

Replacements and Replaceables: Making the Case for Code Variants

Venkatesh Vinayakarao¹, Rahul Purandare¹, Sumit Keswani¹,
Devika Sondhi¹, and Anita Sarma²

¹IIIT-Delhi, India

²Oregon State University, Corvallis, OR, USA.

June 16, 2020

Abstract

There are often multiple ways to implement the same requirement in source code. Different implementation choices can result in code snippets that are similar, and have been defined in multiple ways: *code clones*, *examples*, *simions* and *variants*. Currently, there is a lack of a consistent and unambiguous definition of such types of code snippets. Here we present a characterization study of code variants – a specific type of code snippets that differ from each other by at least one desired property, within a given code context. We distinguish code variants from other types of redundancies in source code, and demonstrate the significant role that they play: about 25% to 43% of developer discussions (in a set of nine open source projects) were about variants. We characterize different types of variants based on their code context and desired properties. As a demonstration of the possible use of our characterization of code variants, we show how search results can be ranked based on a desired property (e.g., speed of execution).

1 Introduction

“*There’s more than one way to skin a cat*”, a popular English idiom, also holds true for software development. Developers often have to reason among the benefits and drawbacks of different implementation choices before making a selection. These choices can include differences in the speed or complexity of computation, the style of coding, the library used or licensing requirements. Currently, code snippets reflecting different choices are known by different names: variants, clones, simions, idioms, examples; and are often used interchangeably. In this work, we define and characterize code variants, and differentiate them from clones, simions, and examples.

Table 1: Code variants are discussed in the bug descriptions of popular open source projects.

Project	Discussion
Apache Math [Bug# MATH-901]	<i>One of the reasons this variant is faster is because it is less accurate, which may not be acceptable for commons-math.</i>
Apache Math [Bug# MATH-1293]	<i>While the jury is still out, I made another variant of the patch ...</i>

Code variants represent an alternative implementation of a code snippet, where each alternative provides the same functionality, but has different properties that make some better suited to the overall project requirements.

Developers routinely need to analyze existing code, find better reuse alternatives, and look to develop high-quality code that meets some desired properties. For example, developers in Apache Commons Math discuss a variant in the FastMath library, which is faster, but is so because it is less accurate (Table 1 presents such an exchange in the issue tracker). The discussion involves issues about specific implementation choices, the library, the underlying algorithm, and license restrictions. Such discussions on selection of a specific variant are common.

Finding variants, however, is nontrivial. Developers search for variants over the web, in online discussion forums, and issue trackers. For example, when investigating Stack Overflow¹ (SO), a popular programming Q&A site containing 13M questions and 20M answers, we found in a random sampling of 300 posts, 100 variants that discussed 34 topics and spanned 13 programming languages. Similarly, we found that 25% to 43% of developer discussions are about variants in our dataset. Our dataset comprised of fifteen open source projects in five different programming languages, with varying sizes and domains. However, searching for variants is a fragmented, ad hoc process that can be time-consuming.

Mining for variants is an intrinsically difficult problem because of the following reasons. First, there is no consensus about what variants are. Research points to different types of semantic similarities [1, 2], but there are inconsistencies in their definitions (see Section 2). Second, Juergens et al. [3] show that behaviorally similar code can vary significantly in structure, making it difficult to match source code patterns. To the best of our knowledge, the concept of variants has not been discussed in the past and hence, it becomes important to understand what variants are and about its characteristics. Moreover, to build useful variant mining tools, we need to understand the code contexts and the desired properties that make a specific variant better than the rest. Our work attempts to draw some essential characteristics of variants which we believe could contribute towards automating variant mining which can help in the

¹<http://stackoverflow.com/tour>




Figure 1: Reuse in source code happens at different levels by size for different reasons. This work focuses on code variants.

development process.

To the best of our knowledge, ours is the first study in understanding variants and their characteristics. Our key contributions are as follows:

- We define and discuss how variants differ from semantic clones, simions and examples.
- We study variants and characterize their code contexts and desired properties.
- We analyze the source code and developer discussions in fifteen open source projects and show that variants exist and they matter to developers.
- As an application of our variant characterization study, we propose a systematic approach to mine variants using data from discussion forums.

2 Background and Terminology

Reuse happens in source code at several levels. This phenomenon is captured in Figure 1. Reuse starts at token-level, and occurs at line level too, as demonstrated by the naturalness of software [4]. Idioms, clones, simions and code variants represent reuse over multiple lines of code. Also, we observe that reuse grows beyond programs into applications and product variants. This work focuses on reuse at the level of multiple lines of code.

Code variants are closely related to two major types of code snippets: 1) Semantically similar code snippets (clones, idioms and simions), and 2) Code examples. Each of these types has its similarities and dissimilarities with variants.

Before we compare variants with semantic clones, simions, idioms and program examples, we need a definition for them. We modify Muralidharan et al.’s [5] definition of variants as given below (and further discussed in Section 3):

Definition 1 A code variant represents an alternative implementation for a given code snippet under a specific context in which one of the two implementation choices must score better on at least one desired property over the other.

This emphasizes on *code context* and *desired properties* of code snippets. We adapt Kirke’s [6] definition of method context to define code context as follows:

Definition 2 Code context of a code snippet describes the fitment of the snippet inside the larger project. It captures the intent, dependencies of the surrounding code, input to the snippet, the output from the snippet and the states in which the system may get into.

In a variant pair, one variant is said to be *preferred* over the other if it has at least one *desired property* when compared with the other within the given context. As an example, Table 1 shows a trade-off between speed of execution and accuracy. In this context, accuracy is more important, and hence a variant which is more accurate is preferred. Therefore, we define properties as follows:

Definition 3 Multiple implementation choices may satisfy the requirements in the given code context. Yet, each implementation choice (a code snippet) differs from one another in ways that could be either functional or para-functional or both. These qualities of code snippets that serve as differentiators are referred to as *properties*.

Factors such as lack of knowledge may lead developers to use less optimal choices initially and later discover snippets with more desired properties. For example, a recursive version of factorial is seldom used in real-world projects especially if the return value cannot hold large values. A hard-coded version avoids recomputation. Table 2 shows these two versions of factorial implementation. If a developer goes for the less optimal choice inadvertently, it may eventually get changed during the review or maintenance process. Thus, the knowledge of variants at development time can help developers write better programs.

Next, we investigate the existing literature to differentiate variants from clones, simions, idioms and examples.

2.1 Clones

Clones were originally defined as redundant snippets introduced due to a copy and paste activity [7]. They can be semantic or structural. Semantic clones appear in different varieties, such as wide-miss clones [8], interleaved clones [9], and high-level concept clones [8]. Gabel et al.’s [9] definition of semantic clones is as follows: *Two disjoint, possibly non-contiguous sequences of program syntax $S1$ and $S2$ are semantic code clones if and only if $S1$ and $S2$ are syntactic code clones or $\rho(S1)$ is isomorphic to $\rho(S2)$* . Here, ρ is a Program Dependence Graph (PDG) based transformation function. PDG captures control and data dependency in code snippets and abstracts away other syntactic details. Elva and Leavens [10] define semantic clones as functionally identical code fragments.

Table 2: Two variants of factorial program are shown below. Note that (a) is a text book example for recursion. Since, the factorial value grows very large quickly, integer type can hold only up to 12!. Hence (b) is preferred for real-world projects since it avoids re-computation.

(a) Recur- sive	(b) Real-World
<pre>public static int fact(int x) { if (x==1 x==0) return 1; return fact(x-1) * x; }</pre>	<pre>public int factorial(int n) { switch(n) { case 0: return 1; ... case 12: return 479001600; default : throw new IllegalArgumentException();} }</pre>

Ira Baxter defines clones as segments of code that are similar according to some (typically lexical) definition of similarity [11].

Syntactic clones [11, 12] are of three types. Type-1 clones are exact copies. Type-2 clones are copies where only the identifier names and variable types are changed. They are otherwise structurally similar. There is no single accepted definition of Type-3 clones [11, 12]. One definition of Type-3 clone is based on the Levenshtein distance between the pair of snippets which quantifies the minimum number of additions and deletions of tokens to transform one snippet to other.

2.2 Idioms and Simions

Keivanloo et al. [12] and Juergens et al. [3] indicate that code similarities go beyond these definitions of clone types. Allamanis and Sutton [1] define a code idiom as “*a syntactic fragment that recurs across software projects and serves a single semantic purpose*”. They claim that programmers use the term idiomatic to refer to code that is used repetitively. Idioms have both syntactic and semantic similarity.

Juergens et al. [3] call the snippets that are behaviorally similar as *Simions*. Simions need not originate from copy and paste activity. Even though simions have independent origins, they have been described in the context of redundant code with the intent of identification and refactoring.

2.3 Programming by Example

Developers seldom read the entire documentation before they start. They learn from code snippets on the web or other projects [13, 14]. *Code examples are small source code fragments whose purpose is to illustrate how a programming language construct, an API, or a specific function or method works* [15]. Examples play a significant role in comprehension, reuse, and bug-fixing [16]. As

Table 3: Fundamental differences exist between semantic clones, code examples and code variants.

	Semantic Clones	Code Examples	Code Variants
Definition	Code snippets with no difference in properties of interest within the given code context. Therefore, one snippet can replace the other.	Code snippets with an instructive property against an information need.	Code variants represent alternative implementations suitable for a specific code context in which one variant must have some desired properties over the other.
Differences	Clones are necessarily semantically similar and have no differences in desired properties. Hence, these snippets are redundant. Some amount of structural similarity is also assumed in cases where PDG based definitions are followed.	Neither semantic nor syntactic similarity warranted. Provides instructive value as in the usage of API or how to implement, and so on.	Semantically similar but has different desired properties.
Example	Two sorting implementations of same worst-case complexity where that is the only quality that matters to developers.	Any API usage tutorial. For example, in Java, the code snippets describing the usage of Arrays.sort feature.	Various sorting implementations with different time complexities are variants, if speed is the only desired property.

as a result, several researchers have explored locating [1, 17], selecting [18] and analyzing [16] examples.

3 Code Variants

Use of the term “variant” is quite popular in the development community. Table 4 shows the number of occurrences of the term “variant” in the same group of projects that Gabel et al. selected for semantic clones study. Not only developers but literature too provides several evidences of the use of the term “variant” [19, 5]. Variants are used in at least three major contexts: 1) Code variants (focus of this paper), 2) Program or product variants (as in product lines and application variants) [20], and 3) Configuration variants [5] (as in tuning a product or product configuration).

Variants are neither clones, nor simions, nor idioms Code variants are similar to clones in the sense that both are functionally similar set of code snippets. Variants differ from clones for the reasons of purpose and properties in the code context. Variants are always discussed with the intent of *bringing in* code snippets with desired properties. Clones are discussed in the context of refactoring. Developers *clean up* clones to promote reuse. Existence of a clone is considered as a bad smell. There is no difference between the desired properties present in the clone instances.

Gabel’s definition of clone does not capture these aspects of semantic clones. We suggest the following definition:

Definition 4 *Let P be the set of n desired properties $\{p_1, p_2, \dots, p_n\}$. Let $score_{p_i}(\phi)$ be a function computing the strength of code snippet ϕ over any property p_i . Code snippets ν_1 and ν_2 are **clones** if neither ν_1 nor ν_2 score over each other on any property of interest and thus ν_1 and ν_2 can replace each other in the given code context. In other words, $\forall_i score_{p_i}(\nu_1) = score_{p_i}(\nu_2)$.*

This definition does not depend on structural similarity at all, and instead focuses on desired properties in a code context. This emphasis helps us differentiate variants from simions and idioms as well. Simions and idioms are semantically similar irrespective of the code context.

Unlike clones, a stronger variant scores over a weaker variant at least on one desired property. Hence, we rephrase variants defined in Section 2 as follows:

Definition 5 *Code snippets ν_1 and ν_2 are **variants** if there is at least one property of interest by which ν_1 is better than ν_2 , or ν_2 is better than ν_1 in the given code context. Both ν_1 and ν_2 should be acceptable in the current code context i.e., $\exists_i score_{p_i}(\nu_1) \neq score_{p_i}(\nu_2)$.*

Variants are not examples Unlike other types, code examples need not be always similar in behavior. For instance, examples could be instructive to explain API usage in a variety of functionally different snippets.

Variants are not bug-fixes or enhancements Let us assume ν_2 is an enhancement sought over ν_1 . Even though an enhancement may add some desired property to the existing code, we observe that the intent has changed. ν_1 and ν_2 would have been variants in the earlier code context when ν_1 was under development; however, in the new context, ν_2 alone is acceptable and ν_1 does not fit. Same applies to bug-fixes as well. A buggy-snippet does not meet the expectations of the code context, and hence is no more a candidate for being a variant.

Variants are not mutants Mutants are artificially changed code used to assess the quality of test cases. They do not fit into the given code context. Hence, mutants are also not code variants.

In summary, similarities and differences exist among semantic clones, idioms,




Figure 2: In a given code context, if two code snippets ν_1 and ν_2 show the same desired properties, ν_1 and ν_2 are clones as shown in (a). If ν_2 has even more desired properties while having all desired properties as in ν_1 , as shown in (b), ν_2 is a stronger variant of ν_1 . Conversely, ν_1 is a weaker variant of ν_2 . As shown in (c), ν_1 and ν_2 are complex variants if they only have a few desired properties in common.

simions, and code examples, in terms of syntactic and semantic distance, purpose, properties, origin and connotation. Due to lack of clear definitions, the terms semantic clones, examples, and variants are often used interchangeably or even incorrectly in literature and developer discussions. In Table 3, we summarize this discussion by defining, differentiating, and exemplifying each of these terms. Note that none of these types are defined or discussed along with code context or desired properties. Thus, code variants are different from all of these types.

3.1 Types of Variants

From the perspective of mining variants, developers need to decide if a particular variant is better or worse or even incomparable to the rest. To aid this activity, we introduce two categories of variants, namely, simple and complex variants.

Definition 6 We refer to a set of variants as *simple* if for any pair of variants (ν_1, ν_2) in the set, one variant scores not less than the other ($\forall_i \text{score}_{p_i}(\nu_1) \geq \text{score}_{p_i}(\nu_2)$ or $\forall_i \text{score}_{p_i}(\nu_2) \geq \text{score}_{p_i}(\nu_1)$) for all desired properties in a specific code context. Also, recall that $\exists_i \text{score}_{p_i}(\nu_1) \neq \text{score}_{p_i}(\nu_2)$ if ν_1 and ν_2 are variants.

If ν_1 and ν_2 are simple variants in the given code context, and ν_2 is stronger than ν_1 , we mean that ν_2 scores over ν_1 on all desired properties (Figure 2 (b)). Recall that by the definition of variants, there is at least one desired property by which ν_1 is better than ν_2 . In practice, we may find that most efficient solutions may suffer from issues such as readability and licensing, and hence may not be better on all desired properties. It is possible that ν_1 and ν_2 may be equal in the number of desired properties they satisfy but still are not clones because the sets of properties each of them satisfy may have differences.

Strict Partial Order The relation over the strength of code variants (represented by the symbol ‘>’) is a strict partial order over the set \mathcal{V} of variants.

Table 4: The term “variants” is often used by developers. For the Gabel’s dataset, we explored the bug repository and found several discussions around variants.

System	Size (in MLoc)	“Variants” usage
GIMP	.78	109
GTK	.88	324
MySQL	1.13	166
Postgresql	.74	262
Linux Kernel	7.31	286

In other words, $\nu_1 > \nu_1$ cannot hold (irreflexivity) since we need at least one property by which the snippet being compared with should differ to be called as a variant. $\nu_2 > \nu_1$ indicates that ν_2 is a stronger variant of ν_1 , and $\nu_1 > \nu_2$ cannot hold (antisymmetry), and $\nu_2 \neq \nu_1$ (irreflexivity). In addition, if we have ν_3 such that $\nu_3 > \nu_2$, then $\nu_3 > \nu_1$ (transitivity).

In the case of complex variants, it might be possible for developers to apply a weight function to choose a specific complex variant as a strong variant. Without weights or additional such preference information, it will be unclear to developers which variant to select (Figure 2 (c)). Figure 2 (a) shows the case where ν_1 and ν_2 have the same properties, no more or no less and thus they become clones.

Definition 7 We call ν_1 and ν_2 as **complex variants** if ν_1 scores over ν_2 for some desired properties, and ν_2 scores over ν_1 for some other desired properties. More formally, $\exists_{i,j} (\text{score}_{p_i}(\nu_1) > \text{score}_{p_i}(\nu_2)) \wedge (\text{score}_{p_j}(\nu_1) < \text{score}_{p_j}(\nu_2)) \wedge (i \neq j)$.

As an example, consider internet traffic monitoring APIs such as Fiddler and Titanium. A discussion on SO² suggests that Titanium is preferred over Fiddler given the licensing constraints. Hence Titanium is a stronger variant than Fiddler. In this case, Titanium is also a simple variant of Fiddler since it is easier to choose the prior over the latter. As another example for simple variant, an $O(n \log n)$ solution is accepted to be better than $O(n^2)$ solution in the context where worst-case time complexity is the desired property. Consider the use of Multinomial³ and Bernoulli⁴ Naive Bayes implementation choices. There are different contexts in which either of them is preferred. Thus they become complex variants in the context where we need additional information to categorically select one over the other. The three snippets in Figure 3 (a) are examples for clones with respect to speed of execution since all are recursive in nature and have no further complexity.

To the best of our knowledge, none of the existing tools focus on differences in desired properties for the given code context. In this sense, they are not

²<https://tinyurl.com/gr66vje>

³<https://tinyurl.com/gulbmb8>

⁴<https://issues.apache.org/jira/browse/SPARK-489>,
<https://issues.apache.org/jira/browse/OPENNLP-777>

variant mining tools. Next, we discuss the types of code contexts and desired properties in the following sub-sections.

3.2 Code Context

While desired properties distinguish variants, the code context relates them together. Code context acts like candidate variant filtering criteria. The code context captures different aspects such as users’ concerns, and development and execution contexts. As defined in Section 2, code context description comprises of one or more of the following: intent, dependencies, input/output and state. We arrived at this taxonomy based on a literature survey of 11 related research papers [21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31] and our experiments on nine open source projects (See Section 4).

Intent Programs are products of human desire to solve specific problems or accomplish well-defined tasks. Hence, an understanding of the problem being solved plays an important part in identifying or recommending variants that “fit” the purpose. Purpose includes functional and para-functional requirements. For example, “computing factorial”, “implementing little endian algorithm” are examples of intent. Intent specification is a hard problem [21] which goes beyond just naming and describing the problem using natural language phrases. Nguyen et al. [22] relate execution context and intent. Their thesis is that the intent can be captured using the API usage patterns in the current code. Ponzanelli et al. [23] show that fully qualified name of the code elements, current code, custom API types, and method names are strong pointers to intent.

Dependencies Often, implementation is constrained to a specific programming language, certain pre-built libraries, or components. Search for variants must honor these constraints. Constraints may also include structural elements such as methods or classes as in Java. We refer to such constraints as dependencies. For example, a REST API for financial data may be provided by multiple providers which become variant choices. In this context, we assume that non-REST APIs are not sought by the developers. Robillard [24] claim that neglecting such dependencies may lead to *low-quality modifications*. They discuss structural dependencies in the scope of *program elements* and mention *methods* and *fields* as examples.

Input/Output Input and output examples are used as context to search [25] and synthesize [26] source code. Nix argues that the problem of synthesizing expressions mapping given a set of inputs to the given set of outputs (in the sub-context of repetitive text editing) is NP-Hard. Programming-by-Example community shows steps taken in this direction with string transformation [27]. In summary, a decade old research in this area has produced solutions for text

Table 5: Examples of developer discussions taken from open source projects describing the desired properties (of different types) in code variants.

Type	Issue	Developer Discussion
Algorithmic	Eclipse Bug 384730	<i>There are already some implementation of this algorithm. However, most of them are pretty complex and slow. I would like to contribute a smaller and simpler version compatible with the ZEST layout engine.</i>
Algorithmic	Eclipse Bug 409427	<i>We should replace the existing xml format with the much denser and faster loading EMF BinaryResource implementation.</i>
Resource-oriented	HTTPClient- 1267	<i>Does httpcomponents not support transient cookies? If so we cannot use this library for session login, since most session-protected sites use this.</i>
Resource-oriented	Eclipse Bug 293637	<i>Ribbon must be licensed by each adopter. If Eclipse will provide Ribbon, than every RCP application with Ribbon must be licensed. This violates EPL.</i>
Diction	Eclipse Bug 196585	<i>It is better to use the setter methods on the model classes (e.g. TracWikiPageVersion) than having constructors with many parameters. That way the order of the parameters does not get mixed up and the code is easier to refactor and to read.</i>
Diction	Eclipse Bug 338065	<i>Our coding conventions currently demand to declare all method parameters as final in order to prevent parameter assignments. Meanwhile, parameter assignments can effectively be revealed by the Eclipse tooling and by tools like FindBugs on the CI server.</i>

editing and spreadsheet processing; however, synthesizing large sized programs remains a challenge [28].

State Often, developers complain of a specific state that the system gets into. For example, in Zope⁵, a developer states, “*For huge transactions ZEO spends a long time (in the order of minutes) in the call to “vote”. This makes it irresponsible for other request...*”. Current context of the code under execution includes the snapshot of its variables, the line under execution, and the resources available at that time for the program [29]. This definition of context is used heavily in debugging [29], program repair [30], and real-time updates [31] to software systems. These systems use a variety of techniques such as automata and logic for capturing and representing the context.

3.3 Desired Properties

As discussed in Definition 3, desired properties distinguish variants. Desired properties in a variant can be classified into broadly three groups: a) Algorith-

⁵<https://bugs.launchpad.net/zodb/+bug/143274>

mic, b) Resource-Oriented and c) Diction.

Algorithmic Properties Algorithms play a significant part in computation and their properties are well studied [32]. Developers seek efficient algorithms to make their code score on para-functional attributes such as security [33], accuracy [34], readability [15] and scalability [9]. Sridhara et al. [35] discuss the importance of identifying high-level algorithmic steps in source code. Patterns, signatures and structures are limited in their ability to detect algorithms in source code [36]. Many reuse techniques [37, 38] that work at function level or code fragment level focus on semantic similarity and ignore the variability across variants. Mishne et al. [36] extract concept graphs to represent algorithmic information. They use a finite list of concept types, such as loop, assign and block, in their representation. Yet, this model shows good results for code based on C language.

Resource-oriented Properties For reasons such as licensing [39, 40], certain libraries, components, sub-systems, interfaces, and services are considered better or relevant. This property has nothing to do with the syntax or semantics of the code snippet. Instead, it is about the extraneous (non-code) elements associated with the snippet, such as the legal constraints, and trust factors. Long [41] observes that many third-party libraries are no longer actively maintained. He calls this the *used car fiasco*. He brings up more issues in reuse, such as *One size fits all* and *Of course it's reusable*. Moreno et al. [15] discuss the effort to reuse the code snippet.

Diction Properties Diction refers to the *style of speaking or writing as dependent upon choice of words*⁶. Some developers may prefer **for** over **while** to code a loop. Resulting code is semantically the same. Naming conventions may contribute to the ranking of one variant over the other [42]. We call such variants as diction variants. Diction variants cover all non-algorithmic and non-resource-oriented properties, such as patterns, refactoring needs, conventions and style. Often, programming language libraries give multiple ways to implement the same functionality within the same resource and algorithmic constraints. Syntactic sugars [43] are classic examples for this type of variants.

Certain syntactic choices have distinguished benefit over the other. A recursive version of factorial is rarely used in practical scenarios. A memoization approach avoids recomputation and is desired especially when large inputs values are bounded so that a four byte variable such as Java *int* can hold the result. Yet, the role of diction variants have been largely ignored by the research community. In academic context, most plagiarism tools depend on these differences to avoid marking student works as duplicate.

Absence of one or more of these properties leads to low-quality code snippet for which developers seek replacement. This absence may introduce faults, bad smells or sub-optimal code. Table 5 shows real developer discussions from

⁶<http://www.dictionary.com/browse/diction>

Eclipse and HTTPClient projects. We have mapped these discussions to one of the three types of properties discussed.

4 Empirical Evaluation

Here we present the results of our experimental studies structured around a few Research Questions (RQ). The experiment data is shared at the project website⁷.

We use a dataset of fifteen open source projects (Table 6). When making the project selection, we wanted to accommodate a variety of programming languages, domains of use, and project sizes. Three of the projects in our dataset are written in Java, three in C, and the remaining three are in Python. Java is a statically typed and compiled object-oriented language whereas C is not object-oriented, but is static (weakly) typed. Python is a dynamically (strongly) typed and interpreted object-oriented language.

Allamanis et al. [44] evaluated idioms on popular Java projects that were selected based on their z-score in GitHub. We pick Atmosphere and Hibernate from that list. To ensure a variety in the project domains, we replaced one project from the list of Java projects used by Allamanis et al., [44] with Apache Math which is an algorithmically sensitive library. Three of these projects, GIMP, GTK+ and MySQL, written in C, were used by Gabel et al. [9] in their work on semantic clones. Roy et al. [45] in their study of code clones used the Python projects: Plone, SCons, and Zope, which we also select. The projects in our dataset vary from 68K (Atmosphere) to 1130K (MySQL) lines of code, and include disparate domains such as mathematics, databases and editors.

RQ 1: Do variants matter?

Before exploring different properties of variants in depth, we investigated about the relevance of variants and about their nature of presence in the development community. Here we seek to understand how often developers make choices about variants in their projects.

We took 100 random defects in each of the fifteen projects in our dataset. When we analyzed the discussion in these defects, we found 277 variants as shown in Table 7. This indicates that variants exist in abundance. Developers actively compare and seek variants. Table 8 shows a sample of such discussions. Discussions were categorized as about variants if they discussed different forms of implementation or were about implementation choices. We found that variant discussions exist in **25%** (in Zope) to **43%** (in Apache Math) of bug reports in these projects. Discussions about simple variants occurred 17% times on average, as compared to 14% for complex variants. Discussions about simple variants were uncomplicated, where a developer proposed a variant in a patch and was approved. Discussions about complex variants were more involved, going over the pros and cons with respect to the desired properties of the variants.

⁷<http://variants.usite.pro/index.html>

Table 6: Projects dataset used to evaluate if variants exist in real open source projects. Diversity in programming languages (Java, C and Python), size and domain characterizes these projects.

#	Project	PL	LoC	Domain
1	Apache Math	Java	375K	Mathematics
2	Atmosphere	Java	68K	Client-Server
3	Hibernate ORM	Java	930K	Domain Model persistence
4	Gimp	C	780K	Image Manipulation
5	GTK+	C	880K	UI Widget Toolkit
6	MySQL	C	1130K	Database
7	Plone	Python	74K	Content Management
8	SCons	Python	228K	Build Tool
9	Zope	Python	272K	Web Application Server

Table 7: Volume of variant discussions in open source projects depends on the project domain. Alg_v , RO_v and D_v represent algorithmic, resource-oriented and diction variants respectively. Dominating types are in bold.

Project	Simple			Complex			Sum
	Alg_v	RO_v	D_v	Alg_v	RO_v	D_v	
ApacheMath	4	1	8	11	7	12	43
Atmosphere	2	7	8	6	2	5	30
Hibernate	5	3	10	2	4	5	29
GIMP	11	2	9	7	3	3	35
GTK+	6	3	7	9	3	4	32
MySQL	11	1	8	6	4	1	31
Plone	2	2	10	4	3	5	26
SCons	8	0	6	3	2	7	26
Zope	10	1	2	6	3	3	25
Average	7	2	8	6	3	5	31

Overall, *31% bugs discussed variants*. Hence, we conclude that variants play a significant role in software development. This serves as a strong motivation to further explore variant characteristics. We discuss these in the following RQs.

RQ 2: How are variants distributed across the types of desired properties?

As developers are constantly looking for variants to reuse, they may look out for certain properties that make one variant preferred over others for the given project requirements. We were interested in knowing the distribution of these

Table 8: Developer discussions on implementation choices is common in SO. Some questions from SO in which multiple implementation choices were discussed are given below along with related programming language and the reference ID to look up the discussion item.

Ref	Tag	Question
7519283	sql	Fastest way to check if a character is a digit?
1102891	java	How to check if a String is numeric in Java
53513	python	Best way to check if a list is empty
392022	bash	Best way to kill all child processes

variants across different properties. We looked up the discussions talking of variants in the fifteen open source projects to analyze the distribution of these variants with respect to their types: algorithmic, resource-oriented, and diction.

Algorithmic variants were typically discussed because new algorithms were sought in new requirements, enhancements or bug reports. Table 7 shows that there were on average 7% simple and 6% complex algorithmic variants. Thus, on an average, 13% algorithmic variants were found in these bug reports. Only 5% variants (2% simple and 3% complex) belonged to the resource-oriented class. They are discussed only when there is a concern or conflict such as licensing, library compatibility and coding conventions. These concerns can be seen in tracking bugs which call for changes at multiple places while keeping the bug count to just one. In two of the projects (MySQL and Apache Math), we found only one resource-oriented variant in the 100 randomly chosen defects. MySQL had only one simple resource-oriented variant. This does not mean that there are very few resource-oriented variants in these projects. It only suggests that the distribution is skewed towards other types of variants. Moreover, constraints such as licensing, library compatibility and coding conventions are expected to be known to seasoned developers, which could be another reason for a relatively lower count of resource-oriented variants discussions when compared with the count based on other properties discussed. A time-boxed, focused search for an hour resulted in 20 resource-oriented MySQL variants, which shows that discussions on these resource-oriented variants exist, but are less frequent. On average, 13% of discussions were about diction variants (8% simple and 5% complex). This makes a case for the use of type-1 to type-3 clone detection tools as they are useful in this space (diction variants).

Drawing such characteristics could be useful for developers in mining variants based on the properties best suited for the project requirements. For instance, developers looking at optimising program function involving data clustering may be looking for algorithmic variants of clustering. This led us to RQ 4 which studies if these desired properties among variants are dependent on the domain of the project.

Further, we found that 41% of the intents lead to algorithmic variants (Ta-

Table 9: Variants are often sought and their desired properties are discussed in bug reports of real-world projects.

Project	Intent	Variant type	Language	#Bug
Ruby	Reverse	Algorithmic	C	12744
Apache Commons Math	Sine	Resource-Oriented	Java	MATH-901
CiviCRM	Maps API	Resource-Oriented	PHP	CRM-524
Moodle	charting	Resource-Oriented	Javascript	MDL-55090
ReactOS	Convert	Resource-Oriented	C	CORE-5849
Apache Lens	Query	Algorithmic	SQL	LENS-1028
Hadoop Map Reduce	Get NameNode URI	Diction	Java	MAP-REDUCE-6483

ble 7). These were largely due to differences in requirements of performance or accuracy of results. For example, in Ruby project, an implementation referred to as `reverse.each_char` that uses `rb_enc_left_char_head` covers all encodings to scan a string backwards⁸. Overall, 18% of intents led to resource-oriented variants. These were largely due to limitations and constraints (such as license) on the reuse of third-party libraries. For example, porting existing code from a different programming language was cited as an issue in some of the bug reports. Finally, diction variants were also a significant part of developer discussions. About 41% of existing variants belonged to this type. Typical diction variants discussions were about better coding conventions and styles.

RQ 3: How are variants distributed across the context types?

Recall that we identified intent, input/output, dependencies and state as the four major constituents of code context. To study the characteristics based on which one could classify one code as a variant of another it is essential to spot the common attributes among variants. To find such attributes we investigate the variant discussions. Further, to validate if the identified four characteristics indeed capture the code context of variants, we classify the contexts described along with the variant discussions in our dataset into these four constituents. Table 11 gives an essence of the characteristics based on which the contexts were classified. We were successful in classifying all the contexts into these

⁸<https://bugs.ruby-lang.org/issues/12744>

Table 10: Code contexts are primarily described using one or more of these four types. We map bug descriptions containing variants to these types.

#	Project	Intent	Dep.	I/O	State
1	Apache Math (Java)	33	11	2	0
2	Atmosphere (Java)	24	12	2	1
3	Hibernate ORM (Java)	18	9	1	3
4	Gimp (C)	16	6	3	10
5	GTK+ (C)	23	4	4	8
6	MySQL (C)	21	3	9	9
7	Plone (Python)	19	7	8	6
8	SCons (Python)	17	10	0	3
9	Zope (Python)	20	7	4	10

four constituents. We randomly selected 100 of these defects and cross-verified to ensure inter-rator agreement. We found no ambiguity in the understanding of context boundaries and thus the categorization was found to be accurate. Table 10 shows the breakdown per project. We found that intent dominates (in 191 out of 343 contexts found i.e., 56%) across the bug reports as most reports discussed were built upon the underlying functionality of code variants. Dependency, input/output and state covered 20%, 10% and 15% of contexts in bug reports, respectively. Identifying the nature of similarities could help developers filter out candidates for variants.

RQ 4: Are variants domain-dependent?

We investigated whether the nature of the project (domain) is connected to the existence or importance of the variants. We observed that the distribution of the type of variants changes heavily across projects. Dominant types are highlighted (in bold) in Table 7. Algorithmic variants dominated in GIMP. As a project about image editing, the code snippets discussed were on intents, such as finding optimal palettes and usage of image masks that were algorithmic. Atmosphere is a client server framework, and it discussed several (30% i.e., 9 out of 30) resource-oriented variants. Diction variants appeared in large numbers (58% i.e., 15 out of 26) in Plone as it talked about removing deprecated exception handling syntax and labels that were not in sorted order in the code. Some sample discussions supporting such observations are shown in Table 12. From these instances, we gather that domain characteristics influence reuse in different ways, such as library usage, licensing concerns, algorithmic depth, performance requirements, and so on. This leads to an uneven distribution of variants across types.

Table 11: Variant discussions from the dataset capturing the code contexts.

Code Context	Bug ID	Excerpts from the Discussion	Essence of the Discussion w.r.t the context type
Intent	MATH-785	S1: The ContinuedFraction calculation can underflow in the evaluate method, similar to the overflow case already dealt with. S2: The evaluation of the continued fraction has been changed to the modified Lentz-Thompson algorithm which does not suffer from underflow/overflow problems as the original implementation.	Variants have a common intent here which is to compute continued fraction.
State	Atmosphere-2037	S1: <i>the _disconnect method is not being called for Android Chrome when closing the browser though the android apps panel. Only in case of changing the URL on the browser, that method is called.</i> S2: <i>It works with Firefox desktop, Firefox mobile, chrome desktop, etc.</i>	State of disconnect method invocation on closing different browsers being discussed.
Input/Output	MATH-1143	S1: <i>A DerivativeStructure and UnivariateDifferentiableFunction are great tools if one needs to investigate the whole function but are not convenient if one just needs derivative in a given point.</i> S2: <i>Give the derivatives in the “natural” order, which is in increasing order when you have one parameter and high order derivatives, and in parameters order when you have only first order derivatives for all parameters.</i>	Common output of derivative, across all variants.
Dependency	MATH-1098	<i>As we will certainly not add a dependency to another library, we could start with our own set of annotations</i>	Two variants of a function where both do not have a dependency on another library.

Table 12: Samples from the dataset showing the influence of domain on distribution of variants.

Code Context	Domain	Bug ID	Excerpts from the Discussion	Essence of the Discussion w.r.t the domain
GIMP (Algorithmic)	Image Manipulation	GIMP-66257	<i>S1:When the Gimp converts an image to indexed mode using “Generate Optimal Palette” and there are more colors in the RGB image than the palette can hold, then all colors in the image are changed. This is not a problem in some cases (and it is even desirable most of the time because the averaged colors give better results) but in other cases this can lead to bad results if some of the colors had to keep a precise value. S2:One relatively painless method/hack would be to add yet another pass over the image data which snaps all of the colours in the final palette back to their closest in-image colour before doing the final remap or dither.</i>	Precision being discussed across variants of converting image to indexed mode.
Atmosphere (Resource-Oriented)	Client-Server	Atmosphere-888	<i>Add timeout support for WebSocket to prevent thread waiting indefinitely</i>	A variant to better manage threads(resources).

Plone (Diction)	Content Man- agement	Plone-732	<i>The toolbar displays the "xx days ago" information during the loading of a page as ISO date time string and then turns it into "xx days ago" after the complete loading of the HTML. This is an visual artifact and confuses the eye especially when the page is loaded over a slow line or takes some time for rendering. The toolbar should only display "xx days ago" and not the ISO string at all. Perhaps the toolbar or this particular toolbar items should be hidden by default and made visible after moment.js</i>	Alter the content displayed.
--------------------	----------------------------	-----------	--	------------------------------------

RQ 5: Are variants language dependent?

We analyzed if programming languages affected the volume of variants (see Table 6 and 7). It was interesting to see if constructs of a given programming language impacted implementation choices. Automated techniques for variant mining will be harder to devise if the volume of variants are programming language dependent.

Volume of variants in Java projects ranged from 29% to 43%. C projects ranged from 31% to 35% (See Sum column in Table 7). Python has the least occurrence of variants of 25% to 26%. Python projects were considerably smaller as compared to Java and C projects. Hence, that could be one reason for reduced discussion on variants. Overall, we do not see any conclusive evidence to show that programming languages affect the existence of variants.

5 Application of Variant Characterization

As an application of the usage of characterization of variants, we investigate ranking of code search results based on "desired properties" in which a user

is interested. Search engines typically rank results based on relevance of query terms. Recently, researchers have focused on the diversity of search results based on structures and semantics of source code [46]. Our hypothesis is that diversity of search results based on a set of desired properties, is helpful.

Problem Statement We formulate the problem of ranking code snippets for a query phrase based on the strength of desired properties. Our objective is to compute a $score_P(\phi, q)$, where P denotes the set of desired properties and q is a query phrase. Each property p_i is represented using natural language terms for a specific code snippet ϕ . The $score_P(\phi, q)$ gives the strength of the variants. Recall that two code snippets with no difference in desired properties should have the same $score_P(\phi, q)$ as mentioned in Definition 4.

Building a knowledgebase of snippets and their properties We observed that SO posts contain discussions on desired properties of code snippets. Our approach started with the construction of a list of known desired properties. For each property, we collected stemmed synonyms ($p_i, syn(p_i)$) and antonyms ($p_i, ant(p_i)$). We used the Snowball stemmer [47] for this purpose. From each SO post, we collected the code snippets, and then computed the term frequencies ($tf(.)$) of synonyms and antonyms. After this exercise, we ended up with a knowledgebase of triples $(\phi, p_i, score)$, where score (see Equation 1) is the strength of desired property p_i demonstrated by the code snippet ϕ . This score can be used to boost the ranking of snippets based on desired properties while using ranking algorithms like BM25 [48].

$$score_{p_i}(\phi) = \sum_{s \in syn(p_i)} tf(s) - \sum_{a \in ant(p_i)} tf(a) \quad (1)$$

Knowledgebase compression We computed scores for each snippet per post. Duplicate code snippets across posts carried different scores for the same property since the terms surrounding them were different. As the frequency of snippets increase, the average of the scores converge.

There are a variety of approaches to compute structural similarity of source code [49, 50, 51]. We used the de-duplication approach explained in [52] for its simplicity and effectiveness for method-level similarity computation. In this approach, code ϕ is transformed into a set of predefined terms ψ to represent programming structures. This approach traverses through the code snippet, and collects loops, operators and conditionals in the same sequence as they appear in the snippet. Assignments, returns, method declarations, comments and rest of the code except loops, conditionals and expressions, are ignored. For instance, the transformed snippet for the first snippet in Figure 3(a) is `if<= if* if-`. This denotes that the operators `<=`, `*` and `-` are used inside a conditional statement. The structural similarity between two code snippets ϕ_1 and ϕ_2 is:

$$similarity(\phi_1, \phi_2) = \frac{|\psi_1 \cap \psi_2|}{\max\{|\psi_1|, |\psi_2|\}} \quad (2)$$

As we merge duplicate snippets, we aggregate their scores as $(\phi, p_i, \sum(score))$. Snippets that are exactly the same will have a similarity score of 1, and snippets with no common transformed structures will have a similarity score of 0.

```

public float factorial
(float num){
  if(num<=1)
    return 1;
  else
    return num *
      factorial(num-1);
}

public int factorial
(int input){
  if(input==0)
    return 1;
  else
    return input*
      this.factorial
        (input-1);
}

int fact(int n) {
  if ( n==1)
    return 1;
  return
    fact (n-1) * n;
}

```

(a) Top-3 results from CodeExchange for query “factorial” before applying our ranking algorithm.

```

public static double
factorial(int s) {
  if (s < 0 || s > 17)
    return Double.NaN;
  double[] a = { 1.0,
    1.0, 2.0, 6.0, 24.0,
    ...
    355687428096000.0 };
  return a[s];
}

public BigInteger
asBigInteger()
{
  BigInteger result =
    BigInteger.ONE;
  for (int i = 2; i <= n;
    ++i) result = result.
    multiply( BigInteger.
      valueOf(i));
  return result;
}

public int factorial
(int n) {
  if (n < 0) throw new
    CustomException();
  if (n == 0)
    return 1;
  if (n == 1)
    return 1;
  return n *
    factorial(n - 1);
}

```

(b) Top-3 results after ranking on the desired property of “speed of execution” for the query “factorial”.

Figure 3: Comparison of results before and after r-ranking.

Ranking search results Given a set of search results, we matched each code snippet in it with the knowledgebase. We used the same Equation 2 to find similarity of code snippets. To compute the score on multiple properties, we computed the score for each property (p_i) and added them to get the overall score, $(\phi, P, \sum(score))$. Figure 3 shows the top three results before and after

Table 13: Average precision before (P_1) and after (P_2) the application of re-ranker for 10 queries for the desired attribute of “speed of execution”.

Query	P_1	P_2	Query	P_1	P_2
factorial	0.1	0.7	read from file	0.3	0.6
substring search	0.1	0.6	parse xml	0.2	0.6
finding duplicate	0.2	0.5	calculate mean	0.1	0.4
matrix multiplication	0.1	0.2	iterate over list	0.1	0.4
serialization	0.1	0.6	write log file	0.4	0.5

applying our ranking algorithm. Notice that the first result is a switch-case based implementation, which is argued to be faster in SO. Apache Math uses this implementation for factorial.

Structural Heterogeneity Results can be improved by dropping all but the first result that look structurally similar. CodeExchange [46], a code search engine, shows 16 recursive factorial results out of the first 20 results. Since we are only interested in diversity with respect to desired properties, we used the approach detailed in [52] to drop similar results at each rank. Figure 3 shows the results after application of structural heterogeneity.

Evaluation To evaluate this approach, we took the top-10 results of a search on “factorial” in CodeExchange, and compared it with results from our ranking approach for the desired property, “speed of execution”, over the same underlying content. We chose speed of execution as it plays an important role in software development, especially in mobile environments [53]. We took the top-10 results after applying our ranking algorithm on the same content as CodeExchange (code snippets from GitHub), for each query and compared them with the top-10 results of CodeExchange. To decide if a code snippet is a variant, we used Definition 4 and Definition 5. We compared the search results ranking using Mean Average Precision (MAP) over ten queries as shown in Table 13. Our approach improves the MAP from 0.17 to 0.51.

Discussion Our approach is a proof of concept for mining variants. Precision of our approach depends on the precision of several sub-steps, such as similarity computation, de-duplication, synonyms and antonyms computation. We assume that if a desired property is strongly demonstrated in a code snippet, that property will be discussed in the SO post. This assumption may not hold true in all posts. Also, the synonyms and antonyms that show up in the post may refer to code snippets other than the one added to the answer. Further, we have ignored negative qualifiers, and hence our results are not always accurate. For example, “not efficient” is considered as “efficient”. Although this problem can be solved, our focus here is to demonstrate a simple feasible ranking approach, we leave the natural language processing problem for future work. Other limitations to this proof of concept are posts that carry multiple code

snippets, in which case all snippets get the same score. Further, some of these code snippets might not even implement the given topic. For instance, an SO post on calculating averages that is concerned about exceeding the datatype double’s limits had a code snippet to explain double representation⁹.

While these limitations exist, our objective is to show that ranking based on desired properties of code snippet is feasible even with a simple approach. Figure 3 compares our results with the default output of CodeExchange.

More applications We believe this characterization of variants opens up several new applications. For instance, code completion can be sensitive to desired properties. Learning API usage can also benefit from this fundamental idea of looking at code snippets along with the desired properties. Often, when a better code is shown to the developers, they are able to understand the flaws in the existing code. Thus, this approach can help in bug detection as well.

6 Related Work

To the best of our knowledge, this is the first characterization study on code variants. However, there is work on independent aspects related to variants such as automatically mining relevant code [54, 55, 14], and calculating similarity measures [52, 36] on source code snippets. There is also work on extracting code contexts [22, 23]. The studies on code clones [56, 7] are also close to our work. However, none of these works consider code variants as the first-class entity and investigate mining, clones, or context from that perspective.

Searching/Mining Relevant Code Snippets Researchers have largely ignored the search for “better” source code snippets to replace an existing snippet. This requires ranking of snippets based on properties of interest to the developers. Existing research related to searching source code snippets focus on handling partial programs [57], indexing [56], refining queries [58], and such aspects of building a search system. Majority of existing research which compares code snippets is limited to structural [37] and semantic similarity [11].

Extracting Code Context Importance of understanding and extracting code context is highlighted in several existing works [59, 23] as discussed in Section 3.2. However, none of these papers study code context in detail. Our work builds upon their findings and provides a thorough review of different types of code contexts.

Studies on redundancies in source code Hindle et al. [4] observe that source code being a human product is repetitive. Juergens et al. [3] claim that semantically similar code taken from various sources can be syntactically

⁹<http://stackoverflow.com/questions/1930454/what-is-a-good-solution-for-calculating-an-average-where-the-sum-of-all-values-e>

heterogeneous. In another study [7], they also report that inconsistent changes to clones lead to maintenance issues. We leverage these studies, but find that semantic clones have inconsistent definition and code variants are conceptually different from them.

Diversity in search results Ambiguous queries are shown to demand diversity in search results [60]. Martie and van der Hoek [46] explore diversity in code search results. Our work differs from them because of the computation of diversity based on desired attributes. We believe that our work complements the existing work on computing diversity in code search results. Developers benefit from avoiding looking at semantically similar search results.

7 Threats to Validity

Internal Validity Studying code variants on open source projects poses several challenges. We have used nine open source projects for evaluation. We strived to control variability by selecting 100 random bugs each from multiple projects. Other projects and domains might give different numbers. We have selected each project from a different domain to mitigate this threat. It is possible that these bugs are not representative of these projects. Also, these projects may not be representative of the entire set of open source projects. Choosing multiple projects was an attempt to reduce this kind of bias. However, we have ensured that there is no bias in terms of project size or popularity within our data set.

External Validity Our results may not generalize to all types of code (for instance, scripting or functional). For evaluation, we have taken a mix of Java, C and Python projects popularly used in clone research. Hence, our work applies to high-level imperative languages at the least. Our results for variant mining depends on the discussion forum data. These results may not generalize to variants that are domain specific implementations with inadequate developer discussions.

8 Conclusion

Code variants are very different from other structural and semantically similar code snippets such as code clones. They appear frequently in source code and developer discussions. Therefore, understanding code variants is important for software development and maintenance. Currently, there are inconsistent definitions of code snippets that are similar. To the best of our knowledge, we present the first study to characterize and distinguish code variants from other types of code that are similar to each other. In this work, we define code variants, classify them as simple and complex, and categorize them into three main types: algorithmic, diction and resource-oriented.

With the availability of code from open source projects and discussion forums, developers are increasingly turning to online sources for finding implementation that match specific desired behavior. Further, with “big code” becoming accessible to the research community, an understanding of code variants and their characteristics can help build tools that provide automated code (variant) searches, including functionality such as, ranking code search results, mining source code variants, and recommender systems for variants. As a demonstration of the possible use of our characterization of code variants, we show how search results from a code search engine can be reordered based on a desired (implementation) property that a developer might have (e.g., speed of execution). This constitutes a first step towards an open research area where code variants can be mined from heterogenous sources (e.g., project code, issue tracker discussions, Q&A forums) by using the code context and the desired behavior that a developer is looking for.

9 Acknowledgments

References

- [1] Miltiadis Allamanis and Charles Sutton. Mining idioms from source code. In *Proceedings of the 22Nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, FSE 2014, pages 472–483, New York, NY, USA, 2014. ACM.
- [2] Mark Gabel and Zhendong Su. A study of the uniqueness of source code. In *Proceedings of the Eighteenth ACM SIGSOFT International Symposium on Foundations of Software Engineering*, FSE ’10, pages 147–156, NY, USA, 2010. ACM.
- [3] Elmar Juergens, Florian Deissenboeck, and Benjamin Hummel. Code similarities beyond copy & paste. In *Proceedings of the 2010 14th European Conference on Software Maintenance and Reengineering*, CSMR ’10, pages 78–87, Washington, DC, USA, 2010. IEEE Computer Society.
- [4] Abram Hindle, Earl T. Barr, Zhendong Su, Mark Gabel, and Premkumar Devanbu. On the naturalness of software. In *Proceedings of the 34th International Conference on Software Engineering*, ICSE ’12, pages 837–847, Piscataway, NJ, USA, 2012. ACM.
- [5] Saurav Muralidharan, Amit Roy, Mary Hall, Michael Garland, and Piyush Rai. Architecture-adaptive code variant tuning. In *Proceedings of the Twenty-First International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS ’16, pages 325–338, New York, NY, USA, 2016. ACM.
- [6] Jens Krinke. Effects of context on program slicing. *J. Syst. Softw.*, 79(9):1249–1260, September 2006.

- [7] Elmar Juergens, Florian Deissenboeck, Benjamin Hummel, and Stefan Wagner. Do code clones matter? In *Proceedings of the 31st International Conference on Software Engineering*, ICSE '09, pages 485–495, Washington, DC, USA, 2009. IEEE Computer Society.
- [8] Andrian Marcus and Jonathan I. Maletic. Identification of high-level concept clones in source code. In *Proceedings of the 16th IEEE International Conference on Automated Software Engineering*, ASE '01, pages 107–, Washington, DC, USA, 2001. IEEE Computer Society.
- [9] Mark Gabel, Lingxiao Jiang, and Zhendong Su. Scalable detection of semantic clones. In *Proceedings of the 30th International Conference on Software Engineering*, ICSE '08, pages 321–330, New York, NY, USA, 2008. ACM.
- [10] Rochelle Elva and Gary T. Leavens. JSCTracker: A semantic clone detection tool for Java code. Technical Report CS-TR-12-04, Computer Science, University of Central Florida, Orlando, Florida, March 2012.
- [11] Rainer Koschke. Survey of research on software clones. In *Duplication, Redundancy, and Similarity in Software*, 23.07. - 26.07.2006, 2006.
- [12] Iman Keivanloo, Chanchal K. Roy, and Juergen Rilling. Sebyte: A semantic clone detection tool for intermediate languages. In *Program Comprehension (ICPC), 2012 IEEE 20th International Conference on*, ICPC '12, pages 247–249, June 2012.
- [13] Oleksandr Panchenko, Hasso Plattner, and Alexander Zeier. What do developers search for in source code and why. In *Proceedings of the 3rd International Workshop on Search-Driven Development: Users, Infrastructure, Tools, and Evaluation*, SUITE '11, pages 33–36, New York, NY, USA, 2011. ACM.
- [14] Naiyana Sahavechaphan and Kajal Claypool. Xsnippet: Mining for sample code. In *Proceedings of the 21st Annual ACM SIGPLAN Conference on Object-oriented Programming Systems, Languages, and Applications*, OOPSLA '06, pages 413–430, New York, NY, USA, 2006. ACM.
- [15] Laura Moreno, Gabriele Bavota, Massimiliano Di Penta, Rocco Oliveto, and Andrian Marcus. How can i use this method? In *Proceedings of the 37th International Conference on Software Engineering - Volume 1*, ICSE '15, pages 880–890, Piscataway, NJ, USA, 2015. ACM.
- [16] Jonathan Sillito, Frank Maurer, Seyed Mehdi Nasehi, and Chris Burns. What makes a good code example?: A study of programming q&a in stackoverflow. In *Proceedings of the 2012 IEEE International Conference on Software Maintenance*, ICSM '12, pages 25–34, Washington, DC, USA, 2012. IEEE Computer Society.

- [17] Joel Brandt, Mira Dontcheva, Marcos Weskamp, and Scott R. Klemmer. Example-centric programming: Integrating web search into the development environment. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 513–522, New York, NY, USA, 2010. ACM.
- [18] Collin McMillan, Denys Poshyvanyk, and Mark Grechanik. Recommending source code examples via api call usages and documentation. In *Proceedings of the 2Nd International Workshop on Recommendation Systems for Software Engineering*, RSSE '10, pages 21–25, New York, NY, USA, 2010. ACM.
- [19] Babak Salamat, Todd Jackson, Andreas Gal, and Michael Franz. Orchestra: Intrusion detection using parallel execution and monitoring of program variants in user-space. In *Proceedings of the 4th ACM European Conference on Computer Systems*, EuroSys '09, pages 33–46, New York, NY, USA, 2009. ACM.
- [20] Julia Rubin and Marsha Chechik. A framework for managing cloned product variants. In *Proceedings of the 2013 International Conference on Software Engineering*, ICSE '13, pages 1233–1236, Piscataway, NJ, USA, 2013. ACM.
- [21] Nancy G. Leveson. Intent specifications: An approach to building human-centered specifications. *IEEE Trans. Softw. Eng.*, 26(1):15–35, January 2000.
- [22] Anh Tuan Nguyen, Tung Thanh Nguyen, Hoan Anh Nguyen, Ahmed Tamrawi, Hung Viet Nguyen, Jafar Al-Kofahi, and Tien N. Nguyen. Graph-based pattern-oriented, context-sensitive source code completion. In *Proceedings of the 34th International Conference on Software Engineering*, ICSE '12, pages 69–79, Piscataway, NJ, USA, 2012. ACM.
- [23] Luca Ponzanelli, Gabriele Bavota, Massimiliano Di Penta, Rocco Oliveto, and Michele Lanza. Mining stackoverflow to turn the ide into a self-confident programming prompter. In *Proceedings of the 11th Working Conference on Mining Software Repositories*, MSR 2014, pages 102–111, New York, NY, USA, 2014. ACM.
- [24] Martin P. Robillard. Topology analysis of software dependencies. *ACM Trans. Softw. Eng. Methodol.*, 17(4):18:1–18:36, 2008.
- [25] Kathryn T. Stolee, Sebastian Elbaum, and Daniel Dobos. Solving the search for source code. *ACM Trans. Softw. Eng. Methodol.*, 23(3):26:1–26:45, June 2014.
- [26] Robert P. Nix. Editing by example. *ACM Trans. Program. Lang. Syst.*, 7(4):600–621, October 1985.

- [27] Rishabh Singh and Sumit Gulwani. Learning semantic string transformations from examples. *Proc. VLDB Endow.*, 5(8):740–751, April 2012.
- [28] Martin C. Rinard. Example-driven program synthesis for end-user programming: Technical perspective. *Commun. ACM*, 55(8):96–96, August 2012.
- [29] Tom Bergan, Dan Grossman, and Luis Ceze. Symbolic execution of multi-threaded programs from arbitrary program contexts. In *Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages & Applications, OOPSLA '14*, pages 491–506, New York, NY, USA, 2014. ACM.
- [30] Aritra Dhar, Rahul Purandare, Mohan Dhawan, and Suresh Rangaswamy. Clotho: Saving programs from malformed strings and incorrect string-handling. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering, ESEC/FSE 2015*, pages 555–566, New York, NY, USA, 2015. ACM.
- [31] Jérémy Buisson and Fabien Dagnat. Recaml: Execution state as the cornerstone of reconfigurations. *SIGPLAN Not.*, 45(9):27–38, September 2010.
- [32] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition, 2009.
- [33] Hossain Shahriar and Mohammad Zulkernine. Mitigating program security vulnerabilities: Approaches and challenges. *ACM Comput. Surv.*, 44(3):11:1–11:46, June 2012.
- [34] Lingxiao Jiang, Ghassan Mishserghi, Zhendong Su, and Stephane Glondu. Deckard: Scalable and accurate tree-based detection of code clones. In *Proceedings of the 29th International Conference on Software Engineering, ICSE '07*, pages 96–105, Washington, DC, USA, 2007. IEEE Computer Society.
- [35] Giriprasad Sridhara, Lori Pollock, and K. Vijay-Shanker. Automatically detecting and describing high level actions within methods. In *Proceedings of the 33rd International Conference on Software Engineering, ICSE '11*, pages 101–110, New York, NY, USA, 2011. ACM.
- [36] Gilad Mishne and Maarten de Rijke. Source code retrieval using conceptual similarity. In *Coupling Approaches, Coupling Media and Coupling Languages for Information Retrieval, RIAO '04*, pages 539–554, Paris, France, France, 2004. ACM.
- [37] Sushil K. Bajracharya, Joel Ossher, and Cristina V. Lopes. Leveraging usage similarity for effective retrieval of examples in code repositories. In *Proceedings of the Eighteenth ACM SIGSOFT International Symposium on*

Foundations of Software Engineering, FSE '10, pages 157–166, New York, NY, USA, 2010. ACM.

- [38] Shaowei Wang, David Lo, and Lingxiao Jiang. Active code search: Incorporating user feedback to improve code search relevance. In *Proceedings of the 29th ACM/IEEE International Conference on Automated Software Engineering*, ASE '14, pages 677–682, New York, NY, USA, 2014. ACM.
- [39] Yuhao Wu, Yuki Manabe, Tetsuya Kanda, Daniel M. German, and Katsuro Inoue. A method to detect license inconsistencies in large-scale open source projects. In *Proceedings of the 12th Working Conference on Mining Software Repositories*, MSR '15, pages 324–333, Piscataway, NJ, USA, 2015. ACM.
- [40] Thomas A. Alspaugh, Hazeline U. Asuncion, and Walt Scacchi. Analyzing software licenses in open architecture software systems. In *Proceedings of the 2009 ICSE Workshop on Emerging Trends in Free/Libre/Open Source Software Research and Development*, FLOSS '09, pages 54–57, Washington, DC, USA, 2009. IEEE Computer Society.
- [41] John Long. Software reuse antipatterns. *SIGSOFT Softw. Eng. Notes*, 26(4):68–76, July 2001.
- [42] Miltiadis Allamanis, Earl T. Barr, Christian Bird, and Charles Sutton. Suggesting accurate method and class names. In *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering*, ESEC/FSE 2015, pages 38–49, New York, NY, USA, 2015. ACM.
- [43] Justin Pombrio and Shriram Krishnamurthi. Resugaring: Lifting evaluation sequences through syntactic sugar. *SIGPLAN Not.*, 49(6):361–371, June 2014.
- [44] Miltiadis Allamanis and Charles Sutton. Mining source code repositories at massive scale using language modeling. In *Proceedings of the 10th Working Conference on Mining Software Repositories*, MSR '13, pages 207–216, Piscataway, NJ, USA, 2013. ACM.
- [45] Chanchal K. Roy and James R. Cordy. Are scripting languages really different? In *Proceedings of the 4th International Workshop on Software Clones*, IWSC '10, pages 17–24, New York, NY, USA, 2010. ACM.
- [46] Lee Martie and André van der Hoek. Sameness: An experiment in code search. In *Proceedings of the 12th Working Conference on Mining Software Repositories*, MSR '15, pages 76–87, Piscataway, NJ, USA, 2015. IEEE Press.
- [47] The english (porter2) stemming algorithm. <http://snowball.tartarus.org/algorithms/english/stemmer.html>, Accessed 02-02-2017.

- [48] Stephen Robertson, Hugo Zaragoza, and Michael Taylor. Simple bm25 extension to multiple weighted fields. In *Proceedings of the Thirteenth ACM International Conference on Information and Knowledge Management, CIKM '04*, pages 42–49, New York, NY, USA, 2004. ACM.
- [49] Yoshiki Higo and Shinji Kusumoto. How should we measure functional sameness from program source code? an exploratory study on java methods. In *Proceedings of the 22Nd ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE 2014*, pages 294–305, New York, NY, USA, 2014. ACM.
- [50] Fang-Hsiang Su, Jonathan Bell, Kenneth Harvey, Simha Sethumadhavan, Gail Kaiser, and Tony Jebara. Code relatives: Detecting similarly behaving software. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE 2016*, pages 702–714, New York, NY, USA, 2016. ACM.
- [51] Saul Schleimer, Daniel S. Wilkerson, and Alex Aiken. Winnowing: Local algorithms for document fingerprinting. In *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data, SIGMOD '03*, pages 76–85, New York, NY, USA, 2003. ACM.
- [52] Venkatesh Vinayakara, Rahul Purandare, and Aditya V. Nori. Structurally heterogeneous source code examples from unstructured knowledge sources. In *Proceedings of the 2015 Workshop on Partial Evaluation and Program Manipulation, PEPM '15*, pages 21–26, New York, NY, USA, 2015. ACM.
- [53] Yepang Liu, Chang Xu, and Shing-Chi Cheung. Characterizing and detecting performance bugs for smartphone applications. In *Proceedings of the 36th International Conference on Software Engineering, ICSE 2014*, pages 1013–1024, New York, NY, USA, 2014. ACM.
- [54] Iman Keivanloo, Juergen Rilling, and Ying Zou. Spotting working code examples. In *Proceedings of the 36th International Conference on Software Engineering, ICSE '14*, pages 664–675, New York, NY, USA, 2014. ACM.
- [55] Valerio Terragni, Yepang Liu, and Shing-Chi Cheung. Csnippex: Automated synthesis of compilable code snippets from q&a sites. In *Proceedings of the 25th International Symposium on Software Testing and Analysis, ISSTA 2016*, pages 118–129, New York, NY, USA, 2016. ACM.
- [56] Mu-Woong Lee, Jong-Won Roh, Seung-won Hwang, and Sunghun Kim. Instant code clone search. In *Proceedings of the Eighteenth ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE '10*, pages 167–176, New York, NY, USA, 2010. ACM.
- [57] Alon Mishne, Sharon Shoham, and Eran Yahav. Typestate-based semantic code search over partial programs. In *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and*

Applications, OOPSLA '12, pages 997–1016, New York, NY, USA, 2012. ACM.

- [58] Sonia Haiduc, Giuseppe De Rosa, Gabriele Bavota, Rocco Oliveto, Andrea De Lucia, and Andrian Marcus. Query quality prediction and reformulation for source code search: The refoqus tool. In *Proceedings of the 2013 International Conference on Software Engineering, ICSE '13*, pages 1307–1310, Piscataway, NJ, USA, 2013. ACM.
- [59] Michael Poulin. A task of behaviour modelling in execution context. In *Proceedings of the 2014 Workshop on Behaviour Modelling-Foundations and Applications, BM-FA '14*, pages 8:1–8:3, New York, NY, USA, 2014. ACM.
- [60] Michael J. Welch, Junghoo Cho, and Christopher Olston. Search result diversity for informational queries. In *Proceedings of the 20th International Conference on World Wide Web, WWW '11*, pages 237–246, New York, NY, USA, 2011. ACM.