RESEARCH ARTICLE - EMPIRICAL

Software: Evolution and Process WILEY

On the relation between architectural smells and source code changes

Darius Sas[1] [ID] | Paris Avgeriou[1] | Ilaria Pigazzini[2] | Francesca Arcelli Fontana[2]

[1]Bernoulli Institute for Mathematics, Computer Science, and Artificial Intelligence, University of Groningen, Groningen, The Netherlands

[2]Department of Informatics, Systems, and Communications, University of Milano-Bicocca, Milan, Italy

Correspondence
Darius Sas, University Of Groningen, University of Groningen Faculty of Science and Engineering (FSE), Bernoulliborg, Nijenborgh 9, Groningen 9747 AG, The Netherlands.
Email: d.d.sas@rug.nl

Abstract

Although architectural smells are one of the most studied type of architectural technical debt, their impact on maintenance effort has not been thoroughly investigated. Studying this impact would help to understand how much technical debt interest is being paid due to the existence of architecture smells and how this interest can be calculated. This work is a first attempt to address this issue by investigating the relation between architecture smells and source code changes. Specifically, we study whether the *frequency* and *size* of changes are correlated with the presence of a selected set of architectural smells. We detect architectural smells using the Arcan tool, which detects architectural smells by building a dependency graph of the system analyzed and then looking for the typical structures of the architectural smells. The findings, based on a case study of 31 open-source Java systems, show that 87% of the analyzed commits present more changes in artifacts with at least one smell, and the likelihood of changing increases with the number of smells. Moreover, there is also evidence to confirm that change frequency increases after the introduction of a smell and that the size of changes is also larger in smelly artifacts. These findings hold true especially in Medium–Large and Large artifacts.

KEYWORDS

architectural smells, architectural technical debt, empirical study, software repository mining, technical debt, technical debt interest

1 | INTRODUCTION

Architectural smells (ASs) are defined as "commonly-used (although not always intentional) architectural decisions that negatively impact system quality."[1] ASs manifest themselves in the system as undesired dependencies, unbalanced distribution of responsibilities, and excessive coupling between components and in many other forms that break one or more software design principles and good practices, ultimately affecting maintainability and evolvability.[2] We note that the presence of AS does not always inevitably indicate that there is a problem, but it points to places in the system's architecture that should be further analyzed.[2] ASs are considered as a type of architectural technical debt (ATD), as they (may) result in increased complexity and "can make future changes more costly or impossible."[3] The interest of the research community in AS has grown exponentially over the past years: According to a systematic mapping study by Verdecchia et al,[4] they are one of the most studied types of ATD.

Research work on AS has ranged from broad studies that define new smell types and study their evolution over time[5–7] to more specific ones that focus on a particular architecture style (e.g., AS in systems built with Model-View-Controller [MVC] or Microservices[8]). Few studies, however,

have extensively investigated the impact of AS on maintenance effort. Although ASs are considered detrimental to software maintenance, forcing developers to pay high technical debt interest,[1] there is little empirical evidence to explore and confirm this phenomenon. Although there has been research on the impact of code smells on maintenance effort, ASs seem completely independent from code smells[9] and arguably more severe.

This study addresses this gap by exploring the impact of a specific set of AS on maintenance effort in terms of the actual changes made by developers to the source code. Specifically, we compare the *frequency* and *size* of changes between source code artifacts affected and not affected by ASs. We perform the comparison both by controlling for the size of the artifacts and without any control for size, to eliminate size as a confounding factor. We consider *change frequency*, that is, the number of times an artifact was changed across multiple versions, and *change size*, that is, the number of lines of code (LOC) added, deleted, and modified,[*] as proxies of the effort spent, based on previous work: Change frequency is a factor that was found to affect maintenance effort,[10,11] whereas change size (also referred to as code churn or Total Amount of Changes (TACH)—see Section 4) was used to estimate the effort in previous studies.[12,13] This can give an *indication* of how much technical debt interest (rather than the actual interest per se) is paid by developers due to the presence of the detected smells (not all changes entail paying interest—see Threats to Validity section). Furthermore, our findings can be used towards building a model to calculate, based on actual changes, the ATD interest[3] paid when maintaining artifacts affected by AS.

The ASs considered in this study are Cyclic Dependency (CD), Hub-Like Dependency (HL), Unstable Dependency (UD), and God Component (GC).[2,5,14] We selected to study these smells as they are some of the most prominent architecture smells, and there already exist tools that provide their automatic detection.

The novel contributions of this study are (1) the vast majority of related work examines code smells, while we focus on ASs, which were found to be independent from code smells;[9] (2) we study four different types of AS, and only CD was previously investigated by other studies, whereas the other three were overlooked; (3) we provide a new, interesting view of how AS affect artifacts before and after the introduction (RQ2).

The rest of the paper is structured as follows: Section 2 summarizes similar work from the literature; Section 3 describes in detail the goals, research questions, and the selected projects of this study; Section 4 reports the data collected as well as the collection process; Section 5 presents the data analysis procedures; Section 6 reports and examines the obtained results; Section 7 discusses our interpretation of the results and compares them with similar findings from the literature; Section 8 enunciates the threats to the validity of this study; and finally, Section 9 concludes the paper and considers possible future work.

2 | RELATED WORK

2.1 | Impact of ASs

In a recent work, Le et al[15] defined a set of six ASs based on an automated reverse architecture model extraction. Next, they investigated whether files affected by ASs (i.e., smelly files) are more likely to have issues (extracted from issue-tracking systems) associated to them than clean files. Additionally, they also checked if smelly files are more change-prone than clean files. The case study was performed on eight different open-source Java systems, and the results confirmed that smelly files are more fault- and change-prone in the eight systems analyzed. Contrary to the work of Le et al,[15] in our work, we investigate a different set of ASs based on concrete software artifacts, rather than on architectural recovery views; we use 31 projects, rather than eight, and we measure several facets of change-proneness (not only the number of commits a file has changed), using a well-established suite of metrics.

Oyetoyan et al[6] have studied the relation between CD and the change frequency of the affected classes near them on twelve Java open-source systems. They investigated both general CD between classes and special kinds of CD (e.g., cycles that contain both parents and children classes, abbreviated as STK, and cycles across branches of the package containment tree) that have been conjectured to be particularly undesirable. Their results show that the presence of cycles does increase the change frequency of the classes affected and of the neighbor classes, but this is not true for classes affected by STK cycles in most of the systems considered. Moreover, their findings also suggest that classes belonging to cycles spanning across branches of the package containment tree (the tree of the packages) do not exhibit a higher correlation with change frequency. Our work differs from this study in the following aspects: We investigate four types of smells, including CD, both at class and package level; our data include more systems and more commits per system; and we use multiple well-established metrics to measure change.

2.2 | Impact of antipatterns, design patterns, and design smells

Khom et al[16] investigated the effect of antipatterns (classes that embody poor design choices and stem in-between design and implementation[16]), on class and change proneness. More specifically, the authors investigated whether classes participating in antipatterns have a higher likelihood

---

*See Section 4 for a full description.

than others to change or be involved in issues documenting faults. Their study focused on four open-source Java systems and a total of 54 releases. Their findings confirmed that classes participating in antipatterns are more change prone than others. Specifically, the *MessageChain* antipattern has been found to consistently have the greatest impact on change proneness across all the four systems analyzed. The impact of the other antipatterns largely depends on the studied system. Concerning fault proneness, the results are very similar to what was observed for change proneness.

Another work on antipatterns and their relation with changes and faults was published by Jaafar et al.[17] In their work, rather than focusing on problematic classes as previous studies, they focused on classes that depend upon classes affected by antipatterns and/or participate in design patterns. Their work focused on six design patterns and 10 antipatterns detected throughout 39 releases of four systems. The findings indicate that classes having dependencies with antipatterns are more prone to fault, whereas this is not always true for classes with dependencies with design patterns. Additionally, the findings also show that classes depending upon antipatterns are more prone to logic faults and structural changes, whereas classes depending on design patterns are more prone to code addition and syntax faults.

Sharma et al[18] conducted an empirical study to investigate the relationship between design and AS in C# projects. Their distinction between design and architecture smells corresponds to our distinction between class-level and package-level smells. In their work, Sharma et al. studied the existence of a possible correlation between design and architectural smells to determine whether they capture the same concept. More specifically, they focused on collocation and causation, by investigating the temporal relationship between design and architectural smells to find out whether one type of smell causes the other. The results of the analysis show evidence of the individuality and uniqueness of design smells with respect to architectural smells.

2.3 | Impact of code smells

Aniche et al[19] have studied the impact of code smells on change and fault-proneness in MVC architectures prior to performing a qualitative analysis involving the developers of the 120 projects they considered. The projects were automatically extracted from GitHub, and the authors defined a set of smells specifically tailored for the MVC architecture by surveying 53 developers. The results concerning change and fault proneness show that classes affected by smells are more prone to change than non-smelly classes; traditional smells seem to have a stronger negative impact, although when controlling for size the difference is less marked on change proneness. No impact was observed on fault proneness when controlling for size for both MVC-specific and traditional smells.

Another study on code smells and change-proneness was done by Khomh et al.[20] In their work, the authors study the impact of 29 code smells on change proneness in two open-source Java projects. More precisely, they investigate whether smelly classes are more change-prone, how the number of smells influences this aspect, and differences in this impact between the different smell types. Their findings show that smelly classes are in fact more change prone in both projects analyzed. Additionally, they also show that a higher number of smells often imply a high change proneness. They also found that *HasChildren*, *MessageChainClass*, *NotAbstract*, and *NotComplex* smell types have the highest change proneness, but this is heavily project-dependant.

This study, in contrast, focuses on ASs, and as it was found in a previous study, ASs are independent from code smells.[9] Moreover, ASs, contrary to code smells, affect multiple classes and/or packages, have complex structures (e.g., dependencies among the affected components), and require large refactorings in order to be removed.[2] This means that research on code smells is not applicable to ASs, and the only similarity with code smells in this regard is that each type of AS needs to be investigated individually.

3 | CASE STUDY DESIGN

The present study is designed and reported following the guidelines published by Runeson et al.[21] Specifically, the case study design follows an embedded multiple-case format: multiple cases, each having numerous units of analysis, as shown in Figure 1. The individual source code files and packages analyzed for a given project constitute the units of analysis; the projects represent the cases. The domain of the project (e.g., web service and database) is the context, containing one or multiple cases.

3.1 | Terminology

In the next sections, we will use the term *change frequency* to indicate the number of times an artifact undergoes any kind of change in a given number of commits. For example, if an artifact changes in 3 commits out of the 100 considered, its change frequency is .03.

The term *change size* refers to the sum of the number of LOC added, deleted, and/or modified to/from an artifact in a single given commit. This is commonly referred to as *code churn*. A formal definition of how we measure changes is provided in Section 4.
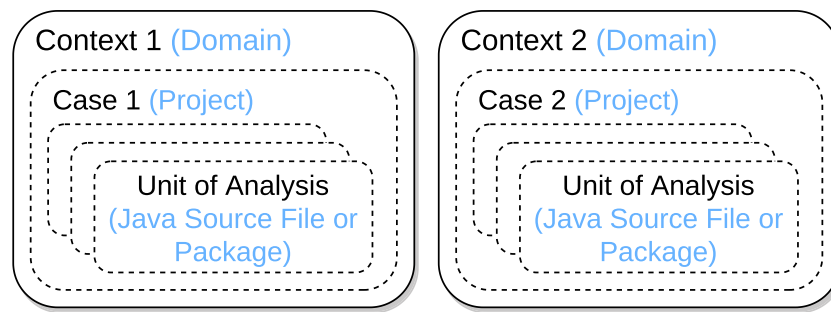
FIGURE 1 The case study design using Runeson et al.'s representation[21]

Finally, we note that we use the terms *commit* and *version* interchangeably. Additionally, the term *release* is used when a certain commit/version is explicitly packaged and tagged for public release.

3.2 | Goal and research questions

The goal of this study is to understand the impact of ASs on source code changes. Using the Goal-Question-Metric approach,[22] the goal can be formulated as follows:

> **Analyse** changes in source code artifacts **for the purpose of** understanding the impact of architectural smells **with respect to** the frequency and size of those changes **from the point of view of** software developers and architects **in the context of** open source Java software systems.

By (Java) source code artifacts, we mean both source code files (classes) and source code packages.
The goal can be broken down into three main research questions, as follows.

RQ1 Do classes and packages with smells change more frequently than classes and packages without smells?
    RQ1a Do different smell types have a different impact on frequency of change?
    RQ1b Does the number of smells have a different impact on frequency of change?

We ask this question to shed some light on the actual relationship between the existence of ASs and the *change frequency* of classes and packages. Such a relationship, in case it exists, confirms that ASs' presence correlates with increased maintenance effort, with respect to the frequency of changes of the affected artfacts.

The two subquestions, RQ1a and RQ1b, further explore the connection between ASs and change frequency by looking at how different smell types and multiple smells correlate to changes.

RQ2 What is the difference in the change frequency of an artifact before and after a smell is introduced?
    This question aims at identifying whether the introduction of a smell impacts the *change frequency* of a certain artifact. More precisely, it provides insights on whether the presence of the smell can be related to an increased change frequency in an affected artifact. Theoretically, one would expect that the introduction of a smell leads to an increase in the change frequency in (at least some of) the artifacts affected by the smell. Finally, the results of this research question, in case we do find evidence of such an increase, will strengthen the findings of RQ1.

RQ3 Is the size of the changes in source code artifacts affected by smells, larger than in nonaffected artifacts?

This question focuses on the magnitude, or *size,* of the changes made (in terms of added, deleted, and changed LOC) to the artifacts that are affected by smells. Theoretically, these artifacts should exhibit bigger changes (thus more complex ones) because working on an suboptimal design is harder and thus requires changing more LOC to be maintained. Bigger changes, in most scenarios (e.g., fixing bugs, adding features, and refactoring), mean developers have spent more time to implement them, resulting in a higher amount of interest paid.[12,13]

We emphasize that, with these research questions we are **not seeking** to establish causality between smells and changes by any means, but rather we aim at investigating correlations. This is further explained in the Section 7.

Finally, a replication package, containing the protocol, the data, the R scripts, and a collection of 14 plots that visualize the data, is available online[†].

3.3 | Analyzed projects

To conduct our study, we selected the 31 projects listed in Table 1. The inclusion criteria used during the selection of the projects were as follows:

1. Nontrivial Java projects with at least 10.000 LOC in the last commit;
2. Actively maintained and used by the community (the Contributors page on GitHub should show a consistently active development[‡]);
3. At least 3 years of active development on GitHub (or similar sites).

During the selection process, we also strove to diversify the domains of the included systems as much as possible, as indicated in Table 1, as well as to increase as much as possible the period of analysis taken into consideration. To this end, our dataset contains 31 projects, with an average period of analysis of 11.5 years, a maximum of 22.1 years, and a minimum of 3.5 years with an average of 126.8 commits analyzed per project. Figure 2 reports the distribution of the total number of LOC of the commits analyzed for each project.

4 | DATA COLLECTION

For every system S listed in Table 1, we analyzed one commit (or version) v every 4 weeks, from the first commit available in the repository to the latest on the main branch (either master or trunk). We selected a 4-week-long interval between each commit because we wanted to ensure that the change-related metrics we selected were calculated at a meaningful level of granularity, allowing enough files to change from one commit to the next one. Such custom intervals were used in similar contexts by previous studies.[23–25] Additionally, a fixed interval between commits avoids the introduction of bias and ensures the results are consistent across the different release rates of our projects.[24] The selection of the 4-week-long interval is further discussed in the Threats to Validity section. The change-related data were extracted using `git diff` between each pair of consecutive commits. The period of analysis started from the first ever commit available on the repository to the last one available as of May 2021. Next, as part of our data cleaning process, we removed the commits with no changes at the beginning and ending of a project, as these entail inactive leading and trailing periods.

For every artifact x, namely, class c or package p, in each commit v, we collected the following *independent* variables: (1) a boolean variable denoting whether x was affected by ASs or not, (2) four boolean variables indicating whether a certain type of smell affects x, (3) and an integer variable counting the total number of smell instances per smell type that affected it. We also measure for every artifact x the changes in the system using a well-established suite of metrics provided by Elish and Al-Khiaty[26] —these are the *dependant* variables in our study:

1. Change Has Occurred (CHO). This metric is the basis for calculating the *change frequency* of an artifact. *CHO* measures whether a class c, or package p, has changed or not in the current commit v with respect to the previous commit $v - 1$ in the dataset:

$$CHO_v(c) = \begin{cases} 1 & \text{if } c \text{ has changed in } v \\ 0 & \text{otherwise} \end{cases} \qquad CHO_v(p) = \bigvee_{c \in p}^{p} CHO(c).$$

Note that to calculate *CHO* for a package p (i.e., right-most formula), we do a binary sum (i.e., binary *OR*) between all the elements c directly contained in p.

2. Percentage of Commits a Class has Changed (PCCC). This metric computes the *change frequency* of an artifact using *CHO* and is represented as a percentage to normalize it. The metric was described and used in previous studies[27,28] and is basically the *FRCH* metric defined by Elish and Al-Khiaty[26] but normalized as a percentage.

$$PCCC_b^a(x) = \frac{\sum_{v=a}^{b} CHO_r(x)}{b - a} \times 100,$$

TABLE 1 Demographics of the projects analyzed in this study

Category	Project	# Commits analyzed	First commit	Last commit	KLOC 1st–last commit	Description
Data storage and management	accumulo	99	December 23, 2011	November 1, 2019	193–237	Data Storage System
	calcite	81	November 23, 2014	May 21, 2021	109–187	Dynamic Data Management
	cassandra	136	April 10, 2009	November 4, 2019	36–178	Distrib. NoSQL database
	chukwa	73	October 31, 2008	April 1, 2019	8–31	Data Collection
	jackrabbit	155	December 24, 2006	November 4, 2019	94–241	Content Repository
	jackson	93	February 6, 2012	November 5, 2019	31–59	Data Binding Library
Web engines and web tools	httpcomp.	126	February 9, 2006	October 3, 2019	0–33	HTTP Toolset
	jspwiki	186	August 25, 2001	November 1, 2019	1–32	Wiki Engine
	retrofit	51	June 1, 2015	June 18, 2020	3–10	Android HTTP client
	spring-boot	47	October 31, 2017	May 31, 2021	91–143	Spring-based project manager
	struts	158	April 24, 2006	November 4, 2019	24–41	Web Apps Framework
Search engines	elasticsearch	49	July 16, 2015	March 26, 2019	295–614	Search engine
	jena	95	June 1, 2012	November 17, 2019	209–348	Semantic Web
	lucene	173	October 20, 2001	August 3, 2015	5–453	Search Engine
	tika	144	August 17, 2007	November 2, 2019	2–63	Content Analysis Toolkit
Development tools	ant-ivy	130	July 15, 2005	November 2, 2019	10–42	Dependency Manager
	jenkins	186	December 17, 2006	May 29, 2021	14–125	Automation server
	jgit	123	November 17, 2009	November 17, 2019	18–113	Java implementation of Git
	selenium	132	February 16, 2011	May 30, 2021	2–53	Automation web libraries
	testng	147	September 21, 2006	October 21, 2019	13–59	Testing Framework
Document manipulation	pdfbox	137	August 17, 2008	November 3, 2019	26–82	PDF Library
	poi	206	March 24, 2002	November 1, 2019	19–99	MS Office API
	xerces2	189	January 3, 2000	July 15, 2019	36–116	XML Library
JDBC drivers	druid	99	June 23, 2011	November 3, 2019	42–85	Alibaba JDBC Library
	pgjdbc	211	September 29, 1997	November 4, 2019	2–30	PostgreSQL JDBC driver
Networking and messaging	activemq	161	January 25, 2006	November 6, 2019	61–177	Message Server
	mina	120	March 22, 2005	June 18, 2019	6–28	Network Framework
Game engine	libgdx	139	April 22, 2010	May 30, 2021	23–222	Game engine
Data binding	fastjson	104	September 14, 2011	April 5, 2021	12–41	Alibaba JSON data mapper
	gson	99	September 27, 2008	May 14, 2021	6–10	Google JSON data mapper
Utility	guava	83	January 11, 2010	July 25, 2019	33–117	Google Core Library

Note: Dates refer to the period of analysis taken into consideration, not age of the system. Additionally, the categories are only indicative.
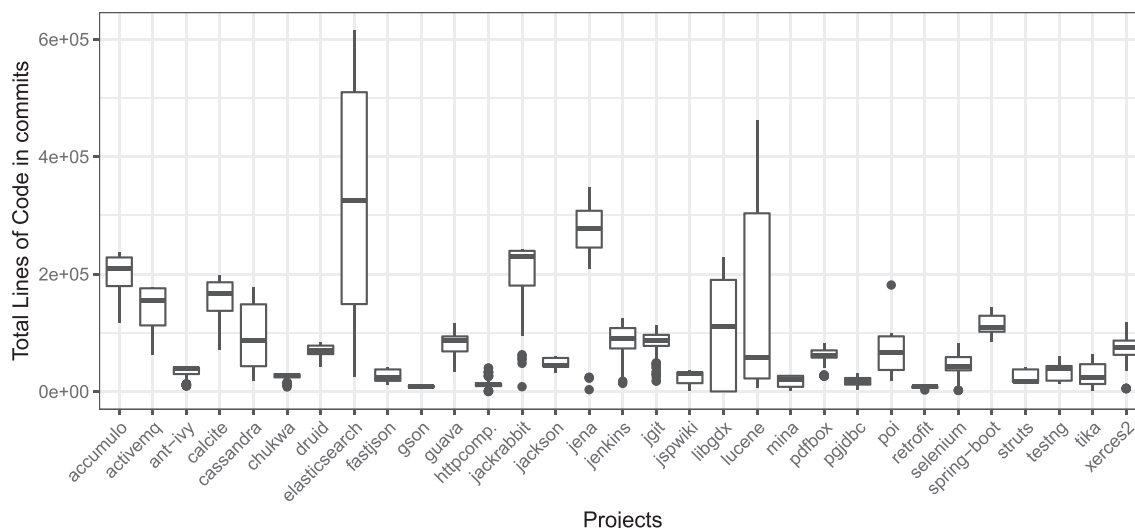
FIGURE 2 The distribution of the total number of lines of code of each version for each project

where v is the commit for which CHO is computed and $[a, b]$ is the interval of commit indexes considered. Intuitively, this metric counts the number of commits where an artifact has undergone changes and divides it by the number of commits in the period considered.

3. TACH. Also called *change size*, or *code churn*, is the sum of added LOC (*NAL*), deleted lines (*NDL*), and twice the changed lines (*NCL*) because the last commit[26] for a given class c or package p:

$$TACH(c) = NAL(c) + NDL(c) + 2 \times NCL(c) \quad TACH(p) = \sum_{c \in p}^{p} TACH(c).$$

The calculation of *TACH* for packages is simply the sum of *TACH* for each class c directly contained in p.

To collect the data, we used a combination of two tools: Arcan[5] and ASTracker.[14] Arcan collected the artifacts affected by ASs in each selected commit in the history of the systems directly from the source code files. The output of Arcan is a graph file containing the dependency network of the commit analyzed, including the smells detected. The algorithms used to detect ASs are explained in detail by Fontana et al. in their paper.[5] The detection is based on the software design principles reported by Martin et al[29] and Martin Lippert.[2] In short, CD is detected using a Depth-First Search algorithm that visits all the nodes in the dependency graph while checking, which were already visited. UD is detected using Martin's Instability metric:[30] If the *majority* of a package's dependencies are less stable than itself, then it is marked as an UD smell. HL is detected by simply looking at the number of incoming and outgoing dependencies a certain artifact has: If the sum of these dependencies surpasses a certain system-based threshold, then the artifact is marked as a hub. Finally, GC is detected using an automatically calculated variable threshold[31] using the distribution of the total amount of LOC of the packages in a benchmark of over 100 systems; the packages in the analyzed system are then compared with this threshold and the artifacts surpassing it are marked as GCs.§

Arcan's results were validated in different studies. A first validation of the results of Arcan was performed on two open-source projects with a precision of 100%.[5] Next, the results of Arcan were also validated in an industrial setting by two different studies: first on industrial C/C++ projects obtaining 50% precision[32] and then on industrial Java projects obtaining 70% precision.[33] The precision metric was chosen as the main indicator of Arcan's performance because the true positive rate was found to be the main concern for developers during the mentioned studies.

The second tool we used, ASTracker, computed the abovementioned change metrics and identified the elements affected by each smell. ASTracker's main feature is to track ASs from one version to the next (i.e., link the same instances detected in two adjacent versions), but for this study, it was only used to calculate the change metrics as stated above. To guarantee the correctness of the implementation of the change metrics, we used thorough unit testing.

§See https://fse.studenttheses.ub.rug.nl/19603/ for more details.

At last, the Peregrine high-performance computing cluster, offered by the University of Groningen, provided the computational power necessary to carry out the whole data collection process.

5 | DATA ANALYSIS

5.1 | Controlling for size

Changes to source code files are intuitively more frequent in files of greater size (i.e., more LOC). In fact, source code size has been empirically found to interfere with the actual findings in several cases.[34,35] Thus, source code size is a confounding factor in our analysis that could skew the results unpredictably and obfuscate the impact of smells on change frequency and size. To mitigate this threat, as already mentioned in the Sections 1 and 2, the data analysis will include controlling for size. Specifically, we will analyze the data both by considering all artifacts (without controlling for source code size) and by grouping the artifacts (either classes or packages) into four size groups, based on their effective LOC. This way we can compare how smells impact files of similar size. The groups are defined as follows: Small $= [1, Q_1)$, Medium–Small (M–Small) $= [Q_1, Q_2)$, Medium–Large (M–Large) $= [Q_2, Q_3)$, and Large $= [Q_3, Q_4]$, where Q_1, Q_2, Q_3, and Q_4 are the first, second, third, and fourth quartiles, respectively, of the distribution of the LOC of classes (or packages, when working with smells affecting packages) in a given project. This means that these values differ for each project. Table 2 shows the quartiles of the LOC distribution in the whole dataset.

This approach was proposed by Aniche et al. in a previous study.[19] We adopted it as it allows us to compare smelly and non-smelly artifacts with comparable size. This method guarantees that all four groups have the same number of files, which is an important prerequisite to ensure that the results of the study are not skewed. Indeed, if we were to partition the files, for example, with a range of 45 LOC per group, the resulting small group (0–45 LOC) would have 200K+ artifacts, whereas the others just a few thousands. This imbalance would greatly affect the outcome.

5.2 | RQ1—Do classes and packages with smells change more frequently than classes and packages without smells?

For this RQ, we statistically analyze the significance of the association between changes in affected and nonaffected artifacts. The Fisher's exact test of independence[36] is performed on two categorical variables: In our case, these variables are CHO and whether this artifact is affected by a smell. The input to the test is a contingency table where all the possible values of the two (categorical) variables are listed on the rows and columns of the table, respectively. The null and alternative hypotheses of the tests (one test for each 4-month-long period considered) are as follows:

- Null hypothesis H_0^{RQ1}: Artifacts affected by smells are _equally_ likely to be subject to changes than artifacts not affected by smells ($\pi_1 = \pi_2$)
- Alt. hypothesis H_1^{RQ1}: Artifacts affected by smells are _more_ likely to be subject to changes than artifacts not affected by smells ($\pi_1 > \pi_2$)

where π_1 and π_2 represent the proportions of the two categories with respect to the overall population.

To ensure the test is _supported_, we need to make sure that the proportions in the contingency tables used to run the tests are not excessively unbalanced towards one category. Contingency tables are likely to be unbalanced if the time period is too small because only limited changes can happen in a certain amount of time and that time cannot be enough to determine whether the correlation is present or not. In other words, given that there are more nonchanging files than changing files, a period of 1 month is likely to be insufficient for enough files to change. Thus, we aggregated our data to a 4-month granularity (rather than 1month); this is approximately the average release rate we mined from the Git tags of our projects. We call these "versions" _pseudo-releases_. Thus, for each pseudo-release _v_, we test for the null hypothesis, namely, whether there is no statistical difference in the proportions of changes for artifacts affected and not affected by ASs. This analysis will include all types of smells, both at class and package level, detected by Arcan.

TABLE 2 Distribution of the lines of code metric in classes and packages in the whole dataset

	0%	25% (Q_1)	50% (Q_2)	75% (Q_3)	100% (Q_4)
Class	1	10	27	77	14,990
Package	1	549	1340	2976	59,074

The next step is to compare the percentages of pseudo-releases that do show a significant difference (accepting H_1^{RQ1}) and the pseudo-releases that do not show any significant difference (accepting H_0^{RQ1}), which will allow us to answer RQ1. Note that we opted to perform one test per commit per project, rather than one test per project, to ensure that the imbalance in changes detected is not the result of a few change hot-spots throughout the history of the system but rather a more constant phenomenon. The confidence level used for this test and all the following tests is equal to $\alpha = .05$.

5.2.1 | RQ1a—Do different smell types have a different impact on frequency of change?

In order to answer RQ1a, we used a logistic regression model.[36] This kind of model allows to predict the value of a binary dependant variable given a set of multiple independent variables. Moreover, it can be exploited to compute the effect size between the dependant variable and each independent variable, to identify which variable influences the outcome. In this case, we chose the CHO metric (see Section 4) as dependant variable and the number of smell instances of each smell type t as independent variables.

The hypotheses of this analysis are as follows:

- Null hypothesis H_0^{RQ1a}: The type of smells does not have an impact on the occurrence of changes of artifacts.
- Alt. hypothesis H_1^{RQ1a}: The type of smells does have an impact on the occurrence of changes of artifacts.

The analysis was performed individually for each 4-month commit period, for all projects. Then, for each type of smell, we counted the number of times that the p values obtained by the logistic regression were significant.

5.2.2 | RQ1b—Does the number of smells have a different impact on frequency of change?

For RQ1b, we used the nonparametric Mann–Whitney statistical test to check whether the average number of smells per commit in artifacts that do not change and in artifacts that do change is statistically similar. Formally, we calculate

$$changed(v) = \sum_x^{C_v} \frac{n_v(x)}{|C_v|} \quad unchanged(v) = \sum_x^{U_v} \frac{n_v(x)}{|U_v|}$$

where C_v is the set of artifacts x that changed in commit v, U_v is the set of unchanged artifacts, and $n_v(x)$ counts the number of smells x has in v.

The hypotheses for this analysis are as follows:

- Null hypothesis H_0^{RQ1b}: The number of smells in artifacts that do not change is equal to the number of smells in artifacts that do change ($\mu_{unchanged} = \mu_{changed}$)
- Alt. hypothesis H_1^{RQ1b}: The number of smells in artifacts that do not change is *less* than the number of smells in artifacts that do change ($\mu_{unchanged} < \mu_{changed}$)

with μ representing the mean of the populations (changed and unchanged). Additionally, to further reinforce the findings, we also check whether there is any correlation (using Spearman's ρ) between the number of smells an artifact is affected by and the number of changes or their size.

5.3 | RQ2—What is the difference in the change frequency of an artifact before and after a smell is introduced?

The analysis for this research question will look at the PCCC metric of a certain artifact before and after a smell is introduced in that element. We then aggregate the data per project and perform a Wilcoxon signed-rank test[36] for each project.

Formally, for every artifact x affected by a smell in the lifetime of a system S, we compute

$$d_S(x) = PCCC_{after}(x) - PCCC_{before}(x),$$

where $PCCC_{before}(x) = PCCC_k^i(x)$ and $PCCC_{after}(x) = PCCC_j^k(x)$ with i being the commit index where x first appeared, k where it was first affected by a smell, and j when it was last affected by a smell or the final commit. Artifacts with either a before or after window smaller than 5 commits were filtered out to avoid skewed data. The values assumed by d_S for the selected artifacts from S are used as input for the test.

The hypotheses for this test are as follows:

- Null hypothesis H_0^{RQ2}: The change frequency of artifacts before and after a smell is introduced is the same ($\theta_{d_s} = 0$)
- Alt. hypothesis H_1^{RQ2}: The change frequency of artifacts after a smell is introduced is greater than before ($\theta_{d_s} > 0$)

where θ represents the true median of the underlying population. Additionally, using the Shapiro–Wilk test, we also test for the normality of d_S to ensure we chose the appropriate statistical test.

5.4 | RQ3—Is the size of the changes in source code artifacts affected by smells, larger than in nonaffected artifacts?

For this RQ, we want to investigate if there is a significant difference in the *variance* of the size of the changes in affected versus nonaffected artifacts in each commit analyzed. We look at the variance because the majority of commits have a relatively small change size, whereas few commits (e.g., the pull requests) have a very large change size.

To determine whether there is a significant difference in these two groups (smelly vs. non-smelly), we perform a Brown–Forsythe test for the homogeneity of variance.[36]

For each commit v, we compute the aggregate change size (TACH metric) of changing artifacts by averaging all the change sizes for that commit. Formally,

$$smelly(v) = \sum_x^{A_v} \frac{TACH(x)}{|A_v|} \quad clean(v) = \sum_x^{N_v} \frac{TACH(x)}{|N_v|},$$

where A_v is the set of affected artifacts in a commit v and N_v the non-smelly artifacts. Note that we use the term "clean" to indicate **non-smelly** artifacts for conciseness.

Next, we test the following hypotheses on those two variables for each project:

- Null hypothesis H_0^{RQ3}: The variance in change size is equal in affected and clean artifacts by smells ($\sigma_{smelly}^2 = \sigma_{clean}^2$)
- Alt. hypothesis H_1^{RQ3}: The variance in change size is not equal in affected and clean artifacts ($\sigma_{smelly}^2 \neq \sigma_{clean}^2$)

where σ_{smelly}^2 and σ_{clean}^2 represent the true variance in the underlying populations.

6 | RESULTS

6.1 | Relation between change frequency and smelly artifacts (RQ1)

The results of the Fisher's tests (main question of RQ1) performed on each 4-month period of the 31 systems analyzed are pretty straightforward when not controlling for size. Figure 3 reports in detail the number of 4-month periods (or pseudo-releases) where the null hypothesis was accepted, rejected, or the test was unsupported by the data. The proportion of smelly artifacts that change is consistently higher than non-smelly artifacts that change in 82% of the total 4-month periods analyzed in most of the projects (rejecting H_0^{RQ1}). In other words, **artifacts with smells do change more frequently**. For the remaining 18%, if we remain conservative and assume that the 11.5% of the unsupported tests are accepted, the smelly and non-smelly artifacts are equally likely to change (accepting H_0^{RQ1}). Note that the unsupported tests are usually the ones corresponding to the pseudo-releases in the early phases of the project with a relatively little number of smells and/or changes. We also note that these percentages hold for the majority of the projects, with six exceptions: Elasticsearch, Jena, HTTP-components, Guava, Retrofit, and Spring-boot. These projects exhibit the opposite scenario, with more than 50% of the pseudo-releases featuring changes in non-smelly artifacts (neither accepting nor rejecting H_0^{RQ1}). This can be at least partially explained in all six cases: They either have a very low density of smells (HTTP-
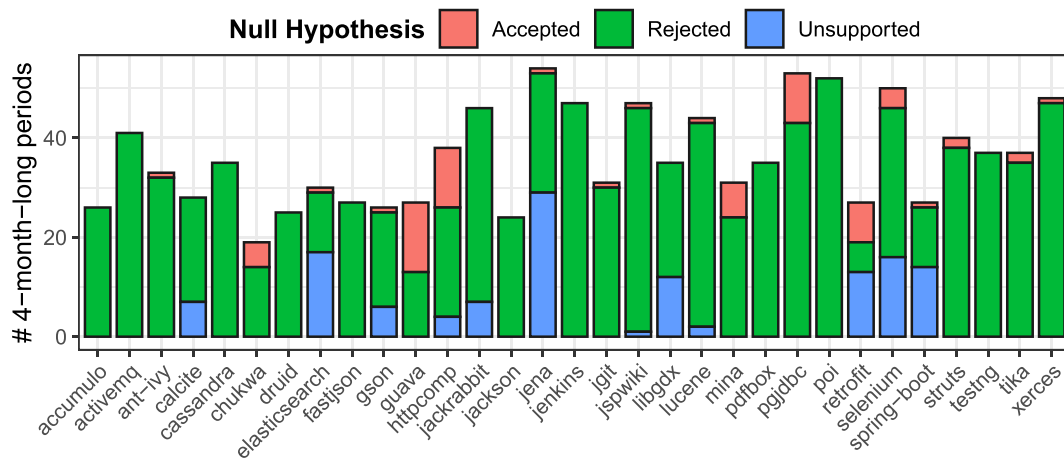
FIGURE 3 Results of the Fisher's test for each project (averages: accepted: 6.5%; rejected: 82.1%; unsupported: 11.4%)

components and Retrofit) or the actual proportion of smelly components that change is lower (up to 10 times) than non-smelly components that change (Elasticsearch, Guava, and Jena).

When adjusted for size (using the LOC—see Section 5.1), the results depicted in Table 3 show that for *Medium–Large* and *Large* artifacts, the null hypothesis H_0^{RQ1} was rejected 66.1% and 78.9% of the times on average across all projects, respectively. For *Small* and *Medium–Small* artifacts, percentages drop to 30.2% and 45.4%, respectively. In total, H_0^{RQ1} was rejected 55.7% of the times and accepted 30.2%, whereas in the remaining 14.1% of times, the analysis was unsupported.

Ultimately, the results controlled for size do not deviate too much from the uncontrolled ones but allow us to discern that **the larger the file, the more an artifact is likely to change if affected by a smell**.

6.1.1 | RQ1a

The aim of answering RQ1a was to understand if the specific type of the smells affecting the artifacts has an impact on the occurrence of changes. Table 4 introduces the results of the multinomial logistic regression model. For each type of smell, it shows the proportion of the 4-month-long periods where the null hypothesis was rejected. A large number of rejected instances mean that the given type of smell has a significant effect on the dependant variable of the regression model, that is, the *occurrence of changes*. The table reports all the statistically significant rates where a variable was considered relevant in the prediction of a change. Each column is a different model calibrated for that size group (or using all files in the case of "Uncontrolled"). We first notice that all the variables exhibit an increase in significance as we look at size groups of larger files. There seems to be no particular smell type, perhaps only excluding HL, that provides a clear contribution to the regression model over the other types.

The results imply that, in most cases, HL is the smell that contributes the most to changes; however, there is no sufficient evidence to affirm that there is a clear distinction between different types of smell. Thus, we conclude that we accept H_0^{RQ1a} and affirm that **there is *no significant difference* in the prediction power of different smell types on source code changes**.

6.1.2 | RQ1b

Furthermore, for RQ1b, we tested whether the number of smells (including 0) affecting an artifact is an important variable contributing to its change frequency. The test results, depicted in Figure 4, show that in all projects, but two (Guava and Pgjdbc), the average number of smells in changing artifacts is statistically higher than in nonchanging artifacts (rejecting H_0^{RQ1b}) when not controlling for size. If we consider the different size groups, the rejection rate is higher in the larger ones, whereas in the *Small* group, we reject H_0^{RQ1b} only twice. Additionally, larger size groups also have a larger Cliff's Delta coefficient with 31 tests having $\delta > .5$ in the *Large* and

TABLE 3 Percentages of rejected H_0^{RQ1} per project by size group (averages: accepted: 33.6%; rejected 59.1%; unsupported: 7.3%)

	% p value <.05			
Project	Small	Med.–Small	Med.–Large	Large
accumulo	38.5	65.4	96.2	100.0
activemq	78.0	65.9	100.0	100.0
ant-ivy	57.6	57.6	90.9	100.0
calcite	67.9	64.3	75.0	71.4
cassandra	74.3	94.3	100.0	100.0
chukwa	0.0	31.6	42.1	63.2
druid	40.0	84.0	96.0	100.0
elasticsearch	30.0	43.3	43.3	43.3
fastjson	7.4	33.3	85.2	88.9
gson	3.8	0.0	30.8	69.2
guava	37.0	0.0	25.9	29.6
httpcomp.	2.6	13.2	36.8	39.5
jackrabbit	15.2	69.6	84.8	84.8
jackson	33.3	95.8	100.0	100.0
jena	9.3	29.6	38.9	44.4
jenkins	44.7	72.3	95.7	100.0
jgit	48.4	77.4	96.8	96.8
jspwiki	31.9	70.2	68.1	89.4
libgdx	8.6	28.6	78.3	100.0
lucene	38.6	47.7	56.8	75.0
mina	25.8	12.9	38.7	67.7
pdfbox	54.3	77.1	94.3	100.0
pgjdbc	24.5	35.8	64.2	79.2
poi	36.5	55.8	92.3	100.0
retrofit	0.0	0.0	3.7	11.1
selenium	18.0	36.0	40.0	62.0
spring	14.8	3.7	25.9	44.4
struts	40.0	32.5	85.0	95.0
testng	5.4	5.4	37.8	100.0
tika	13.5	51.4	70.3	91.9
xerces2	35.4	52.1	56.2	97.9
Avrg.	30.2	45.4	66.1	78.9

TABLE 4 Results of the multinomial logistic regression in percentage of commits a variable was statistically significant in predicting a change (rejecting H_0^{RQ1a})

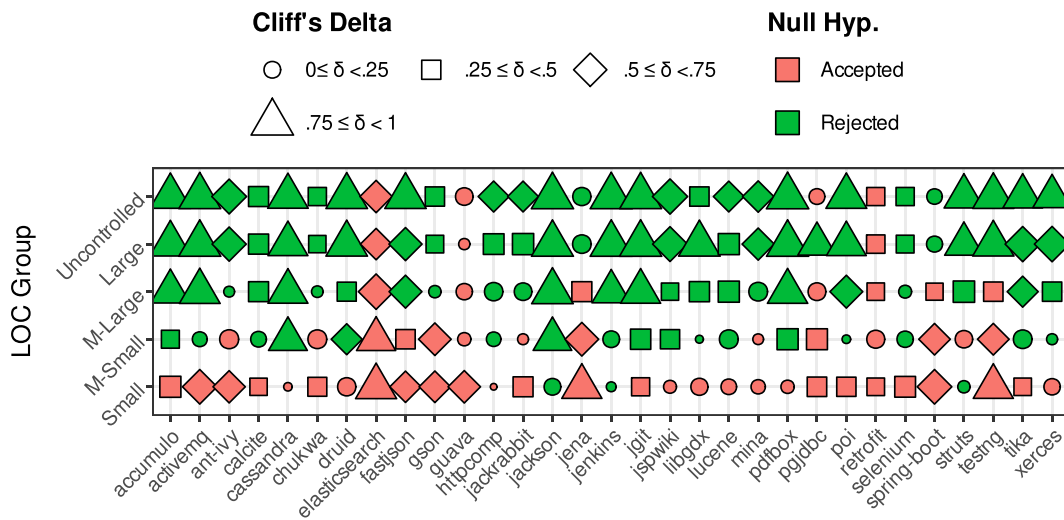|	Commits % when variable is significant				
Variable	Small	Med.–Small	Med.–Large	Large	Uncontrolled
Cyclic Dependency	12.2	14.1	20.2	29.1	39.7
Unstable Dependency	8.3	14.6	20.6	32.4	29.7
Hub-Like Dependency	22.1	22.3	26.7	44.6	51.9
God Component	–	–	–	30.5	38.5

FIGURE 4 Mann–Whitney tests testing H_0^{RQ1b} by size groups
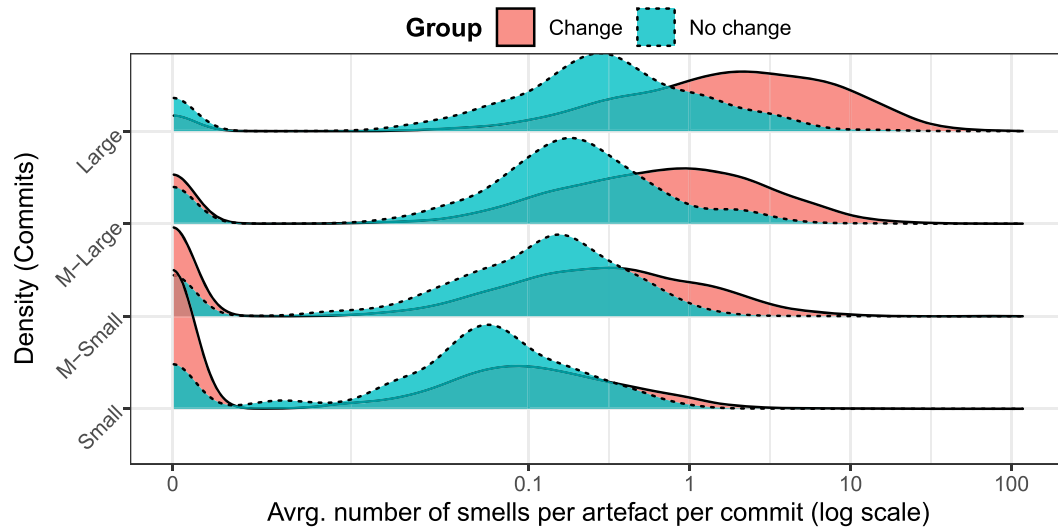

FIGURE 5 Average number of smells in changing/not-changing artifacts in all the analyzed commits by size groups

M–Large groups against the 3 in the M–Small and Small groups, highlighting the difference in magnitude between the values of the two variables tested.

To better grasp the contrast in smell density between changing and nonchanging artifacts in different size groups, Figure 5 plots the density (i.e., # commits) of the (average) number of smells of these two categories. In the figure, one can see how the average number of smells in changing and nonchanging artifacts varies in the commits in our dataset. When no smells affect an artifact (leftmost side of the plot), in several commits, only the smaller files change. But looking at the Medium–Large and Large files, we observe both curves shifting in shape and moving towards the right of the plot, with the changing artifact curve growing larger and distantiating itself from the nonchanging artifacts curve. This means that **for artofacts with more smells, the larger they are, the more likely they are to change**.

Figure 6 shows how the change frequency varies by the number of smells affecting an artifact, with different colors representing different projects. As it can be noted, there is a steep increase in the number of changes as the number of smells increases from 0 to 15, before stabilizing and slightly growing towards the end of the plot.[¶] The Spearman statistical correlation tests show that 8 projects show a strong ($\rho \geq .7$) positive

---

[¶]For reference, the right-most project in yellow-ish/ocra is Cassandra. Per-project plots are available in the replication package.
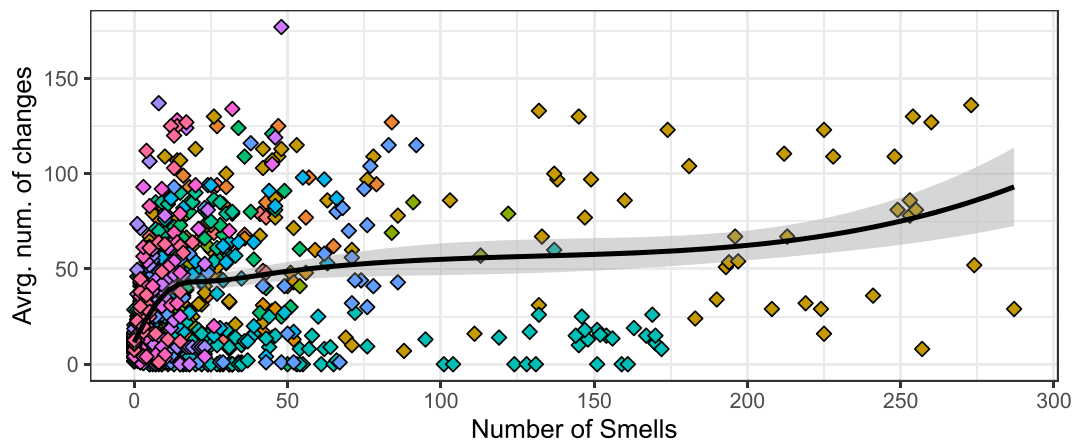
FIGURE 6   Average change frequency by number of smells per project (color-coded). Locally Estimated Scatterplot Smoothing (LOESS) regression curve shows the trend

correlation; 7 show a moderate ($.5 \leq \rho < .7$) positive correlation; 5 projects have a weak ($.3 \leq \rho < .5$) correlation; for 2 projects, there is little-to-no correlation ($\rho < .3$); and for the remaining 9 projects, $p > .05$.

In summary, **the more smells affecting an artifact, the higher the change frequency for that artifact**.

6.2 | Impact on change frequency after the introduction of a smell (RQ2)

Whereas the results of RQ1 hint that changes are more likely, and more frequent, in smelly artifacts, they do not tell us anything about the effects of the introduction of a smell on the change frequency of a particular artifact throughout its lifetime. This particular aspect is considered and tested by RQ2 through a series of Wilcoxon signed-rank tests; this was confirmed to be suitable in this case because most of the projects have their d_S (as defined in Section 5.3) function not normally distributed. Note that we do not control for size for this RQ because this is a temporal analysis; thus, there is no way to establish exactly which size category one artifact belongs to, as the LOC fluctuate over time. The results, presented in Table 5, indicate that for 16 projects (57.1%), there is an increase in the frequency of changes (in the PCCC metric, to be precise) after a smell is introduced. For 12 projects (42.9%) instead, the opposite holds, and more changes happen before the introduction of the smell. Finally, 3 projects did not contained enough samples and were ignored.

We can further inspect the distribution of d_S in Figure 7, where we can see how the distribution of changes to the artifacts is skewed either towards the "before" or "after" the introduction of a smell side, depending on the project. Therefore, we can conclude that, **in some cases, the introduction of an AS has increased the frequency of changes in the affected component**. We offer a potential explanation for the 12 projects (skewed towards the "before" part) that do not conform to this trend in the Section 7, but we would like to note that 9 of the 16 projects for which we rejected the null hypothesis had $n \leq 30$, whereas the accepted ones only had 6. This means that the tests were likely accepted because of an insufficient number of samples.

6.3 | Comparison of magnitude of changes in smelly and non-smelly artifacts (RQ3)

For this last question, when testing the null hypothesis H_0^{RQ3} on all projects, without adjusting for the size of the artifacts, we reject it for all of them – meaning that change size (TACH metric) in artifacts affected by smells has a consistently higher variance than the non-smelly ones. However, when controlling for the size of the affected artifacts, a different picture emerges. The results are presented in Figure 8. Smaller classes and packages do not exhibit this pattern as consistently as the larger ones do; in fact, in *Small* artifacts, we note the opposite in the majority of the projects. For *Large*, *M–Large*, and *M–Small* artifacts, the results are consistent in rejecting the null hypothesis H_0^{RQ3}. This can be further visualized in Figure 9, where the violin plots of the average change size of smelly and non-smelly artifacts in the analyzed commits can be visually compared for each size group. The more elongated the shape of the violin, the larger the variance in the corresponding group. By observing this figure, we note that the change size in smelly artifacts tends to increase (the violin shifts upwards) as the size of the artifacts increases. In stark contrast, the

TABLE 5 Wilcoxon signed-rank results and the sample size (# of artifacts) for the test (H_0^{RQ2} rejected in bold) (total: accepted: 42.9%; rejected: 57.1%)

Project	p value	Null Hyp.	Obs.
accumulo	.50	Accepted	11
activemq	**<.01**	**Rejected**	120
ant-ivy	**<.01**	**Rejected**	64
calcite	**<.01**	**Rejected**	75
cassandra	.30	Accepted	178
chukwa	**<.01**	**Rejected**	17
druid	.61	Accepted	47
elasticsearch	**.04**	**Rejected**	28
fastjson	.83	Accepted	31
gson	.50	Less than 10 obs.	5
guava	**<.01**	**Rejected**	34
httpcomp.	**<.01**	Less than 10 obs.	9
jackrabbit	**<.01**	**Rejected**	82
jackson	.28	Accepted	12
jena	.86	Accepted	11
jenkins	**<.01**	**Rejected**	178
jgit	**<.01**	**Rejected**	42
jspwiki	**<.01**	**Rejected**	59
libgdx	.30	Accepted	98
lucene	**<.01**	**Rejected**	220
mina	**<.01**	**Rejected**	22
pdfbox	**<.01**	**Rejected**	103
pgjdbc	.07	Accepted	23
poi	.25	Accepted	36
retrofit	**.04**	Less than 10 obs.	4
selenium	**.03**	**Rejected**	16
spring-boot	**.02**	**Rejected**	25
struts	.05	Accepted	13
testng	.97	Accepted	15
tika	.17	Accepted	33
xerces2	**<.01**	**Rejected**	86

violin shapes of the non-smelly group are surprisingly similar across the four different groups; this contrast highlights the impact of smells on affected artifacts w.r.t. change size.

Hence, both visual analysis and statistical tests converge to the same conclusion that **smelly artifacts undergo changes of higher magnitude than non-smelly artifacts, especially in larger artifacts**. More precisely, smelly artifacts have an average change size (TACH) across all the projects of 1608, whereas non-smelly artifacts settle at 109. The difference is one order of magnitude higher in smelly artifacts; we note that the smelly artifacts also have a higher variance.

7 | DISCUSSION

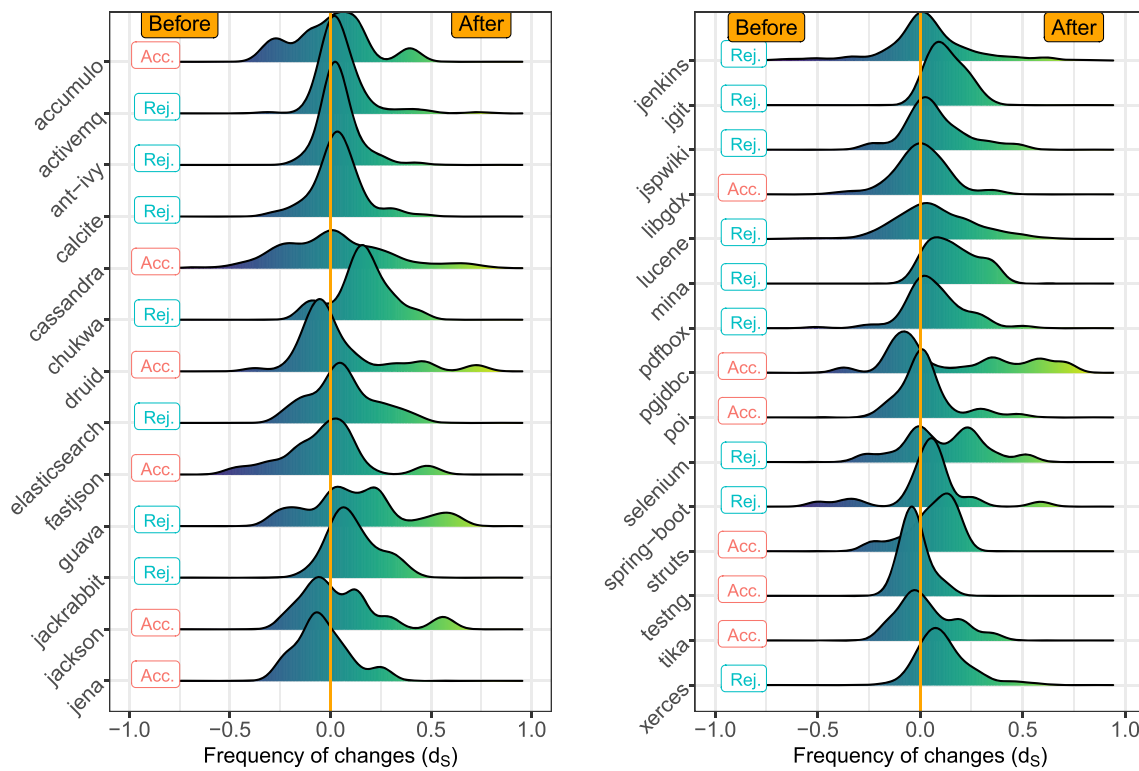In the following section, we discuss and elaborate on the results presented above.

FIGURE 7 Density of Percentage of Commits a Class has Changed (PCCC) before and after the introduction of a smell (d_S function)
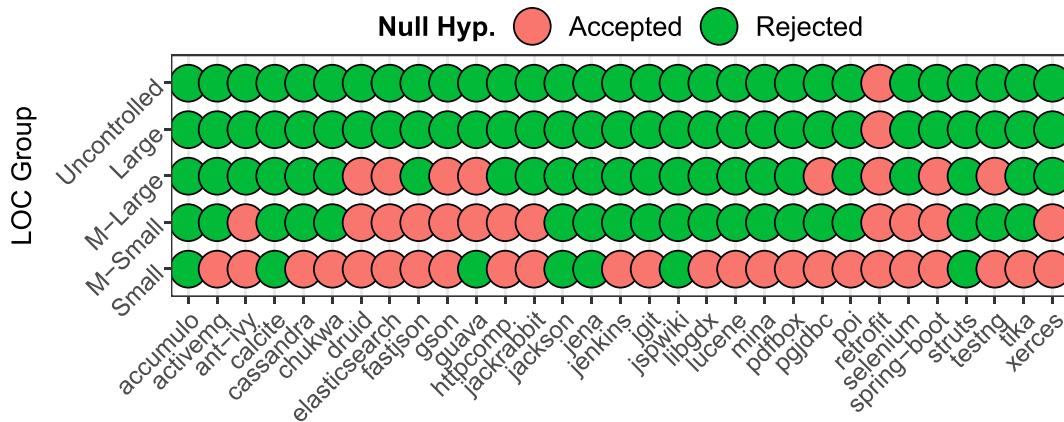

FIGURE 8 Brown–Forsythe tests results by size groups for H_0^3

From the obtained results, we have empirically confirmed that ASs (at least the ones considered for this study), exhibit a *correlation* with the *change frequency* and *change size* of the artifacts they affect and especially the larger artifacts. As stated previously, our goal was to seek and establish *correlation*, rather than *causality*.

The results from RQ1 show that artifacts affected by at least one smell exhibit more changes than artifacts without smells. We also saw that as the number of smells increases, so does the likelihood of the affected artifact to change. Interestingly, we found evidence suggesting that the four different types of smells that we studied affect change frequency in a similar way. Theoretically speaking, the main drawback associated with the UD smell type[5] is an increased likelihood to change caused by a low stability of the artifacts it depends upon.[30] Indeed, while we observed that UD-affected artifacts have an increased change frequency, we also expected them to have a higher change frequency than artifacts affected by the other smell types. However, this is not what we observed, as all four smell types seem to have a similar effect on change, with HL surpassing UD in fact. The HL smell was hypothesized to be quite prone to propagate changes due to its numerous dependencies, which increase the

likelihood that a change propagates to the central component before rippling to the components depending on it (see related work[14] for further information). Given this theoretical description, we would predict that changes may propagate within the structure of HL smells, but we did not expect it to deviate from the other smell types and surpass UD.

The results of RQ2 show an increase in change frequency *after* the introduction of a smell. This finding implies that the introduction of a smell leads to an increase in the effort developers spend on the particular artifact(s) affected by that smell compared with the period of time that artifact was not affected by any smell. Note that change frequency and size were used to estimate the effort spent by a developer in previous studies too.[10,11,37] Nonetheless, it is interesting to note that this result is not valid for all of the projects considered, and in some cases, the opposite situation occurs. We conjecture that this result is highly dependant on how development teams decide to implement new functionality in the system. If developers reuse existing classes, then these classes are likely to require changes for a longer period of time and especially after a smell affects them (as seen from our results). If developers do not reuse existing classes and implement a new functionality in new classes and packages, then old and smelly classes are less likely to be changed, because they serve their purpose as they are without requiring further changes. More generally, in the "old" features of a system, very little maintenance effort is spent on good design and architecture, for example, by refactoring smells; this means that components affected by smells rarely get changed.

A perfect example of this is provided by the project PgJDBC. PgJDBC, among other projects, has a package named "v2," suggesting that functionality for the previous release ("v1") is implemented separately in a different package. The classes implementing the functionality for the previous release are thus no longer extended, and therefore, they no longer change, but they are still kept in the repository in order to support legacy functionality. Assuming that at least some of these classes were affected by smell, the resulting effect is that their change frequency after the smell was introduced is close to zero. Moreover, given that these are open source projects, we cannot assume that they undergo constant development, and changes in the popularity of a project may influence how many pull requests, commits, and changes are performed. Ultimately, these two factors greatly influence the variability in the results obtained from RQ2.

With the obtained results from RQ3, there is additional evidence to support that developers spend *more time* on smelly artifacts, where we noticed a consistently larger change size—again, especially in the larger files. This is especially true in pull requests commits, the type of commits where usually new functionality, or big bug fixes, are introduced. This result becomes even clearer when observing Figure 9: There is a strong contrast between the constant change size in non-smelly artifacts across the four different size groups, versus the increasing change size in smelly artifacts. This clearly shows the spending of extra effort to perform changes in smelly artifacts.

Putting together the results of all the three research questions, we conclude that developers are not only compelled to make *more frequent* changes to smelly artifacts but also to make *larger* changes. Ultimately, if we assume that change is a proxy of the effort spent maintaining the components affected by smells,[10,11] the technical debt interest of those components is increased by two factors: change frequency and change size. An important caveat is that these findings do not include any input from the actual developers; therefore, further research is required in order to understand the full extent to which ASs perturb development activities from the perspective of software practitioners themselves. For the time being, we can conclude that ASs constitute a *high risk*, as their accumulation can increase technical debt interest to unsustainable levels.

A common trend in our results is that smelly *Large* and *Medium–Large* smelly files exhibit statistically different patterns in change frequency and size in contrast to the *Small* and *Medium–Small* groups. It is interesting to explore why this happens mostly in these size groups. The main reason is that GC and HL are defined based on the number of LOC or incoming and outgoing dependencies of the affected artifact. Namely, they
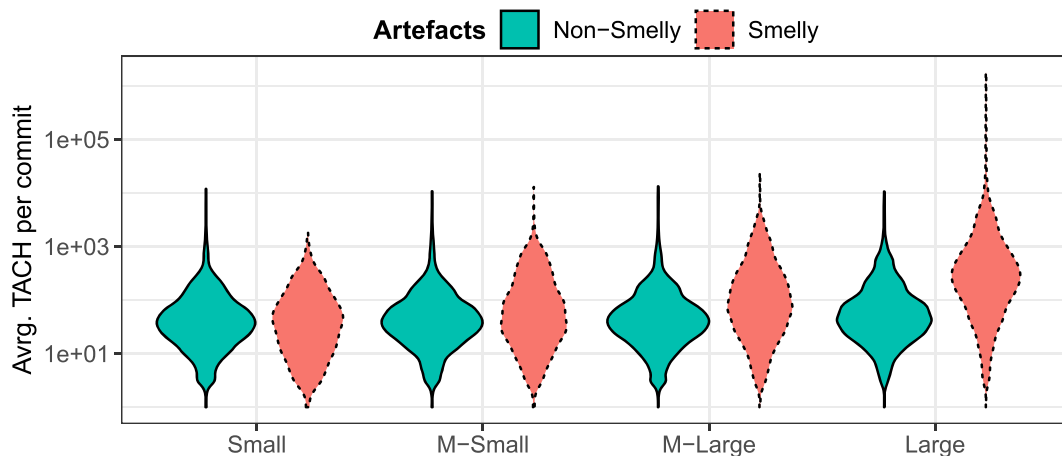

FIGURE 9 Violin plots of the distribution of the average change size (Total Amount of Changes [TACH] metric) in the analyzed commits grouped by smelly and non-smelly artifacts (log scale)

cannot affect small artifacts by definitions. Smaller artifacts can however still be impacted by a HL smell if they depend on the hub (central artifact), because change may propagate to them from the hub. This difference between smaller and larger files has a clear implication for researchers: We advise the development of better prioritization methods for refactoring ASs by prioritizing smells affecting larger artifacts; these are the ones where developers pay the most technical debt interest.

Finally, some of the findings that emerged from this study match what Oyetoyan et al[6] and Le et al[15] have found in their own works. Specifically, our results from RQ1 match what Le et al[15] found; namely, the average number of changes in smelly files is higher than in non-smelly files (see Figure 6). On top of that, we have also shown how the number of changes positively correlates with the number of smells (RQ1b). Answering RQ1a instead, we have noted among others that potentially *not all types of cycles* are impactful on changes. This corroborates what Oyetoyan et al[6] found about CDs; that is, certain types of cycles do not have an impact on changes. However, we did not investigate precisely which category of cycles does so and neither if they affect neighbor artifacts; we consider this future work.

8 | STUDY LIMITATIONS

The identified limitations of this study are described in terms of *reliability*, *external validity*, and *construct validity* as described by Runeson et al.[21] Internal validity was not considered as we did not examine causal relations.[21]

8.1 | Construct validity

Construct validity concerns to what extent this study is measuring what it is claiming to be measuring.[21] To ensure construct validity, we adopted the well-known case study design guidelines provided by Runeson et al[21] and iteratively revised the protocol during the duration of the study. Thus, the data collection and analysis processes were meticulously planned and implemented to ensure that the final results would answer precisely the three main research questions of interest of this study.

One concrete threat to construct validity is the arbitrary selection of the 4-week interval between the analyzed commits. While the selection of this particular interval was computationally convenient (i.e., more commits would pose higher requirements for processing time), in the more active projects, this time interval might have caused the loss of information for frequency-related metrics. For instance, a class might have changed several times during the course of 4 weeks, but we only count it is as one big change. As a result, the coarse-grained frequency data may have impacted the analysis and thus the results of RQ2. Additionally, this interval might clash with the culture of each development team in pushing changes to the central repository and the size of those changes. However, the very selection of this particular interval also partially mitigated this risk, if we compare our study with related work, where most studies[7,15,16] use time in-between releases, which is usually longer and more susceptible to the risks mentioned above.

Similarly, the pseudo-release data aggregation we performed for RQ1 might be incorrect even though it is based on empirical evidence. The problem is that we used the date the Git tag was added, which might not match the official release date of that release. To mitigate this risk, we manually inspected all the dates and ensured they were reasonable and matched the versions' numbering order (e.g., v1.1 comes before v1.2, and their dates match such order). Tags with different release numbers but with the same date were removed.

Another threat to construct validity is related to our use of change frequency and size as indicators of technical debt interest. The same indicators have been used in previous studies,[37,38] as there is no way to directly measure technical debt interest. However, it is important to keep in mind that they are only proxies and the actual interest paid by developers might vary significantly. Thus, assuming that an increase in change frequency and size corresponds to a *direct* increase in technical debt interest paid by the development team while implementing new features, or making changes to the code base, is not always correct. The more frequent and bigger changes required to implement those features may be a result of the inherent difficulty of implementing the features themselves or even other external factors. On the other hand, it is also unlikely that *all* the new features and changes are characterized by inherently difficult elements to design and implement.

8.2 | External validity

This aspect of validity reflects to what extent the results of this study can be fitted to the whole population of projects considered and relatable contexts.

Two threats have been identified in this case. The first one involves the types of projects we selected for our study. While all of them are open source projects, 18 of them are projects from the Apache Foundation, and only 13 are non-Apache projects. The imbalance is caused by the fact that most Apache projects have a very long and consistent history, which made them more likely to be adopted for our analysis. We decided to mitigate this aspect by diversifying as much as possible the application domains of the selected projects. Moreover, we collected our data from

31 projects, considerably more than what had been done by previous, similar studies (i.e., Le et al[7] used 14 projects; Le et al[15] used 8 projects, and Oyetoyan et al[6] used 12 projects), thus strengthening external validity.

The other threat to external validity concerns the ASs we used for our analysis. It is very hard to generalize the results to other ASs, and it is probably not possible to do so with enough confidence for every type of smell. This very much depends on the type of smell and the detection strategy for that smell. Therefore, we cannot claim any generalization of our results to other ASs.

8.3 | Reliability

Reliability is the aspect of validity focusing on the degree to which the data and the analysis depend on the researchers performing them.

All the tools and the data used in this study are freely available online (see related studies and footnote 2) to allow researchers to study or replicate our results using the same data or even a different set of projects.

The intermediary findings and data analysis steps were all inspected and discussed by all the authors of this paper to ensure their reliability. Moreover, similar data collection and analysis techniques have been also used in previous studies on code smells (e.g., Khomh et al.[20]) and ASs (e.g., Le et al.[15]), assuring that it is indeed possible to do this type of analysis for these types of artifacts.

9 | CONCLUSIONS AND FUTURE WORK

The present study has thoroughly investigated the relationship between a set of four ASs and the changes in the affected components. In total, 31 projects, adding up to a total of 360 years of development and over 305 million LOC, were statically analyzed and then statistically tested against our hypotheses.

The main findings of this case study show that (1) artifacts affected by ASs change more frequently than non-smelly artifacts; (2) the type of the smell does not have a significant correlation with changes; (3) the more smells affect an artifact the more likely it is to change; (4) the change frequency of an artifact increases after the introduction of a smell in the majority of the systems; and (5) the size of changes is significantly higher in smelly artifacts than in non-smelly ones. These findings are especially valid for artifacts belonging to the *Medium–Large* and *Large* size groups. We thus concluded that ASs are very likely to be associated with an increase in the technical debt interest developers pay each time they work on artifacts affected by smells.

Given the results obtained from our RQs, it would be interesting to explore how the presence of ASs is perceived by the very developers and architects of a software system. More specifically, a natural continuation of RQ1 is to investigate if practitioners do perceive that affected components are more prone to changes than nonaffected components and whether there is any difference, in this regard, between different types of smells. Concerning RQ2, it would be interesting to explore how the introduction of a smell is perceived, what lead to the introduction of the smell, and whether developers were aware of it. Finally, for RQ3, a possible research direction is to understand whether the difference measured in our study has a perceivable impact by developers and architects; in other words, to study how big changes must be in order to make a difference in the effort perceived.

DATA AVAILABILITY STATEMENT
The data that support the findings of this study are openly available in Zenodo at https://doi.org/10.5281/zenodo.4459015.

ORCID
Darius Sas https://orcid.org/0000-0003-3383-3298

REFERENCES
1. Garcia J, Popescu D, Edwards G, Medvidovic N. Identifying architectural bad smells. In: 2009 13th Proceedings of the European Conference on Software Maintenance and Reengineering, CSMR; 2009; Kaiserslautern, Germany:255-258.
2. Lippert M, Roock S. *Refactoring in Large Software Projects: Performing Complex Restructurings Successfully*: John Wiley & Sons, Inc; 2006. http://eu.wiley.com/WileyCDA/WileyTitle/productCd-0470858923.html

3. Avgeriou P, Kruchten P, Ozkaya I, Seaman C. Managing technical debt in software engineering (Dagstuhl Seminar 16162). *Dagstuhl Rep*. 2016;6(4): 110-138. http://drops.dagstuhl.de/opus/volltexte/2016/6693

4. Verdecchia R, Malavolta I, Lago P. Architectural technical debt identification: the research landscape. In: 2018 ACM/IEEE International Conference on Technical Debt. Gothenburg, Sweden; 2018:11-20. http://www.ivanomalavolta.com/files/papers/TechDebt_2018.pdf

5. Fontana FA, Pigazzini I, Roveda R, Zanoni M. Automatic detection of instability architectural smells. In: Proceedings - 2016 IEEE International Conference on Software Maintenance and Evolution, ICSME 2016; 2016; Raleigh, NC, USA:433-437.

6. Oyetoyan TD, Falleri JR, Dietrich J, Jezek K. Circular dependencies and change-proneness: an empirical study. In: 2015 IEEE 22nd International Conference on Software Analysis, Evolution, and Reengineering, SANER 2015 - Proceedings. Institute of Electrical and Electronics Engineers Inc.; 2015; Montreal, QC, Canada:241-250.

7. Le DM, Behnamghader P, Garcia J, Link D, Shahbazian A, Medvidovic N. An empirical study of architectural change in open-source software systems. In: 2015 IEEE/ACM 12th Working Conference on Mining Software Repositories; 2015; Florence, Italy:235-245.

8. Neri D, Soldani J, Zimmermann O, Brogi A. Design principles, architectural smells and refactorings for microservices: a multivocal review. *SICS Softw-Intens Cyber-Phys Syst*. 2019:35:3-15.

9. Arcelli Fontana F, Lenarduzzi V, Roveda R, Taibi D. The Journal of Systems and Software Are architectural smells independent from code smells? An empirical study. *J Syst Softw*. 2019;154:139-156. https://doi.org/10.1016/j.jss.2019.04.066

10. Sjoberg DIK, Yamashita A, Anda BCD, Mockus A, Dyba T. Quantifying the effect of code smells on maintenance effort. *IEEE Trans Softw Eng*. 2013; 39(8):1144-1156.

11. Olbrich S, Cruzes DS, Basili V, Zazworka N. The evolution and impact of code smells: a case study of two open source systems. In: 2009 3rd International Symposium on Empirical Software Engineering and Measurement, ESEM 2009; 2009; Lake Buena Vista, FL, USA:390-400.

12. El-Emam K. A methodology for validating software product metrics. Technical Report NRC 44142, Nat'l Research Council of Canada; 2000.

13. Mockus A, Votta LG. Identifying reasons for software changes using historic databases. In: Conference on Software Maintenance. IEEE; 2000; San Jose, CA, USA:120-130.

14. Sas D, Avgeriou P, Arcelli Fontana F. Investigating instability architectural smells evolution: an exploratory case study. In: 35th International Conference on Software Maintenance and Evolution. IEEE; 2019; Cleveland, OH, USA:557-567. https://ieeexplore.ieee.org/document/8919109/

15. Le DM, Link D, Shahbazian A, Medvidovic N. An empirical study of architectural decay in open-source software. In: Proceedings - 2018 IEEE 15th International Conference on Software Architecture, ICSA 2018. IEEE; 2018; Seattle, WA, USA:176-185. https://ieeexplore.ieee.org/document/8417151/

16. Khomh F, Penta MD, Guéhéneuc YG, Antoniol G. An exploratory study of the impact of antipatterns on class change- and fault-proneness. *Empir Softw Eng*. 2012;17(3):243-275.

17. Jaafar F, Guéhéneuc YG, Hamel S, Khomh F, Zulkernine M. Evaluating the impact of design pattern and anti-pattern dependencies on changes and faults. *Empir Softw Eng*. 2016;21(3):896-931.

18. Sharma T, Singh P, Spinellis D. An empirical investigation on the relationship between design and architecture smells. *Empir Softw Eng*. 2020;25(5): 4020-4068. https://doi.org/10.1007/s10664-020-09847-2

19. Aniche M, Bavota G, Treude C, Gerosa MA, van Deursen A. Code smells for Model-View-Controller architectures. *Empir Softw Eng*. 2018;23(4):2121-2157.

20. Khomh F, Di Penta M, Guéhéneuc YG. An exploratory study of the impact of code smells on software change-proneness. In: Proceedings - Working Conference on Reverse Engineering, WCRE; 2009; Lille, France:75-84.

21. Runeson P, Höst M, Rainer A, Regnell B. *Case Study Research in Software Engineering - Guidelines and examples*. John Wiley & Sons, Inc; 2012.

22. van Solingen R, Basili V, Caldiera G, Rombach HD. Goal Question Metric (GQM) Approach. *Encyclopedia of Software Engineering*. 2nd edition. Wiley; 2002.

23. Nagappan N, Ball T. Using software dependencies and churn metrics to predict field failures: an empirical case study. In: First International Symposium on Empirical Software Engineering and Measurement (ESEM 2007); 2007; Madrid, Spain:364-373.

24. Kouroshfar E, Mirakhorli M, Bagheri H, Xiao L, Malek S, Cai Y. A study on the role of software architecture in the evolution and quality of software. In: IEEE International Working Conference on Mining Software Repositories. IEEE Computer Society; 2015; Florence, Italy:246-257.

25. Arcelli Fontana F, Avgeriou P, Pigazzini I, Roveda R. A study on architectural smells prediction. In: 2019 45th Euromicro Conference on Software Engineering and Advanced Applications (SEAA); 2019; Kallithea, Greece:333-337.

26. Elish MO, Al-Khiaty MAR. A suite of metrics for quantifying historical changes to predict future change-prone classes in object-oriented software. *J Softw Evol Process*. 2013;25(5):407-437. http://doi.wiley.com/10.1002/smr.1549

27. Arvanitou EM, Ampatzoglou A, Chatzigeorgiou A, Avgeriou P. A Method for Assessing Class Change Proneness. In: Proceedings of the 21st International Conference on Evaluation and Assessment in Software Engineering (EASE'17). Association for Computing Machinery, New York, NY, USA; 2017:186-195. https://doi.org/10.1145/3084226.3084239

28. Zhang J, Sagar S, Shihab E. The evolution of mobile apps: an exploratory study. In: Proceedings of the 2013 International Workshop on Software Development Lifecycle for Mobile (DeMobile 2013). Association for Computing Machinery, New York, NY, USA; 2013:1-8. https://doi.org/10.1145/2501553.2501554

29. Martin RC, Grenning J, Brown S. *Clean Architecture: A Craftsman's Guide to Software Structure and Design*. Prentice Hall; 2018.

30. Martin R. OO Design Quality Metrics. *Qual Eng*. 1994;8(4):537-542.

31. Fontana FA, Ferme V, Zanoni M, Yamashita A. Automatic metric thresholds derivation for code smell detection. In: International Workshop on Emerging Trends in Software Metrics, WETSoM; 2015; Florence, Italy:44-53.

32. Martini A, Fontana FA, Biaggi A, Roveda R. Identifying and prioritizing architectural debt through architectural smells: a case study in a large software company. *European Conference on Software Architecture*. Cham: Springer; 2018:320-335. http://link.springer.com/10.1007/978-3-030-00761-4_21

33. Arcelli Fontana F, Locatelli F, Pigazzini I, Mereghetti P. An Architectural Smell Evaluation in an Industrial Context. (c): 68–74; 2020.

34. El Emam K, Benlarbi S, Goel N, et al. The confounding effect of class size on the validity of object-oriented metrics IEEE Transactions on Software Engineering, 27 (2001). *IEEE Trans Softw Eng*. 2001;27(7):630-650. http://ieeexplore.ieee.org/document/935855/

35. Zhou Y, Leung H, Xu B. Examining the potentially confounding effect of class size on the associations between object-oriented metrics and change-proneness. *IEEE Trans Softw Eng*. 2009;35(5):607-623.

36. Sheskin DJ. *Handbook of Parametric and Nonparametric Statistical Procedures*. 5th ed. Chapman & Hall/CRC; 2007.

37. Nugroho A, Visser J, Kuipers T. An empirical model of technical debt and interest. In: Proceedings - International Conference on Software Engineering. New York, New York, USA: ACM Press; 2011:1-8. http://portal.acm.org/citation.cfm?doid=1985362.1985364

38. Ampatzoglou A, Michailidis A, Sarikyriakidis C, Ampatzoglou A, Chatzigeorgiou A, Avgeriou P. A framework for managing interest in technical debt: An industrial validation. In: Proceedings - International Conference on Software Engineering. Gothenburg, Sweden; 2018:115-124.