic value  $s(t, g)$  for each target goal  $g$ . When two individuals  $t_1$  and  $t_2$  have the same primary  $h(t_1, g) = h(t_2, g)$  value, then the fitness function can break the tie by inspecting the secondary values  $s(t_1, g)$  and  $s(t_2, g)$  in order to decide which one of the two individuals is fitter for evolution.

In order to compute  $s(t, g)$ , EvoMASTER keeps track of the HTTP call  $a_i \in t$  which led to the best value for  $h(t, g)$ , collecting all secondary SQL distances from  $a_i \in t$  [23]. We extended EvoMASTER to also include each computed  $H_c(F)$  NoSQL distance as a secondary objective. Let us call this group of database distances  $D_i$ .

Once EvoMASTER collects the set  $D_i$  of secondary database distances,  $s(t, g)$  can be computed. Let  $k_i^t$  be the number of database (either SQL or NoSQL) commands executed by the test  $t$  at HTTP call  $i$ , where  $k_i^t \geq |D_i|$ , then,  $s(t, g)$  is computed as an average but only if the number  $k_i^t$  of database commands is the same. If when

comparing a test  $t_1$  with a test  $t_2$  we have  $k_i^{t_1} \neq k_j^{t_2}$  (where  $i$  and  $j$  are independent), then we rather prefer and select the one with larger  $k$ , i.e., doing more database operations.

By using  $s(t, g)$  as a secondary objective to minimize, the aim is to reward test cases that get closer to having `find()` operations returning non-empty data sets. However, this is just a secondary objective, as we cannot be sure that having data returned by the database will necessarily have a beneficial impact on achieving coverage of  $g$ .

## 6 NoSQL Data Generation

To achieve a specific testing goal  $g$ , such as covering a branch in the SUT, the database may need to be in a specific state. For instance, the goal might only be covered if a particular `find()` command returns the desired data. Consequently, reaching this goal may require performing a series of HTTP calls on the SUT to prepare the database appropriately. As discussed in [23] for SQL databases, while this approach can also theoretically work for NoSQL databases, it presents several potential challenges:

- (1) Preparing the database might involve executing several HTTP calls on the SUT, which not only complicates the search process but also reduces the clarity of the test.
- (2) In some cases, it may even be infeasible to insert the required data into the database using the available SUT endpoints. This limitation arises when certain parts of the database are effectively “read-only” from the SUT’s perspective, particularly when multiple systems access the same database.
- (3) In contrast to relational databases, NoSQL databases such as MongoDB do not have a precisely defined schemas for each collection. However, the SUT might expect certain attributes in a given collection to correctly transform documents into SUT classes. Therefore, it is required to dynamically infer the expected type of the elements in a MongoDB collection during runtime.

To address these challenges, EvoMASTER includes the capability to write data directly into relational databases [23]. For a test case  $t$  involving multiple HTTP calls to the SUT, EvoMASTER inserts specialized *initialization calls* at the start to configure the relational database state before interacting with the SUT. These initialization calls specify not only the target table but also the values to be inserted into each column. The mutation of these values follows the same rules applied to mutating values in regular HTTP calls within the test case  $t$ .

In order to enable data insertion, we extend EvoMASTER by adding a new type of specialized initialization call for inserting a *document* in a target MongoDB collection  $c$ . As with SQL databases, EvoMASTER will probabilistically add such MongoDB initialization calls whenever a `find()` command fails to return a result in collection  $c$ . Since NoSQL databases allow the storage of documents with varying attributes within the same collection, the insertion of a document will be generally valid, as no schema is defined. However, it is often the case that it is expected a specific data format for the documents stored in a given collection. For example, if a given field in the document has an integer value instead of a string, the mapping from the document to the expected result class will fail when the SUT reads and parses such data. Such failure to comply

**Table 4: Descriptive statistics of the employed REST APIs.** For each SUT, #Endpoints represents the number of declared HTTP endpoints in those APIs, #SourceFiles represents the number of files (i.e., the number of public Java classes) composing such SUT, and #LOCs represents their total number of lines of code (LOC).

SUT	#Endpoints	#SourceFiles	#LOCs
<i>bibliothek</i>	8	33	2176
<i>genome-nexus</i>	23	405	30004
<i>gestaohospital-rest</i>	20	33	3506
<i>ocvn-rest</i>	258	526	45521
<i>reservations-api</i>	7	39	1853
<i>session-service</i>	8	15	1471
<i>Total</i>	324	1051	84531

to an implicit format might lead to missing significant executions in the SUT.

In order to infer the expected data format, whenever a call to `MongoCollection.find()` is detected, we also collect the expected result class. With this information, by means of reflection, our instrumentation derives all the result class' fields and their types. Therefore, whenever EvoMASTER inserts a document in a given collection MongoDB *c*, it probabilistically instantiates a document satisfying the inferred data format.

## 7 Empirical Study

In this paper, we carried out an empirical study to answer the following two research questions:

**RQ1:** How do our novel techniques improve upon existing white-box fuzzing for RESTful APIs?

**RQ2:** How do our results compare to the ones of existing black-box fuzzers?

### 7.1 Artifact Selection

To answer our research questions, we needed a set of REST APIs for experimentation that use a MongoDB database. Our novel techniques are implemented for the JVM (e.g., based on bytecode manipulation). However, this is just a technical detail, as they could be reimplemented for other programming languages as well (e.g., Go and Python). Although REST APIs are widely used in industry, especially for microservice architectures, these two constraints (i.e., use of MongoDB and running on the JVM) reduces the available pool of possible APIs that could be found on open-source repositories such as GitHub.

To address this issue, we use all the JVM REST APIs using MongoDB that are present in the latest version of the EMB corpus [2, 27]. EMB is a curated set of APIs used for experimentation in software testing research for Web APIs.

Table 4 shows some statistics of the selected six APIs used for these experiments. Full details can be found at [2], including source code and links to the original GitHub repositories those APIs were collected from. Note that the reported lines of code (LOCs) is only based on the business logic of these APIs. They do not count all the third-party libraries used in these APIs (which can be millions of

lines of code [27], e.g., for embedded HTTP servers and development frameworks such as Spring).

Using six APIs is less than what done in large tool comparisons such as in [59] (20 APIs) and in [44] (20 APIs). However, six APIs are more than what used when introducing RESTLER [29] (2 APIs) and EvoMASTER [17] (2 APIs), and same as Morest [47] (6 APIs). Still, they are less than what employed for evaluating ARAT-RL [43] (10 APIs) and SCHEMATHESIS [41] (16 APIs).

While tool comparisons of black-box fuzzers are not influenced by the programming language they are written in, or which type of database (if any) they use, our settings are necessarily more restricted (i.e., JVM and MongoDB). Still, our case study size is comparable to current research standards.

### 7.2 Experimental Settings

To answer our research questions, we carried out two different set of experiments on the six APIs discussed in Section 7.1.

In the first set of experiments on *white-box* testing, we compared our novel techniques with a default, base version of EvoMASTER. As far as we know, EvoMASTER is the only existing open-source white-box fuzzers for REST APIs, so no comparison with other tools can be done. In these comparisons, the two techniques are evaluated based on achieved code coverage and detected faults. To reduce bias, these metrics are computed with EvoMASTER's own reporting tooling. Each experiment was repeated 30 times to take into account the randomness of these algorithms. Each fuzzing session was left running for one hour. In total, this took  $2 \times 6 \times 30 \times 1h = 15$  days of computation effort.

In the second set of experiments, we ran four existing black-box fuzzers on these six APIs, using the same experiment settings (i.e., 1-hour fuzzing sessions, with 30 repetitions). These experiments took  $4 \times 6 \times 30 \times 1h = 30$  days of computation effort. The selected fuzzers are EvoMASTER [21], SCHEMATHESIS [41], RESTLER [29] and ARAT-RL [43].

Black-box EvoMASTER was chosen as it makes easier to compare with white-box testing, as being the same tool this reduces possible side-effects of implementation details impacting the comparisons. Moreover, it was the tool that gave best results in [44, 59]. With more than 2 000 stars on GitHub, RESTLER and SCHEMATHESIS seem to be the most popular tools among practitioners. Furthermore, SCHEMATHESIS has shown similar results as EvoMASTER in tool comparisons [59], making it arguably the best black-box fuzzer for REST APIs. Finally, ARAT-RL is a more recent tool (compared to the others), where in an empirical analysis [43] it was shown to give better results than EvoMASTER and RESTLER (as well as better than another fuzzer called Morest [47]). But, ARAT-RL was not compared with SCHEMATHESIS in [43].

Besides these four fuzzers, as discussed in Section 4, there are other REST API fuzzers, such as for example bBoxrt [46], Morest [47], Restest [51], RestCT [58], and RestTestGen [57]. However, arguably, based on existing empirical studies in the literature (e.g., [43, 44, 59]), the selected four fuzzers are a good representation of the current state-of-the-art in *black-box* REST API fuzzing.

For the experiments, we downloaded the latest released versions of these tools, as of the 6th of August 2024. This means we used EvoMASTER version 3.0.0, SCHEMATHESIS version 3.33.3 and RESTLER

version 9.2.4. For ARAT-RL, we did not use its latest published version 0.1 (which is old, from summer 2023), but rather from its latest commit (code 04e22bb). This is because the critical option of specifying for how long to run the fuzzing sessions (e.g., for one hour) was not present in the released version 0.1.

These tools were compared based on achieved code coverage. This was done by instrumenting the APIs with JaCoCo [3], and collecting code coverage results at the end of 1-hour fuzzing sessions. However, this approach has two potential issues.

The first issue is that we are not comparing fault detection. The reason is that each tool report detected faults in their own formats and ways which are not directly comparable. Looking at returned 500 HTTP status codes via a proxy (e.g., as done in [32]), or at stack-traces in the server-side error logs (as done in [43]) have their own set of limitations. Furthermore, they only show one type of faults that these tools can detect (e.g., crashes). But, there are many other kinds of faults that can be detected, e.g., based on schema validation, robustness and security testing [38]. Only focusing on one type of faults would be quite restrictive, and possibly biased. Looking at the total number of self-reported detected faults would be problematic as well, as tools that are equipped with more oracles would be unfairly advantaged. As such, as long as there is no standardized way among tools on how to report detected faults, focusing only on achieved code coverage is a reasonable compromise for comparisons among different tools.

The second issue is that this approach of measuring code coverage during fuzzing (which is the common approach done in all comparisons in the literature, as far as we know) ignores the generated tests. A black-box fuzzer could evaluate hundreds of thousands of HTTP calls within one hour. Generating test suites with hundreds of thousands of HTTP calls would be impractical. Therefore, tools use minimization strategies when outputting test suites, based for example on black-box coverage metrics [50]. However, whether such minimized test suites would still achieve the same level of code coverage is not guaranteed. Comparing tools based on generated tests is a major technical effort, as tests can be generated in different languages (e.g., Python and JUnit), which would then need to be possibly compiled (with all needed third-party libraries) before being executed. Furthermore, unless the size of the generated test suites is taken into account in the comparison metrics, tools that output lots and lots of test cases (which would not be useful for practitioners, as not manageable) could have an unfair advantage.

Our novel techniques, aimed at MongoDB support, are for white-box testing. To compare them with existing black-box fuzzers, we need to use the same metrics, collected with the same tools. Unfortunately, it would not be recommended to run the APIs with JaCoCo while white-box fuzzing them, as the bytecode manipulations of EvoMASTER's instrumentation could conflict with the one of JaCoCo, possibly leading to messed up results. For the sake of these experiments, we run the tests generated instrumented with JaCoCo *after* the one hour fuzzing sessions. Note that EvoMASTER white-box instrumentation is only needed during the fuzzing (e.g., to collect and compute search-based heuristics). It is not needed for when the generated tests are run, so there is no conflict of any kind with JaCoCo. As EvoMASTER's minimization algorithms take into account the achieved code coverage (as that is computed for each test execution), there is no risk of losing code coverage in

the generated test suites (unless there is a fault in EvoMASTER itself). However, how EvoMASTER measures code coverage and how JaCoCo does that is not necessarily the same.

These two sets of experiments took together  $15 + 30 = 45$  days of computation effort. All experiments were run on a HP Z6 G4 Workstation with Intel(R) Xeon(R) Gold 6240R CPU @2.40GHz 2.39GHz, 192 GB RAM, with 64-bit Windows 11.

### 7.3 Experimental Results

Table 5 shows the results of our first set of experiments, where the default version of EvoMASTER (named *Base* here) is compared with our extension for MongoDB support (named *Mongo*). The two configurations are compared following standard guidelines in software engineering research [22]. In particular, pair comparisons per SUT are analyzed with Wilcoxon-Mann-Whitney U-tests, with  $\hat{A}_{12}$  standardized effect-sizes.

Regarding achieved code coverage, in all cases but one (i.e., ocnv) our novel techniques provided better results that are statistically significant, with strong effect-sizes (average effect-size  $\hat{A}_{12} = 0.91$ ). On four out of six APIs, the effect-size was the maximum  $\hat{A}_{12} = 1.00$ . This means that, on these four APIs, every single run out of 30 for *Mongo* gave better results than any run for *Base*.

On the one hand, on session-service the achieved code average increased on average by 18%. On the other hand, the case of ocnv is rather particular. It seems our novel techniques provide some improvement, but 30 runs are not enough to provide enough confidence to claim it with statistical significance. Looking at the source code of ocnv, for most of this API's endpoints, data seems that is mainly written into the MongoDB database, and not read. This could explain why our novel techniques do not provide much benefits for fuzzing an API like this.

Regarding fault detection, based on the results shown in Table 5, our novel techniques can reliably find new faults in three of these APIs (average  $\hat{A}_{12} = 0.70$ ). However, whether faults are found depends entirely on whether such faults exist in the first place. These six APIs are taken as they are, with no manually injected faults. Our novel techniques enable to better explore the execution paths of these APIs (see previously discussed improvements on achieved code coverage), but many faults in REST APIs are typically located in the first validation layer [49]. Therefore, the number of new detectable faults is expected to be lower.

Our novel techniques have a computational cost. Such cost makes each test case evaluation longer. Given the same amount of time, this means that less test cases can be evaluated. If the novel techniques provide no benefit to the search process, then they would have a detrimental effect on the final results. In Table 5 we can see that on average there is a decrease of around 18% in the number of evaluated HTTP calls (note that each test case is composed of one or more HTTP calls towards the tested API). However, this is not as simple as stating that this is due to computational overhead, alone. The reason is that not all test case executions take the same amount of time. A test case that returns immediately due to input validation (e.g., a 400 user-error code) will be faster than a successful call that executes large parts of the API's source code, including calls to databases and external services. An example of this is session-services, where only 2/3 of test cases could



**Table 5: Performance comparisons between the *Base* and *Mongo* configurations, in terms of average (i.e., arithmetic mean) line coverage and average number of detected faults. Results of statistical tests are reported, including p-values and  $\hat{A}_{12}$  effect sizes. For p-values lower than the threshold  $\alpha = 0.05$ , the effect sizes  $\hat{A}_{12}$  are shown in bold. We also report the average number of HTTP calls done during the search, and their scaled difference compared to the *Base* configuration.**

SUT	Line Coverage %				# Detected Faults				# HTTP Calls		
	Base	Mongo	$\hat{A}_{12}$	p-value	Base	Mongo	$\hat{A}_{12}$	p-value	Base	Mongo	Ratio
<i>bibliothek</i>	27.0	31.2	<b>1.00</b>	< 0.001	0.0	2.5	<b>1.00</b>	< 0.001	278002	133799	48.13
<i>genome-nexus</i>	37.3	38.9	<b>0.86</b>	< 0.001	0.0	0.0	0.50	1.000	16576	12180	73.48
<i>gestaohospital-rest</i>	40.1	53.9	<b>1.00</b>	< 0.001	1.0	1.7	<b>0.83</b>	< 0.001	61829	71091	114.98
<i>ocvn-rest</i>	28.9	29.1	0.62	0.114	297.3	299.0	0.53	0.743	100824	100665	99.84
<i>reservations-api</i>	49.1	55.1	<b>1.00</b>	< 0.001	4.0	4.9	<b>0.79</b>	< 0.001	45581	41305	90.62
<i>session-service</i>	57.3	75.7	<b>1.00</b>	< 0.001	8.5	8.7	0.55	0.441	285487	188874	66.16
Average	40.0	47.3	0.91		51.8	52.8	0.70		131383	91319	82.20
Median	38.7	46.4	1.00		2.5	3.7	0.67		81327	85878	82.05

be evaluated within the same amount of time (i.e., 285k vs. 188k HTTP calls). Nevertheless, this API achieved much better coverage.

**RQ1:** We provide strong improvements in terms of code coverage compared to the state-of-the-art REST API white-box fuzzing. The average effect-size is  $\hat{A}_{12} = 0.91$ , with improvements of up to 18%. In terms of fault detection, more faults are found.

Table 6 shows the results of our second set of experiments. In this table, six tools are compared: four black-box fuzzers (i.e., ARAT-RL, EvoMASTER BB, RESTLER and SCHEMATHESIS), white-box EvoMASTER (i.e., EvoMASTER WB), and EvoMASTER extended with our novel techniques (i.e., MONGO). Comparisons are based on achieved code coverage (measured with JaCoCo). For each API, each tool is ranked based on their line coverage, where Rank 1 is best, and Rank 6 is the worst. To verify the difference in performance among the tools, a Friedman Test is carried out.

Table 6 shows some interesting results. First of all, considering the *Average* scores, our novel techniques (i.e., MONGO) gives the best overall results (i.e., the best average Rank 1.5), by a *large* margin in terms of line coverage (i.e.,  $52.7 - 43.8 = +8.9\%$  over the second best tool). The fact that EvoMASTER WB is the second best tool is not surprising. White-box techniques that can exploit internal knowledge of the tested SUT (e.g., runtime line coverage and processed database commands) have a clear advantage over black-box techniques. Surprisingly, ARAT-RL achieved the best results on two APIs (gestaohospital and session-service); we discuss these cases below.

Based on the results in Table 6, with an average code coverage of 58%, ARAT-RL gives better results than the second best tool, which is EvoMASTER BB with average coverage 56.6%. Interestingly, on this API EvoMASTER WB gives significantly worse results than EvoMASTER BB (i.e., 38.9%). To explain this rather odd behavior, we can think of at least two possible conjectures: (1) the fitness function for the evolutionary search of EvoMASTER gives little guidance on this API. As a black-box approach can evaluate many more test cases in the same amount of time (as it does not have to collect coverage results and compute white-box heuristics at each test case execution), it can better explore the search space. (2) There is a fault in EvoMASTER, where the outputted generated JUnit test cases do not contain all the coverage improving test cases evolved during the fuzzing session. Recall that, for EvoMASTER WB and

MONGO, coverage is based on the execution of the generated JUnit test suites, and not on all the calls done during the 1-hour fuzzing sessions. If we were to apply the same approach for ARAT-RL, then its results would be 0% coverage, as that tool prototype does not generate any executable test case (only log summaries).

session-service is another interesting case. If we exclude the results of ARAT-RL, then EvoMASTER WB gives better results than any black-box fuzzer, and MONGO provides strong improvements upon it (i.e., from 58.5% to 80%). Still, this is lower than the 88.3% achieved by ARAT-RL. Fuzzing a REST API poses many variegated challenges [59]. Handling MongoDB interactions is only one of many aspects that a fuzzer needs to deal with. In the particular case of session-service, there is something special there that ARAT-RL is very efficient at dealing with. Without further in-depth analyses, we cannot currently conjecture its nature.

The goal of this empirical study is to compare our novel techniques with the state-of-the-art, demonstrating their value as a research contribution. Comparing existing black-box fuzzers also allows us to examine how our replicated results align with prior studies on the selected APIs. For instance, our findings are consistent with the large-scale experiments reported in [59]: EvoMASTER BB and SCHEMATHESIS achieved the best results, outperforming RESTLER (ARAT-RL had not yet been introduced). In contrast, our results differ from those in [43], which introduced ARAT-RL. In our study, ARAT-RL still outperformed RESTLER but not EvoMASTER BB (note that SCHEMATHESIS was not considered in [43]). Interestingly, on two of the six APIs we analyzed, ARAT-RL produced the best results overall, even surpassing white-box fuzzing. This highlights the limits of conclusions from inter-tool comparisons.

First, the outcome of the comparisons largely depends on the selected APIs used in the experiments. By cherry-picking SUTs, one can show whatever results they want. Trying to be as fair and unbiased as possible in the SUT selection is of paramount importance for sound scientific studies. To provide stronger results, at posteriori we could have excluded from our study gestaohospital and session-service. But, besides being unethical and scientifically fraudulent, we could not have claimed then that we used all SUTs in the EMB corpus that rely on a MongoDB database. Having an established corpus of APIs like EMB is essential for scientific research, to avoid this kind of issues. However, as for any benchmark, it is important that novel techniques are not tailored for a specific

**Table 6: Performance comparisons between the six analysed tools. For each tool, we report the achieved average (out of 30 runs) line coverage (measured with JaCoCo), as well as the minimum and maximum values (out of 30 runs) achieved (shown in squared ‘[]’ parentheses). For p-values lower than the threshold  $\alpha = 0.05$ , the effect sizes  $\hat{A}_{12}$  are shown in bold. We report in ‘()’ parenthesis the rank of each tool (from ‘1’ best, to ‘6’ worst).**

SUT	ARAT-RL	EvoMASTER BB	RESTLER	SCHEMATHESIS	EvoMASTER WB	MONGO
<i>bibliothek</i>	35.2 [35.2,35.2] (5.5)	39.7 [39.7,39.7] (3)	35.2 [35.2,35.2] (5.5)	39.7 [39.7,39.7] (3)	39.7 [39.7,39.7] (3)	<b>46.9 [42.3,48.3] (1)</b>
<i>genome-nexus</i>	12.8 [12.8,12.8] (6)	31.4 [28.6,33.9] (3)	24.4 [24.4,24.6] (5)	27.3 [26.3,28.5] (4)	35.6 [33.5,36.7] (2)	<b>37.4 [31.5,39.4] (1)</b>
<i>gestahospital-rest</i>	<b>58.0 [45.3,65.6] (1)</b>	56.6 [46.6,61.6] (2)	21.5 [21.5,21.5] (6)	51.4 [48.0,56.9] (4)	38.9 [38.7,38.9] (5)	53.7 [50.3,56.6] (3)
<i>ocvn-rest</i>	10.1 [10.1,10.1] (5.5)	27.7 [27.7,27.9] (3)	10.1 [10.1,10.1] (5.5)	27.4 [27.3,27.5] (4)	27.9 [27.8,28.0] (2)	<b>28.2 [27.8,34.3] (1)</b>
<i>reservations-api</i>	26.2 [26.2,26.2] (6)	27.6 [27.6,27.6] (4.5)	27.6 [27.6,27.6] (4.5)	34.3 [30.8,56.3] (3)	62.4 [62.4,62.4] (2)	<b>70.0 [68.8,70.3] (1)</b>
<i>session-service</i>	<b>88.3 [88.1,89.3] (1)</b>	54.1 [54.1,54.1] (4.5)	50.9 [50.9,50.9] (6)	54.1 [54.1,54.1] (4.5)	58.5 [57.9,59.1] (3)	80.0 [76.7,84.9] (2)
Average	38.4 (4.2)	39.5 (3.3)	28.3 (5.4)	39.0 (3.8)	43.8 (2.8)	52.7 (1.5)
Friedman Test	$\chi^2 = 15.421$ , $p\text{-value} = 0.009$					

corpus (i.e., no overfitting). Corpora that are extended each year with new APIs for experimentation is necessary.

Second, ideally, from a scientific standpoint one would like to compare “techniques” and not “tools”. A tool could give better results just because it is more mature and robust. Building an effective fuzzer that can handle all the different aspects of OpenAPI schemas and properties of REST APIs is no trivial task. It requires time and effort. In this regard, older but actively maintained tools such as SCHEMATHESIS and EvoMASTER have a clear advantage. Still, novel techniques can be designed that can be particularly effective on some APIs, like it happened in our experiments with ARAT-RL. To better “sell” our novel results, at posteriori we could have excluded the comparisons with ARAT-RL, and rather use some other tools that seem to give worse results (e.g., based on previously published tool comparisons in the literature such as in [59]). Apart from being unethical, it is our opinion that it would *reduce* the scientific contribution of this paper. Showing that a novel technique works effectively on some SUTs is important, as well as it is important to show where it does not work so good. We integrated our novel techniques in the evolutionary search-based engine of EvoMASTER. But, based on the results obtained by ARAT-RL, it is clear that there is an important research direction in studying how to integrate these novel techniques in reinforcement learning approaches, like done for example in ARAT-RL. This will be future work.

**RQ2:** On average, with six real-world APIs, our novel techniques provide the overall best results. Still, there are important research directions for further improvements, such as integrating our techniques with reinforcement learning approaches.

## 8 Threats to Validity

A threat to internal validity is that the compared tools are using randomized algorithms. All experiments were repeated 30 times, and analyzed with the appropriate statistical tests recommended in the literature [22], such the Wilcoxon-Mann-Whitney U-Test, Vargha-Delaney  $\hat{A}_{12}$  standardized effect-size, and Friedman Test.

For external validity, six real-world APIs were used in our empirical study, chosen from the EMB Corpus [2, 27]. How good our novel techniques would fare on different APIs is not something that can be known for sure without further experiments. Using a larger set of APIs for experimentation is not simple, as we are constrained on the programming language (i.e., must run on the JVM, as we do white-box testing based on EvoMASTER’s bytecode

instrumentation) and type of API (it needs to use MongoDB). Although Java and MongoDB are widely used in industry, finding non-trivial instances of this type of web service on open-source repositories such as GitHub is not so simple. Collaborations with industrial partners to access to more APIs for experimentation will be an important research direction to follow.

## 9 Conclusion

In this paper, we have proposed a series of novel techniques to handle interactions with MongoDB when doing white-box fuzzing of REST APIs. Based on automated bytecode instrumentation, we can automatically intercept and analyze at runtime all calls made by the APIs towards their MongoDB databases. This guides fuzzing toward more effective inputs and enables direct data insertion into MongoDB from test cases.

Our novel approach is evaluated on six real-world APIs, providing better results than existing fuzzers. Compared to the state-of-the-art, large improvements in code coverage are achieved (e.g., up to +18%), as well as being able to detect more faults on average.

Our empirical study also shows some potential directions for further research to obtain better results. For example, on some APIs, reinforcement learning techniques as done in ARAT-RL give better results. How to best integrate our novel techniques with reinforcement learning will be an important research question to address.

## Acknowledgments

This work is funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (EAST project, grant agreement No. 864972), and partially funded by UBACYT-2020 20020190100233BA, PICT-2019-01793. Man Zhang is supported by NSFC (grant agreement No. 62502022).

## Data Availability

EvoMASTER and its new extension presented in this article are open-source on GitHub [1], with each release automatically uploaded to Zenodo for long-term storage (e.g., [28]).

## References

- [1] [n. d.]. EvoMaster. <https://github.com/WebFuzzing/EvoMaster>. Online, Accessed July 28, 2025.
- [2] [n. d.]. EvoMaster Benchmark (EMB). <https://github.com/WebFuzzing/EMB>. Online, Accessed July 28, 2025.

- [3] [n. d.]. JaCoCo. <https://www.jacoco.org/>. Online, Accessed July 28, 2025.
- [4] [n. d.]. KMongo. <https://litote.org/kmongo/>. Online, Accessed July 28, 2025.
- [5] [n. d.]. Micronaut Data MongoDB. <https://micronaut-projects.github.io/micronaut-data/latest/guide/#mongo>. Online, Accessed July 28, 2025.
- [6] [n. d.]. MongoDB. <https://www.mongodb.com/>. Online, Accessed July 28, 2025.
- [7] [n. d.]. MongoDB Java Drivers. <https://www.mongodb.com/docs/drivers/java-drivers/>. Online, Accessed July 28, 2025.
- [8] [n. d.]. Morphia. <https://morphia.dev/>. Online, Accessed July 28, 2025.
- [9] [n. d.]. Query Operators for MongoDB. <https://www.mongodb.com/docs/manual/reference/operator/query/>. Online, Accessed July 28, 2025.
- [10] [n. d.]. RestAssured. <https://github.com/rest-assured/rest-assured>. Online, Accessed July 28, 2025.
- [11] [n. d.]. Spring Boot. <https://spring.io/projects/spring-boot>. Online, Accessed July 28, 2025.
- [12] [n. d.]. Spring Data MongoDB. <https://spring.io/projects/spring-data-mongodb>. Online, Accessed July 28, 2025.
- [13] [n. d.]. SQL. <https://www.iso.org/standard/63555.html>. Online, Accessed July 28, 2025.
- [14] S. Ali, L.C. Briand, H. Hemmati, and R.K. Panesar-Walawege. 2010. A systematic review of the application and empirical investigation of search-based test-case generation. *IEEE Transactions on Software Engineering (TSE)* 36, 6 (2010), 742–762.
- [15] Mohammad Alshraideh and Leonardo Bottaci. 2006. Search-based software test data generation for string data using program-specific search operators. *Software Testing, Verification and Reliability (STVR)* 16, 3 (2006), 175–203.
- [16] A. Arcuri. 2013. It really does matter how you normalize the branch distance in search-based software testing. *Softw. Test., Verif. Reliab.* 23, 2 (2013), 119–147. doi:10.1002/stvr.457
- [17] Andrea Arcuri. 2017. RESTful API Automated Test Case Generation. In *IEEE International Conference on Software Quality, Reliability and Security (QRS)*. IEEE, 9–20.
- [18] Andrea Arcuri. 2018. EvoMaster: Evolutionary Multi-context Automated System Test Generation. In *IEEE International Conference on Software Testing, Verification and Validation (ICST)*. IEEE.
- [19] Andrea Arcuri. 2018. Test suite generation with the Many Independent Objective (MIO) algorithm. *Information and Software Technology* 104 (2018), 195–206.
- [20] Andrea Arcuri. 2019. RESTful API Automated Test Case Generation with EvoMaster. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 28, 1 (2019), 3.
- [21] Andrea Arcuri. 2020. Automated Black-and White-Box Testing of RESTful APIs With EvoMaster. *IEEE Software* 38, 3 (2020), 72–78.
- [22] A. Arcuri and L. Briand. 2014. A Hitchhiker's Guide to Statistical Tests for Assessing Randomized Algorithms in Software Engineering. *Software Testing, Verification and Reliability (STVR)* 24, 3 (2014), 219–250.
- [23] Andrea Arcuri and Juan P Galeotti. 2020. Handling SQL databases in automated system test generation. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 29, 4 (2020), 1–31.
- [24] Andrea Arcuri and Juan P Galeotti. 2021. Enhancing Search-based Testing with Testability Transformations for Existing APIs. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 31, 1 (2021), 1–34.
- [25] Andrea Arcuri, Juan Pablo Galeotti, Bogdan Marculescu, and Man Zhang. 2021. EvoMaster: A Search-Based System Test Generation Tool. *Journal of Open Source Software* 6, 57 (2021), 2153.
- [26] Andrea Arcuri, Man Zhang, and Juan Pablo Galeotti. 2024. Advanced White-Box Heuristics for Search-Based Fuzzing of REST APIs. *ACM Transactions on Software Engineering and Methodology (TOSEM)* (2024). doi:10.1145/3652157
- [27] Andrea Arcuri, Man Zhang, Amid Golmohammadi, Asma Belhadi, Juan P Galeotti, Bogdan Marculescu, and Susruthan Seran. 2023. EMB: A curated corpus of web/enterprise applications and library support for software testing research. In *2023 IEEE Conference on Software Testing, Verification and Validation (ICST)*. IEEE, 433–442.
- [28] Andrea Arcuri, Man Zhang, Susruthan Seran, Asma Belhadi, Juan Pablo Galeotti, Bogdan, Amid Golmohammadi, Onur Duman, Agustina Aldasoro, Philip, Alberto Martín López, Hernan Ghianni, Ömür Şahin, Annibale Panichella, Kyle Niemeyer, and Marcello Maugeri. 2025. *WebFuzzing/EvoMaster: v3.4.0*. doi:10.5281/zenodo.14597412
- [29] Vaggelis Atlidakis, Patrice Godefroid, and Marina Polishchuk. 2019. RESTler: Stateful REST API Fuzzing. In *ACM/IEEE International Conference on Software Engineering (ICSE)*. 748–758.
- [30] A. Baresel and H. Sthamer. 2003. Evolutionary testing of flag conditions. In *Genetic and Evolutionary Computation Conference (GECCO)*. 2442–2454.
- [31] George Candea and Patrice Godefroid. 2019. *Automated Software Test Generation: Some Challenges, Solutions, and Recent Advances*. Springer International Publishing, Cham, 505–531. doi:10.1007/978-3-319-91908-9\_24
- [32] Davide Corradini, Amedeo Zampieri, Michele Pasqua, and Mariano Ceccato. 2021. Restats: A test coverage tool for RESTful APIs. In *2021 IEEE International Conference on Software Maintenance and Evolution (ICSME)*. IEEE, 594–598.
- [33] Roy Thomas Fielding. 2000. *Architectural styles and the design of network-based software architectures*. Ph.D. Dissertation. University of California, Irvine.
- [34] Gordon Fraser and Andrea Arcuri. 2011. EvoSuite: Automatic Test Suite Generation for Object-Oriented Software. In *Proceedings of the 19th ACM SIGSOFT symposium and the 13th European conference on Foundations of software engineering*. 416–419.
- [35] Gordon Fraser and Andrea Arcuri. 2013. Whole Test Suite Generation. *IEEE Transactions on Software Engineering* 39, 2 (2013), 276–291.
- [36] Felix Gessert, Wolfram Wingerath, Steffen Friedrich, and Norbert Ritter. 2017. NoSQL database systems: a survey and decision guidance. *Comput. Sci. Res. Dev.* 32, 3–4 (2017), 353–365. doi:10.1007/S00450-016-0334-3
- [37] Amid Golmohammadi, Man Zhang, and Andrea Arcuri. 2022. Testing RESTful APIs: A Survey. *ACM Transactions on Software Engineering and Methodology* (2022).
- [38] Amid Golmohammadi, Man Zhang, and Andrea Arcuri. 2023. Testing RESTful APIs: A Survey. *ACM Transactions on Software Engineering and Methodology* (aug 2023). doi:10.1145/3617175
- [39] M. Harman, L. Hu, R. Hierons, A. Baresel, and H. Sthamer. 2002. Improving evolutionary testing by flag removal. In *Genetic and Evolutionary Computation Conference (GECCO)*. 1351–1358.
- [40] Mark Harman, S Afshin Mansouri, and Yuanyuan Zhang. 2012. Search-based software engineering: Trends, techniques and applications. *ACM Computing Surveys (CSUR)* 45, 1 (2012), 11.
- [41] Zac Hatfield-Dodds and Dmitry Dygalo. 2022. Deriving Semantics-Aware Fuzzers from Web API Schemas. In *2022 IEEE/ACM 44th International Conference on Software Engineering: Companion Proceedings (ICSE-Companion)*. IEEE, 345–346.
- [42] Pooyan Jamshidi, Claus Pahl, Nabor C Mendonça, James Lewis, and Stefan Tilkov. 2018. Microservices: The journey so far and challenges ahead. *IEEE Software* 35, 3 (2018), 24–35.
- [43] Myeongsoo Kim, Saurabh Sinha, and Alessandro Orso. 2023. Adaptive rest api testing with reinforcement learning. In *2023 38th IEEE/ACM International Conference on Automated Software Engineering (ASE)*. IEEE, 446–458.
- [44] Myeongsoo Kim, Qi Xin, Saurabh Sinha, and Alessandro Orso. 2022. Automated Test Generation for REST APIs: No Time to Rest Yet. In *Proceedings of the 31st ACM SIGSOFT International Symposium on Software Testing and Analysis (Virtual, South Korea) (ISSTA 2022)*. Association for Computing Machinery, New York, NY, USA, 289–301. doi:10.1145/3533767.3534401
- [45] Bogdan Korel. 1990. Automated software test data generation. *IEEE Transactions on software engineering* 16, 8 (1990), 870–879.
- [46] Nuno Laranjeiro, João Agnelo, and Jorge Bernardino. 2021. A black box tool for robustness testing of REST services. *IEEE Access* 9 (2021), 24738–24754.
- [47] Yi Liu, Yuekang Li, Gelei Deng, Yang Liu, Ruiyuan Wan, Runchao Wu, Dandan Ji, Shiheng Xu, and Minli Bao. 2022. Morest: Model-based RESTful API Testing with Execution Feedback. In *ACM/IEEE International Conference on Software Engineering (ICSE)*.
- [48] Ke Mao, Mark Harman, and Yue Jia. 2016. Sapienz: Multi-objective automated testing for android applications. In *ACM Int. Symposium on Software Testing and Analysis (ISSTA)*. ACM, 94–105.
- [49] Bogdan Marculescu, Man Zhang, and Andrea Arcuri. 2022. On the Faults Found in REST APIs by Automated Test Generation. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 31, 3 (2022), 1–43.
- [50] Alberto Martín-López, Sergio Segura, and Antonio Ruiz-Cortés. 2019. Test coverage criteria for RESTful web APIs. In *Proceedings of the 10th ACM SIGSOFT International Workshop on Automating TEST Case Design, Selection, and Evaluation*. 15–21.
- [51] Alberto Martín-López, Sergio Segura, and Antonio Ruiz-Cortés. 2021. RESTest: Automated Black-Box Testing of RESTful Web APIs. In *ACM Int. Symposium on Software Testing and Analysis (ISSTA)*. ACM, 682–685.
- [52] Sam Newman. 2021. *Building microservices*. "O'Reilly Media, Inc."
- [53] Alessandro Orso and Gregg Rothermel. 2014. Software testing: a research travelogue (2000-2014). In *Proceedings of the on Future of Software Engineering, FOSE 2014, Hyderabad, India, May 31 - June 7, 2014*, James D. Herbsleb and Matthew B. Dwyer (Eds.). ACM, 117–132. doi:10.1145/2593882.2593885
- [54] Annibale Panichella, Fitsum Kifetew, and Paolo Tonella. 2018. Automated Test Case Generation as a Many-Objective Optimisation Problem with Dynamic Selection of the Targets. *IEEE Transactions on Software Engineering (TSE)* 44, 2 (2018), 122–158.
- [55] José Miguel Rojas, José Campos, Mattia Vivanti, Gordon Fraser, and Andrea Arcuri. 2015. Combining multiple coverage criteria in search-based unit test generation. In *International Symposium on Search Based Software Engineering*. Springer, 93–108.
- [56] José Miguel Rojas, Gordon Fraser, and Andrea Arcuri. 2016. Seeding strategies in search-based unit test generation. *Software Testing, Verification and Reliability* 26, 5 (2016), 366–401.
- [57] Emanuele Viganis, Michael Dallago, and Mariano Ceccato. 2020. RESTTEST-GEN: Automated Black-Box Testing of RESTful APIs. In *IEEE International Conference on Software Testing, Verification and Validation (ICST)*. IEEE.

- [58] Huayao Wu, Lixin Xu, Xintao Niu, and Changhai Nie. 2022. Combinatorial Testing of RESTful APIs. In *ACM/IEEE International Conference on Software Engineering (ICSE)*.
- [59] Man Zhang and Andrea Arcuri. 2023. Open Problems in Fuzzing RESTful APIs: A Comparison of Tools. *ACM Transactions on Software Engineering and Methodology (TOSEM)* (may 2023). doi:10.1145/3597205