

# How are functionally similar code clones syntactically different? An empirical study and a benchmark

Stefan Wagner, Asim Abdulkhaleq, Ivan Bogicevic, Jan-Peter Ostberg, Jasmin Ramadani

**Background.** Today, redundancy in source code, so-called “clones”, caused by copy&paste can be found reliably using clone detection tools. Redundancy can arise also independently, however, caused not by copy&paste. At present, it is not clear how only *functionally similar clones* (FSC) differ from clones created by copy&paste. Our aim is to understand and categorise the syntactical differences in FSCs that distinguish them from copy&paste clones in a way that helps clone detection research. **Methods.** We conducted an experiment using known functionally similar programs in Java and C from coding contests. We analysed syntactic similarity with traditional detection tools and explored whether concolic clone detection can go beyond syntax. We ran all tools on 2,800 programs and manually categorised the differences in a random sample of 70 program pairs. **Results.** We found no FSCs where complete files were syntactically similar. We could detect a syntactic similarity in a part of the files in < 16 % of the program pairs. Concolic detection found 1 of the FSCs. The differences between program pairs were in the categories algorithm, data structure, OO design, I/O and libraries. We selected 58 pairs for an openly accessible benchmark representing these categories. **Discussion.** The majority of differences between functionally similar clones are beyond the capabilities of current clone detection approaches. Yet, our benchmark can help to drive further clone detection research.

# How Are Functionally Similar Code Clones Syntactically Different?

## An Empirical Study and a Benchmark

Stefan Wagner\*, Asim Abdulkhaleq, Ivan Bogicevic, Jan-Peter Ostberg, and Jasmin Ramadani

Institute of Software Technology, University of Stuttgart, Germany

### Abstract

**Background.** Today, redundancy in source code, so-called “clones”, caused by copy&paste can be found reliably using clone detection tools. Redundancy can arise also independently, however, caused not by copy&paste. At present, it is not clear how only functionally similar clones (FSC) differ from clones created by copy&paste. Our aim is to understand and categorise the syntactical differences in FSCs that distinguish them from copy&paste clones in a way that helps clone detection research.

**Methods.** We conducted an experiment using known functionally similar programs in Java and C from coding contests. We analysed syntactic similarity with traditional detection tools and explored whether concolic clone detection can go beyond syntax. We ran all tools on 2,800 programs and manually categorised the differences in a random sample of 70 program pairs.

**Results.** We found no FSCs where complete files were syntactically similar. We could detect a syntactic similarity in a part of the files in < 16 % of the program pairs. Concolic detection found 1 ‰ of the FSCs. The differences between program pairs were in the categories *algorithm*, *data structure*, *OO design*, *I/O* and *libraries*. We selected 58 pairs for an openly accessible benchmark representing these categories.

**Discussion.** The majority of differences between functionally similar clones are beyond the capabilities of current clone detection approaches. Yet, our benchmark can help to drive further clone detection research.

---

\*Corr. author: Stefan Wagner, Universitätsstr. 38, 70569 Stuttgart, Germany, phone +49 711 685 88455, stefan.wagner@informatik.uni-stuttgart.de

# 1 Introduction

As software is now a key ingredient of current complex systems, the size of software systems is continuously increasing. While software with a code size of several thousand lines has been considered large in the seventies and eighties, we now reach code sizes of tens to hundreds of millions of lines of code. This has strong effects on the complexity and manageability of these systems and, as a result, on the cost of maintaining them.

By abstraction and code generation, modern programming languages and development techniques help to reduce the amount of code we have to understand. Nevertheless, it still tends to be overwhelming. A factor that aggravates the situation is that there is unnecessary code in these huge code bases: unreachable code, never executed code and redundant code. The latter has been of increasing interest in the software engineering research community under the term “cloning”. Especially clones that resulted from copy&paste can now be detected reliably (Rattan et al., 2013). In our own research with companies, we often found rates of redundant code caused by copy&paste in the range of 20 % – 30 % (Wagner, 2013).

One challenging is the detection of functionally similar source code. We will refer to it as *functionally similar clones* (FSCs). FSCs might not have been created by copy&paste but developers independently needed and implemented certain functionalities in their code base. We deliberately go beyond *type-4 clones* (Koschke, 2007) or *simions* (Deissenboeck et al., 2012) which only include functional equivalence. In a refactoring session with the goal to reduce the size of a code base, a developer would still be interested in mostly similar and not only exactly equivalent functionality. Although this problem is in general undecidable, there have been several heuristic efforts (Deissenboeck et al., 2012; Jiang & Su, 2009; Kim et al., 2011; Marcus & Maletic, 2001).

Juergens, Deissenboeck, & Hummel (2010b) showed that traditional clone detection approaches and tools are hardly able to detect functionally equivalent clones because they rely on syntactic similarities.

## 1.1 Problem Statement

So far, the work by Juergens, Deissenboeck, & Hummel (2010b) is the only study investigating the syntactical differences in functionally similar clones. Furthermore, their study is limited: they use only programs implementing a single specification in Java. Therefore, we have no clear understanding of what differences make an individually developed functionally similar clone really different from copy&paste clones. Hence, a realistic, open benchmark for comparing and improving such approaches is also lacking although it is necessary for faster progress in the field (Lakhotia et al., 2003).

## 69 1.2 Research Objectives

70 The objective of this study is to better understand the differences that make  
71 up functionally similar clones to support future research on their detection. In  
72 particular, we want to classify and rate differences and build a representative  
73 benchmark.

## 74 1.3 Contribution

75 We contribute a large-scale quantitative study combined with a qualitative anal-  
76 ysis of the differences. We selected 2,800 Java and C programs which are solu-  
77 tions to the Google Code Jam programming contest and are therefore function-  
78 ally similar. We identified copy&paste clones by using two clone detection tools  
79 (ConQAT and Deckard) to quantify syntactic similarities. We explored how a  
80 type-4 detection tool (CCCD) using cocolic detection performs in detecting the  
81 not syntactically similar FSCs. We created a categorisation of differences be-  
82 tween undetected clones and quantified these categories. Finally, we derived a  
83 benchmark based on real FSCs covering the categories and degrees of differences  
84 which can drive the improvement of clone detection tools.

## 85 1.4 Terminology

86 The most important terminological issue for this study is whether we need a new  
87 term for the phenomenon we investigate. The most established term is *type-4*  
88 *clone*. Yet, the definition in Roy et al. (2009) emphasises that the code fragments  
89 have to perform the *same* computation. We want to emphasise the *similarity*,  
90 however. This is expressed by the term *simion*, introduced in Juergens et al.  
91 (2010b), but that term has not been used in other context. Therefore, we  
92 propose the more general term *functionally similar clone*. Nevertheless, we  
93 believe the community should agree on one of these terms to avoid confusion in  
94 the future.

95 We define *functionally similar clones* as *two code fragments that provide a*  
96 *similar functionality w.r.t a given definition of similarity but can be implemented*  
97 *quite differently*. With this definition, we are based on the general definition of  
98 a code clone: “two code fragments form a clone pair if they are similar enough  
99 according to a given definition of similarity.” (Bellon et al., 2007) Intuitively,  
100 what we are interested in is similar enough so that the clones are interesting for  
101 a developer of the system while changing it. For our FSCs, we expect this to  
102 include similarity in terms of the same output on the same input in *most* cases.  
103 The data types of the input and output should be somehow transformable into  
104 each other. The nonfunctional properties of the code fragments, however, can  
105 be very different.

106 For example, as a developer, I would be interested in the different implemen-  
107 tations of search algorithms in my system also when they have quite different  
108 performance characteristics. I could then decide if the used algorithms and im-  
109 plementations are suitable for the context. Finally, as with other code clones,

the similarity should be such that the clones could potentially be refactored. If the fragments are so different that a refactoring is very elaborate or results in a very complex design, a developer will not be interested in it.

As there is a large diversity in how the further terms around FSCs are used, we provide definitions for the clone types we investigate in this paper. Moreover, we define terms for granularities of the software programs under analysis in Tab. 1.

Table 1: Terminology

<b>Type-1 clone</b>	Similar code fragments except for variation in whitespace, layout and comments (Bellon et al., 2007)
<b>Type-2 clone</b>	Similar code fragments except for variation in identifiers, literals, types, whitespaces layouts and comments (Bellon et al., 2007)
<b>Type-3 clone</b>	Similar code fragments except that some statements may be added or deleted in addition to variation in identifiers, literals, types, whitespaces, layouts or comments (Bellon et al., 2007)
<b>Type-4 clone</b>	Two or more code fragments that perform the same computation but are implemented by different syntactic variants. (Roy et al., 2009)
<b>Functionally similar clone (FSC)</b>	Code fragments that provide a similar functionality w.r.t a given definition of similarity but can be implemented quite differently
<b>Solution file</b>	A single program in one file implementing the solution to a programming problem
<b>Solution set</b>	A set of solution files all solving the the same programming problem
<b>Clone pair</b>	Two solution files from the same solution set which we assume to be functionally similar

The structure of the remainder of the paper follows the guidelines in Jedlitschka & Pfahl (2005).

## 2 Related Work

A *code clone* consists of at least two pieces of code that are similar according to a definition of similarity. Most commonly, **clone detection** approaches look for exact clones (also called *type-1*) and clones with simple changes such as renaming (also called *type-2*). These types of clones are detectable today in an efficient and effective way. Even clones with additional changes (*inconsistent*, *near-miss* or *type-3* clones) can be detected by several detection approaches and tools (Deissenboeck et al., 2008; Jiang et al., 2007a; Kamiya et al., 2002). There are also two surveys (Koschke, 2007; Roy & Cordy, 2007) and a systematic literature review (Rattan et al., 2013) on this topic. Tiarks et al. (2011) investigated in particular type-3 clones and also their differences. They concentrated, however, on differences in code metrics (e.g. fragment size), low level edits (e.g. variable) and abstracted them only slightly (e.g. to type substitution).

Juergens, Deissenboeck, & Hummel (2010b) reported on an experiment to investigate the differences between syntactical/representational and semantic/behavioural similarities of code and the detectability of these similarities. They used a simple student assignment called *email address validator* and also inspect the open-source software *JabRef*. Both of them are in Java. To detect the clones of types 1–3, they used the clone detection tools ConQAT and Deckard. They reviewed the open-source system manually to identify if behaviourally similar code that does not result from copy&paste can be detected and occurs in real-world software. The results indicate that behaviourally similar code of independent origin is highly unlikely to be syntactically similar. They also reported that the existing clone detection approaches cannot identify more than 1 % of such redundancy. We build our work on their study but concentrate on understanding the differences in more detail based on a diverse sample with a larger sample size and different programming languages.

Several researchers have proposed to move away from the concrete syntax to detect what they call **semantic clones**. Marcus & Maletic (2001) used information retrieval techniques on source code to detect semantic similarities. Krinke (2001) proposed to use program dependence graphs (PDG) for abstracting source code. Higo et al. (2011) extended this to an incremental approach. Komondoor & Horwitz (2001) also used PDGs for clone detection and see the possibility to find non-contiguous clones as a main benefit. Gabel, Jiang, & Su (2008) combined the analysis of dependence graphs with abstract syntax trees in the tool Deckard to better scale the approach.

A very different approach to detecting semantic clones comes from Kim et al. (2011) who use static analysis to extract the memory states for each procedure exit point. They could show that they find more semantically similar procedures as clones than previous clone detectors including PDG-based detectors. Finally, Kamiya (2013) proposed to analyse potential executions by analysing method calls on the Bytecode level in Java programs. Nevertheless, the used approach as well as the examples of found semantic clones suggest that the syntactic representation still plays a role and that the clones have been created by copy&paste. These semantic clone detection techniques cannot guarantee that

164 they also find all functionally similar clones as a completely different structure  
165 and memory states can generate similar functionality.

166 Jiang & Su (2009) were the first to comprehensively detect **functionally**  
167 **similar code** by using random tests and comparing the output. Hence, they  
168 were also the first who were able to detect clones without any syntactic similar-  
169 ity. They claim they are able to detect “*functionally equivalent* code fragments,  
170 where functional equivalence is a particular case of semantic equivalence that  
171 concerns the input/output behavior of a piece of code.” They detected a high  
172 number of functionally equivalent clones in a sorting benchmark and the Linux  
173 kernel. Several of the detected clones are dubious, however, as it is not clear  
174 how useful they are. They state: “Assuming the input and output variables  
175 identified by EQMINER for these code fragments are appropriate, such code  
176 fragments are indeed functionally equivalent according to our definition. How-  
177 ever, whether it is really useful to consider them functionally equivalent is still  
178 a question worth of future investigation.”

179 Deissenboeck et al. (2012) followed an analogous approach to Jiang & Su  
180 (2009) to detect functionally similar code fragments in Java systems based on  
181 the fundamental heuristic that two functionally similar code fragments will pro-  
182 duce the same output for the same randomly generated input. They imple-  
183 mented a prototype based on their toolkit ConQAT. The evaluation of the  
184 approach involved 5 open-source systems and an artificial system with indepen-  
185 dent implementations of the same specification in Java. They experienced low  
186 detection results due to the limited capability of the random testing approach.  
187 Furthermore, they mention that the similarities are missed due to chunking,  
188 i.e. if the code fragments perform a similar computation but use different data  
189 structures at their interfaces. They emphasise that further research is required  
190 to understand these issues.

191 CCCD (Krutz & Shihab, 2013) also claims to detect functionally similar  
192 code for C programs based on concolic analysis. Its creators evaluated their  
193 implementation of the approach on the benchmarks mentioned below and found  
194 a 92 % recall even in the type-4 clones in those benchmarks. As the tool is freely  
195 available in a virtual machine, we were able to include it in our experiment.

196 A clear comparison and measurement of the improvement in clone detection  
197 research would require a comprehensive **benchmark**. There have been few  
198 approaches (Lakhotia et al., 2003; Roy et al., 2009; Tempero, 2013) trying to  
199 establish a benchmark but they are either small and artificial or do not contain  
200 (known) FSCs. The only exception is the recent *BigCloneBench* (Svajlenko  
201 et al., 2014) which has a huge number of clones mined from source code reposi-  
202 tories. Yet, they do not classify the types of differences and also state “there is  
203 no consensus on the minimum similarity of a Type-3 clone, so it is difficult to  
204 separate the Type-3 and Type-4 clones”.

## 205 3 Experimental Design

206 To reach our research objectives, we developed a study design based on the idea  
 207 that we investigate sets of programs which we knew to be functionally similar:  
 208 accepted submissions to programming contests. We formulated four research  
 209 questions which we all answer by analysing these programs and corresponding  
 210 detection results. All instrumentation, analysis scripts and results are freely  
 211 available in Wagner et al. (2014).

### 212 3.1 Research Questions

213 As we have independently developed but functionally similar programs, we first  
 214 wanted to establish how much syntactic similarity is in these programs. We can  
 215 investigate this by quantifying the share of type-1–3 clones

216 **RQ 1: What share of independently developed similar programs are**  
 217 **type-1–3 clones?**

218 Then we wanted to understand what is different in clones not of type-1–  
 219 3. This should result in a categorisation and rating of the differences between  
 220 FSCs.

221 **RQ 2: What are the differences between FSC that go beyond type-1–3**  
 222 **clones?**

223 Although we could not fully evaluate type-4 detectors, we wanted at least  
 224 to explore what a modern clone detection approach can achieve on our FSCs.  
 225 This should give us an indication how much more research is needed on those  
 226 detection approaches.

227 **RQ 3: What share of FSC can be detected by a type-4 clone detector?**

228 Finally, to make our results an operational help for clone detection research,  
 229 we wanted to create a representative benchmark from the non-detected clones.

230 **RQ 4: What should a benchmark contain that represents the differ-**  
 231 **ences between FSC?**

### 232 3.2 Hypotheses

233 We define two hypotheses regarding RQ 1. As we investigate the share of  
 234 detectable Type-1–3 clones, we wanted to understand if there are differences  
 235 between the used tools and analysed languages because this might have an  
 236 influence on the generalisability of our results. We formulated the two null  
 237 hypotheses:

238 **H1: There is no difference in the share of detected Type-1–3 clones**  
 239 **between the analysed programming languages.**

240 **H2: There is no difference in the share of detected Type-1–3 clones**  
 241 **between the analysed clone detection tools.**

242 Moreover, in RQ 2, we wanted to understand the characteristics of non-  
 243 detected clone pairs and, therefore, categorised them. In this categorisation,  
 244 we also rated the degree of difference in each category. An ideal categorisa-



tion would have fully orthogonal categories and, hence, categories would not be correlated in the degree of difference:

**H3: There is no correlation between the degrees of difference between categories.**

Furthermore, we could imagine that different programming languages might cause disparately strong differences in certain categories. As this again has an impact on the generalisability of our results, we formulated this null hypotheses:

**H4: There is no difference in the degree of difference between the analysed programming languages.**

### 3.3 Design

The overall study design is a combination of quantitative and qualitative analysis. For the quantitative part of our study, we used a factorial design with two factors (programming language and clone detection tool). As applying the treatments of both factors was mostly automated we could apply almost all factor levels to all study object programs (which we call *solutions*). Only if a detection tool did not support a certain programming language, we would not apply it. We tried to minimise that but to include a contemporary tool, we accepted an unbalanced design. Table 2 shows the factors in our experiment.

Table 2: The factorial design used in this experiment

		Programming language	
		Java	C
<b>Clone detection tool</b>	CCCD	–	✓
	ConQAT	✓	✓
	Deckard	✓	✓

We will describe the programming languages, clone detection tools and corresponding programs under analysis in more detail in the next subsection.

### 3.4 Objects

The general idea of this experiment was that we analyse accepted solutions to programming contests because we know that for a given problem, the solutions must be functionally similar. Therefore, our selection of study objects needed to include clone detection tools we could access and execute as well as solutions in programming languages supported by most of the detection tools.

#### 3.4.1 Clone Detection Tools

Primarily, we needed clone detection tools for detecting type-1–3 clones to investigate with RQ 1 the syntactic similarity of FSCs. We did a literature and web search for available tools.

Many research prototypes were not available or could not be brought to execute correctly. Several tools were not included in the study due their lower performance and scalability or their lack of support for some clone types. CloneDR and CPMiner have lower performance and scalability compared to Deckard (Jiang et al., 2007a). CCFinder has also lower performance than Deckard and does not support type-3 clones (Svajlenko & Roy, 2014).

In the end, we chose two clone detection tools that both can analyse Java and C programs: **ConQAT** (Deissenboeck et al., 2008) and **Deckard** (Jiang et al., 2007a). ConQAT is described in Rattan et al. (2013) as modern, useful and fast open-source clone detector framework. In the studies we mentioned above, Deckard has shown to have good performance and scalability. Hence, both are well established and have been used in previous studies, especially Juergens et al. (2010b). At the time of the study, those were the two tools which were both freely available and were possible to make them work for us.

**ConQAT** is a stable open-source dashboard toolkit also used in industry. It is a general-purpose tool for various kinds of code measurement and analysis. For our experiment, ConQAT offers several specific clone detection configurations for various programming languages including Java, C/C++, C# and Cobol. It has separate detection algorithms for type-1/2 clones and type-3 clones. We employed the latter algorithm. ConQAT has been used in various studies on clone detection (Juergens et al., 2010a, 2009) including the study we build on (Juergens et al., 2010b).

The language-independent clone detection tool **Deckard** works on code in any programming language that has a context-free grammar. Deckard uses an efficient algorithm for identifying similar subtrees and applies it to tree representations of source code. It automatically generates a parse tree builder to build parse trees required by its algorithm. By a similarity parameter it is possible to control whether only type-1/2 clones or type-3 clones are detected. Deckard is a stable tool used in other studies (Gabel et al., 2008; Jiang et al., 2007b) including the study we build on.

To explore the state of type-4 clone detection tools, we also searched for such tools. Most existing tools, however, could not be used. For example, EqMiner (Jiang & Su, 2009) was too tightly coupled with the Linux kernel and MeCC (Kim et al., 2011) could not detect clones across files. Finally, we were able to only include a single type-4 detector.

**CCCD** (Krutz & Shihab, 2013) is a novel clone detection tool that uses concolic analysis as its primary approach to detect code clones. Concolic analysis combines symbolic execution and testing. CCCD detects only clones in programs implemented in C. The concolic analysis allows CCCD to focus on the functionality of a program rather than the syntactic properties. Yet, it has the restriction that it only detects function-level clones.

### 3.4.2 Solution Sets and Solutions

We looked at several programming contests and the availability of the submitted solutions. We found that *Google Code Jam*<sup>1</sup> provided us with the broadest selection of programming languages and the highest numbers of submissions. Google Code Jam is an annual coding contest organised by Google. Several tens of thousands of people participate each year. In seven competition rounds, the programmers have to solve small algorithmic problems within a defined time frame. Although over one hundred different programming languages are used, the majority of the solutions are in C, C++, Java and Python. Most solutions of the participants are freely available on the web.<sup>2</sup>

We define a *solution* as a single code file delivered by one participant during the contest. We define a *solution set* as a set of solutions all solving the same problem. A typical solution set consists of several hundred to several thousand solutions. We can be sure that all solutions in a solution set should be FSCs because they passed the judgement of the programming contest. This is also implemented as automated tests.<sup>3</sup> Even if there are differences in the programs, e.g. in the result representation, these are instances of similarity instead of equivalence.

We selected 14 out of 27 problem statements of the Google Code Jam 2014. For every problem we randomly chose 100 solutions in Java and 100 solutions in C from sets of several hundreds to several thousands of solutions. Table 3 shows a summary of the size of the chosen solution sets. Hence, on average a C solution has a length of 46 LOC and a Java solution of 94 LOC.

Table 3: Summary of the solution sets

	#No. Sets	#Files/Set	#Procedures	LOC
C	14	100	2,908	64,826
Java	14	100	8,303	131,398

In Table 4, we detail the size of the selected solution sets. The solution sets differ in size but the means all lie between 33 and 133 LOC per solution.

## 3.5 Data Collection Procedure

### 3.5.1 Preparation of Programs Under Analysis

We implemented an instrumentation which automatically downloaded the solutions from the website, sampled the solution sets and solutions and normalised the file names. The instrumentation is freely available as Java programs in

<sup>1</sup><https://code.google.com/codejam/>

<sup>2</sup><http://www.go-hero.net/jam/14/>

<sup>3</sup><https://code.google.com/codejam/quickstart.html>

Table 4: Information on the solution sets

Java				C			
Set	#Files	LOC	#Proc.	Set	#Files	LOC	#Proc.
1	100	11,366	823	1	100	3,917	233
2	100	7,825	523	2	100	3,706	167
3	100	10,624	575	3	100	4,750	265
4	100	6,766	473	4	100	3,928	219
5	100	7,986	585	5	100	4,067	187
6	100	10,137	611	6	100	6,840	166
7	100	13,300	869	7	100	4,701	263
8	100	8,568	614	8	100	4,679	176
9	100	8,580	717	9	100	6,831	227
10	100	9,092	459	10	100	4,063	159
11	100	8,536	584	11	100	4,624	266
12	100	11,412	648	12	100	3,574	163
13	100	9,436	465	13	100	3,335	168
14	100	7,770	357	14	100	5,811	249

Wagner et al. (2014). Every downloaded solution consisted of a single source code file.

### 3.5.2 Configuration of Clone Detection Tools

We installed ConQAT, Deckard and CCCD and configured the tools with a common set of parameters. As far as the parameters between the tools were related to each other, we tried to set the same values based on the configuration in Juergens et al. (2010b). We set the parameters conservatively so that the tools find potentially more clones as we would normally consider valid clones. This ensured that we do not reject our null hypotheses because of configurations. We based the minimal clone length on the parameters for practical use mentioned in Juergens et al. (2010b) and experimented with lower values. Finally, we decide to make 6 the threshold, because lower values created many findings which we did not consider a valid clone. For the gap clone ratio, we also tried different values, but decided to use 0.3 as it was close to the recommended value, generated reasonable clones and clones similar to the ones detected by CCCD with the Levenshtein score of 35. All the detailed configurations are available in Wagner et al. (2014).

For CCCD, we only included clone pairs which have a Levenshtein similarity score below 35 as advised in Krutz & Shihab (2013) where the score is calculated for the concolic output for each function. The detection failed for 4 of the solutions of our sample set. We had to exclude them from further analysis.

### 3.5.3 Executing Clone Detection Tools

We manually executed the clone detection tools for every solution set. ConQAT generated an XML file for every solution set containing a list of found clone classes and clones. Deckard and CCCD generate similar CSV files. Our instrumentation tool parsed all these result files and generated reports in a unified format. The reports are tables in which both rows and columns represent the solutions. The content of the table shows the lowest detected clone type between two files. Additionally, our tool calculated all shares of syntactic similarity as described in the next section and wrote the values into several CSV files for further statistical analysis. We also wrote all the detected clones into several large CSV files. Altogether, the tools reported more than 9,300 clones within the Java solutions and more than 22,400 clones within the C solutions.

## 3.6 Analysis Procedure

### 3.6.1 Share of Syntactic Similarity (RQ 1)

All solutions in a solution set solve the same programming problem and were accepted by Google Code Jam. Hence, their functionality can only differ slightly and, therefore, they are functionally similar. To understand how much of this similarity is expressed in syntactic similarity, we calculate the share of FSCs which are also type-1-2 or type-1-3 clones.

Inspired by Juergens et al. (2010b), we distinguish partial and full syntactic similarity. The *share of full syntactic similarity* is the ratio of clone pairs where all but a defined number of the statements of the solutions of the pair were detected as a clone in relation to all clone pairs. We set the threshold of this to a maximum of 16 lines of code difference within a clone pair, which leads to ratios of 5 % to 33 % of difference based on the solutions' lines of code.

$$\text{Share of full syntactic similarity} = \frac{|\text{Found full clone pairs}|}{|\text{All clone pairs}|} \quad (1)$$

Because we expected the share of full syntactic similarity to be low, we wanted to check whether there are at least some parts with syntactic similarity. It would give traditional clone detection tools a chance to hint at the FSC. Furthermore, it allowed us to inspect more closely later on what was not detected as a clone. We called the ratio *share of partial syntactic similarity*.

$$\text{Share of partial syntactic similarity} = \frac{|\text{Found partial clone pairs}|}{|\text{All clone pairs}|} \quad (2)$$

For a more differentiated analysis, we calculated two different shares each representing certain types of clones. We first computed the share for type-1-2 clones. This means we only need to accept exact copies, reformatting and renaming. Then, we determined the shares for type-1-3 clones which includes type-1-2 and adds the additional capability to tolerate smaller changes.

In ConQAT and Deckard, we can differentiate between type-1/2 clones and type-3 clones by configuration or result, respectively. In ConQAT, clones with a gap of 0 are type-1/2 clones. In Deckard, analysis results with a similarity of 1 are type-1/2 clones. The others are type-3 clones. The instrument tooling described in Sec. 3.5 directly calculated the various numbers. We computed means per clone type and programming language.

For a further statistical understanding and to answer the hypotheses H1–H4, we did statistical hypotheses tests. For answering H1 and H2, we performed an analysis of variance (ANOVA) on the recall data with the two factors *programming language* and *detection tool*. We tested the hypotheses at the 0.05 level. All analyses implemented in R together with the data are available in Wagner et al. (2014).

The combined descriptive statistics and hypothesis testing results answered RQ 1.

### 3.6.2 Classifying Differences (RQ 2)

For the categorisation of the differences of FSCs that were not syntactically similar, we took a random sample of these clone pairs. As we had overall 69,300 clone pairs for Java and C, we needed to restrict the sample for a manual analysis. We found in an initial classification (see also Sec. 3.7) that a sample of 0.5 % per language and fully/partially different clone pairs is sufficient for finding repeating categories and getting a quantitative impression of the numbers of clone pairs in each category. With larger samples, the categories just kept repeating. Therefore, we took a sample of 2 % of the syntactically different clone pairs: 70 pairs each of the fully and partially different clone pairs (35 C and 35 Java).

The set of *fully syntactically different clone pairs* is the set of all pairs in all solution sets minus any pair detected by any of the type-1–3 detection. We apply random sampling to get pairs for further analysis: First, we randomly selected one of the solution sets in a language. Second, we randomly selected a solution file in the solution set and checked if it was detected by Deckard or ConQAT. If it was detected, we would discard it and select a new one. Third, we randomly picked a second solution file, checked again if it was detected and discard it if it was.

The set of *partially syntactically different clone pairs* is then the superset of all partially different clone pairs minus the superset of all fully different clone pairs. From that set, we randomly selected clone pairs from all partially different pairs of a programming language and checked if it was fully different. If that was the case, we would discard it and take a new random pair. We found their analysis to be useful to understand also smaller syntactic differences.

We then employed qualitative analysis. We manually classified the characteristics in the clone pairs that differed and, thereby, led to being not detected as type-1–3 clone. This classification work was done in pairs of researchers in three day-long workshops in the same room. It helped us to discuss the categories and keep them consistent. The result is a set of categories of characteristics that

446 describe the differences. We added quantitative analysis to it by also counting  
447 how many of the sampled clone pairs have characteristics of the found types.

448 After the creation of the categories we also assessed the degree of difference  
449 (high, medium, or low) per category. From the discussion of the categories, we  
450 discovered that this gave us a comprehensive yet precise way to assign clone  
451 pairs to the categories. Furthermore, it gave us additional possibilities for a  
452 quantified analysis. First, we wanted to understand better how we categorised  
453 and assessed the degrees of difference as well as answer H3. Therefore, we per-  
454 formed correlation analysis on them. We chose Kendall's tau as the correlation  
455 coefficient and tested all correlations on the 0.05 level.

456 For answering H4, we performed a multivariate analysis of variance (MANO-  
457 VA) which allows more than one dependent variable to be used. Here, our  
458 dependent variables are the degrees of difference and the independent variable is  
459 the programming language. In this analysis, we have a balanced design because  
460 we ignored the category *OO design* which was only applicable to Java programs.  
461 We use the Pillar-Bartlett statistic for evaluating statistical significance. We  
462 checked H4 also on the 0.05 level.

463 These categories with frequencies as well as the results of the hypothesis  
464 tests answered RQ 2.

### 465 3.6.3 Running a Type-4 Detector (RQ 3)

466 As this part of the study is only for exploratory purposes, we focused on the  
467 recall of CCCD in the FSCs. As all solutions contain a *main* function, we  
468 expected it to find each *main*-pair as clone. We calculate the recall as the  
469 number of detected clone pairs by the sum of all clone pairs. A perfect clone  
470 detection tool would detect all solutions from a solution set as clones.

### 471 3.6.4 Creating a Benchmark (RQ 4)

472 After the categorisation to answer RQ 2, we had a clear picture of the various  
473 differences. Therefore, we could select representative examples of each difference  
474 for each programming language and put them into our new benchmark. To check  
475 that the clone pairs cannot be detected by the tools, we run the tools again on  
476 the benchmark. If one of the tools still detected a clone, we would replace the  
477 clone pair by another representative example until no clones are detected.

478 We created the benchmark by choosing clone pairs consisting of two source  
479 code files out of the same solution set. The two files therefore solve the same  
480 problem. We selected three pairs where the difference between the files belong  
481 to that category for each of the categories we created by answering RQ 2. We  
482 chose three pairs for all of the three levels of difference. The other categories  
483 of the pairs are low or non-existent so that the pair is focused on one cate-  
484 gory. Additionally, we added one extra clone pair with extreme differences in  
485 all categories.

486 Preferably, we would provide the source code of the chosen solutions directly  
487 all in one place. Yet, the copyright of these solutions remains with their authors.

Therefore, we provide source files following the same structure as the original files but not violating the copyright.

A final set of clone pairs that are not detected as full clones by any of the tools constitutes the benchmark and answered RQ 4.

### 3.7 Validity Procedure

To avoid selection bias, we performed random sampling where possible. We randomly selected the solution sets and solutions that we use as study objects. In addition, before we manually analysed the category of syntactically different clone pairs, we chose random samples of clone pairs.

To avoid errors in our results, we manually checked for false positives and clone types with samples of clones in the solution sets. Furthermore, by working in pairs during all manual work, we controlled each other and detected problems quickly. Overall, the manual inspection of 70 clone pairs for RQ 2 also was a means to detect problems in the detection tools or our instrumentation.

For the manual categorisation, we started by categorising 30 syntactically different clone pairs to freely create the categories of undetected clone pairs. Afterwards, we discussed the results among all researchers to come to a unified and agreed categorisation. The actual categorisation of clone pairs was then performed on a fresh sample. Additionally, we performed an independent categorisation of a sample of 10 categorised clone pairs and calculated the inter-rater agreement using Cohen's kappa.

## 4 Analysis and Results

We structure the analysis and results along our research questions. All quantitative and qualitative results are also available in Wagner et al. (2014).

### 4.1 Share of Syntactic Similarity (RQ 1)

We summarised the results of the calculated shares for fully and partially syntactically similar clone pairs in Tab. 5. We divided the results by programming languages, detection tools and detected clone types. The results differ quite strongly from tool to tool but only slightly between the programming languages. The average syntactic similarities and the standard deviations (SD) are all very low. ConQAT detects more full and partial clones in clone pairs.

Table 6 shows the ANOVA results which we need for answering hypotheses H1 and H2. As our experiment is unbalanced, we use the Type II sum of squares. This is possible because we found no significant interactions between the factors in any of the ANOVA results.

The results give us no single evaluation of the hypotheses H1 and H2. We have to differentiate between partial and full syntactic similarity. For the partial similarity, we consistently see a significant difference in the variation in the detection tools but not in the programming languages. Hence, for partial clone



Table 5: Full and partial syntactic similarity (in %)

Lang.	Tool	Partially similar				Fully similar			
		Type 1-2		Type 1-3		Type 1-2		Type 1-3	
		Mean	SD	Mean	SD	Mean	SD	Mean	SD
Java	ConQAT	6.36	0.05	11.53	0.07	0.00	0.00	0.00	0.00
	Deckard	0.33	0.00	0.87	0.01	0.00	0.00	0.00	0.00
	<b>Mean</b>	3.35	0.03	6.11	0.04	0.00	0.00	0.00	0.00
C	ConQAT	5.24	0.09	11.48	0.13	1.30	0.00	1.73	0.00
	Deckard	0.28	0.00	1.44	0.01	0.01	0.00	0.01	0.00
	<b>Mean</b>	1.82	0.00	4.32	0.06	0.47	0.00	0.58	0.00
<b>Grand mean</b>		2.45	0.04	5.07	0.04	0.26	0.00	0.35	0.00

similarity, we corroborate H1 that there is no difference in recall between the analysed programming languages. Yet, we reject H2 in favour of the alternative hypothesis that there is a difference in the similarity share between different tools. For full similarity, we reject H1 in favour of the alternative hypothesis that there is a difference between the analysed programming languages. Instead, we accept H2 that there is no difference between the analysed detection tools.

How can we interpret these results? *The overall interpretation is that the share of syntactic similarity in FSCs is very small.* There seem to be many possibilities to implement a solution for the same problem with very different syntactic characteristics. When we only look at the full syntactic similarity, the results are negligible. Both tools detect none in Java and only few clone pairs for C. Hence, the difference between the tools is marginal. The difference is significant between C and Java, however, because we found no full clone pairs in Java. As we saw in a manual inspection, the full detection is easier in C if the developers implement the whole solution in one main function.

For partial syntactic similarity, we get higher results but still stay below 12 %. Hence, for almost 90 % of the clone pairs, we do not even detect smaller similarities. We have no significant difference between the languages but the tools. ConQAT has far higher results than Deckard in the type-1-3 clones. The distinct detection algorithms seem to make a difference here. For the further analysis, we accept an FSC as syntactically similar if one of the tools detected it.

## 4.2 Categories of Differences (RQ 2)

Initially, we created 18 detailed categories. In our qualitative analysis and discussions, we finally reduced them to five main categories of characteristics describing the differences between the solutions in a clone pair. The five categories are *algorithm*, *data structure*, *object-oriented design*, *input/output* and *libraries*.

Table 6: ANOVA results for variation in recalls (Type II sum of squares, \* denotes a significant result)

Partial type 1-2	Sum of Squares	F value	Pr(>F)	
Language	0.0005	0.2352	0.6294	
Tool	0.0491	12.2603	$3 \cdot 10^{-5}$	*
Partial type 1-3				
Language	0.0010	0.0210	0.8853	
Tool	0.1884	20.5846	$1 \cdot 10^{-7}$	*
Full type 1-2				
Language	$1 \cdot 10^{-7}$	7.8185	0.0072	*
Tool	$2 \cdot 10^{-8}$	1.1566	0.2871	
Full type 1-3				
Language	$2 \cdot 10^{-7}$	7.7757	0.0074	*
Tool	$5 \cdot 10^{-8}$	1.9439	0.1692	

We could assign each of the 18 initial categories there and realised that we can assign them to different *degrees of difference*. Therefore, we ended up with a categorisation including an ordinal quantification of the degree of difference with the levels *low*, *medium* and *high*. The overall categorisation is shown in Fig. 1. The centre of the dimensions would be a type-1 clone. The further out we go on each dimension, the larger the difference.




Figure 1: The categories of characteristics of differences between clone pairs

To make the categories and degrees of difference clearer, we give examples of characteristics in clone pairs in Tab. 7 that led us to classify them in the specific degrees of difference. The guiding principle was how much effort (in terms of

edit operations) it would be to get from one solution to the other.

Table 7: Examples for the levels in the degree of difference per category

<b>Algorithm</b>	<i>low</i>	Only syntactic variations
	<i>medium</i>	Similarity in the control structure but different method structure
	<i>high</i>	No similarity
<b>Data structure</b>	<i>low</i>	Different data types, e.g. int – long
	<i>medium</i>	Related data types with different interface, e.g. array vs. List
	<i>high</i>	Standard data types vs. own data classes or structs
<b>OO design</b>	<i>low</i>	Only one/few static methods vs. object creation
	<i>medium</i>	Only one/few static methods vs. data classes or several methods
	<i>high</i>	Only one/few static methods vs. several classes with methods
<b>Library</b>	<i>low</i>	Different imported/included but not used libraries
	<i>medium</i>	Few different libraries or static vs. non-static import
	<i>high</i>	Many different or strongly different libraries
<b>I/O</b>	<i>low</i>	Writing to file vs. console with similar library
	<i>medium</i>	Strongly different library, e.g. Scanner vs. FileReader
	<i>high</i>	Strongly different library and writing to file vs. console

The two main aspects in any program are its *algorithms* and its *data structures*. This is reflected in our two main categories. Our corresponding degrees of difference reflect that there might be algorithms that are almost identical with e.g. only a *switch* instead of nested *if* statements up to completely different solutions, e.g. iterative vs. recursive. Similarly, in data structures, we can have very simple type substitutions which change the behaviour but are still functionally very similar (e.g. from *int* to *long*) but also completely user-defined data types with strong differences.

Related to data structures is the category *OO design*. We made this a separate category because it only applies to OO languages and it had a particular kind of occurrence in the programs we inspected. Some developers tended to write Java programs like there were no object-oriented features while others created several classes and used their objects.

As our programming environments and languages are more and more defined by available *libraries*, this was also reflected in the differences between solutions. If one developer of a solution knew about a library with existing functionality needed, and the other developer implemented it herself, this created code that

looks strongly different but can have similar functionality.

Finally, maybe a category that arose because the programming contest did not specify if the input and output should come from a console or a file was the usage of *I/O*. Nevertheless, we think that this might also be transferable to other FSCs and contexts because we might be interested in functionally similar code even if one program writes the output on a network socket while the other writes into a file.

Table 8 shows descriptive statistics for the categories in our sample of undetected clone pairs. The column *Share* shows the ratio of clone pairs with a degree of difference higher than 0 in relation to all clone pairs in that language. The median and median absolute deviation (MAD) give the central tendency and dispersion of the degrees in that category. For that, we encoded *no difference* = 0, *low* = 1, *medium* = 2 and *high* = 3.

Table 8: Descriptive statistics of degrees of difference over categories and programming languages

Lang.	Category	Share	Median	MAD
Java	Algorithm	96 %	3	0.0
	Libraries	86 %	1	1.5
	I/O	83 %	2	1.5
	Data structure	72 %	1	1.5
	OO design	71 %	1	1.5
C	Algorithm	76 %	2	1.5
	Libraries	73 %	1	1.5
	Data structure	66 %	1	1.5
	I/O	38 %	0	0.0
Total	Algorithm	86 %	2	1.5
	Libraries	79 %	1	1.5
	OO design	71 %	1	1.5
	Data structure	69 %	1	1.5
	I/O	60 %	1	1.5

All categories occur in the majority of clone pairs. The categories *algorithm* and *libraries* even occur in about three quarters of the clone pairs. The occurrence of categories is consistently smaller in C than in Java. The medians are mostly low but with a rather large deviation. Only *input/output* in C has a median of 0. This is consistent with our observation during the manual inspection that I/O is done similarly in the C programs.

For evaluating H3, we calculated Kendall's correlation coefficients for all combinations of categories. The results are shown in Tab. 9. The statistical tests for these correlations showed significant results for all the coefficients.

Therefore, we need to reject H3 in favour of the alternative hypothesis that there are correlations between the degrees of difference between different categories.

Table 9: Correlation matrix with Kendall's correlation coefficient for the category degrees (all are significant)

	Algo.	Data struct.	OO design	I/O	Libraries
Algorithm	1.00	0.38	0.44	0.15	0.31
Data struct.	0.38	1.00	0.26	0.25	0.21
OO design	0.44	0.26	1.00	0.29	0.39
I/O	0.15	0.25	0.29	1.00	0.27
Libraries	0.31	0.21	0.39	0.27	1.00

Finally, for evaluating H4, we show the results of the MANOVA in Tab. 10. We can reject H4 in favour of the alternative hypothesis that there is a difference between the degrees of difference between the analysed programming languages. This is consistent with the impression from the descriptive statistics in Tab. 8.

Table 10: MANOVA results for variation in degree of differences (Type I sum of squares, \* denotes a significant result)

	Pillai-Bartlett	approx. F	Pr(>F)
Language	0.1513	6.0196	0.0002 *

In summary, we interpret these results such that *there are differences in FSC pairs in their algorithms, data structures, input/output and used libraries. In Java, there are also differences in the object-oriented design.* On average, *these differences are mostly small but the variance is high.* Hence, we believe that with advances in clone detectors for tolerating the smaller differences, there could be large progress in the detection of FSCs. Yet, there will still be many medium to large differences. We also saw that the programming languages vary in the characteristics of undetected difference. Therefore, it might be easier to overcome those differences in non-object-oriented languages, such as C, than in object-oriented languages which offer even more possibilities to express solutions for the same problem. Yet, we were impressed by the variety in implementing solutions in both languages during our manual inspections.

Our categories are significantly correlated with each other. This can mean that there might be other, independent categories with less correlation. Nevertheless, we believe the categories are useful because they describe major code aspects in a way that is intuitively understandable to most programmers. It would be difficult to avoid correlations altogether. For example, a vastly different data structure will always lead to a very different algorithm.

### 627 4.3 Type-4 Detection (RQ 3)

628 Table 11 shows the recall of fully and partially detected clone pairs in our sample.  
 629 CCCD has a considerable recall for partial clones in the clone pairs of about  
 630 16 %. It does, however, detect almost none of the clone pairs a full clones.  
 631 The overlap with ConQAT and Deckard, and therefore type-1–3 clones, is tiny  
 632 (0.05 % of the recall).

Table 11: Full and partial clone recall means over solution sets for CCCD (in %)

	Mean	SD
Partial	16.03	0.07
Full	0.10	0.00

633 We interpret this result such that also contemporary type-4 detection tools  
 634 have still problems detecting real-world FSCs and to handle the differences we  
 635 identified in RQ 2.

### 636 4.4 Benchmark (RQ 4)

637 The number of study objects used in our analysis is quite high. As described  
 638 above, we examined 1,400 Java files and 1,400 C files. For many demonstrations  
 639 and clone detection tool analyses a much smaller file set is sufficient. We call  
 640 this smaller set of files *benchmark*.

641 The first half of the benchmark we provide consists of 29 clone pairs. For  
 642 Java, we include 16 clone pairs. The set of clone pairs we provide for C is  
 643 structured in exactly the same way as the Java samples except that we do not  
 644 have the three clone pairs that differ only in object-oriented design. Therefore,  
 645 we do not have 16 samples here but 13 which make the 29 clone pairs for both  
 646 languages.

647 Figure 2 shows a rating of an example clone pair in the benchmark set where  
 648 the two files only differ significantly in the kind of input/output, but not in the  
 649 other categories. Figure 3 shows how such a high difference in the category  
 650 input/output could look like between two solutions in a solution set.

651 We provide this distribution of clone pairs for both partial clones and full  
 652 clones. Hence, the total number of clone pairs within the benchmark is 58.  
 653 Figure 4 shows an overview of the structure of the whole benchmark set. This  
 654 structure enables developers of a clone detection tool to test their tool easily as  
 655 well analyse the nature of the clones found and not found by a tool.

656 Our benchmark provides several advantages to the research community.  
 657 First, *developers of a clone detection tool can easily test their tool with the*  
 658 *source files as input*. They can see whether their tool detects the clones or they  
 659 can analyse why it did not. Second, the clones in the benchmark are easily  
 660 understandable examples for the categories we created. Third, *the clones in*




Figure 2: Example category rating of a clone pair in the benchmark set

```
public static void main(String[] args) {
    reader = new BufferedReader(
        new FileReader("A-large.in"));
    writer = new PrintWriter("a.out");
}

public static void main(String[] args) {
    File file = new File(System.in);
    try (Scanner scanner = new Scanner(
        new FileReader(file))) {
        File out = new File(System.out);
        try (PrintWriter writer = new ...
```

Figure 3: Example of a high difference in the category Input/Output

the benchmark were not built artificially; the solutions were implemented independently by at least two persons during the Code Jam contest. Despite our modifications to avoid copyright problems, neither changing structure nor algorithm, the code clones are more realistic than fully artificial copies where one file is modified as part of a study.

## 5 Threats to Validity

We analyse the validity threats for this study following common guidelines for empirical studies (Wohlin et al., 2012; Yin, 2003).

### 5.1 Conclusion Validity

As most of our measurements and calculations were performed automatically, the threats to conclusion validity are low. For the corresponding hypothesis tests, we checked all necessary assumptions. Only the classification and rating of the degree of difference is done manually and, hence, could be unreliable.

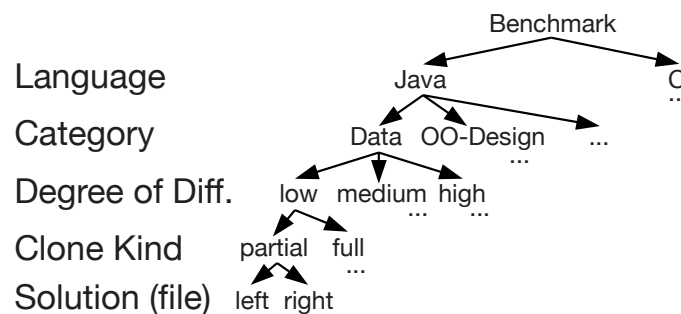


Figure 4: Structure of the benchmark set (overview)

674 We worked in pairs to reduce this threat. Furthermore, one of the researchers  
 675 performed an independent classification of a random sample of 10 clone pairs  
 676 to compare the results. We calculated Cohen's kappa for the categories of  
 677 differences between clone pairs as presented in Table 12.

Table 12: Kappa values for difference categories

Category	Kappa
Data structures	0.41
OO design	0.35
Algorithms	0.47
Libraries	0.36
Input/Output	0.47

678 We interpret the kappa results according to the classification by Landis &  
 679 Koch (1977). Hence, our results are a moderate agreement between the cate-  
 680 gories: data structures, algorithms and input/output. For the categories object-  
 681 oriented design and libraries we have a fair agreement. We consider this to be  
 682 reliable enough for our investigations.

## 683 5.2 Internal Validity

684 There is the threat that the implementation of our instrumentation tooling may  
 685 contain faults and, therefore, compute incorrect results for the detected clones  
 686 and recalls. We reduced this threat inherently by the manual inspections done  
 687 to answer RQ 2 and independently to investigate the type-4 clones.

688 A further threat to internal validity is that we took our solution sets from  
 689 Google Code Jam. We cannot be sure that all the published solutions of the  
 690 Code Jam within a solution set are actually functionally similar. We rely on  
 691 the fact that the organisers of the Code Jam must have checked the solutions to



rank them. Furthermore, we assume to have noticed in the manual inspections if there were solutions in a solution set with highly differing functionality.

### 5.3 Construct Validity

To fully understand the effectiveness of a clone detection approach, we need to measure precision as well as recall. In our study, we could not measure precision directly because of the large sample size. We checked for false positives during the manual inspections and noted only few rather short clones. Our minimal clone length is below recommended thresholds. This is a conservative approach to the problem. By that, we will find more clones than in an industrial approach. We decided to use this threshold to be sure that we cover all the interesting clone pairs that would be lost due to variation in the precision of the tools.

There is a threat because we count each clone pair only once. In partial clones, one clone pair might contain a type-2 as well as a type-3 partial clone. In those cases, we decided that the lower type – the easier detection – should be recorded. Hence, the assignment to the types might be imprecise. We accept this threat as it has no major implication for the conclusions of the experiment.

### 5.4 External Validity

There is also a threat to external validity in our usage of solutions from Google Code Jam. The submitted programs might not represent industrial software very well. Participants had a time limit set for turning in their solutions. Furthermore, the programming problems contained mostly reading data, performing some calculations on it and writing data. This might impact the method structure within the solutions. This threat reduces the generalisability of our results. Yet, we expect that other, more complex software will introduce new kinds of difference categories (e.g. differences in GUI code) and only extend but not contradict our results.

For the study, we chose three well-known and stable clone detection tools. Two of them analyse Java and C programs detecting type-1 to type-3 clones. The third one detects type 4 clones and supports only programs written in C and only finds clones in complete functions. Overall, we are confident that these tools represent the available detection tools well.

## 6 Conclusions and Future Work

In this paper, we investigated the characteristics of clones not created by copy&-paste. We base our study on Juergens et al. (2010b), but this is the first study with programs implementing different specifications in diverse programming languages including CCCD as concolic clone detector for type-4 clones. We found that a full syntactic similarity was detected in less than 1 % of clone pairs. Even partial syntactic similarity was only visible in less than 12 %. The concolic approach of CCCD can detect FSCs without syntactic similarity as type-4 clones.

731 Yet, a full detection was only possible in 0.1 % of clone pairs.

732 Our categorisation of the differences of clone pairs that were not syntactically  
733 similar showed that usually several characteristics make up these differences. On  
734 average, however, the differences were mostly small. Hence, we believe there is  
735 a huge opportunity to get a large improvement in the detection capabilities of  
736 type-4 detectors even with small improvements in tolerating additional differ-  
737 ences. We provide a carefully selected benchmark with programs representing  
738 real FSCs. We hope it will help the research community to make these improve-  
739 ments.

## 740 6.1 Relation to Existing Evidence

741 We can most directly relate our results to Juergens, Deissenboeck, & Hummel  
742 (2010b). We support their findings that using type-1–3 detectors, below 1 % is  
743 fully and below 10 % is partially detected. We can add that with the type-4  
744 detection of CCCD, the partial clone recall can reach 16 %. They introduce  
745 categories which were derived from other sources but not created them with a  
746 systematic qualitative analysis. Yet, there are similarities in the categories.

747 Their category *syntactic variation* covers “if different concrete syntax con-  
748 structs are used to express equivalent abstract syntax”. We categorised this  
749 as small algorithm difference. Their category *organisational variation* “occurs  
750 if the same algorithm is realized using different partitionings or hierarchies of  
751 statements or variables”. We categorise these differences as a medium algorithm  
752 difference. Their category *delocalisation* “occurs since the order of statements  
753 that are independent of each other can vary arbitrarily between code fragments”  
754 is covered as difference in algorithm in our categorisation. Their category *gen-  
755 eralisation* “comprises differences in the level of generalization” which we would  
756 cover under *object-oriented design*. They also introduce *unnecessary code* as  
757 category with the example of a debug statement. We did not come across such  
758 code in our sample but could see it as a potential addition.

759 Finally, they clump together *different data structure and algorithm* which  
760 we categorised into separate categories. We would categorise these variations  
761 as either data structure or algorithm differences with probably a high degree  
762 of difference. They found that 93 % of their clone pairs had a variation in  
763 the category *different data structure or algorithm*. We cannot directly support  
764 this value but the tendency. We found that 91 % of the inspected clone pairs  
765 had a difference at least in either *algorithm* or *data structure* and especially for  
766 algorithm the difference was on average large.

767 Tiarks, Koschke, & Falke (2011) created a categorisation for differences in  
768 type-3 clones. Therefore, their focus was on classifying syntactic differences  
769 that probably hail from independent evolution of initially copied code. Yet, the  
770 larger the differences, the more their categories are similar to ours. For example,  
771 they abstract edit operations to *type substitution* or *different algorithms*. We  
772 believe, however, that our categorisation is more useful for FSCs and to improve  
773 clone detection tools along its lines.

## 774 6.2 Impact

775 Independently developed FSCs have very little syntactic similarity. Therefore,  
776 type-1–3 clone detectors will not be able to find them. Newer approaches,  
777 such as CCCD, can find FSCs but their effectiveness still seems limited. Hence  
778 more research in approaches more independent of syntactic representations is  
779 necessary. We will need to find ways to transfer the positive results of Jiang  
780 & Su (2009) with the Linux kernel to other languages and environments while  
781 overcoming the challenges in such dynamic detections as discussed, for example,  
782 in Deissenboeck et al. (2012). We hope our benchmark will contribute to this.

## 783 6.3 Limitations

784 The major limitation of our study is that we did not use a wide variety of  
785 types of programs that exist in practice. The programs from Google Code Jam  
786 all solve structurally similar problems, for example, without any GUI code.  
787 We expect, however, that such further differences would rather decrease the  
788 syntactic similarity even more. The categories might have to be extended to  
789 cover these further differences. Nevertheless, the investigated programs were all  
790 developed by different programmers and are not artificial.

791 Furthermore, we had to concentrate on three clone detectors and two pro-  
792 gramming languages. Other tools and languages might change our results but  
793 we are confident that our selection is representative of a large class of detectors  
794 and programming languages.

## 795 6.4 Future Work

796 We plan to investigate the differences between the tools and the detected clone  
797 pairs of different types in more detail. In particular, we would like to work  
798 with researchers who have built type-4 detectors to test them against our clone  
799 database and to inspect the found and not found clones.

## 800 Acknowledgment

801 The authors would like to thank Benjamin Hummel, Lingxiao Jiang and Daniel  
802 Krutz for their help in getting their tools to work and Kornelia Kuhle for feed-  
803 back on the text.

# References

- 805 Bellon, S., Koschke, R., Antoniol, G., Krinke, J., & Merlo, E. (2007). Compar-  
806 ison and evaluation of clone detection tools. *IEEE Transactions on Software*  
807 *Engineering*, 33(9), 577–591.
- 808 Deissenboeck, F., Heinemann, L., Hummel, B., & Wagner, S. (2012). Challenges  
809 of the dynamic detection of functionally similar code fragments. In *Proc. 16th*  
810 *European Conference on Software Maintenance and Reengineering (CSMR)*,  
811 (pp. 299–308). IEEE.
- 812 Deissenboeck, F., Juergens, E., Hummel, B., Wagner, S., Mas y Parareda, B., &  
813 Pizka, M. (2008). Tool support for continuous quality control. *IEEE Software*,  
814 25(5), 60–67.
- 815 Gabel, M., Jiang, L., & Su, Z. (2008). Scalable detection of semantic clones.  
816 In *Proc. 30th International Conference on Software Engineering (ICSE '08)*,  
817 (pp. 321–330). ACM.
- 818 Higo, Y., Yasushi, U., Nishino, M., & Kusumoto, S. (2011). Incremental code  
819 clone detection: A PDG-based approach. In *Proc. 18th Working Conference*  
820 *on Reverse Engineering (WCRE)*. IEEE.
- 821 Jedlitschka, A., & Pfahl, D. (2005). Reporting guidelines for controlled ex-  
822 periments in software engineering. In *Proc. 4th International Symposium on*  
823 *Empirical Software Engineering (ISESE)*. IEEE.
- 824 Jiang, L., Misherghi, G., Su, Z., & Glondou, S. (2007a). Deckard: Scalable  
825 and accurate tree-based detection of code clones. In *Proc. 29th International*  
826 *Conference on Software Engineering (ICSE)*, (pp. 96–105). IEEE.
- 827 Jiang, L., & Su, Z. (2009). Automatic mining of functionally equivalent code  
828 fragments via random testing. In *Proc. Eighteenth International Symposium*  
829 *on Software Testing and Analysis (ISSTA'09)*, (pp. 81–92). ACM.
- 830 Jiang, L., Su, Z., & Chiu, E. (2007b). Context-based detection of clone-related  
831 bugs. In *Proc. 6th Joint Meeting of the European Software Engineering Con-*  
832 *ference and the ACM SIGSOFT Symposium on The Foundations of Software*  
833 *Engineering (ESEC/FSE)*, (pp. 55–64). ACM.
- 834 Juergens, E., Deissenboeck, F., Feilkas, M., Hummel, B., Schaetz, B., Wagner,  
835 S., Domann, C., & Streit, J. (2010a). Can clone detection support quality  
836 assessments of requirements specifications? In *Proc. ACM/IEEE 32nd Inter-*  
837 *national Conference on Software Engineering (ICSE'10)*, (pp. 79–88). ACM.
- 838 Juergens, E., Deissenboeck, F., & Hummel, B. (2010b). Code similarities beyond  
839 copy & paste. In *Proc. 14th European Conference on Software Maintenance*  
840 *and Reengineering (CSMR)*, (pp. 78–87). IEEE.

- 841 Juergens, E., Deissenboeck, F., Hummel, B., & Wagner, S. (2009). Do code  
842 clones matter? In *Proc. 31st International Conference on Software Engineer-*  
843 *ing (ICSE'09)*, (pp. 485–495). IEEE.
- 844 Kamiya, T. (2013). Agec: An execution-semantic clone detection. In *Proc. 21st*  
845 *IEEE International Conference on Program Comprehension (ICPC)*. IEEE.
- 846 Kamiya, T., Kusumoto, S., & Inoue, K. (2002). CCFinder: A multilinguistic  
847 token-based code clone detection system for large scale source code. *IEEE*  
848 *Transactions on Software Engineering*, 28(7), 654–670.
- 849 Kim, H., Jung, Y., Kim, S., & Yi, K. (2011). MeCC: Memory comparison-  
850 based clone detector. In *Proc. 33rd International Conference on Software*  
851 *Engineering (ICSE '11)*, (pp. 301–310). ACM.
- 852 Komondoor, R., & Horwitz, S. (2001). Using slicing to identify duplication  
853 in source code. In *Proc. 8th International Symposium on Static Analysis*  
854 *(SAS'01)*, (pp. 40–56). Springer.
- 855 Koschke, R. (2007). Survey of research on software clones. In *Dagstuhl Seminar*  
856 *Proc. Duplication, Redundancy, and Similarity in Software*.
- 857 Krinke, J. (2001). Identifying similar code with program dependence graphs.  
858 In *Proc. Eighth Working Conference on Reverse Engineering (WCRE)*, (pp.  
859 301–309). IEEE.
- 860 Krutz, D. E., & Shihab, E. (2013). CCCD: Concolic code clone detection. In  
861 *Proc. 20th Working Conference on Reverse Engineering (WCRE'13)*. IEEE.
- 862 Lakhotia, A., Li, J., Walenstein, A., & Yang, Y. (2003). Towards a clone detec-  
863 tion benchmark suite and results archive. *11th IEEE International Workshop*  
864 *on Program Comprehension*, (pp. 285–286).
- 865 Landis, R. J., & Koch, G. G. (1977). The measurement of observer agreement  
866 for categorical data. *Biometrics*, 33(1), 159–74.
- 867 Marcus, A., & Maletic, J. I. (2001). Identification of high-level concept clones  
868 in source code. In *Proc. 16th Annual International Conference on Automated*  
869 *Software Engineering (ASE 2001)*, (pp. 107–114). IEEE.
- 870 Rattan, D., Bhatia, R., & Singh, M. (2013). Software clone detection: A sys-  
871 tematic review. *Information and Software Technology*, 55(7), 1165–1199.
- 872 Roy, C. K., & Cordy, J. R. (2007). A survey on software clone detection research.  
873 Tech. Rep. 2007-541, Queen's University, Kingston, Canada.
- 874 Roy, C. K., Cordy, J. R., & Koschke, R. (2009). Comparison and evaluation of  
875 code clone detection techniques and tools: A qualitative approach. *Science*  
876 *of Computer Programming*, 74(7), 470–495.

- 877 Svajlenko, J., Islam, J. F., Keivanloo, I., Roy, C. K., & Mia, M. M. (2014). To-  
878 wards a big data curated benchmark of inter-project code clones. In *Proc. In-*  
879 *ternational Conference on Software Maintenance and Evolution (ICSME'14)*,  
880 (pp. 476–480). IEEE.
- 881 Svajlenko, J., & Roy, C. K. (2014). Evaluating Modern Clone Detection Tools.  
882 In *Proceedings of the 30th International Conference on Software Maintenance*  
883 *and Evolution*, (pp. 321–330). IEEE.
- 884 Tempero, E. (2013). Towards a curated collection of code clones. In *Proc. 7th*  
885 *International Workshop on Software Clones (IWSC'13)*, (pp. 53–59). IEEE.
- 886 Tiarks, R., Koschke, R., & Falke, R. (2011). An extended assessment of type-3  
887 clones as detected by state-of-the-art tools. *Software Quality Journal*, 19(2),  
888 295–331.
- 889 Wagner, S. (2013). *Software Product Quality Control*. Springer.
- 890 Wagner, S., Abdulkhaleq, A., Bogicevic, I., Ostberg, J.-P., & Ramadani, J.  
891 (2014). Detection of functionally similar code clones: Data, analysis software,  
892 benchmark. DOI 10.5281/zenodo.12646.
- 893 Wohlin, C., Runeson, P., Höst, M., Ohlsson, M. C., Regnell, B., & Wesslén, A.  
894 (2012). *Experimentation in Software Engineering*. Springer.
- 895 Yin, R. K. (2003). *Case Study Research: Design and Methods*. Applied Social  
896 Research Methods. SAGE Publications.