

GeoRocket: A scalable and cloud-based data store for big geospatial files

Michel Krämer

Fraunhofer Institute for Computer Graphics Research IGD, 64283 Darmstadt, Germany

Abstract

We present GeoRocket, a software for the management of very large geospatial datasets in the cloud. GeoRocket employs a novel way to handle arbitrarily large datasets by splitting them into chunks that are processed individually. The software has a modern reactive architecture and makes use of existing services including Elasticsearch and storage back ends such as MongoDB or Amazon S3. GeoRocket is schema-agnostic and supports a wide range of heterogeneous geospatial file formats. It is also format-preserving and does not alter imported data in any way. The main benefits of GeoRocket are its performance, scalability, and usability, which make it suitable for a number of scientific and commercial use cases dealing with very high data volumes, complex datasets, and high velocity (Big Data). GeoRocket also provides many opportunities for further research in the area of geospatial data management.

Keywords: Geospatial Data, Cloud, Big Data, Distributed Computing, Databases

1. Motivation and significance

The global data volume is growing continuously. By the year 2025, it will have reached 163 zettabytes (or 163 trillion gigabytes) [1]. The main drivers of this data growth are mobile phones, autonomous cars, satellites, and other devices with built-in location sensors [2]. Data collected by these devices can be located in time and place [3] and is called *spatiotemporal data* (or *geospatial data*, *geodata*). Kitchin & McArdle recognise geospatial data as *Big Data* characterised by its volume, variety, and velocity [4]. This means geospatial data sets are typically large, heterogeneous, and acquired in a

Email address: michel.kraemer@igd.fraunhofer.de (Michel Krämer)

URL: <https://georocket.io> (Michel Krämer)

10 short amount of time. Earth observation satellites, airborne laser scanners,
11 and terrestrial mobile mapping systems, for example, record hundreds of
12 thousands of samples per second [5] and produce a few GiB of data up to
13 several TiB in a couple of hours [6]. With the growing data volume, users
14 face new challenges as their current computer systems lack storage space and
15 computational power. At the same time, they require new software solutions
16 capable of handling such data.

17 In our previous work, we investigated the possibilities of using the cloud
18 and microservice architectures to process large amounts of heterogeneous
19 geospatial data [7, 8]. We focused on use cases from various domains such as
20 land management, urban planning, and marine applications [9, 10] where we
21 could show that geospatial data can be of great value given there is sufficient
22 computational power, enough storage resources, and suitable software.

23 To complement this, we now explore new ways to store, index, and query
24 big geospatial data in a scalable, efficient, and inexpensive manner. We
25 developed a novel software solution called *GeoRocket* enabling users to store
26 large amounts of geospatial vector data and to access, analyse, and share
27 it in a distributed environment—in our case the cloud. The key properties
28 of *GeoRocket* are its scalability, the indexing functionalities, as well as the
29 modular architecture and lightweight interfaces. At the same time, it is
30 schema-agnostic and format-preserving and provides users with a pragmatic
31 way to store data.

32 With *GeoRocket*, we pursue a novel path that differentiates it from ex-
33 isting software solutions:

- 34 • *PostGIS* [11] is an extension to *PostgreSQL* [12] that provides a low-
35 level interface for application developers to store and analyse geospa-
36 tial vector data in a traditional, relational database. In contrast,
37 *GeoRocket* is a high-level data store that makes use of other storage
38 technologies (see Section 2). The geospatial entities in *GeoRocket* are
39 semantic features and not geometries. *GeoRocket* is not a relational
40 database. It also employs its own query language instead of SQL (see
41 Section 2.3).
- 42 • *GeoServer* [13] and *Deegree* [14] are storage solutions for geospatial data
43 that have a long history. They have a monolithic architecture and use
44 a traditional client/server approach. *GeoRocket* has a modern reactive
45 and distributed architecture. It has been designed to run in the cloud
46 and to harness the possibilities in terms of performance, scalability, and
47 cost-effectiveness.

- 48 • *3DCityDB* [15] is a database that is specifically made to store CityGML
49 files describing 3D city models [16]. CityGML is an application schema
50 of the Geography Markup Language (GML) [17], which is itself based
51 on XML. In contrast to 3DCityDB, GeoRocket supports multiple file
52 formats and is schema-agnostic, so that it can handle CityGML, but
53 also GML, or even arbitrary XML files. Besides, 3DCityDB has again
54 a monolithic architecture.
- 55 • *rasdaman* [18] is a storage and analytics solution for large geospatial
56 raster data. It runs in a distributed environment and has been specifi-
57 cally designed for Big Data. The main difference to GeoRocket is the
58 type of the data stored. GeoRocket supports vector data and *rasdaman*
59 is made for raster data.
- 60 • *Cesium ion* [19] is a commercial solution to host massive 3D datasets
61 in the cloud and stream them efficiently for 3D web visualisation in the
62 browser with the JavaScript framework Cesium [20]. Data uploaded to
63 Cesium ion will be processed, optimized, and tiled to improve stream-
64 ing and visualisation performance. The original data is not accessible.
65 In contrast, GeoRocket allows access to the unmodified data. It also
66 supports 2D as well as 3D data (albeit the supported file formats differ).
67 It is a more generic solution that can be used in various applications,
68 whereas Cesium ion focuses on 3D web visualisation. Both solutions
69 complement each other and can be used in tandem.

70 The main contribution of this paper is the novel way to handle arbitrarily
71 large data sets (see Section 2). We describe GeoRocket’s event-driven archi-
72 tecture and our approach to importing and indexing. We also show samples
73 of GeoRocket’s query language to demonstrate how GeoRocket can be used.
74 Section 3 describes an illustrative example where we used our software in a
75 real-world application. GeoRocket provides potential for further research and
76 commercial exploitation, which we discuss in Section 4. The paper finishes
77 with conclusions in Section 5.

78 2. Software description

79 Figure 1 shows a generic overview of a GeoRocket deployment. The main
80 components are the GeoRocket server, the storage back end, and the index.
81 The server is responsible for importing and exporting geospatial files. The
82 actual data is kept in the storage back end. GeoRocket supports multiple
83 back ends such as Amazon S3 [21], MongoDB [22], distributed file systems,
84 or the local file system (typically used for testing purposes). In addition to




Figure 1: Overview of GeoRocket server, clients, and back-end services

85 the storage back end, GeoRocket keeps an inverted index about information
 86 found inside imported files. With this, users can search a large data set and
 87 extract the parts that are relevant to their use case. The index is main-
 88 tained by the Open-Source framework Elasticsearch [23]. The processes of
 89 importing, indexing, and querying are described in Section 2.1.

90 In addition to the server, the storage back end, and the index, there is a
 91 command-line interface called GeoRocket CLI. It allows users to import and
 92 export files, as well as to manage *tags* and *properties* (see list of definitions
 93 below). The GeoRocket server also has an HTTP interface that can be used
 94 by other clients.

95 Before going into detail about the architecture of GeoRocket, we define
 96 a number of commonly used terms.

97 **Chunk** A chunk represents a geospatial object (also called *feature*) within
 98 an imported file. For example, in a CityGML file containing a 3D city
 99 model, a chunk represents a building (specified by a `cityObjectMember`
 100 element). Analogously, in a GeoJSON file [24], a chunk is a feature in
 101 a feature collection. During import, GeoRocket splits geospatial files
 102 into individual chunks (see Section 2.1) and saves them in its storage
 103 back end.

104 **Layer** A layer is a user-defined label (or folder, or directory) that can be
 105 used to structure a large collection of chunks in GeoRocket’s storage
 106 back end. A chunk is always put into exactly one layer. If the user
 107 doesn’t define one during import, the chunk will be put into the root
 108 layer called ‘/’. Layers are structured hierarchically and parent layers
 109 always include all chunks of their sub-layers.

110 **Property** Properties are user-defined key-value pairs that can be attached
111 to one or more chunks. Keys are unique.

112 **Tag** A tag is a user-defined label that can be attached freely to one or more
113 chunks to structure data. Basically, a tag is a property with no value.

114 **Metadata** A metadata object includes user-defined tags and properties, as
115 well as other automatically derived information (e.g. the imported file's
116 spatial reference system).

117 **Indexed attribute** Indexed attributes are key-value pairs that GeoRocket
118 detects during import. Unlike properties, they are not user-defined
119 but directly extracted from the imported file (e.g. CityGML generic
120 attributes or GeoJSON properties).

121 Note that chunks, layers, and indexed attributes are immutable. If a
122 geospatial feature should be changed—i.e. if its attributes or geometry should
123 be modified or if it should be moved from one layer to another—the feature
124 has to be deleted and a modified one has to be imported again. User-defined
125 metadata such as properties and tags, however, can be changed later.

126 This is also one of the reasons why we developed a new query language
127 instead of using SQL. Joins and updates would be too complex or impossible
128 to implement, in particular since GeoRocket is no relational database as
129 described above.

130 *2.1. Software Architecture*

131 GeoRocket has been implemented with Vert.x, an Open-Source toolkit for
132 building reactive applications [25]. Its architecture consists of so-called *ver-*
133 *ticles*, which are independent components that communicate with each other
134 by sending messages through an event bus. The software design adheres to
135 the reactive manifesto [26]. GeoRocket is responsive, resilient, elastic, and
136 message-driven. That means it is able to respond in a timely manner even
137 under high load, and it provides good scalability in terms of data volume and
138 number of users/parallel requests. At the same time it is fault-tolerant and
139 can quickly recover from failures. Responsiveness, scalability, and elasticity
140 are the result of both the event-driven architecture design based on Vert.x
141 and the novel approach to importing and indexing. Fault-tolerance was im-
142 plemented with patterns such as isolation, asynchronous timeouts, fail-fast,
143 and retries. We refer to Nygard for more information on this topic [27].

144 Note that these properties allow GeoRocket to be deployed to the cloud
145 and to make use of the benefits offered by it. Individual verticles (or multiple
146 instances of GeoRocket) can be deployed to distributed virtual machines (or

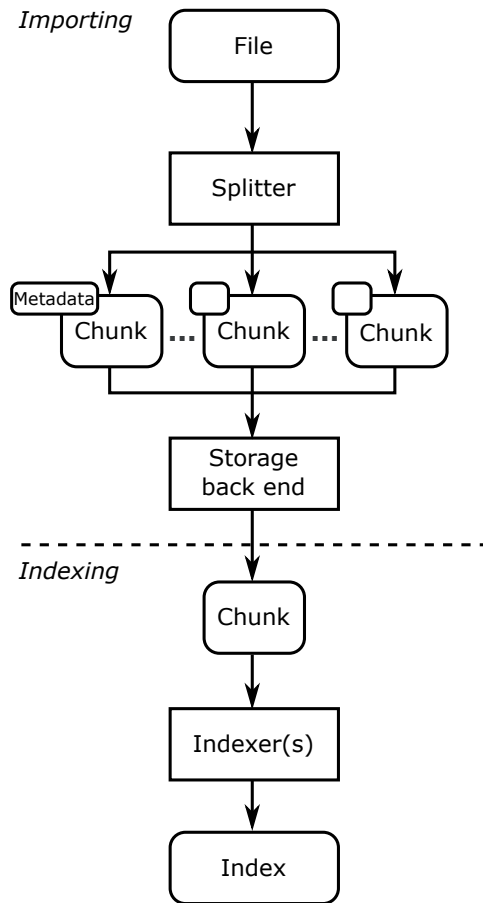


Figure 2: Importing and indexing a geospatial file: The file is split into chunks, which are in turn sent to the storage back end. After that, the indexing process runs asynchronously.

147 even containers) in the cloud to achieve high performance, reliability, and
 148 scalability. The event-driven architecture with loosely coupled verticles (or
 149 instances) allows GeoRocket to scale elastically on demand, which can help
 150 optimise resource usage and ultimately reduce operating costs.

151 Figure 2 depicts the process of importing and indexing a geospatial file
 152 and how data flows between verticles. In order to be able to process arbitrary
 153 data volumes, GeoRocket uses a novel streaming approach that applies the
 154 divide-and-conquer paradigm. At the beginning, the geospatial file is divided
 155 into individual chunks by the *Splitter* verticle. The *Splitter* also attaches
 156 user-defined metadata objects to each chunk. The chunks are saved in the
 157 configured storage back end. When all chunks have been written to the back
 158 end, the *importing* process is finished.




Figure 3: Exporting a file from GeoRocket: the merger retrieves chunk IDs from the indexer and merges matching chunks from the storage back end.

159 As soon as the first chunk has been written, the *indexing* process is started
 160 asynchronously. The *Indexer* verticle reads every imported chunk from the
 161 storage back end and looks for known patterns such as attributes, geometries,
 162 or bounding boxes. For this, it uses lightweight stream-based parsing and
 163 regular expressions. This approach is faster and more scalable than loading
 164 the chunk completely into memory and interpreting it semantically. It also
 165 helps GeoRocket interpret geospatial data in a schema-agnostic manner. Af-
 166 ter parsing, the Indexer saves the extracted information into the index. Note
 167 that there can be more than one Indexer, each of them responsible for a cer-
 168 tain kind of pattern. This allows GeoRocket to be extended with pluggable
 169 Indexers and to support indexing of heterogeneous data sets.

170 The process of querying and exporting files from GeoRocket is depicted
 171 in Figure 3. The main component in this diagram is the *Merger* verticle.
 172 It sends a query (see Section 2.3) to the Indexers, which in turn search
 173 the index for chunk IDs matching the query’s criteria. The chunk IDs are
 174 then sent back to the Merger, which in turn loads matching chunks from the
 175 storage back end. These chunks are joined to a valid output file that is finally
 176 rendered to the client.

177 2.2. Software Functionalities

178 From the software description above, we can derive the following key
 179 functionalities of our software:

- 180 • GeoRocket has been *designed for the cloud*. It has a distributed ar-
181 chitecture consisting of independent components (verticles) that can
182 be deployed redundantly to achieve *scalability, performance, and fault-*
183 *tolerance*. This is in contrast to existing alternative software products
184 that usually have a monolithic architecture.
- 185 • It is backed by the Open-Source framework Elasticsearch that allows
186 *indexing and querying of very large data sets* in flexible ways and with
187 high performance. Elasticsearch is itself designed to run in a distributed
188 environment and fits well to the architecture of GeoRocket.
- 189 • GeoRocket is *schema-agnostic*, which means it does not require specific
190 data schemas to work properly. Instead, it tries to identify common
191 patterns in large heterogeneous data sets and uses the extracted infor-
192 mation for indexing.
- 193 • Since GeoRocket’s data store is immutable, the software is *format-*
194 *preserving*. This means, every imported file can later, during export, be
195 reconstructed as it was (apart from possible minor whitespace changes
196 between chunks).

197 2.3. Sample queries

198 In this section, we demonstrate how the query language of GeoRocket
199 can be used to retrieve data. The structure of the language is lightweight.
200 It consists of terms, logical operators, and comparison operators. Terms can
201 be simple strings, dates, or bounding boxes (spatial areas defined by four
202 numbers minimum X, minimum Y, maximum X, and maximum Y). You
203 can also use the logical operators **AND**, **OR**, and **NOT**. The following example
204 retrieves all chunks located inside the given bounding box *and* containing the
205 string **Berlin** (e.g. as a value in one of the indexed attributes or properties)
206 or that are labelled with the tag **Berlin**:

```
207 AND (13.378, 52.515, 13.380, 52.517 Berlin)
```

208 You can also use comparison operators to constrain a term to a certain
209 indexed attribute or property. The following complex example combines
210 the logical operator **AND** with the comparison operators **EQ** (equals) and **GTE**
211 (greater than or equal to) to search for chunks that (a) lie inside a given
212 bounding box, (b) whose indexed attribute or property **name** equals **Berlin**,
213 and (c) whose indexed attribute or property **importedDate** is greater than or
214 equal to **2018-02-13** (i.e. chunks that have been imported on or after this
215 date):


```

216 AND(13.378,52.515,13.380,52.517 EQ(name Berlin)
217     GTE(importedDate 2018-02-13))

```

218 Note that, in this example, `name` and `importedDate` are either user-defined
219 (if they are properties) or their existence depends on the imported data
220 (in case they are indexed attributes). They are not created by default by
221 GeoRocket.

222 A more detailed description of GeoRocket’s query language can be found
223 in the user documentation [28].

224 3. Illustrative Example

225 In this section, we describe a real-world use case that demonstrates how
226 GeoRocket can be used to store a very large geospatial dataset in a public
227 cloud and to keep it up to date. The use case involves a data set of 3D build-
228 ing models provided by the German federal state of North Rhine-Westphalia
229 (Land NRW). The dataset is in the CityGML format (Level of Detail 2) and
230 is licensed under the dl-de/by-2-0 (Datenlizenz Deutschland - Namensnen-
231 nung - Version 2.0, www.govdata.de/dl-de/by-2-0). It can be downloaded
232 from www.opengeodata.nrw.de/produkte/geobasis/3d-gm/3d-gm_lod2/.

233 In order to demonstrate how this dataset can be kept up to date, we set
234 up a GeoRocket cluster using *Amazon Web Services (AWS)*. Table 1 shows
235 the EC2 instances we created and their configuration.

Description	Type	vCPUs	RAM	Volume size
1 × GeoRocket 1.3.0	c5.xlarge	4	8 GiB	40 GiB
3 × Elasticsearch 6.4.0	m5.2xlarge	8	32 GiB	100 GiB
1 × MongoDB 4.0.2	m5.large	2	8 GiB	100 GiB

Table 1: Instances of our GeoRocket cluster on AWS

236 Our cluster consisted of five instances (1 for GeoRocket, 3 for Elastic-
237 search, and 1 for MongoDB) running in the AWS region eu-central-1b (Frank-
238 furt). The volumes mounted into the instances were SSDs provided by the
239 Amazon Elastic Block Store (EBS). All instances were running the Ubuntu
240 16.04 LTS AMI (Amazon Machine Image). We deployed and provisioned
241 them using the Infrastructure-as-Code (IoC) tool Terraform [29].

242 After setting up the cluster, we imported the complete dataset with
243 GeoRocket’s command-line application (CLI). The dataset had a total size
244 of 224.3 GB split up into 35,022 files. Since the CLI uses GZIP compression
245 during upload, only 23.7 GB had to be transferred. We also recorded the
246 space usage on our EC2 instances. The MongoDB database was 69.7 GiB

247 large. The size of the sharded Elasticsearch index was 36.7 GiB. By de-
248 fault, Elasticsearch creates one replica of each index, so the total size of the
249 Elasticsearch storage was 73.4 GiB distributed over the three EC2 instances.
250 MongoDB and Elasticsearch use Snappy and LZ4 compression respectively,
251 which is the reason why the used space was lower than the total size of our
252 dataset. Both MongoDB and Elasticsearch contained entries for 10,529,668
253 chunks, which means there were this many geospatial objects in the dataset.

254 In order to demonstrate how such a large dataset can be managed with
255 GeoRocket, we performed a workflow that is realistic and regularly happens
256 in this or a similar way in municipalities or federal agencies. The dataset
257 contained buildings in level of detail 2 (LoD2), which means they were only
258 represented by wall and roof geometries. Suppose the dataset should be up-
259 dated and more detailed building models should be added: for the purpose
260 of city marketing, the LoD2 models of the popular shopping street ‘Schilder-
261 gasse’ in Cologne should be replaced by highly detailed geometries. Further
262 suppose that old objects should not be removed from the dataset but kept
263 for historical reasons.

264 First, we used the command-line application to mark the buildings in the
265 Schildergasse in Cologne as outdated by adding a property `deleted` denoting
266 the date when the buildings were replaced. Since the dataset contained xAL
267 2.0 addresses [30], we were able to use the terms `Schildergasse` and `Köln`
268 (German for Cologne) in our command.¹

```
269 georocket property set -props deleted:2018-09-13 \  
270     AND(Schildergasse Köln)
```

271 We then imported the new buildings:

```
272 georocket import Schildergasse_update.gml
```

273 After this, we were able to download the complete city model of Cologne
274 excluding the old models of the Schildergasse with the following query:

```
275 georocket search AND(NOT(LTE(deleted 2018-09-13)) Köln)
```

276 This query matches all objects from Cologne but not those that have a
277 `deleted` property whose value is less than or equal to 2018-09-13.

278 All operations performed very fast. Setting the property was finished in
279 a few milliseconds. The new file was only a few MiB large and importing
280 it was a matter of seconds. The latency for downloading the complete city
281 model of Cologne was again very low (a few milliseconds).

¹If you run the CLI on a Unix shell such as bash, you need to escape parentheses with backslashes. The Windows Command Prompt, on the other hand, does not require them. For the sake of readability, we omitted the backslashes here.

282 If the municipality or federal agency wanted to actually delete old data
283 from the dataset on a regular basis for the sake of housekeeping (e.g. at the
284 end of every year), they could use the following command:

```
285 georocket delete LT(deleted 2018)
```

286 This command would remove all objects from the dataset that have been
287 marked as deleted in 2017 or earlier. GeoRocket is able to automatically
288 parse dates in ISO format and compare their values accordingly.

289 4. Impact

290 As mentioned in Section 1, geospatial data is increasingly becoming larger
291 and more complex. Users are faced with new problems related to data volume
292 and heterogeneity, as well as the speed with which data is acquired. One of
293 the aims of developing GeoRocket was to make it possible to analyse and
294 share such data by leveraging the possibilities of the cloud. This has opened
295 up a number of new research directions and potential for changing the way
296 users and companies work with Big Geo Data.

297 Firstly, we are currently working on developing novel visual analysis meth-
298 ods for large geospatial data sets. In the research project DataBio funded
299 by the European Commission (grant agreement No 732064), we are using
300 GeoRocket as a store for data from the agricultural domain. Based on this,
301 we are developing a visual tool to interactively explore the data set and to
302 perform analyses and aggregations.

303 We are also using GeoRocket in the area of Smart City Clouds where
304 it enables users to access and share the large amounts of information col-
305 lected in a Smart City for the first time for use cases such as urban planning
306 or traffic management. In this respect, we have extended GeoRocket with
307 means to encrypt data in the cloud while keeping the possibility to search it
308 using Searchable Symmetric Encryption (SSE) [31]. We have also explored
309 the possibility to share geospatial data in a secure way in a Smart City Cloud
310 for applications related to security [32]. Furthermore, we discussed the pos-
311 sibility to use GeoRocket in the area of processing of large geospatial data
312 for use cases such as land monitoring or urban planning [7, pp. 48–49 and
313 163–164].

314 Besides the research opportunities, we believe GeoRocket also benefits
315 users and companies. As mentioned in Section 1, there are existing prod-
316 ucts that can manage geospatial data, but they typically have a monolithic
317 software architecture and are supposed to run in a traditional client/server
318 setting. GeoRocket, on the other hand, has been designed to run in the
319 cloud and to leverage its possibilities, not only in terms of performance and

320 scalability but also cost-effectiveness. Deploying GeoRocket to the cloud can
321 be much less expensive for users and companies than maintaining dedicated
322 on-premise hardware. This particularly applies to public administrations or
323 small and medium enterprises.

324 5. Conclusions

325 In this paper, we presented GeoRocket, a scalable and cloud-based data
326 store for geospatial files. We compared it to existing products and described
327 its architecture and query language. We also presented an illustrative exam-
328 ple showing how GeoRocket can be used in a real-world application. Finally,
329 we discussed the impact of our software with regards to opportunities for
330 scientific research and commercial exploitation.

331 There is an ongoing paradigm shift in Computer Sciences towards Cloud
332 Computing. This particularly applies to Geoinformatics, which has only
333 started to make successful use of the cloud. GeoRocket is one of the first
334 applications that is specifically designed to manage geospatial data and to
335 run in the cloud. In this paper, we were able to show that our novel ap-
336 proach to data handling based on splitting files into chunks and indexing
337 them individually has many benefits regarding performance, scalability, and
338 usability. Since GeoRocket is schema-agnostic, it supports a wide range
339 of geospatial datasets and can be used in multiple applications. It is also
340 format-preserving and avoids information loss that typically happens when
341 you have to transform data between different models. These properties make
342 GeoRocket superior to existing solutions.

343 The illustrative example presented in the paper only scratches the surface
344 of what is possible with GeoRocket. In Section 4, we already mentioned
345 briefly that we are also working on visual analysis methods and secure data
346 storage. In the future, we will focus on novel approaches to increase the
347 use of the large amounts of data managed by GeoRocket through Big Data
348 methods and Visual Analytics. We will also further improve the performance
349 of GeoRocket and explore its use for time series and other spatiotemporal
350 data.

351 References

- 352 [1] D. Reinsel, J. Gantz, J. Rydning, Data age 2025 - the evolution of data
353 to life-critical, Tech. rep., An IDC White Paper, Sponsored by Seagate
354 (2017).

- 355 [2] M. F. Goodchild, Citizens as sensors: the world of volunteered
356 geography, *GeoJournal* 69 (4) (2007) 211–221. doi:10.1007/
357 s10708-007-9111-y.
- 358 [3] R. R. Vatsavai, A. Ganguly, V. Chandola, A. Stefanidis, S. Klasky,
359 S. Shekhar, Spatiotemporal data mining in the era of big spatial data:
360 Algorithms and applications, in: *Proceedings of the 1st ACM SIGSPA-*
361 *TIAL International Workshop on Analytics for Big Geospatial Data*,
362 ACM, 2012, pp. 1–10. doi:10.1145/2447481.2447482.
- 363 [4] R. Kitchin, G. McArdle, What makes Big Data, Big Data? Exploring
364 the ontological characteristics of 26 datasets, *Big Data & Society* 3 (1)
365 (2016) 1–10. doi:10.1177/2053951716631130.
- 366 [5] C. Cahalane, T. McCarthy, C. P. McElhinney, MIMIC: Mobile Mapping
367 Point Density Calculator, in: *Proceedings of the 3rd International Con-*
368 *ference on Computing for Geospatial Research and Applications*, ACM,
369 2012, pp. 15:1–15:9. doi:10.1145/2345316.2345335.
- 370 [6] N. Paparoditis, J. P. Papelard, B. Cannelle, A. Devaux, B. Soheilian,
371 N. David, E. Houzay, Stereopolis II: A multi-purpose and multi-sensor
372 3D mobile mapping system for street visualisation and 3D metrology,
373 *Revue française de photogrammétrie et de télédétection* 200 (1) (2012)
374 69–79.
- 375 [7] M. Krämer, A microservice architecture for the processing of large
376 geospatial data in the cloud, Ph.D. thesis, Technische Universität Darm-
377 stadt (2018). doi:10.13140/RG.2.2.30034.66248.
- 378 [8] M. Krämer, I. Senner, A modular software architecture for processing
379 of big geospatial data in the cloud, *Computers & Graphics* 49 (2015)
380 69–81. doi:10.1016/j.cag.2015.02.005.
- 381 [9] J. Böhm, M. Bredif, T. Gierlinger, M. Krämer, R. Lindenbergh, K. Liu,
382 F. Michel, B. Sirmacek, The IQmulus Urban Showcase: Automatic Tree
383 Classification and Identification in Huge Mobile Mapping Point Clouds,
384 *ISPRS - International Archives of the Photogrammetry, Remote Sensing*
385 *and Spatial Information Sciences XLI-B3* (2016) 301–307. doi:10.5194/
386 isprs-archives-XLI-B3-301-2016.
- 387 [10] M. Belényesi, D. Kristóf, IQmulus public project deliverable D1.2.3 -
388 Revised User Requirements, Tech. rep. (2014).
- 389 [11] PostGIS (2019). <https://postgis.net/> [accessed 25 October 2019].

- 390 [12] PostgreSQL (2019). <https://www.postgresql.org/> [accessed 25 Oc-
391 tober 2019].
- 392 [13] GeoServer (2019). <http://geoserver.org/> [accessed 25 October 2019].
- 393 [14] Deegree (2019). <https://www.deegree.org/> [accessed 25 October
394 2019].
- 395 [15] 3DCityDB (2019). <https://www.3dcitydb.org/> [accessed 25 October
396 2019].
- 397 [16] G. Gröger, T. H. Kolbe, C. Nagel, K.-H. Häfele (eds.), OGC City Ge-
398 ography Markup Language (CityGML) Encoding Standard 2.0, Tech.
399 rep., Open Geospatial Consortium (2012).
400 URL <http://www.opengeospatial.org/standards/citygml>
- 401 [17] C. Portele (ed.), OpenGIS Geography Markup Language (GML) Encod-
402 ing Standard 3.2.1, Tech. rep., Open Geospatial Consortium (2007).
403 URL <http://www.opengeospatial.org/standards/gml>
- 404 [18] rasdaman (2019). <https://www.rasdaman.com/> [accessed 25 October
405 2019].
- 406 [19] Cesium ion (2019). <https://cesium.com/cesium-ion/> [accessed 25
407 October 2019].
- 408 [20] Cesiumjs (2019). <https://cesiumjs.org/> [accessed 25 October 2019].
- 409 [21] Amazon S3 (2019). <https://aws.amazon.com/s3/> [accessed 25 Octo-
410 ber 2019].
- 411 [22] MongoDB (2019). <https://www.mongodb.com/> [accessed 25 October
412 2019].
- 413 [23] Elasticsearch (2019). [https://www.elastic.co/de/products/
414 elasticsearch](https://www.elastic.co/de/products/elasticsearch) [accessed 25 October 2019].
- 415 [24] H. Butler., M. Daly, A. Doyle, S. Gillies, S. Hagen, T. Schaub, RFC 7946
416 – The GeoJSON Format, Tech. rep., Internet Engineering Task Force
417 (IETF) (2016).
418 URL <https://tools.ietf.org/html/rfc7946>
- 419 [25] Vert.x (2019). <https://vertx.io/> [accessed 25 October 2019].

- 420 [26] J. Bonér, D. Farley, R. Kuhn, M. Thompson. The reactive manifesto v2.0
421 (2014). <https://www.reactivemanifesto.org/> [accessed 25 October
422 2019].
- 423 [27] M. T. Nygard, Release It! Design and Deploy Production-Ready Soft-
424 ware, Pragmatic Bookshelf, 2007.
- 425 [28] M. Krämer. GeoRocket user documentation (2018). [https:
426 //georocket.io/docs/user-documentation/](https://georocket.io/docs/user-documentation/) [accessed 25 October
427 2019].
- 428 [29] Terraform (2019). <https://www.terraform.io/> [accessed 25 October
429 2019].
- 430 [30] Extensible Address Language (xAL) Standard Description Document
431 for W3C DTD/Schema – Version 2.0, Tech. rep., Organization for the
432 Advancement of Structured Information Standards OASIS (2002).
433 URL [https://www.oasis-open.org/committees/tc_home.php?wg_
434 abbrev=ciq](https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=ciq)
- 435 [31] B. Hiemenz, M. Krämer, Dynamic searchable symmetric encryption for
436 storing geospatial data in the cloud, International Journal of Information
437 Security (2018) Accepted, to be published.
- 438 [32] M. Krämer, S. Frese, A. Kuijper, Implementing secure applications in
439 smart city clouds using microservices, Future Generation Computer Sys-
440 tems (2019) Submitted, under review.

441 **Required Metadata**

442 **Current code version**

Nr.	Code metadata description	Please fill in this column
C1	Current code version	Git SHA ace352a
C2	Permanent link to code/repository used for this code version	https://github.com/georocket/georocket/tree/v1.3.0
C3	Legal Code License	Apache-2.0
C4	Code versioning system used	git
C5	Software code languages, tools, and services used	Java, Vert.x, Elasticsearch
C6	Compilation requirements, operating environments & dependencies	JDK 8
C7	If available Link to developer documentation/manual	https://georocket.io/docs/
C8	Support email for questions	michel.kraemer@igd.fraunhofer.de

Table 2: Code metadata (mandatory)

443 **Current executable software version**

Nr.	(Executable) software metadata description	Please fill in this column
S1	Current software version	1.3.0
S2	Permanent link to executables of this version	https://github.com/georocket/georocket/releases/tag/v1.3.0
S3	Legal Software License	Apache-2.0
S4	Computing platforms/Operating Systems	Linux, macOS, Windows
S5	Installation requirements & dependencies	Java 8
S6	If available, link to user manual - if formally published include a reference to the publication in the reference list	https://georocket.io/docs/user-documentation/1.3.0
S7	Support email for questions	michel.kraemer@igd.fraunhofer.de

Table 3: Software metadata (optional)