

HBSNIFF: A Static Analysis Tool for Java Hibernate Object-Relational Mapping Code Smell Detection

Zijie Huang, Zhiqing Shao*, Guisheng Fan*, Huiqun Yu, Kang Yang, Ziyi Zhou

Department of Computer Science and Engineering, East China University of Science and Technology, Shanghai 200237, China

Abstract

Code smells are symptoms of sub-optimal software design and implementation choices. Detection tools were actively developed for general code smell related to coupling and cohesion issues, but such tools cannot capture domain-specific problems. In this work, we fill the gap in data persistence and query code quality by proposing HBSNIFF, i.e., a static analysis tool for detecting 14 code smells as well as 4 mapping metrics in Java Hibernate Object-Relational Mapping (ORM) codes. HBSNIFF is tested, documented, and manually validated. It also generates readable and customizable reports for every project. Moreover, it is beneficial to Mining Software Repository (MSR) research requiring large-scale analysis since project compilation is not needed for detection.

Keywords: Code Smell, Object-Relational Mapping, Hibernate, Static Analysis, Object-Oriented Programming

1. Introduction

Code smells (*i.e.*, symptoms of sub-optimal design and implementation choices [1]) are related to determining factors of software quality such as change- and error-proneness [2]. Compared with software defects, code smells

*Corresponding Authors.

Email addresses: hzj@mail.ecust.edu.cn (Zijie Huang), zshao@ecust.edu.cn (Zhiqing Shao), gxfan@ecust.edu.cn (Guisheng Fan), yhq@ecust.edu.cn (Huiqun Yu), 15921709583@163.com (Kang Yang), zhouziyi@mail.ecust.edu.cn (Ziyi Zhou)

5 are more likely to be underestimated, and they have long life-cycles, which
6 means they may impose negative effects in the long run [2]. Thus, code smell
7 detectors were actively developed to capture severe issues that may hinder
8 maintainability. Recently, researchers found that general smell detection
9 tools were not able to capture more specific problems related to domain-
10 specific code [3], and domain-specific smells (*e.g.*, for data persistence [4, 5,
11 6, 7]) were attracting more attention from researchers and practitioners.

12 To facilitate Object-Oriented Programming (OOP), practitioners use
13 Object-Relational Mapping (ORM) frameworks which bridge database and
14 application by filling the gap of data mapping and persistence [4, 8]. Despite
15 their flexibility and capability, there exist various challenges in the applica-
16 tion of ORM [8] including the metamorphic class and table inheritance [9, 10],
17 the inconsistency in data structure [11], and the uncontrollable propagation
18 of relational data retrieval [12]. Consequently, ORM usage is regarded as
19 a double-edged sword [4, 8] or even an anti-pattern [13]. However, recent
20 studies suggested ORM need not affect performance if used properly [4], and
21 practitioners need more static analysis tool support to help them with de-
22 velopment [3]. In response, related studies [13, 14] outlined several smells
23 and refactoring strategies to cope with them. In most cases, they either did
24 not provide tools, or the tools were early prototypes and requires project
25 compiling, which is not ideal [15, 16] for large-scale analysis over real-world
26 systems. Thus, we fill the gap by proposing a static analysis tool called HB-
27 SNIFF (HiBermate Sniffer) for ORM code smell detection. Similar to related
28 studies [3, 12, 13, 14, 17, 18, 19], we use the trending Java HIBERNATE¹
29 framework as the context of our implementation.

30 The contributions of our work include:

- 31 • We implement a HIBERNATE code quality static analysis tool called
32 HBSNIFF for detecting 14 code smells and calculating 4 mapping metrics.
- 33 • Compared with the state-of-the-art dynamic analysis tool [14], we im-
34 plement more code smells, and our tool supports the detection of projects
35 using Java version greater than 1.7.
- 36 • We manually validate our tool on 5 open-source projects and 1 commer-
37 cial project, and we prove the high impact of 7 out of 8 performance smells
38 in a case study.

39 The highlights of HBSNIFF are:

¹<https://hibernate.org/>

40 • Compilation is not required for the evaluated project.
41 • Test cases and documentations for every smell detector and metric are
42 included.

43 • The tool generates a customizable and readable report for every project.
44 • The code is open-sourced on GITHUB.

45 The rest of this paper is organized as follows. In Section 2 we introduce
46 the problems, background, and summarize related work. Section 3 presents
47 the architecture and the implemented smells, while Section 4 outlines the
48 evaluation results, the advantage of our tool, as well as further research
49 opportunities. In Section 5 we present illustrative examples of our tool.
50 Finally, Section 6 concludes the paper and describes future research.

51 2. Problems and Background

52 In this section, we introduce the background and related work of ORM
53 code quality analysis as well as code smells.

54 2.1. Java ORM and Hibernate Query Language (HQL)

Listing 1: Example of An ORM Entity (User.java)

```
55 package com.example.blog.models;  
56 import javax.persistence.*; import java.util.*;  
57 import java.io.*;          import lombok.Data;  
58 @Data @Entity @Table("users")  
59 public class User implements Serializable {  
60     @Id @Column  
61     private Integer id;  
62     @Column(unique = true)  
63     private String email;  
64     @Column  
65     private String password;  
66     @Column  
67     private String name;  
68  
69     @OneToMany(fetch = FetchType.LAZY, mappedBy = "user")  
70     private List<Post> posts = new ArrayList<>();  
71  
72     public User() {} // No-Arg Constructor  
73 }
```

74 Java ORM frameworks implement the Java Persistence API (JPA²) in
75 order to map the tables, columns, and relationships of trending RDBMS
76 (Relational DataBase Management Systems, *e.g.*, MySQL³) to the OOP-
77 driven classes, attributes, and inheritance [8]. Listing 1 shows an example
78 of an ORM entity class. Classes annotated with `@Entity` indicates that it
79 is an entity in the ORM context, while the `@Table` annotation specifies its
80 corresponding RDBMS data table. Moreover, `@Id` could be used to specify
81 the unique identifier (in most cases, the corresponding field of the Primary
82 Key of RDBMS data table), and relational annotations like `@ManyToOne`,
83 `@OneToMany` could be used to describe relationships between entities with
84 `FetchType` (*e.g.*, `EAGER`, `LAZY`) specified to determine whether data should
85 be fetched in advance or on demand.

Listing 2: Example of A User Query (Method Extracted from UserRepository.java)

```
86  
87 public List<User> getUserByEmail(String queryEmail){  
88     EntityManager em = connection.createEntityManager();  
89     String hql = "from User u where u.email=:email";  
90  
91     Query query = em.createQuery(hql, User.class);  
92     query.setParameter("email", queryEmail);  
93  
94     List<User> results = query.getResultList();  
95     return results;  
96 }
```

97 Listing 2 demonstrates an example of an HQL query using the entity class
98 of Listing 1, which is a method to retrieve users by their email addresses.
99 The HQL template including the `email` parameter is defined in the `String`
100 variable called `hql`, and it is later populated by the method parameter called
101 `queryEmail`. Finally, the populated HQL query is executed by HIBERNATE,
102 and the results are returned in a `List` collection of the `User` entities. HQL
103 is a SQL-like query language designed for ORM [13, 20]. HQL could either
104 be generated by ORM or specified by developers. Then, ORM will translate
105 HQL to executable SQLs for RDBMS. Later, the results of the queries will
106 be processed by ORM and presented with the form of OOP in the context of

²<https://www.jcp.org/en/jsr/detail?id=338>

³<https://www.mysql.com/>

107 JVM (Java Virtual Machine). We focus on the human-written HQLs in this
108 work.

109 2.2. ORM Code Quality Analysis and Optimization

110 Since ORM relies heavily on automatic relational data retrieval and map-
111 ping, performance issues are primary concerns of practitioners. Procaccianti
112 *et al.* [21] found that ORM approaches can introduce a 70% increase in exe-
113 cution time, and they significantly increased energy consumption. However,
114 there existed conflicting opinions indicating ORM need not affect perfor-
115 mance if used properly [4]. Thus, researchers intended to find best prac-
116 tices to ensure the appropriate usage of ORM. Chen *et al.* [22] proposed
117 a cache optimizers for HIBERNATE-based web applications, which improved
118 the throughput up to 138% comparing to default caching strategy. Singh *et al.* [23]
119 *et al.* [23] exploited multi-objective genetic algorithms to improve ORM per-
120 formance by optimizing ORM configurations. Lorenz *et al.* [9] compared the
121 performance of 3 different mapping strategies and provided visualizations of
122 radar charts as conclusions. Chen *et al.* [3] conducted an empirical study
123 on ORM code changes, they found that ORM codes were frequently modi-
124 fied, and such modifications lacked static analysis tool support. Meurice *et al.*
125 *et al.* [11] proposed a detection approach to address inconsistencies in ORM
126 code after database schema change. Nazário *et al.* [24] proposed a devel-
127 opment framework to solve 12 problems with mapping problems as well as
128 their consequences.

129 2.3. Domain-Specific Code Smells

130 The state-of-the-arts of code smell detection tools mainly focused on gen-
131 eral code smells such as coupling, cohesion, and complexity issues [25], such
132 as Feature Envy, Spaghetti Code, and Complex Class. However, since recent
133 work revealed there may be “other fish in the sea [1, 26]” that may impose
134 uncaptured impact to software maintainability, various domain-specific code
135 smell detection tools were proposed. For example, in the domains related
136 to our work, Aniche *et al.* outlined various smells of model, view, controller
137 codes in web applications [27], and the quality of Structured Query Language
138 (SQL) queries [5, 6, 7] were also discussed.

139 In terms of ORM smells that we implement, Holder *et al.* [10] proposed
140 a metric suite to measure ORM mapping code complexity. Silva *et al.* [14]
141 proposed a set of rules to check if HIBERNATE entity codes follow JPA spec-
142 ifications. Loli *et al.* [13] summarized ORM code smells proposed in prior

143 studies [3, 12, 17, 18, 19] as well as in grey literature, and they investigated
 144 the agreement of developers towards the definition of smells. Results showed
 145 most developers agree with the definitions and severity.

146 3. Software Framework

147 3.1. Software Architecture

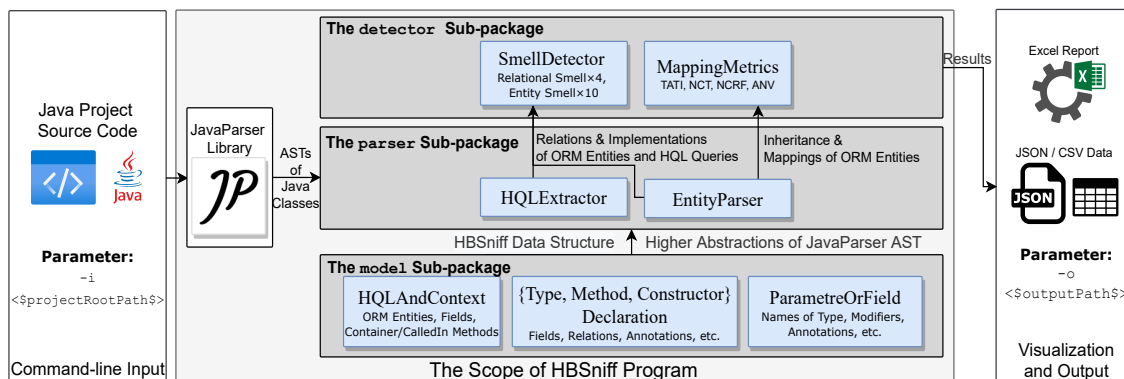


Figure 1: The general architecture of HBSNIFF.

148 HBSNIFF is a Java project using MAVEN⁴ as dependency management
 149 and building tool which consists of 3 major modules (sub-packages), *i.e.*,
 150 `model`, `parser`, and `detector`. The general architecture is depicted in Fig-
 151 ure 1. The light gray colored box represents the scope of HBSNIFF code,
 152 while the dark gray colored boxes are sub-packages of HBSNIFF. The blue-
 153 colored boxes are classes in the sub-packages, and the white-colored boxes
 154 are external libraries and resources.

155 First, users specify the path of the project to detect as input (the `-i`
 156 parameter) and the path of the output directory (the `-o` parameter). Then,
 157 the program constructs higher level abstractions (whose data models are
 158 available in the `model` sub-package) in the `parser` sub-package using JAVA-
 159 PARSER⁵. Meanwhile, it locates HQL in the context of the method and
 160 class containing the method call. Moreover, it finds the method that calls
 161 the method containing HQL (`CalledIn` methods in Fig. 1). Afterwards,

⁴<https://maven.apache.org/>

⁵<https://javaparser.org/>

162 the generated `models` will be used to populate the context of code smell
163 `detectors`. Finally, the detection and metric evaluation will be performed,
164 and the results will be converted to EXCEL reports as well as `csv` and `JSON`
165 data. The demonstrations of the outputs are available in Section 5.

166 The implementation of the 10 smell detectors and 4 mapping metrics fol-
167 lows their definitions, which will be described in the next Subsection. Each
168 detector class extends a class called `SmellDetector` which provides standard
169 interface methods and basic capabilities (*e.g.*, loading all available HIBER-
170 NATE entities). Since we only have 4 closely related `MappingMetrics` to
171 detect, they are now implemented in a single file. We will exploit a strategy
172 similar to `SmellDetector` if more metrics are developed in the future.

173 Apart from the detectors, we also implement 2 `parsers` of HQL and HI-
174 BERNATE entities to avoid operating directly on the lower-level `JAVAPARSER`
175 ASTs, which may be more challenging to comprehend and maintain. The
176 `parsers` are designed to retrieve relations, inheritance, mappings, and other
177 implementation details from the original `JAVAPARSER` ASTs. We introduce
178 the 2 `parsers` in the next paragraphs.

179 The `EntityParser` consists of several methods aiming to parse Java
180 classes to the `JAVAPARSER CompilationUnits`, and convert
181 `CompilationUnit` to HBSNIFF-defined `TypeDeclaration` in the `model` pack-
182 age. Each `TypeDeclaration` refers to a Java class (including HIBERNATE
183 entities), which contains its nested fields, its relations with nested types, its
184 class- and method-wide annotations, its constructors, and so on. Meanwhile,
185 HBSNIFF also defines the abstractions of methods and constructors.

186 The `HQLExtractor` locates the `createQuery` method call which indicates
187 potential HQL usage [28], and generates `HQLAndContext` objects containing
188 the HQL query `String` and its corresponding context, *e.g.*, the signature of
189 the method containing HQL (*i.e.*, container method of `HQLAndContext` in
190 Fig. 1), the available types and their fields (presented in the
191 `ParametreOrField` objects from the `model` sub-package) of the entities in
192 the FROM phrase of the HQL query, the methods which call the container
193 of HQL (*i.e.*, `CalledIn` methods of `HQLAndContext` in Fig. 1), and so on.

194 3.2. Inter-Entity Relational Smells *Detected*

195 Relational smells summarized in [13] are inter-entity smells caused by
196 inappropriate usage of data retrieval strategies in entity relationships. Some
197 of the relational smells are related to the N+1 performance issue. The
198 N+1 problem occurs when an application retrieves a parent entity from the

199 database, and then loops through a collection field of the entity containing
200 N other entities. HIBERNATE may generate a query for every iteration to
201 retrieve smelly entities, which means we call to the database recurrently. In
202 total, the application will call the database once for every row (*i.e.*, for N
203 times) returned by the original query, and the plus one refers to the original
204 query. This problem could lead to performance issue if the size of N is
205 large. However, reducing N is not acceptable since we may need that large
206 amount of data. The appropriate way to address it is to correctly configure
207 the relationships between entities and the strategies of data fetching. For
208 example, using the LAZY FetchType with specified batchSize for on-demand
209 batch retrieval. The relational smells are listed as follows.

210 **(1) Eager Fetch [12, 17, 18]:** HIBERNATE preloads all fields annotated
211 with the EAGER FetchType when the data class is initialized, even if some
212 of them will never be accessed. To avoid performance issues (*i.e.*, retrieving
213 too much data in advance), data should be retrieved on demand.

214 **(2) Lacking Join Fetch [12, 18]:** Fields annotated with EAGER
215 FetchType should be joined by `join fetch` in HQL to be retrieved through
216 one query using `join`. Otherwise, such fields would be retrieved by N additional
217 queries if the parent object is initialized, resulting in the N+1 problem.

218 **(3) One-By-One [12, 17]:** A collection annotated with @OneToMany
219 or @ManyToMany using LAZY FetchType will be fetched one-by-one in every
220 loop iteration. @BatchSize should be involved to load data on demand and
221 in batch.

222 **(4) Missing ManyToOne [19]:** Using @OneToMany annotation in a field
223 without @ManyToOne presented on the other side of the relationship may also
224 lead to the N+1 problem.

225 3.3. Intra-Entity and Application Smells *Detected*

226 Entity smells are caused by inappropriate definition or application of
227 entity fields and methods. **Smells (5) to (13)** are summarized in [14], while
228 **Smell (14)** is described in [13]. The entity smells are listed as follows.

229 **(5) Collection Field:** Collection fields should use `Set` instead of `List`
230 due to performance concern, *e.g.*, an insertion after deletion in a `List` may
231 cause HIBERNATE to remove all the entities and re-insert them.

232 **(6) Final Entity:** Using `final` classes as entities would disable the proxy
233 functionality of HIBERNATE to enhance the performance of lazy loading.
234 Thus, the LAZY FetchType will fall back to EAGER, and no warning or error

235 message will be thrown by HIBERNATE. As a result, the performance will be
236 harmed silently.

237 **(7) Missing No Argument Constructor:** A no argument constructor
238 should be implemented for HIBERNATE to generate an entity object using
239 reflection, otherwise HIBERNATE will use Java reflection to initialize entities,
240 which will consume more resources. Moreover, if HIBERNATE is used as a
241 provider of JPA, it will throw an `org.hibernate.InstantiationException`,
242 and the application may crash since it may not be able to handle it.

243 **(8) Missing Identifier:** Identifier field should be specified to uniquely
244 determine an entity. Otherwise, comparators of objects may be confused
245 when dealing with data objects containing identical data from different rows.

246 **(9) Missing Equals Method:** The default `equals` method compares
247 the reference of objects, which is not ideal for comparing entities, especially
248 for collection-related operations. The lack of an appropriate `equals` method
249 will cause failure in reconnecting the detached entities, which may cause data
250 persistence problems such as duplication.

251 **(10) Missing hashCode Method:** `hashCode` is vital for collections
252 such as `HashSet`s to determine equivalent entities. The consequence of this
253 smell is similar to Missing Equals Method.

254 **(11) Using Identifier in Equals or hashCode Methods:** The iden-
255 tifier should not be used in `equals` and `hashCode` since all transient objects
256 may be equal because their identifiers could be null. The consequence of this
257 smell is similar to Missing Equals Method.

258 **(12) Not Serializable:** Entities which would leave the domain of JVM
259 (*i.e.*, detached for data export) should implement the `Serializable` inter-
260 face. Otherwise the serialization may fail, and Java will throw an exception
261 called `java.io.NotSerializableException`.

262 **(13) Missing Accessor Methods:** Although HIBERNATE does not
263 require accessor methods, JPA specification recommends implementing pub-
264 licly visible getters and setters to access and update private fields.

265 **(14) Local Pagination:** Built-in pagination of ORM should be used to
266 fetch the data of each page instead of fetching all data and locally split them
267 for pagination.

268 3.4. Mapping Metrics *Implemented*

269 The 4 Mapping metrics [10] are designed to evaluate data redundancy
270 and performance of entities related to inheritance. Thresholds to identify a
271 smell should be investigated further.

272 **Table Accesses for Type Identification (TATI):** The number of
273 tables needed to identify the requested type of entity. Higher TATI indicates
274 more queries will be executed to construct an object of the entity.

275 **Number of Corresponding Tables (NCT):** The number of tables
276 that contain data of an entity, which measures object retrieval performance.
277 The impact of higher NCT is similar to higher TATI.

278 **Number of Corresponding Relational Fields (NCRF):** The num-
279 ber of relational fields in all tables that correspond to each non-inherited
280 non-key field of an entity, which measures change propagation. Higher NCRF
281 indicates more queries will be executed to change data of an entity since there
282 exists data redundancy in terms of the relational fields.

283 **Additional Null Values (ANV):** The number of null values in the
284 row of union superclasses, which measures the data redundancy. Higher
285 ANV indicates more storage space is used to store null values since entities
286 affected by such a problem are likely to be stored with other entities (with
287 nonidentical fields) in the same table.

288 4. Implementation and Empirical Results

289 HBSNIFF could be executed as a command-line program under JDK
290 (Java Development Kit) version 1.8 and above with a line of command, *e.g.*,
291 `java -jar HBSniff-1.6.8.jar -i <projectRootPath> -o`
292 `<outputPath>`. The tool is shipped together with unit tests for all smells
293 and metrics implemented and documentations for both developers and users.
294 However, since it is also a tool for practical usage, we test it on real-world
295 projects.

296 4.1. Manual Validation of Detection Results

297 We perform smell detection over 5 open-source projects and 1 commercial
298 project (CP). The brief introduction of the projects are listed in Table 1. Prior
299 study [14] constructed a dataset of 77 projects for evaluation, however, the
300 authors found most of them were toy and example projects. To generate
301 a more practical dataset, we pick 3 non-toy projects from [3, 14] having
302 actual purpose and functionality. We also randomly pick the 4th and 5th
303 project by locating `createQuery` method calls using GITHUB search to find
304 projects performing potential HQL execution. The 6th project is used to
305 confirm if our tool can be used in a more realistic scenario. The smells
306 are manually validated by the 1st and the 5th author independently. The

Project	Purpose	Entities	I%	R%
WeixinMultiPlatform [14]	Content management system.	26	100.00	38.46
Jpa-issuetracker [14]	Development issue tracker.	6	100.00	16.67
Broadleaf Commerce [3]	E-commerce framework.	162	100.00	71.60
Devproof Portal ⁶	Blogging platform.	22	100.00	72.73
2ndInvesta ⁷	Invest management.	16	100.00	62.50
CP (Commercial)	Order processing.	27	100.00	77.78

Table 1: Projects analyzed. I% refers to entities affected by intra-entity smells. R% refers to entities affected by relational smells.

307 detection is all accurate and there is no missing case. The analyzed results
308 show that HIBERNATE-based projects are heavily affected by ORM smells,
309 which indicates the need of analyzing empirically the impact of these smells
310 to the quality of software in large-scale.

311 However, we fail to find an appropriate project for assessing the 4 metrics.
312 Nevertheless, we implement the HIBERNATE-based examples of the original
313 paper [10], which is also available with database generation code (ddl) in
314 the example folder in our source code. Note that since recent versions of
315 HIBERNATE do not support mixed inheritance strategy⁸, our implementation
316 is slightly different from the original paper. Finally, the 4 metrics are verified
317 by unit tests and manual evaluation of the sample project ⁹.

318 4.2. Impact of ORM Performance Smells

319 The smells we detect are either maintainability-related (*i.e.*, **smells (8)-**
320 **(13)**) or performance-related. Since the evaluation of maintainability impact
321 relies on empirical studies over large-scale dataset, it is not within the scope
322 of our work. However, we are still interested in the significance of the imple-
323 mented performance smells, which could be measured by benchmark. Thus,
324 we measure the impact of the **smells (1)-(7)** and **(14)** by retrieving or edit-
325 ing relational data in different scales (*i.e.*, 50, 500, and 50,000 instances of
326 relational entities).

327 We construct 3 entities (*e.g.*, A, B, C) for every smell. Entity A is the
328 parent, while B is a child of A, and C is a child of B. The parent entities

⁶<https://github.com/devproof/portal>

⁷<https://github.com/2ndStack/2ndInvesta>

⁸Mixing inheritance is not allowed. <https://hibernate.atlassian.net/browse/HHH-7181>

⁹<https://tinyurl.com/2as46jx8>

Smell	10 × 5 Instances			100 × 5 Instances			10000 × 5 Instances		
	Clean	Smelly	Impact	Clean	Smelly	Impact	Clean	Smelly	Impact
(1)	202	242	0.20	209	1457	5.97	304	130462	428.15
(2)	134	247	0.84	161	1469	8.12	692	123464	177.42
(3)	120	243	1.03	275	1474	4.36	2027	123190	59.77
(4)	122	240	0.97	268	1461	4.45	2013	124036	60.62
(5)	131	271	1.07	255	1561	5.12	507	140353	275.83
(6)	207	507	1.45	211	2743	12.00	344	139279	403.88
(7)	248	250	0.01	1484	1496	0.01	123465	125355	0.02
(14)	194	269	0.39	350	590	0.69	583	2791	3.79

Table 2: The performance of smelly and clean entities. The unit of the performance in the Clean and Smell columns is millisecond.

own the relations. Except for the foreign key field (`parent_id`), the entities have only 1 additional `MYSQL VARCHAR(255)` column called `name` mapped as a corresponding `String` field in `JAVA`. Every instance of Entity B owns 5 instances of entity C as children, and we alter the number of the children of entity A (*i.e.*, instances of entity B) in $\{10, 100, 10,000\}$ to measure the impact of smells to software systems with different data scale. We present a 3-level relation because we intend to recover a more practical scenario by using multi-level nested relation.

To ensure a comparable and fair performance score, each test is performed 10 times, and we re-execute the program rather than performing the queries in loops since `HIBERNATE` built-in cache cannot be disabled, and it may underestimate the impact. Afterwards, we present the medians of the performance data in Table 2. The clean column refers to the entity which is not affected by the smell concerned, while the smelly column refers to the entities affected by the smells. The impact column measures the differences between the two performance divided by the performance of clean entities.

The hardware environment of this experiment is consistent, *i.e.*, AMD Ryzen 7 4800H CPU, DDR4 16GB 2666MHz RAM, and 512GB SSD. We use `MYSQL` as the database. All queries are performed on `HIBERNATE` 5.4.31 and Java 11. To clarify, we do not test the performance impact of the combinations of smells. For example, in the evaluation of the **pagination smell (14)**, we implement `LAZY FETCH` with `BatchSize` specified (and thus the entity is not affected by the prior smells), and we retrieve the first page since it is more likely to be the most visited page. However, **smell (7) Miss-**

353 **ing No Argument Constructor** is an exception since we intend to find
354 out the impact of reflection-based constructor generation in more entities,
355 and thus we test this smell on **EAGER Fetch** entities.

356 From the impact column, we can easily conclude that performance smell
357 could impose more impact if more relational entities present in queries.
358 Meanwhile, except for **smell (7) Missing No Argument Constructor**,
359 almost all smells cause significant decline in performance if there are more
360 than 500 instances of relational entities. Moreover, even in the case of 50 in-
361 stances of instances of relational entities, the [decline of performance caused](#)
362 [by the smells](#) may also reach 20% to 145%. Thus, we believe such smells are
363 worth refactoring.

364 4.3. Comparison with Related Tools

365 **Difference with Respect to [14].** Prior study [14] proposed a design
366 rule checker which is capable of detecting 9 out of 10 intra-entity smells
367 (except for Local Pagination). However, the checker requires the analyzed
368 project to compile, and its upstream parser (**DESIGNWIZARD**¹⁰) provide full
369 support only for the class files compiled with JDK version less than or equal to
370 1.7. JDK 1.7 is no longer supported¹¹ by its manufacturer since April 2015.
371 Compared with **DESIGNWIZARD**, **JAVAPARSER** is a static analysis based
372 Java AST parser that updates weekly and supports the language features up
373 to the latest Java 15. Moreover, we compile 2 projects in the datasets in Table
374 1 to verify the results, and we fix some issues in its implementation, *e.g.*, we
375 can detect the cases of using identifier annotations in accessor methods, calls
376 of parent methods by **super** in **equals** and **hashCode**, and we do not treat
377 missing **equals** and **hashCode** as an occurrence of **smell (11)** since default
378 methods compare object references instead of attributes, and so on.

379 **The 3 Unimplemented Data Usage Smells Mentioned in Section**
380 **4.5 of [13].** The original source of the 3 smells [18] used static analysis to
381 locate method calls of queries, and analyzed the accessed data using dynamic
382 analysis. We do not implement them since static analysis is not able to profile
383 execution. However, we will extend our work to find trials of redundant
384 usage, *e.g.*, locating **findAll** and **update** operations of fetching a whole
385 entity or table. To achieve this goal, a large-scale empirical analysis should

¹⁰<https://github.com/joaoarthurbm/designwizard>

¹¹End of Public Updates Notice. https://java.com/en/download/help/java_7.html

386 be conducted to capture different forms of entity update. Moreover, we may
387 propose new data usage smells, which is not within the scope of this work.

388 *4.4. Remarks on Implementation*

389 **Exclusion of Controversial Smells.** We allow users to exclude every
390 smell in command-line parameter `-e` in case they do not perceive them as
391 real problems, *e.g.*, the impact of missing no argument constructor is almost
392 negligible.

393 **Drawbacks of Static Analysis.** Static analysis has its unavoidable
394 drawback since run-time information is not available. To cope with them,
395 we need to specifically implement solutions for every up-mentioned case to
396 detect the application of third-party plugins. For example, we cover the
397 usage of libraries such as LOMBOK¹² and APACHE COMMONS¹³ to generate
398 `equals`, `hashCode`, and accessor (`getter`, `setter`) methods. We may not
399 be able to detect similar usage if practitioners use other libraries.

400 **Detecting Pagination Smell.** We locate the `setMaxResults` or
401 `setFirstResult` method calls in parent methods with HQL appearance, and
402 we analyze if the code component that called the method defines any `Integer`
403 or `Long` object whose name contains “page” or “limit”. This complies with
404 the sample code of [13], which may be impractical, and may be improved in
405 future work.

406 *4.5. Research Opportunities*

407 Since RDBMS and software applications are different systems, there still
408 exist gaps (also known as “impedance mismatch” [29]) between them. ORM
409 aimed to address this problem, but they introduced new code quality issues
410 and more complexity. For practitioners, the quality of ORM code and the
411 appropriate application of ORM should always be considered to make more
412 reasonable refactoring plans. For researchers, more empirical studies could
413 be made to measure the impact of data persistence code quality to software
414 maintainability and reliability. To these ends, the research opportunities
415 brought by our tool include but not limited to:

- 416 • Evaluating ORM smells in large-scale datasets (*e.g.*, GitHub Mirrors
417 [30]) to study their occurrence, impact, and interactions with other code
418 smells and architectural issues;

¹²<https://projectlombok.org/>

¹³<https://commons.apache.org/>

- 419 • Integrating our command-line tool into the process of Continuous In-
- 420 tegration to generate reports after every development iteration of software
- 421 code, or extending our code to build an IDE-based detection tool such as
- 422 JDEODORANT [25] to provide just-in-time support for practitioners;
- 423 • Using our static analysis code base to develop detection methods for
- 424 new ORM-related smells;
- 425 • Revealing the losses and gains of ORM detaching. For example, can
- 426 we improve code maintainability and query performance by transferring to
- 427 “lightweight” data persistence solutions such as semiautomatic ORM frame-
- 428 works (*e.g.*, MyBatis¹⁴) and native SQL queries with manual data class map-
- 429 ping?

430 5. Illustrative Example

431 Figure 2 illustrates the exported `xls` report of the analyzed commercial
 432 project. Undetected smells are not presented. Fields in orange represents
 433 smelly, and texts in these fields are corresponding comments (*e.g.*, affected
 entity attributes). Light green fields refer to clean entities.

Class Name / Smell	CollectionField	Eager Fetch	Lacking Join Fetch	MissingEquals	MissingGetterSetter	MissingHashCode	NotSerializable	One-By-One	UsingIdInHashCode OrEquals
AbstractModuleConfiguration		state country isoCo			Missing Getter of <id>				
AdminModuleImpl	sections								Using ID <id> from eq
AdminPermissionImpl	qualifiedEntitiesalChick				Missing Getter of <typ				
AdminPermissionQualifiedEntityImpl		adminPermission							
AdminRoleImpl									
AdminSectionImpl	permissions	module			Missing Getter of <use				
AdminUserAttributeImpl		adminUser							Using ID <id> from eq
AdminUserImpl		overrideSandBox							
BandedPriceFulfillmentOptionImpl	bands							bands	
BandedWeightFulfillmentOptionImpl	bands							bands	
BankAccountPaymentImpl									
BroadleafCurrencyImpl					Missing Getter of <del				Using ID <id> from eq
BundleOrderItemFeePriceImpl		bundleOrderItem			Missing Getter of <arr				Using ID <id> from eq
BundleOrderItemImpl	discreteOrderItemsbusku productBundle							discreteOrderItems	Using ID <id> from eq
CandidateFulfillmentGroupOfferImpl		fulfillmentGroup offer			Missing Getter of <dis				Using ID <id> from eq
CandidateItemOfferImpl		orderItem offer			Missing Getter of <dis				Using ID <id> from eq
CandidateOrderOfferImpl		order offer			Missing Getter of <dis				Using ID <id> from eq
CatalogImpl	siteXrefssites				Missing Getter of <arc				Using ID <id> from eq
CategoryAttributeImpl		category							Using ID <id> from eq
CategoryExcludedS									

Figure 2: A snapshot of the generated EXCEL report for the [Broadleaf Commerce](#) project.

434

¹⁴<https://mybatis.org/mybatis-3/>

435 Listing 3 is an example of the JSON output of HBSNIFF. The available
436 information include names of smells (the `name` field), file paths (the `file`
437 `field`), names of the classes (the `className` field), positions of the detected
438 smells in the source code (the `position` field), and comments of the smells
439 (the `comment` field). Such data is also available in the `csv` output.

Listing 3: Example of The JSON output for the portal project

```
440 {  
441   "Article.java": [  
442     {  
443       "name": "UsingIdInHashCodeOrEquals",  
444       "file": "...", // File Path  
445       "position": "(line 36,col 1)-(line 193,col 1)",  
446       "className": "Article",  
447       "comment": "Using ID <id> from equals.  
448         Using ID <id> from hashCode. "  
449     },  
450     ... // Other Smells  
451   ],  
452   "Other Hibernate Entities":[  
453     ... // Other Smells  
454   ], ...  
455 }, ...  
456 }  
457 }
```

458 6. Conclusions and Future Work

459 We presented a static analysis-based Java HIBERNATE ORM code smell
460 detection tool called HBSNIFF which is capable for evaluating 14 smells and
461 4 mapping metrics in uncompiled Java project source codes. Moreover, we
462 conducted unit tests and manual verification for the detectors and metrics to
463 ensure the reliability of our implementation. We also evaluated the impact
464 of the implemented performance smells, and we suggested refactoring them
465 as soon as possible since most of them could greatly impact performance.

466 Future work includes: (1) proposing new ORM smells and improve the
467 existing implementations, (2) extending our scope to Python ORMs, and (3)
468 assessing the impact of ORM smells to architecture degradation and software
469 maintainability.

470 Acknowledgements

471 This work was partially supported by the National Natural Science Foun-
472 dation of China under Grant No. 61772200, and the Natural Science Foun-
473 dation of Shanghai under Grant No. 21ZR1416300.

474 References

- 475 [1] F. Palomba, A. Panichella, A. Zaidman, R. Oliveto, A. De Lucia, The
476 scent of a smell: An extensive comparison between textual and struc-
477 tural smells, *IEEE Transactions on Software Engineering* 44 (10) (2018)
478 977–1000.
- 479 [2] F. Palomba, G. Bavota, M. Di Penta, F. Fasano, R. Oliveto, A. De Lucia,
480 On the diffuseness and the impact on maintainability of code smells:
481 A large scale empirical investigation, *Empirical Software Engineering*
482 23 (3) (2018) 1188–1221.
- 483 [3] T.-H. Chen, W. Shang, J. Yang, A. E. Hassan, M. W. Godfrey,
484 M. Nasser, P. Flora, An empirical study on the practice of maintaining
485 object-relational mapping code in Java systems, in: *Proc. 13th Inter-
486 national Conference on Mining Software Repositories (MSR)*, 2016, p.
487 165–176.
- 488 [4] G. Vial, Lessons in persisting object data using object-relational map-
489 ping, *IEEE Software* 36 (6) (2019) 43–52.
- 490 [5] C. Nagy, A. Cleve, A static code smell detector for SQL queries embed-
491 ded in Java code, in: *Proc. IEEE 17th International Working Conference
492 on Source Code Analysis and Manipulation (SCAM)*, 2017, pp. 147–152.
- 493 [6] B. A. Muse, M. M. Rahman, C. Nagy, A. Cleve, F. Khomh, G. Antoniol,
494 On the prevalence, impact, and evolution of SQL code smells in data-
495 intensive systems, in: *Proc. of the 17th International Conference on
496 Mining Software Repositories (MSR)*, 2020, p. 327–338.
- 497 [7] F. Gonçalves de Almeida Filho, A. D. Forte Martins, T. da Silva Vin-
498 unto, J. M. Monteiro, Í. Pereira de Sousa, J. de Castro Machado,
499 L. Souza Rocha, Prevalence of bad smells in PL/SQL projects, in: *Proc.
500 IEEE/ACM 27th International Conference on Program Comprehension
501 (ICPC)*, 2019, pp. 116–121.

- 502 [8] A. Torres, R. Galante, M. S. Pimenta, A. J. B. Martins, Twenty years
503 of object-relational mapping: A survey on patterns, solutions, and their
504 implications on application design, *Information Software Technology* 82
505 (2017) 1–18.
- 506 [9] M. Lorenz, J.-P. Rudolph, G. Hesse, M. Uflacker, H. Plattner, Object-
507 relational mapping revisited: A quantitative study on the impact of
508 database technology on O/R mapping strategies, in: *Proc. 50th Hawaii*
509 *International Conference on System Sciences (HICSS)*, 2017, pp. 4877–
510 4886.
- 511 [10] S. Holder, J. Buchan, S. G. MacDonell, Towards a metrics suite for
512 object-relational mappings, in: *Proc. 1st International Workshop on*
513 *Model-Based Software and Data Integration (MBSDI)*, 2008, pp. 43–54.
- 514 [11] L. Meurice, C. Nagy, A. Cleve, Detecting and preventing program in-
515 consistencies under database schema evolution, in: *Proc. IEEE 16th*
516 *International Conference on Software Quality, Reliability and Security*
517 *(QRS)*, 2016, pp. 262–273.
- 518 [12] T.-H. Chen, W. Shang, Z. M. Jiang, A. E. Hassan, M. N. Nasser,
519 P. Flora, Detecting performance anti-patterns for applications developed
520 using object-relational mapping, in: *Proc. 36th International Conference*
521 *on Software Engineering (ICSE)*, 2014, pp. 1001–1012.
- 522 [13] S. Loli, L. Teixeira, B. Cartaxo, A catalog of object-relational mapping
523 code smells for Java, in: *Proc. 34th Brazilian Symposium on Software*
524 *Engineering (SBES)*, 2020, pp. 82–91.
- 525 [14] T. M. Silva, D. Serey, J. C. A. de Figueiredo, J. Brunet, Automated
526 design tests to check hibernate design recommendations, in: *Proc. 33th*
527 *Brazilian Symposium on Software Engineering (SBES)*, 2019, pp. 94–
528 103.
- 529 [15] V. Lenarduzzi, V. Nikkola, N. Saarimäki, D. Taibi, Does code quality
530 affect pull request acceptance? An empirical study, *Journal of Systems*
531 *and Software* 171 (2021) 110806.
- 532 [16] M. Tufano, F. Palomba, G. Bavota, M. Di Penta, R. Oliveto, A. De Lu-
533 cia, D. Poshyvanyk, There and back again: Can you compile that snap-
534 shot?, *Journal of Software: Evolution and Process* 29 (4) (2017) e1838.

- 535 [17] T.-H. Chen, Improving the quality of large-scale database-centric soft-
536 ware systems by analyzing database access code, in: Proc. 31st IEEE
537 International Conference on Data Engineering Workshops (ICDEW),
538 2015, pp. 245–249.
- 539 [18] T.-H. Chen, W. Shang, Z. M. Jiang, A. E. Hassan, M. Nasser, P. Flora,
540 Finding and evaluating the performance impact of redundant data ac-
541 cess for applications that are developed using object-relational mapping
542 frameworks, *IEEE Transactions on Software Engineering* 42 (12) (2016)
543 1148–1161.
- 544 [19] P. Węgrzynowicz, Performance antipatterns of one to many association
545 in hibernate, in: Proc. 2013 Federated Conference on Computer Science
546 and Information Systems (FedCSIS), 2013, pp. 1475–1481.
- 547 [20] L. Meurice, C. Nagy, A. Cleve, Static analysis of dynamic database usage
548 in Java systems, in: Proc. 28th International Conference on Advanced
549 Information Systems Engineering (CAiSE), pp. 491–506.
- 550 [21] G. Procaccianti, P. Lago, W. Diesveld, Energy efficiency of ORM ap-
551 proaches: An empirical evaluation, in: Proc. 10th ACM/IEEE Interna-
552 tional Symposium on Empirical Software Engineering and Measurement
553 (ESEM), ACM, 2016, pp. 36:1–36:10.
- 554 [22] T.-H. Chen, W. Shang, A. E. Hassan, M. N. Nasser, P. Flora, Cacheop-
555 timizer: Helping developers configure caching frameworks for hibernate-
556 based database-centric web applications, in: Proc. 24th ACM SIG-
557 SOFT International Symposium on Foundations of Software Engineering
558 (FSE), 2016, pp. 666–677.
- 559 [23] R. Singh, C. Bezemer, W. Shang, A. E. Hassan, Optimizing the
560 performance-related configurations of object-relational mapping frame-
561 works using a multi-objective genetic algorithm, in: Proc. 7th
562 ACM/SPEC International Conference on Performance Engineering
563 (ICPE), 2016, pp. 309–320.
- 564 [24] M. F. C. Nazário, E. Guerra, R. Bonifácio, G. Pinto, Detecting and
565 reporting object-relational mapping problems: An industrial report, in:
566 Proc. ACM/IEEE 13th International Symposium on Empirical Software
567 Engineering and Measurement (ESEM), 2019, pp. 1–6.

- 568 [25] N. Tsantalis, T. Chaikalis, A. Chatzigeorgiou, Ten years of JDeodorant:
569 Lessons learned from the hunt for smells, in: Proc. IEEE 25th Inter-
570 national Conference on Software Analysis, Evolution and Reengineering
571 (SANER), 2018, pp. 4–14.
- 572 [26] F. Palomba, D. A. Tamburri, F. Arcelli Fontana, R. Oliveto, A. Zaid-
573 man, A. Serebrenik, Beyond technical aspects: How do community
574 smells influence the intensity of code smells?, IEEE Transactions on
575 Software Engineering 47 (1) (2021) 108–129.
- 576 [27] M. F. Aniche, G. Bavota, C. Treude, M. A. Gerosa, A. van Deursen,
577 Code smells for model-view-controller architectures, Empirical Software
578 Engineering 23 (4) (2018) 2121–2157.
- 579 [28] C. Nagy, L. Meurice, A. Cleve, Where was this SQL query executed?
580 A static concept location approach, in: Proc. IEEE 22nd Interna-
581 tional Conference on Software Analysis, Evolution, and Reengineering
582 (SANER), 2015, pp. 580–584.
- 583 [29] W. R. Cook, R. Greene, P. Linskey, E. Meijer, K. Rugg, C. Russell,
584 B. Walker, C. Wittig, Objects and databases: State of the union in 2006,
585 in: Companion to the 21st ACM SIGPLAN Symposium on Object-
586 Oriented Programming Systems, Languages, and Applications (OOP-
587 SLA), 2006, p. 926–928.
- 588 [30] R. Dyer, H. A. Nguyen, H. Rajan, T. N. Nguyen, Boa: Ultra-large-
589 scale software repository and source-code mining, ACM Transactions
590 on Software Engineering and Methodology 25 (1) (2015) 7:1–7:34.

591 **Required Metadata**

592 **Current executable software version**

Nr.	(executable) Software metadata description	Please fill in this column
S1	Current software version	v1.6.8
S2	Permanent link to executables of this version	<i>https : //github.com/HBSniff/HBSniff/releases/tag/v1.6.8</i>
S3	Legal Software License	GPL
S4	Computing platform/Operating System	Linux, OS X, Microsoft Windows.
S5	Installation requirements & dependencies	JDK 8.0
S6	If available, link to user manual - if formally published include a reference to the publication in the reference list	https://hbsniff.github.io/
S7	Support email for questions	hzj@mail.ecust.edu.cn

Table 3: Software metadata (optional)

593 **Current code version**

Nr.	Code metadata description	Please fill in this column
C1	Current code version	v1.6.8
C2	Permanent link to code/repository used of this code version	<i>https : //github.com/HBSniff/HBSniff</i>
C3	Legal Code License	GPL
C4	Code versioning system used	git
C5	Software code languages, tools, and services used	Java
C6	Compilation requirements, operating environments & dependencies	JDK 8.0, Maven 5
C7	If available Link to developer documentation/manual	https://hbsniff.github.io/
C8	Support email for questions	hzj@mail.ecust.edu.cn

Table 4: Code metadata (mandatory)