MossAD: Defeating Software Plagiarism Detection

BREANNA DEVORE-MCDONALD and EMERY D. BERGER, Univ. of Mass. Amherst, USA

Automatic software plagiarism detection tools are widely used in educational settings to ensure that submitted work was not copied. These tools have grown in use together with the rise in enrollments in computer science programs and the widespread availability of code on-line. Educators rely on the robustness of plagiarism detection tools; the working assumption is that the effort required to evade detection is as high as that required to actually do the assigned work.

This paper shows this is not the case. It presents an entirely automatic program transformation approach, MossAD, that defeats popular software plagiarism detection tools. MossAD comprises a framework that couples techniques inspired by genetic programming with domain-specific knowledge to effectively undermine plagiarism detectors. MossAD is effective at defeating four plagiarism detectors, including Moss and JPlag. MossAD is both fast and effective: it can, in minutes, generate modified versions of programs that are likely to escape detection. More insidiously, because of its non-deterministic approach, MossAD can, from a single program, generate *dozens* of variants, which are classified as no more suspicious than legitimate assignments. A detailed study of MossAD across a corpus of real student assignments demonstrates its efficacy at evading detection. A user study shows that graduate student assistants consistently rate MossAD-generated code as just as readable as authentic student code. This work motivates the need for both research on more robust plagiarism detection tools and greater integration of naturally plagiarism-resistant methodologies like code review into computer science education.

CCS Concepts: • **Software and its engineering** → *Genetic programming; Collaboration in software development*; • **Social and professional topics** → **Computer science education**.

Additional Key Words and Phrases: cryptography and security, computers and society, neural and evolutionary computing, programming languages

ACM Reference Format:

Breanna Devore-McDonald and Emery D. Berger. 2020. MossAD: Defeating Software Plagiarism Detection. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 138 (November 2020), 28 pages. <https://doi.org/10.1145/3428206>

1 INTRODUCTION

Plagiarism in programming courses is unfortunately widespread. Universities report that 10–70% of their students have cheated on coding assignments [Baron 2017; Bidgood and Merrill 2017; EAB 2017; Murray 2010]. Two factors combine to effectively incentivize software plagiarism:

- **Homework solutions are often available on-line.**

The rise of open source hosting sites and programming-oriented Q&A sites has had numerous benefits, but also offers temptations for misuse. Open source hosting sites like GitHub often contain full assignment solutions and code, in addition to solutions to typical undergraduate projects [Jue 2014; McMillan 2015]. Sites like Stack Overflow offer programmers abundant resources for coding assistance and tips, which can also provide solutions to coursework; the

Authors' address: Breanna Devore-McDonald, bdevorem@cs.umass.edu; Emery D. Berger, emery@cs.umass.edu, CICS, Univ. of Mass. Amherst, 140 Governors Drive, Amherst, Massachusetts, 01002, USA.

This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2020 Copyright held by the owner/author(s).

2475-1421/2020/11-ART138

<https://doi.org/10.1145/3428206>

FAQ for Stack Overflow specifically states that “it is okay to ask about homework.” [Coehoorn 2019].

- **Enrollment pressure in computer science makes manual inspection impractical.**

The rise in computer science enrollments in recent years is well-documented. The CRA reported in 2017 that there was a 185% “surge” in undergraduate program enrollments since 2006 [Camp et al. 2017]. This rise in enrollments has led to increased class sizes and faculty workloads, and shortages in the number of faculty, instructors, and teaching assistants [National Academies of Sciences, Engineering, and Medicine 2018, pages 5–6]. These factors combine to make manual inspection of assignments impractical in many cases. In a survey of computer science faculty conducted by the authors (Section 3.1), roughly half of the respondents said they perform little to no manual inspection of solutions. Respondents cited ever-growing class sizes (from 150 to nearly 1,000 students) as a reason for this practice.

Plagiarism thus represents an excellent (short-term) cost-benefit approach for students. Students can easily search for and copy and paste solutions to assignments, minimizing their effort. The absence of oversight due to large class sizes further reduces the risk of detection [Yan et al. 2018].

Software plagiarism detection tools like Moss [Aiken 2020; Schleimer et al. 2003] can change this equation by raising the risk of detection, thus making plagiarism less attractive. For software plagiarism detection to be effective, the cost of defeating it must be high, and the risk of detection must be low. Ideally, the difficulty of evading detection must be as high as the effort required to actually do the assignment itself; this feature is precisely what educators rely on, and what designers of plagiarism detectors intend. The author of one of the plagiarism detectors we analyze here (Sherlock [Joy and Luck 1999]) states that “if (students) can plagiarise and avoid detection by Sherlock (or JPlag or MOSS) then they can program to a good standard already.” [Joy 2020]

Existing software plagiarism detectors work primarily by identifying suspiciously-similar code segments via a combination of tokenization and hashing or directly comparing strings [Chen et al. 2004; Liu et al. 2006; Nichols et al. 2019; Prechelt et al. 2002; Schleimer et al. 2003]. For example, Moss’s tokenization effectively undoes the effect of trivial changes like modifying formatting, editing comments, and renaming variables or functions. It also identifies chunks of code that are structurally similar across pairs of assignments, and computes a score corresponding to the percentage of matches. (Section 2 describes Moss’s algorithms in detail.) Moss’s approach generally makes it difficult for students to manually alter code in a way that would elude detection [Bowyer and Hall 1999].

Accordingly, software plagiarism detectors—especially Moss—have seen widespread and growing adoption. The move to online courses due to the coronavirus pandemic has spurred a further increase [Aiken 2020]. Moss has also recently been integrated into Gradescope, a widely used suite of grading tools [Forbes 2018; Gradescope 2019; Singh et al. 2017].

1.1 MOSSAD

This paper presents MOSSAD, a program transformation framework inspired by genetic programming [Le Goues et al. 2011] that defeats popular software plagiarism detection tools, including Moss and JPlag. MOSSAD can automatically produce plagiarized assignments that escape detection; that is, the resulting similarity scores between the original and plagiarized version of the code are indistinguishable from those of different (legitimate) assignments.

MOSSAD thus defies the conventional wisdom that defeating plagiarism detection is difficult or requires significant programming ability. The techniques that underlie MOSSAD could be implemented manually, relying on only the most basic understanding of programming language principles, letting them evade detection by both plagiarism detectors and some degree of manual

inspection. The results of this paper both indicate the need for research into more robust plagiarism detection techniques and highlight the fact that common practice needs to be revisited to account for this kind of plagiarism.

Section 2 first provides a detailed technical description of Moss's algorithms, which are key to understanding the attack that MOSSAD embodies. Section 3 presents our threat model, informed by surveys of faculty, and performs a security analysis of Moss's algorithms; Section 4 then describes the implementation of MOSSAD and how it leverages the discovered vulnerabilities.

Section 5 empirically demonstrates MOSSAD's real-world effectiveness by conducting all experiments on a suite of actual classroom projects from multiple institutions. The evaluation shows that MOSSAD can take a single input program and generate a variant that Moss, JPlag [Prechelt et al. 2002], and Sherlock [Joy and Luck 1999] all consider no more similar than legitimate programs. More insidiously, because of its non-deterministic nature, MOSSAD can also generate *dozens* of variants from a single input, each of which is considered unsuspecting. We conduct a user study of graduate student assistants trained as teaching assistants and find that they rate MOSSAD-generated code as just as readable as authentic student code, making even manual detection unlikely. Section 6 demonstrates MOSSAD's superiority to approaches like code obfuscation.

Section 7 describes a *partial* countermeasure to MOSSAD, applicable to compiled languages like C and C++. Finally, Section 8 and 9 discuss related work and conclude.

1.2 Contributions

In sum, this paper makes the following contributions:

- It presents a security analysis of Moss and identifies a key vulnerability;
- It describes MOSSAD, an automated program transformation framework that targets this vulnerability;
- It presents a detailed empirical evaluation demonstrating MOSSAD's effectiveness on actual student code, highlighting its ability to produce low similarity scores and thus effectively eluding detection by Moss, JPlag, and Sherlock.

1.3 Ethical Considerations

The goal of this work is to demonstrate that existing software plagiarism detectors can be systematically and effectively defeated, to raise questions to educators about the efficacy and robustness of popular plagiarism detection tools, and to motivate research on improving plagiarism detection systems to make them more robust to attack. To ensure that this publication provides substantial advance warning for instructors to take this threat into account, and following standard practices for vulnerability disclosures, we plan to embargo the release of MOSSAD. However, in the interest of advancing science, we plan to make MOSSAD available to researchers upon request.

Responsible disclosure: The authors disclosed the MOSSAD attacks to the author of Moss, as well as the authors of Sherlock and JPlag, over 120 days prior to the publication of this work. As far as the authors are aware, there have not been any substantial changes to any of these tools since this disclosure. In response to the disclosure, the authors had conversations with both Alex Aiken (of Moss) and Mike Joy (of Sherlock) regarding previous attacks against each of the respective tools and discussion of MOSSAD [Aiken 2018; Joy 2020]. The authors did not receive a response from the authors of JPlag.

Privacy concerns: Throughout this study, we ensured that all of our experiments meet community ethical standards. This work was deemed exempt by our institutional review board (IRB), since the data we use presents no more than minimal risk to our subjects. All student assignment data we use was anonymized and assigned random identification numbers before we received it;

we therefore cannot infer the authors of the code from any of our experiments. Student assignment code will not be published. Email addresses obtained from professors surveyed were solely used to validate their identity.

2 MOSS BACKGROUND

Before delving into how MOSSAD operates, we first provide a detailed description of the operation of the Moss software plagiarism detector, which is necessary to understand why MOSSAD's attack is effective. While we also evaluate MOSSAD against another plagiarism detector, JPlag [Prechelt et al. 2002] (Section 5.7.1), we focus our attention on Moss for the following reasons:

- **Moss is popular.** To the best of our knowledge, Moss is currently the most widely used software plagiarism detection system. Aiken reports that the Moss web service has roughly 300K current Moss accounts, with 50K–100K new Moss accounts per year [Aiken 2018], and it is now integrated with Gradescope.
- **Moss is general.** At the time of writing, Moss supports 24 different programming languages, making it to our knowledge by far the most widely-applicable plagiarism detector.
- **Moss represents the state of the art.** Empirically, we found Moss to be the most effective of the tools we tested. Its design is similar to that employed by other leading plagiarism detectors.

As Section 1 explains, Moss is an automatic system for detecting the similarity of software provided as a free web service, and it is widely used to detect plagiarism of classroom assignments. Moss measures the similarity between every combination of two input source programs (that is, all $O(n^2)$ pairs) and computes the similarity score as the percentage of lines matched between each pair. Beyond the above-cited characteristics, Moss has the following properties that make it attractive as a plagiarism detector:

- **Whitespace and comment insensitivity:** Moss ignores changes in whitespace, capitalization, or text in comments.
- **Noise suppression:** Moss exhibits a low false positive rate by reporting only large chunks of copied code. For example, the appearance of individual tokens like `int` in two files would not be reported as instances of plagiarism.
- **Position independence:** Moss can find plagiarized code regardless of its placement in the program.

2.1 Normalization

To provide whitespace insensitivity, Moss performs a *normalization step* before it moves on to actually computing similarity. Depending on the language of the input files, Moss eliminates irrelevant features that should not distinguish documents, including semantically-irrelevant whitespace or comments. This normalization step also consists of renaming all identifiers to the same value; defeating variable renaming as a means of evading detection. The input files are then tokenized into a normal form across all input languages, and then sent as input to Moss's *fingerprinting engine*.

2.2 Fingerprinting

Moss's fingerprinting engine ensures both noise suppression and position independence. It operates on the normal form that the previous step produces; the same engine is used for all supported source languages. The overarching goal of the fingerprinting engine is to prepare the inputs for efficient similarity assessment by transforming the source files into small sets of values called *fingerprints*. The algorithm aims to produce the fingerprints with enough identifying information as to not sacrifice correctness when assessing similarity.

As a first step, Moss divides the normalized program into a sequence of overlapping k -grams, a contiguous substring of length k , where k is an internally-defined *noise threshold*. Then, Moss computes a hash for each k -gram, resulting in a sequence of multiple hashes representing the input (the fingerprints are chosen from this sequence). The value of k is particularly crucial in the trade-off between noise suppression (false positive rate) and sensitivity: large k would increase confidence in similarity matches, since it represents more source code; however, large k would not allow Moss to detect shorter code matches. As a result, the size of k has a direct impact on positional independence, since Moss is unable to detect relocation of substrings with a length shorter than k . Though the value of k is unknown to the user, it can be obtained straightforwardly through a brute force attack.

After generating the sequence of hashes of k -grams, Moss then performs its unique *winnowing* algorithm, which creates fingerprints by selecting a subset of the hashes to represent the entire input. The first step of winnowing is to group the hashes into overlapping windows of length w , where w is the difference between a fixed *guarantee threshold* t and the *noise threshold* k , plus one. The guarantee threshold is the minimum substring length, in terms of tokens, that is to be detected when compared for similarity (intuitively, k must be no greater than t , otherwise substrings of t length would not be detectable). Like the noise threshold, the guarantee threshold is an unknown value specified in the Moss fingerprinting engine.

Since the window size is defined as $t - k + 1$, each substring of size t has at least one corresponding window such that each hash in the window would be detected as similar, if another input file contained that substring. As a result, to scale down the set of fingerprints for an input, the next step of winnowing is to choose a single hash from each window to be a fingerprint. The algorithm for this selection is straightforward: Moss chooses the *minimum* hash value of each window. If the hash has already been chosen as a fingerprint in the preceding window, then it will not be chosen again in the following window (essentially skipping successive windows with the same minimum hash value). If the minimum hash value appears more than once in a window, Moss chooses the rightmost one. This algorithm actually further scales the set of fingerprints: the minimum hash value of a window is most likely the minimum hash value of the next window (recall that the windows overlap, similar to the k -grams), and therefore results in far fewer fingerprints than the total number of w -sized windows.

The engine keeps track of the positional information of the fingerprints, so that the source code of matches can be retrieved later for reporting purposes. The engine adds the fingerprints it chooses for each file along with the positional information to a database. Then, each file is fingerprinted a second time and queried in the fingerprint database for positional information, and the engine returns the set of all matching fingerprints across the input files. Rather than outputting the raw fingerprints to the user, Moss returns the source code found from querying the fingerprint database, such that the code matches are easily viewable by a human user. Moss users can also specify that they want Moss to exclude boilerplate or template code, in which case the fingerprints of the boilerplate code are ignored if matched.

2.3 Example

Figure 1 provides an example of Moss's normalization and fingerprinting algorithms in action. Figure 1a presents an example consisting of two lines of valid C code. We omit headers and function definitions for brevity and clarity. Figure 1b presents an example set of tokens that could result from smart-tokenizing that code. Figure 1c shows potential hashes of these tokens, assuming a window of length 3 (3-gram model). Figure 1d shows the first stage in the winnowing algorithm, where the k -gram hashes are further grouped into windows of size $w = t - k + 1$, which is 3 in this case. Finally, Figure 1e shows the result of finishing Moss's winnowing algorithm. As previously

```
int hello = 0;
return hello;
```

(a) A simple C program with headers and functions removed.

```
TYP_INT ID EQ NUM SEMI
RET ID SEMI
```

(b) The example set of tokens after smart tokenization of the C code in Figure 1a.

```
30 15 56 83 71 10
```

(c) Representative hashes of the tokens from Figure 1b, assuming $k = 3$.

```
(30 15 56) (15 56 83)
(56 83 71) (83 71 10)
```

(d) Windows of hashes of length $w = t - k + 1 = 3$

```
15 56 10
```

(e) Post-*winning*: keep the smallest hash in each k -gram, dropping consecutive duplicates, producing the fingerprints of the program.

Fig. 1. Moss Algorithm. An example run of performing Moss's smart tokenization and fingerprinting algorithms, assuming that the guarantee threshold t is 5 tokens and the noise threshold k is 3.

noted, winnowing creates a smaller set of hashes to represent the entire document by choosing the smallest hash from each window of size w .

3 MOSS SECURITY ANALYSIS

This section performs a security analysis of Moss. First, we describe our threat model, which was informed by the results of two surveys of computer science faculty. Second, we verify empirically that Moss operates as described in Section 2, which was based on the paper describing it [Schleimer et al. 2003]. Finally, we identify a key weakness of Moss's algorithms, which MossAD exploits.

3.1 Developing the Threat Model

To inform a threat model (that is, establishing the power and limitations of the adversary), we conducted two surveys on two distinct social media platforms to understand the grading practices in computer science courses, including their usage of software plagiarism detectors.

We first conducted a poll on Twitter asking about the degree of manual inspection performed on code. This poll asked respondents to optionally provide further information regarding class sizes and time spent on manual code inspection. Of the respondents ($N=128$), 65% said that they manually inspect nearly 0% of student assignments. The additional comments to the poll ($N=12$) reported that most student code does not get manually inspected; specifically, the code that *does* get inspected are assignments flagged with high similarity scores from Moss or other plagiarism detectors. (We did not ask but believe that even when manual inspection is performed, it is not conducted across every N^2 pair of assignments.)

We conducted a second, more detailed survey using Google Forms advertised on Facebook and Twitter. Of the respondents to this survey of faculty members ($N=50$ responses, 35 unique institutions), 65% report that they manually inspect nearly all of their assignments and 19% report a manual inspection rate of nearly 0%. This result strongly suggests that the populations sampled by the two polls are different. Aggregating the results across the two surveys yields an average of 52% of the respondents inspecting nearly no student solutions.

Of the respondents to the second survey, 58% of those reporting high manual inspection rates also reported that they used Moss to inform their decision of which assignments to inspect by either sorting based on Moss scores or finding cliques from the Moss output. On the other hand, 85% of respondents reporting low manual inspection rates limit their spot checks to the high matches reported by Moss. Informed by the results of these surveys, we developed the following threat model:

Threat Model

- Students are incentivized to minimize the time spent on assignments while maximizing their grade. Therefore, students will not spend more time attempting to defeat the plagiarism detector than actually doing the assignment.
- Students have arbitrary access to peers' assignment solutions or solutions found online. Students may use these solutions as the basis for their plagiarized assignment solution.
- In addition, all solutions are also available to the instructor; attacks that depend on outsourcing (hiring someone to write an entirely new program that is not available online) are to our knowledge beyond the scope of existing plagiarism detection tools, since the original program is unavailable.
- Course staff relies primarily on software plagiarism detection, beyond the lightweight manual inspection described below.
- Course staff will only examine student code for spot checking or for those assignments that receive high similarity scores with another assignment. Students will thus aim to achieve the lowest similarity score possible in order to avoid potential detection, but must also take care in that the submitted code retains readability; that is, submitting clearly obfuscated code would be an unacceptable risk in the case of a spot check.

The implications of our threat model are that the greatest risk to current practices is a system that can automatically and rapidly transform one source program into a multitude of variants that each result in low Moss scores, while maintaining readability (as we show in Section 5, this is exactly what MOSSAD achieves).

3.2 Investigating Moss's Implementation

With our threat model in hand, we next evaluate whether the current implementation of Moss reflects its description in the paper describing it [Schleimer et al. 2003] or if it has materially changed, since an attack based on the algorithm described in the paper could plausibly fail against a new implementation.

As Moss's source code is not publicly available, we could not manually inspect it. Instead, we perform a series of experiments to evaluate its behavior. Our experiments involve repeatedly transforming differently-sized regions of code via a series of program transformations, including code rearrangement and identifier renaming, and observing their impact on Moss's reports.

We quickly were able to verify that noise suppression behaved as expected; that is, Moss's window is sufficiently large to exclude individual lines like `int x`; as constituting plagiarism.

The next experiment was to examine if Moss's whitespace insensitivity worked as described. To do this, we took two copies of the same program and edited one with comments, and changed the names of identifiers. We performed this experiment automatically for 30 files in our dataset (discussed in Section 5.1) and compared the files with Moss. This experiment confirmed that whitespace insensitivity worked as expected: the Moss scores before and after the alterations did not differ.

We then examined Moss's position independence. To do this, for the 30 files, we created copies of files and rearranged the functions. After this, we compared the files with Moss to see if the similarity score had changed. In most cases, the similarity score did not change; in the few cases in which it did, it did not go below an 85% match. We expect the change in Moss score is a product of Moss's scoring method, rather than its similarity-detection algorithm. All of the code was still detected as similar, but the code matches were separated into multiple sections. This division into sections appears to affect the percentage of similarity computed, though not to a large degree.

Finally, we performed the same experiment, but at a finer grain. By rearranging code *within* a function, we could examine how fine of a granularity Moss can detect plagiarism—that is, approximately how many lines of code correspond to a window of tokens. We tested this by initially rearranging a single line of code, and then increasing the number of lines that we rearranged by one at each iteration.

We found that Moss could not detect plagiarism when just one or two consecutive lines of code were rearranged. However, Moss could consistently detect rearranged plagiarism for 3 and 4 consecutive lines of code. Regions of greater than 5 lines of code were always detected. As a result, we infer that the number of tokens used as the length of the hashing window in Moss corresponds to between 3 and 4 lines of code. This result is not precise, as the number of tokens per line in a file can vary, but we found this to be the case consistently within our dataset.

Result: We find that Moss's current implementation closely resembles the original description. We confirm that Moss performs as expected: it correctly assigns high similarity scores to code that has been altered with static transformations. Crucially, we observe behavior that is consistent with the hashing and winnowing approach described in the original paper; it is this aspect of Moss's algorithm that we identify as vulnerable to attack.

3.3 Key Observation: Hash Disruption

Recall that Moss hashes program tokens as k -grams (that is, k is the window size for hashing tokens). As mentioned earlier, this value is a key parameter for Moss: a window size must be tuned so it is large enough to reduce noise (false positives), yet small enough to detect meaningful code matches (true positives).

We hypothesized that Moss's algorithm suffers from the following vulnerability, which we refer to as the *hash disruption hypothesis*: **inserting a single token in a window of length k should cause Moss to fail to detect any similarity**.

To validate this hypothesis, we needed the value of k , which is not made public and which may be different for each programming language Moss supports. As described above, we reverse engineered a rough value for k by inserting statements at different strides. With our approximate k in hand, we were able to verify that, as anticipated, a single change within the window leads to a different hash, since Moss no longer treats such code fragments as similar.

Figure 2 illustrates the hash disruption vulnerability in simplified form. For the purposes of illustration, we omit the effect of winnowing, which only minimizes the number of hashes that are actually compared across files. The hash disruption vulnerability described strikes at the heart of Moss's algorithm; winnowing actually exacerbates its impact.

Recall the code in Figure 1a. To effect hash disruption, we simply add an additional line of code between the two original lines. The new line of code creates a Boolean variable. Following the same method of tokenization as used in Figure 1, Figure 2b presents a tokenization of the new C code. Then, Figure 2c presents a hash of the tokens, again using the same method of hashing as used previously. This figure shows that there are now more hashes than shown in Figure 1c, which is as expected as there are now more tokens to be hashed. Hash 83 from Figure 1b has been disrupted and is no longer present in the new set of hashes, leading to a failed match.


```
int hello = 0;
bool nothing = true;
return hello;
```

(a) To demonstrate Moss's vulnerability, we first add a benign, semantics-preserving addition to the code from Figure 1a.

```
TYP_INT ID EQ NUM SEMI
TYP_BOOL ID EQ BOOL SEMI
RET ID SEMI
```

(b) The example set of tokens after smart tokenization. Noticeably, identifiers are tokenized to the same token.

```
30 15 56 12 45 39 97 62 80 71 9
```

(c) Potential hashes of these tokens. Notice that the hash value 83 (shown in Figure 1b) has been eliminated.

Fig. 2. Deterministically defeating Moss via *hash disruption*. A demonstration of *hash disruption*, which exposes a key vulnerability of Moss's fingerprinting algorithm. After a benign addition to a series of statements, one of the hashes created by Moss's fingerprinting algorithm (83) is no longer present. MOSSAD leverages this vulnerability to drastically reduce the effectiveness of Moss's matching algorithm.

3.4 Other Plagiarism Detection Tools

In addition to Moss, we aimed to collect all of the related systems described in Section 8. This effort was largely unsuccessful, with three exceptions: JPlag [Prechelt et al. 2002], Sherlock [Joy and Luck 1999], and Fett [Nichols et al. 2019]. JPlag and Sherlock are the only tools, aside from Moss, our survey respondents reported using. Other related systems are described in Section 8.

The algorithms used by both JPlag and Sherlock are similar in effect to Moss. Like Moss, JPlag and Sherlock detect software similarity. Both also employ tokenization and normalization of the input file. However, rather than relying on hashing and winnowing to build the file's fingerprints, JPlag and Sherlock both use Running Karp-Rabin Greedy String Tiling to identify code matches and produce a similarity score. This approach has the same effect as Moss's core algorithm (excluding winnowing).

On the other hand, the algorithm used by Fett is unlike that of Moss. Rather than relying on tokenization, Fett uses parse trees with weighted nodes in order to determine what needs to be filtered. For the scoring stage of the algorithm, unlike the hashing and fingerprinting schemes of Moss, Fett uses the Smith-Waterman algorithm to align and identify common subsequences, while penalizing gaps [Smith et al. 1981]. Although there are no "fingerprints" as employed by Moss, JPlag, and Sherlock, and therefore no hashes to be disrupted, MOSSAD has a disruptive effect on Fett's algorithm nonetheless. MOSSAD disrupts Fett's sequences by regularly placing important nodes (such as useless assignments and empty conditionals) at regular intervals, leading to large gap penalties.

In section 5.7, we evaluate the effectiveness of MOSSAD on JPlag, Sherlock, and Fett. As we demonstrate, while MOSSAD targets Moss's fingerprinting algorithm, it in effect implicitly targets the algorithms used by these systems as well, as they are all subject to the same disruption vulnerability.

4 MOSSAD

The hash disruption phenomenon that Section 3 describes forms the groundwork for defeating detection: foil Moss's algorithm by systematically introducing benign (semantics-preserving) statements within windows. This approach is the point of departure for MOSSAD.

Fig. 3. MossAD system overview. MossAD is a program transformation framework inspired by genetic programming that transforms a file into semantics-preserving variants that defeat Moss’s detection algorithms (Section 4).

MossAD is a program transformation framework that fully automates this approach. MossAD presents the worst-case with respect to the threat model developed in Section 3.1: it is not only entirely automated (requiring little to no effort) and fast (taking minutes, as we describe in Section 5), but also can produce numerous variants that each escape detection by Moss, even with quite low detection thresholds. Figure 3 presents an overview of the system. Figure 4 shows an example of running MossAD on an input function with a target Moss score of 10%. Figure 4c shows that Moss considers the original and the MossAD-processed variant to have no matches.

Given an input file that compiles without errors, MossAD uses techniques inspired by genetic programming to transform it into semantically-equivalent variants. Given a user-provided target similarity score (e.g., 25%), MossAD will attempt to produce a variant that achieves a Moss score less than or equal to the target score when compared to the original input file. The current MossAD prototype only supports C/C++ as input programming languages, but the technique described here is directly applicable to any languages that can be statically compiled.

4.1 Initialization

The first step of the MossAD algorithm is the initialization phase. MossAD compiles the input file (with optimization) to an intermediate form. For C, MossAD compiles the files with Clang to object code. If compilation succeeds, the initialization phase then creates a copy of the input source code, which becomes the first variant. Initialization also seeds the *pool*, which initially consists of every line from the user input file, for use by the generation phase.

4.2 Generation

The generation phase performs selections and mutations, and reinitializes the pool for the next generation. The first module of generation is selection. Selection chooses one line of code from the pool and performs naïve semantic checks to speed up the subsequent generation module. These checks include ensuring that the selected line is a complete statement (e.g., ends with a semicolon for C) and will not change the semantics of the program (e.g., is a print statement). Once a selection is chosen, the selection engine sends the chosen line to the next generation module: mutation.

```
int NchooseK(int n, int k){
    if(k==0)return 1;
    if(n==k)return 1;
    else
        return NchooseK(n-1, k-1) +
               NchooseK(n-1, k);
}
```

(a) A student's solution to part of an assignment in our dataset. This function computes $\binom{n}{k}$ in C.

```
int NchooseK(int n, int k){
    if(k==0)return 1;
    int NchooseK(int n, int k);
    if(n==k)return 1;
    else
        return NchooseK(n-1, k-1) +
               NchooseK(n-1, k);
    int i = 0;
}
```

(b) A result of applying MossAD to this code.

```
No matches were found in your submission.
```

(c) Moss's output from comparing these programs.

Fig. 4. MossAD in Action. An example of MossAD's output, using the code in 4a as the input file and 10% as the target Moss score. By injecting just two lines of semantics-preserving code, MossAD successfully defeats Moss.

4.3 Mutation

The mutation engine performs mutations by inserting the selection into the file and ensuring that the original file's semantics are preserved. To do this, the mutation engine first creates a temporary variant (which is a copy of the current variant) to be mutated. Mutation will randomly insert the selected line from selection into the temporary variant and perform a sequence of checks.

In each generation, there are two outcomes for variants. The first is a successful mutation, where the selection engine produces a semantics-preserving statement. If the variant is not semantically equivalent to the original (which is checked as described in Section 4.5), then mutation discards the temporary variant and restarts the selection module with the same pool for regeneration.

Note that a mutation is considered successful whether or not the target Moss score has been reached, since it almost always requires multiple successful mutations to reach the target.

4.4 Score Checking

The second check mutation performs is testing if the target Moss score has been reached by querying Moss and scraping the web data from the resulting URL. If the target has been achieved, then the current generation will be the last and the temporary variant will be outputted by MossAD.

If the target Moss score check fails, the mutation module adds this successful mutation to the pool to increase the odds of inclusion in future variants of lines that have been proven to be successful. It then returns to generation with the temporary variant as the new variant.

4.5 Program Equivalence

To determine program equivalence, the current prototype of MossAD simply directly compares the intermediate representation of programs after compiling both with a high level of optimization. It directly *diffs* the resulting object code after compiling with the `-O3` flag. MossAD uses this method for a variety of reasons: it is sound, fast, and is straightforward to implement. This approach to equivalence checking naturally limits the mutations to those that would be compiled away, such as dead code and redundant assignments.

This approach could, with additional engineering, be extended to JIT-compiled interpreted languages by modifying the JIT-compiler to emit generated object code. The key engineering obstacles to overcome would be non-determinism in the JIT compilers themselves, such as when they trigger optimization based on elapsed time rather than number of instructions executed, or use other, pseudorandom sampling approaches. Additionally, inlining across system library boundaries would need to be disabled, as it could result in spurious matches.

Leveraging more advanced and general approaches for determining program equivalence, such as parameterized program equivalence [Kundu et al. 2009; Roychoudhury et al. 2000] and program equivalence using context-free grammars [Rosen 1972], would open the door to a wider range of mutations that could be applied. In particular, two variants could produce different code while retaining semantic equivalence. Another possible but unsound approach would be to generate a large number of test cases via fuzzing, and using identical behavior on those test cases as a proxy for equivalence.

4.6 Entropy

To increase the likelihood of generating successful mutations, MossAD includes a feature that allows users to make additions to the pool for selection, which we refer to as *entropy*. Users can add their own code to an external file that MossAD will use as an addition to the initial pool of the user input file. In addition, if entropy collection is enabled, MossAD will also store all successful mutations to the entropy file during mutation. This approach aims to increase the odds of successful mutations in future uses of MossAD for other assignments.

For our experiments (Section 5), we manually assembled a single entropy file based on our observations of common C statements. These include common variables and basic mathematical operations on them. For example, we included declarations and assignments to variables named `count`, `i`, `j`, `k`, `ret`, and `status`. We found that these are extremely frequent in public C codebases, such as Linux and Redis, making them generally unsuspecting. The resulting entropy file contains approximately 30 lines of these statements.

5 MOSSAD EVALUATION

Our evaluation focuses on the following research questions:

- (1) **RQ1 (Unsuspectiousness)**: Can MossAD produce code variants that yield unsuspectiously low scores typical of non-plagiarized code?
- (2) **RQ2 (Readability)**: Does MossAD produce code that is as no less readable (and thus no more suspicious) than legitimate programs?
- (3) **RQ3 (Effect of Input Size)**: How does the size of input programs affect MossAD's efficacy?
- (4) **RQ4 (Mass Plagiarism)**: Can MossAD enable *mass plagiarism* by producing numerous mutually dissimilar results from the same input programs?
- (5) **RQ5 (Performance)**: Is the execution time required by MossAD low enough to make it practical?
- (6) **RQ6 (Generality)**: Is MossAD equally effective at defeating other plagiarism detectors?

5.1 Methodology

Dataset. To answer all of these research questions, we evaluate MossAD with course assignments provided by faculty from several universities; all were previously stripped of personally identifiable information. Our dataset consists of three homework projects from undergraduate courses, all in C. Project 1 required students to implement a bubblesort over an array of pointers to strings; solutions range from 11–101 LOC (all figures computed by `sloccount`, which excludes comments and empty

Fig. 5. Across a suite of actual student assignments, MOSSAD consistently produces variants that result in unsuspicious Moss scores. Light gray indicates legitimate assignments, while dark gray indicates MOSSAD-produced variants. (§5.2)

lines). Project 2 required students to implement a program that converts parts of images (stored in a format akin to PPM) from color to grayscale; solutions range from 19–53 LOC. Both Project 1 and 2 include 135 student solutions. Project 3 required students to use recursion to implement a function computing $\binom{n}{k}$. This project includes 120 student solutions, ranging from 14–116 LOC.

Suspiciousness Threshold. Evaluating MOSSAD’s ability to produce unsuspiciously low scores depends on the choice of a threshold of similarity that constitutes suspiciousness. That is, programs that are below this threshold of similarity will not be viewed as being likely to constitute plagiarism, and thus will be less likely to be manually inspected. Note that, for large classes, Moss does not report the similarity of every file, instead reporting only the 250 most similar pairs, in descending order of similarity.

For our evaluation, we empirically set a threshold that would result in low false positive rates across our dataset. We examined the Moss scores for each project in our dataset and found that the average similarity score of assignments pairs (in the top 250 results) was 26.2%, with a bootstrapped 95% confidence interval of [25.3%, 27.2%] (10,000 iterations); see Table 1 for complete statistics. We manually inspected files in the high end of that range and determined that these did not appear to constitute instances of plagiarism. In addition, survey respondents indicated that the average threshold they use for deeming Moss scores suspicious is 30%. Informed by these two results, we conservatively set our threshold for suspicion for experiments to 25%, just below the bottom of the confidence interval.

5.2 RQ1: Unsuspiciousness

The most important metric of performance for a system aimed at defeating plagiarism is its effectiveness at producing unsuspicious variants; that is, the similarity of transformed programs to the original version is reported as below the suspiciousness threshold we derive above (25%). To measure MOSSAD’s effectiveness, we apply MOSSAD to files from three C data sets and used Moss to compare the generated variants against the entire body of student code.

We perform the same set of experiments across the projects in our dataset. For each, we randomly select five assignments to be used as inputs to MOSSAD. We run MOSSAD with a target Moss score of 25% together with the entropy file described in Section 4.6. After generating these MOSSAD variants, we add these five additional files, and then ran Moss across the entire corpus.

MOSSAD performed similarly across all of the projects; Figure 5 presents these results. As previously mentioned, Moss returns a list of the highest scoring pairs of assignments; our graphs present the scores of these pairs. For each pair consisting of at least one MOSSAD variant, the corresponding

Table 1. Summary of effectiveness results for MOSSAD (§4) and “ablated” variants (§5.8). Lowest values are shown in **boldface**. All of the variants produce code that is usually below the average similarity of non-plagiarized student code. MOSSAD always yields the lowest similarity scores. *Individual plagiarism* denotes an attempt to plagiarize from exactly one source program; *mass plagiarism* denotes an attempt to produce multiple variants from a single source program. Note that MOSSAD_{DET} is ineffective for mass plagiarism, since its transformations are deterministic, as all generated versions would be identical (§5.8).

Code type	Moss similarity scores			
	Min	Max	Avg	95% CI
<i>Individual plagiarism</i>				
MOSSAD	7%	24%	15.4%	[13.6% , 17.2%]
MOSSAD _{DET} (§5.8.1)	6%	31%	16.7%	[13.8%, 19.7%]
MOSSAD _{NONDET} (§5.8.2)	6%	33%	19.6%	[14.5%, 24.7%]
<i>Mass plagiarism</i>				
MOSSAD (§5.5)	5%	27%	12.1%	[11.5% , 12.7%]
MOSSAD _{NONDET} (§5.8.2)	15%	46%	24.2%	[23.4%, 25.0%]
Legitimate student code	6%	98%	26.2%	[25.3%, 27.2%]

bins are shaded in dark gray, while for each pair that consists solely of authentic student code, the corresponding bins are shaded in light gray.

Table 1 summarizes the average similarity scores and their bootstrapped 95% confidence intervals. For every project, the MOSSAD variants achieved similarity scores lower than the average Moss score for the 250 highest scoring pairs of authentic programs. MOSSAD yielded variants with average similarity scores of 15.4%, while the average similarity of legitimate code was 26.2%. As with the legitimate student code, we used the bootstrap to compute a 95% confidence interval around the means; the confidence interval for MOSSAD variants was [13.6%, 17.2%].

RQ1 (Unsuspectiousness): *MOSSAD generates variants that are unsuspectious: on average, Moss computes their similarity to the original versions to be just 15%, well below the average similarity of authentic programs.*

5.3 RQ2: Readability of MOSSAD-Generated Code

While generating unsuspectious variants is important, it is not enough to avoid the risk of detection in case of spot checks or routine manual inspection. We examine MOSSAD’s effectiveness at producing variants that are just as readable as authentic student code via an experiment with teaching assistants in-training. For the purposes of this evaluation, the term “readable” is used roughly to mean that the code looks like legitimate student code to graders.

For the user study, the participants ($N=30$) were Computer Science PhD students enrolled in a mandatory teacher training course, required before they can serve as Teaching Assistants at their home university. Participants were in the first to fourth years of their PhD.

To simulate the effect of manually grading MOSSAD code in the real world, participants were asked to grade a series of student assignments consisting of a mix of authentic student code, MOSSAD-generated code, and obfuscated code. Each participant was given assignments in random order. All assignments were anonymized. We deliberately did not provide Moss similarity scores; because MOSSAD generates extremely low scores, we did not want to risk biasing students against detecting plagiarism for these cases.

The participants were provided the prompt for the assignment along with a rubric that contained three core elements that needed to be scored on a four-point scale (1-4, higher is better): *Accuracy*, *Program Design*, and *Code Readability*. This grading rubric was derived from several publicly available grading rubrics for Computer Science courses. Each participant was given 20 minutes to grade 10 assignments in order to mimic the average time a grader actually spends manually inspecting code, as reported by survey respondents.

The rubric provided guidance for each score for each task and asked the participants to write a few sentences on their overall impression of the code, including anything they believed the instructor should know. The participants were given time to familiarize themselves with the rubric before they were given assignments to grade. For the Accuracy component, participants were given a simulated autograder score for each assignment. For Program Design, participants were asked to judge the structures used and the overall design of algorithms. Finally, for Readability, participants were instructed to examine the documentation and indentation of the code, as well as if the identifiers used were appropriate for the task and if the code was overall easy to read.

As in all of the experiments in this evaluation, the MOSSAD code presented to participants contained only additional trivial statements and declarations, all of which optimize away during compilation (i.e., they are semantics-preserving).

Results: Using Hedges' G on the two populations of authentic student code scores and MOSSAD code scores, we find the effect size for both Program Design and Code Readability to be low (0.15 and 0.11, respectively). That is, the effect of using MOSSAD on programs results in little impact on their assessed program design or readability. Participants were asked to provide a few sentences in feedback for each assignment and to relay notes to the instructor if they felt the need; none reported anything out of the ordinary, let alone suspicious code or plagiarism.

However, using Hedges' G with the populations of authentic student code scores and obfuscated code scores, we discovered very different, extreme results. In this case, the effect of obfuscation was quite large on both Program Design (1.7) and Code Readability (1.1), empirically substantiating our expectation that obfuscated code looks markedly different from unobfuscated code, as anticipated in Section 6.1.

RQ2 (Readability): *MOSSAD generates variants that are essentially indistinguishable from authentic programs in terms of readability.*

5.4 RQ3: Effect of Input Size on MOSSAD

To assess the impact of input size (that is, the length of input programs) on MOSSAD, we evaluate MOSSAD's effectiveness at generating variants from input files of varying sizes, with and without an entropy file.

Throughout the experiments, input file size was observed to have an effect on both Moss scores and MOSSAD outputs. With Moss, for small enough files, it appears that the window size for hashing is simply too large, since very minimal disruptions are needed to reach a 0% Moss score. It is intuitive that smaller files need fewer disruptions; however, the number of disruptions does not follow a linear scale for achieving low Moss scores. We observe that the number of disruptions for small files to achieve low scores is not proportional to the number of disruptions for large files.

The length of program inputs has a variety of effects on MOSSAD. We first consider these assuming MOSSAD does not use its entropy file feature. For small files, MOSSAD performs poorly without an entropy file. Although smaller files need a proportionally lower amount of insertions to disrupt the amount of windows, the number of successful mutation possibilities are extremely low. Small files generally have fewer semantics-preserving lines of code, so it limits the possible lines that can be re-inserted by MOSSAD. Smaller files also have limited locations for insertions such that

Fig. 6. Mass plagiarism with MossAD yields unsuspicious Moss scores. Because MossAD is non-deterministic, it can generate multiple variants from the same input source code. Using 5 different base files as input to MossAD and creating 30 variants per base file, we find that the highest Moss score of any pair of files is 27%, not significantly larger than the 25% suspiciousness threshold.

the insertions do not affect the output of the program. On the other hand, for larger files, there are more mutation possibilities: more possible statements that can be inserted, and more locations where such statements can be inserted.

We empirically evaluate the length of input files necessary for MossAD to run properly without the need for an entropy file. For this experiment, we sorted the programs from our data sets by length (LOC) and ran MossAD on each assignment in descending order until MossAD consistently timed out: we define these terms below.

Since MossAD is nondeterministic, we used each assignment as an input four times. We set our timeout threshold at five minutes. If MossAD could not find a single successful mutation within that time, or if the variant ever reaches a length of $2.5\times$ the length of the input file without achieving a Moss score that meets the threshold, it would then timeout. We deem “consistent timeouts” as those that timeout for half or more of the iterations, for three unique assignments in a row. Using these methods, we found that, on average, MossAD is not effective on inputs of 35 LOC or less without an entropy file.

When MossAD is run with an entropy file, the system produces target-achieving outputs regardless of input size. An entropy file increases the probability of generating successful mutations, by definition, which is independent of input size. In fact, the experiment to evaluate this is exactly the experiment shown previously in Figure 5, in which MossAD consistently produces variants that achieve Moss scores below 25%.

RQ3 (Effect of Input Size): MossAD is effective on programs above 35 LOC; for shorter programs, it is effective when incorporating an entropy file consisting of generic lines of program source code.

5.5 RQ4: Mass Plagiarism with MossAD

To demonstrate the effects of mass plagiarism with MossAD, our third experiment consisted of generating multiple MossAD-disguised files from the same input file and using Moss to assess the similarity scores between said variants. Using five randomly-selected student assignments from our data sets as input, we generated 30 MossAD variants per base file. Since MossAD is nondeterministic, this process created 30 mutually-dissimilar files per base file. Separately, each set of 30 variants

Fig. 7. MossAD consistently takes less than 5 minutes to create a variant. 95% of Project 3 iterations complete within 7 minutes, the longest of all three data sets.

was assigned similarity scores by Moss; these scores are shown in Figure 6. Recall that Moss only shows the top 250 highest-scoring pairs, therefore, for each base file there are 250 corresponding data points (1250 in total).

Since MossAD is completely nondeterministic, the odds of an insertion hit between files is supremely low, resulting in overall low Moss scores. As shown in Figure 6, 90% of the top 250 pairs score below 26% similarity for each unique base file, which is only 1% higher than the threshold of suspicion. The highest Moss score was 27% and the average score was 12.1%. These results are summarized in Table 1.

RQ4 (Mass Plagiarism): *MossAD is able to generate large numbers of variants from a base file that Moss scores as non-suspicious, enabling mass plagiarism.*

5.6 RQ5: MossAD: Performance

The performance of MossAD primarily depends on the latency and threshold of Moss, which is a function of the number of jobs at any particular moment. In this section, we examine the run time performance of MossAD for each of our three data sets. For each data set, three randomly selected student assignments are used as inputs to MossAD, along with the default Moss score of 25% we established in Section 5.1, and each input is evaluated with MossAD ten times.

Figure 7 plots the distribution of execution times for all 30 iterations for each project. As shown in the graph, Project 3 inputs produced the longest run times with 95% completing within 7.5 minutes. Inputs from Projects 1 and 2 resulted in very similar run times for almost all of their iterations, with 95% of the iterations completing within approximately 4 minutes for Project 1 and 5.5 minutes for Project 2. These run times are nondeterministic and can change depending on how many jobs are running simultaneously on Moss. That said, these times are in general far less than the time required to actually do the homework assignments.

RQ5 (Performance): *MossAD is able to generate variants with low similarity scores in under 10 minutes, far less than the amount of time generally allotted to complete assignments.*

Fig. 8. MOSSAD effectiveness versus JPlag [Prechelt et al. 2002], another widely used software plagiarism detector. MOSSAD-produced variants produce even lower similarity scores with JPlag than with Moss, with one outlier at 38% in Project 3. For each dataset, we omit pairs scoring 0%, which constitute the vast majority of results.

5.7 RQ6: MOSSAD vs. Other Plagiarism Detectors

To address the question of MOSSAD’s effectiveness against other software plagiarism detection tools, aimed to collect all of the systems discussed in Section 8. Unfortunately, this effort was largely unsuccessful, as almost none of these systems are available for use. We reached out to the authors of the cited papers, and were unsuccessful in obtaining any of their systems, with three exceptions: JPlag [Prechelt et al. 2002], Sherlock [Joy and Luck 1999], and Fett [Nichols et al. 2019]. We focus our evaluation on JPlag, since it was the second-most popular tool cited by our survey respondents. Only one respondent reported using Sherlock, and none reported using Fett.

5.7.1 JPlag. Like Moss, JPlag is a system for detecting similarity of software; both systems use tokenization, but instead of hashing, JPlag uses an approach based on string tiling to identify matches. From our survey results, we found that JPlag was the second most common software plagiarism tool, though it was a distant second to Moss.

To examine MOSSAD’s effectiveness on JPlag, we performed the same experiment as outlined in Section 5.2, replacing Moss with JPlag as the similarity detector. We refer to MOSSAD augmented with JPlag as MOSSAD-JPLAG.

Unlike Moss, which returns the top 250 highest scoring pairs of results, JPlag returns *all* possible pairwise combinations of input assignments. As a result, the MOSSAD-JPLAG experiments result in many more observable scores. We repeat all experiments with JPlag with the same threshold used in the previous experiments with Moss (25%).

The results from executing JPlag on all three projects are shown in Figure 8. The vast majority of the pairs (> 8000) scored 0%; we omit these results from these graphs for clarity. As before, the scores for pairs of assignments with at least one MOSSAD-JPLAG file are shown in dark gray, and the scores for authentic student assignment pairs are shown in light gray.

Table 2 summarizes our results. For all three projects, JPlag resulted in average similarity scores that were substantially higher than Moss. Just as with MOSSAD, MOSSAD-JPLAG was successful at producing results whose similarity scores are noticeably lower than the bulk of the distribution of legitimate programs, even with the 0% matches excluded.

MOSSAD-JPLAG produces variants that are below the 25% suspiciousness threshold across the suite, with one exception: a case where MOSSAD-JPLAG produced a variant that scored 38% with JPlag. We note that this value crosses a suspiciousness threshold that we derived for Moss. However, given that JPlag’s scores are much higher (the average similarity of the top 250 matches was 76.9%, this threshold is probably too low for JPlag. That is, we expect users of JPlag would stop inspecting at far higher thresholds. Even including the 0% matches, the average JPlag score assigned

Table 2. Summary of effectiveness results against three other plagiarism detectors: JPlag (with the complete JPlag output and only the top 250 pairs, for ease of comparison with Moss), Sherlock, and Fett.

Code type	Min	Max	Avg	95% CI
JPlag similarity scores				
<i>all matches</i>				
<i>MOSSAD-JPLAG</i>	0%	38%	1.1%	[0.9%, 1.3%]
<i>Legitimate Student Code</i>	0%	100%	28.2%	[27.7%, 28.7%]
<i>top 250 matches</i>				
<i>Legitimate Student Code</i>	30%	100%	76.9%	[74.7%, 78.3%]
Sherlock similarity scores				
<i>MOSSAD</i>	1%	61%	6.2%	[5.7%, 6.8%]
<i>Legitimate Student Code</i>	2%	100%	24.1%	[23.8%, 24.4%]
Fett similarity scores				
<i>MOSSAD</i>	0.6%	37.6%	5.9%	[4.7%, 7.4%]
<i>Legitimate Student Code</i>	2.3%	100%	45%	[38.6%, 50.4%]

to legitimate student code is 28%, whereas the average score assigned to MOSSAD-JPLAG scores is 1%.

For Project 1, the average JPlag scores hover around 30%, and the MOSSAD-JPLAG scores center around a mean of 12% (Figure 8a). One interesting result of using JPlag on this project is a higher percentage of 100% matches within the dataset without MOSSAD-JPLAG variants. In fact, 174 files across the 3 datasets received 100% similarity scores, and 536 scored above 97%. We randomly chose a subset (5) of the 100%-scoring pairs to spot check in order to assess the true positive rate. From these pairs, we found that approximately half exhibit high similarity in terms of algorithm and text, which we believe could be reasonably flagged as suspiciously similar. The other half of the files, although similar, had different algorithms and were not clearly plagiarized.

For Project 2, the JPlag results do not exhibit a clear drop off between the average student assignment similarity score and the outliers that we observed in the MOSSAD experiments using Moss (Figure 8b). Instead, there are three clusters of scores at 30%, 65%, and 90%. The mean MOSSAD-JPLAG score is 10%.

Finally, for Project 3, the JPlag scores follow a heavy-tailed distribution (Figure 8c). Here, the MOSSAD-JPLAG scores are not significantly lower than the mean score of the legitimate solutions, approximately 20%.

5.7.2 Sherlock. Sherlock [Joy and Luck 1999] is the third-most popular similarity detection system that is publicly available; it was cited by one survey respondent. It is a general-purpose academic plagiarism detector with support for both source code as well as natural language. For source code, it translates each source file into three *modes*: the original document, normalized (whitespace and comments are removed), and tokenized, where the code is parsed into tokens according to their basic purpose (similar to Moss’s normalization pass).

For each mode, Sherlock compares every pair of assignments and calculates a similarity score based on the total number of lines. However, Sherlock only displays to the user the sum of the total mode scores.

Using MOSSAD-generated variants for each of our three datasets, we compared the entire corpus augmented with the MOSSAD code using Sherlock. Since Sherlock reports a sum of the scores for each mode, we normalize its output simply by dividing by 3 (the system recommended leaving the mode settings untouched when using C).

MOSSAD is generally as effective against Sherlock as against the other plagiarism detectors we examine. While in one case, a MOSSAD assignment receives a score of 61%, the average for MOSSAD-generated code is 6.2%, well below the average score for legitimate student code (24%); the 95% confidence interval for MOSSAD code is [5.7%, 6.8%]. Table 2 provides a summary of these statistics.

5.7.3 Fett. Fett [Nichols et al. 2019] is a recent system aimed at addressing the drawbacks of Moss, JPlag, and other related plagiarism detection tools. Although none of our survey respondents reported using the tool, Fett is the only additional publicly-available plagiarism detection tool the authors could find. In fact, one of the major pillars of Fett is that it is open-source and easily accessible, along with being language-agnostic and highly accurate [Nichols et al. 2019].

Fett linearizes the resulting parse tree generated by an ANTLR parser via postorder traversal [Yan et al. 2018]. It then sorts the functions of the program by size, prunes the sequences, groups parse nodes into equivalence classes, and assigns weights to remaining nodes based on their class. Then, Fett uses the Smith-Waterman algorithm for sequence alignment scoring, penalizing gaps. Fett generates a matrix of similarity scores in CSV format.

Performing the same experiment described in Section 5.7.2, using the same C dataset and MOSSAD-generated variants, we found that MOSSAD is generally as effective against Fett as the other systems we evaluate.

The largest score assigned to a pair of files, where at least one file in the pair is a MOSSAD variant, is 37%. The average for MOSSAD-generated code is 5.9%, well below the average score for legitimate student code (45%); the 95% confidence interval for MOSSAD code is [4.7%, 7.4%]. These statistics are summarized in Table 2.

Although MOSSAD in its current form is able to completely undermine Fett for C datasets, MOSSAD would need to be altered slightly in order to have the same impact with Java code. For Java, Fett removes all expressions [Nichols et al. 2019], which eliminates the kind of code that MOSSAD inserts by default. A straightforward workaround is to replace the expressions that MOSSAD inserts with alternative pieces of code that would not be eliminated, including empty if-statements or calls to functions with no side effects.

RQ6 (Generality): *MOSSAD is effective against other plagiarism detectors (JPlag, Sherlock, and Fett), consistently achieving low similarity scores.*

5.8 Ablation Study

To isolate the effects of MOSSAD's algorithmic approaches, we conduct an ablation study. We constructed two variants of MOSSAD that we evaluate here. Note that, unlike MOSSAD, neither of these rely on Moss itself. Both of these approaches also directly perform hash disruption by targeting fingerprint windows, while MOSSAD does so implicitly.

- **Deterministic insertion of the same kind of statement.** MOSSAD_{DET} is an entirely deterministic variant that attempts to insert benign lines of code (fresh variable declarations) within each fingerprint window. Unlike MOSSAD, MOSSAD_{DET} always produces the same output from the same input program, running the risk of collision if assignments are plagiarized from the same input. This determinism also precludes mass plagiarism based on MOSSAD_{DET}.
- **Non-deterministic insertion of randomly-chosen statements.** MOSSAD_{NONDET} is a non-deterministic variant that randomly selects locations within each fingerprint window, and inserts randomly chosen lines from the program. This non-determinism reduces the likelihood that two programs will collide, though we expect that its effectiveness for mass plagiarism would be limited.

Fig. 9. Ablation Study: $\text{MOSSAD}_{\text{DET}}$ is also successful at defeating Moss’s detection. The dark gray bars denote a match between at least one $\text{MOSSAD}_{\text{DET}}$ file, and the light gray bars denote purely authentic code matches. As Section 5.8 describes, $\text{MOSSAD}_{\text{DET}}$ is strictly less powerful than MossAD and its deterministic nature makes mass plagiarism impossible.

Fig. 10. Ablation Study: $\text{MOSSAD}_{\text{NONDET}}$ ’s non-deterministic approach also defeats Moss. $\text{MOSSAD}_{\text{NONDET}}$ also occasionally leads to multiple variants producing high similarity scores with each other, but less frequently than its deterministic counterpart $\text{MOSSAD}_{\text{DET}}$.

Methodology. To examine the effects of these variants, we use five randomly selected files from each of our three datasets, and generate variants with both $\text{MOSSAD}_{\text{DET}}$ and $\text{MOSSAD}_{\text{NONDET}}$. We separately augment each dataset with the five corresponding plagiarized files and use Moss to produce similarity scores for each dataset.

Fingerprint Windows. To perform this ablation study, we need to know a single parameter: the length of the fingerprint windows. As Section 3.3 notes, this parameter is not publicly available, and needed to be reverse-engineered. In fact, the actual parameter we need is size of the fingerprint window in terms of the number of lines of code, which we refer to as w . However, different lines of code naturally lead to different numbers of tokens.

We performed extensive experimentation across our suite of student solutions and found that we could reliably perform hash disruption by inserting a variable declaration between every three to four lines of code. We therefore use $w = 4$ lines of code as the window size for these experiments.

5.8.1 $\text{MOSSAD}_{\text{DET}}$. We examine the effectiveness of $\text{MOSSAD}_{\text{DET}}$ at producing code variants that achieve low similarity scores when compared with the original base file by using our system to perform its transformations on five randomly selected files from each data set. After generating these variants, we use Moss to score the similarity of each entire dataset. Figure 9 presents these results. The dark gray bars denote comparisons that were made with at least one $\text{MOSSAD}_{\text{DET}}$ variant; most scored below the threshold of 25%.

Fig. 11. Mass plagiarism with $\text{MOSSAD}_{\text{NONDET}}$ frequently yields suspicious Moss scores. Unlike MossAD, $\text{MOSSAD}_{\text{NONDET}}$ produces variants with suspicious scores in some cases roughly 80% of the time.

The deterministic nature of $\text{MOSSAD}_{\text{DET}}$ introduces risk associated with high Moss scores of multiple $\text{MOSSAD}_{\text{DET}}$ variants. This risk is visible in Figure 9b: the dark gray bars that scored higher than 25% are actually comparisons made between two $\text{MOSSAD}_{\text{DET}}$ files. Since the additions that $\text{MOSSAD}_{\text{DET}}$ makes to an input file are the same, it can cause Moss to detect additional, synthetic similarity. Since multiple $\text{MOSSAD}_{\text{DET}}$ variants are all generated from the same base file, two attempts to plagiarize from the same source code will result in Moss scores of 100%.

5.8.2 $\text{MOSSAD}_{\text{NONDET}}$. The $\text{MOSSAD}_{\text{NONDET}}$ variant works by iterating through all non-overlapping windows of size w in consecutive order, trying to insert a randomly selected line from the source program into a randomly selected position within that window. If the insertion produces a compilation error or the object code does not match the original, the system will repeat this process $w \log(w)$ times. If all attempts fail, $\text{MOSSAD}_{\text{NONDET}}$ moves on to the next window.

By using the source file itself as the source of code to insert, $\text{MOSSAD}_{\text{NONDET}}$ adds greater non-determinism that lowers the probability that the same statement was added to the same window. When combined with the nondeterminism introduced by randomly choosing an insertion location within each window, we expect $\text{MOSSAD}_{\text{NONDET}}$ to lower the number of matches of hashes across multiple variants generated from the same base file. Figure 10 presents the results.

$\text{MOSSAD}_{\text{NONDET}}$ consistently scores below the threshold similarity score of 25% in each project; however, for Projects 1 and 2 the scores center around 25%, with each score between 6% and 30%. Since $\text{MOSSAD}_{\text{NONDET}}$ is not fully nondeterministic, the resulting Moss scores of its outputs are shown to occasionally under-perform MossAD; this is because it has limited areas in which to insert mutations, resulting in decreased chances of generating a mutation that is successful in compilation and in object code comparison.

Because $\text{MOSSAD}_{\text{NONDET}}$ is non-deterministic, it should in principle be applicable to mass plagiarism. We perform the same experiment described in Section 5.5. As Figure 11 shows, it can produce variants with suspicious scores ($> 25\%$) as much as 80% of the time. That said, 90% of the pairs of files score a similarity score below 40%. While this similarity is generally low, it significantly underperforms MossAD in this scenario (maximum similarity: 27%).

5.8.3 Discussion. The major difference between MossAD and the ablated versions is the level of determinism. MossAD is fully nondeterministic; that is, it is free to place its inserted lines anywhere, giving it a higher likelihood of disrupting hash windows. The two ablated versions can at most insert one line of code every so many lines of code. Since lines of code may not actually align with

windows (which are actually in units of tokens, not lines of code), the ablated versions may fail to disrupt hashes that MOSSAD can disrupt.

Ablation Study: *While MOSSAD_{DET} and MOSSAD_{NONDET} are generally effective at undermining Moss, MOSSAD dominates both of them, especially for mass plagiarism.*

5.9 Threats to Validity

In this section, we examine the threats to validity of the evaluation of MOSSAD.

5.9.1 MOSSAD File Sizes. Due to the nature of code insertion, MOSSAD files are always longer than the base file, though the current implementation limits the number of consecutive insertions to ensure that files do not grow too large. For the datasets used throughout the experiments described earlier in this section, the number of lines of code ranges from 10 to 100 (which in and of itself presents another threat to validity); longer code produced by MOSSAD attacks would not necessarily be an easily-identifiable pattern since this (legitimate) variance is already present in the dataset. However, the length of MOSSAD files may be a characteristic that could identify suspiciousness for other datasets in which we have not examined.

5.9.2 Populations. The two surveys used to inform the evaluation have very different results; as previously noted, the surveys themselves were different. The Twitter-only survey was brief and could be answered with a single click, while the second survey required data entry into a Google Form. The latter may have been a barrier to participation. Although the two populations disclosed varying degrees of manual inspection, the user study presented in RQ2 5.3 suggests that even when code is manually inspected, MOSSAD-generated code escapes detection.

For the user study presented in RQ2 5.3, the population is taken from a single institution, an R1 public university, and may not be representative. While it is a reasonable size ($N=30$), a larger sample size may result in different conclusions. Additionally, changes to the study, such as providing a different rubric or grading metric or removing the time limit and letting graders choose the amount of time to spend on manual inspection, could lead to different results.

5.9.3 Moss Latency. For the performance evaluation of MOSSAD in RQ5 5.6, the performance of MOSSAD is heavily dependent on the latency and thresholds of the deployed Moss system, and is subject to variance due to dynamic load. The point of this experiment was to examine MOSSAD performance in the real world. These experiments could inadvertently have used Moss during times of low load; higher load could increase the average time taken for MOSSAD and consequently could have had an impact on the real world results.

6 OBFUSCATION

This section explores the possibility of an alternative approach to MOSSAD: code obfuscation. Intuitively, this approach may seem appealing. By definition, the goal of an obfuscator is to transform a program so that it is not as legible as the original program. This effect appears to hold promise as a means of thwarting software plagiarism detectors. However, we find this is not the case.

We tested a number of publicly-available obfuscators. We found that many of the obfuscations are ones that Moss trivially defeats, such as whitespace removal. We discuss here the obfuscator that was most successful at defeating Moss's detection; however, as we describe below, we find that it suffers from several significant drawbacks that make it ineffective as an attack on software plagiarism detection.

(a) Obfuscation can defeat Moss's similarity detection algorithms, though not reliably.

```
#include <stdio.h>
int o_4e62450cf2d977e80969369f248d3c4a(int o_f934f4cf272bdacbc5c98186815088
o_b549e9d021435e399fbb534b3e11507c, int o_674f832bb3c1f80a6271b9e4bf4d176
); int main() { int o_5be56d8559c61299842126a04c5f2201
o_1c35b3d13f20f38ec54dd6a1540693a7, o_f3e0c8e85198544ae3341829f51ec99f
o_777ae22b54bca1c620015ef082e8865, o_510ef4b5d4c78d394465c5625c814f33
o_7d9e6ab545109f750f135dab568400a, o_19daf6597629f5fc80157f2d25c9132d;
("x25""d\045d\x25""d\045d\x25""d\045d\x25""d
, o_5be56d8559c61299842126a04c5f2201, o_1c35b3d13f20f38ec54dd6a1540693a
, o_f3e0c8e85198544ae3341829f51ec99f, o_777ae22b54bca1c620015ef082e886
, o_510ef4b5d4c78d394465c5625c814f33, o_7d9e6ab545109f750f135dab568400
, o_19daf6597629f5fc80157f2d25c9132d); printf
("x50""3\012\x64""x0A\045d\x0A""2\0655\x0A""
, o_510ef4b5d4c78d394465c5625c814f33, o_7d9e6ab545109f750f135dab568400a)
```

(b) Obfuscation produces highly unreadable code.

Fig. 12. Obfuscators can defeat Moss's similarity detection algorithms, but at a cost. Picheta is the most effective of the obfuscators we evaluated at producing low similarity scores. However, it is not consistently reliable, frequently producing highly suspicious matches. It suffers from two other major drawbacks: (1) it is deterministic, so code plagiarized from the same source will yield a high match, and (2) it produces obviously-obfuscated code.

The obfuscator we evaluate here is an online obfuscator for C/C++, whose stated intended use is to protect against software piracy, reverse engineering, and tampering [Picheta 2020a,b]. This obfuscator has no name; we refer to it here as *Picheta* after its author. Picheta performs three major transformations: *data*, where data is transformed into different radices or turned into a static expression; *lexical*, where all identifiers are hashed; and *control*, where the semantics of the program are slightly modified, including dead code insertion. Of the three obfuscators we found that were at least moderately successful in undermining detection by Moss (Stunnix [stunnix.com 2020], Tigress [Collberg 2020], and Picheta), we found that Picheta was the most successful in producing low Moss similarity scores.

6.1 Evaluation

Using Picheta to plagiarize source files from our datasets, we examined its effectiveness as a means to avoid plagiarism detection. For this experiment, we randomly selected 10 files from the first dataset as inputs to Picheta. The resulting ten obfuscated files were added back to the original set of assignments to act as plagiarized files; the resulting corpus was compared using Moss.

Figure 12a graphs the Moss scores of the top 250 highest-scoring pairs of files. We also performed this experiment using the two additional data sets; the results were similar, so we omit those graphs. We find that Picheta is moderately effective, but roughly 25% of the time, it produces highly suspicious outputs, with resulting Moss similarity scores substantially higher than our suspiciousness threshold (40%–98%).

Like all other obfuscators we examined, Picheta is also deterministic. Thus, plagiarism from the same base file using Picheta would result in the exact same files, leading to likely discovery. Interestingly, as shown in Figure 12a (denoted by *Double Picheta*) many of the matches with high similarity scores are actually artifacts of using the obfuscator itself, further limiting its effectiveness, as obfuscated code is self-similar even when the input programs differ. Beyond these drawbacks, the code produced by the obfuscator is immediately suspicious, as shown in Figure 12b.

7 COUNTERMEASURES

This section describes a limited countermeasure to the MOSSAD line of attacks. The current prototype implementation of MOSSAD relies on introducing semantics-preserving code in an extremely

conservative manner: directly comparing generated object files. The code that MOSSAD introduces is unreachable, redundant, or dead.

We therefore evaluated the effectiveness of the following countermeasure: first, compile all code with a high optimization level (e.g., `-O3`) and output *assembly code* (via the `-S` flag). Second, submit the assembly code to Moss, indicating it should use one of the supported ISAs as its input language (Moss currently supports both x86 and MIPS). We found that for each project, this countermeasure undoes the effect of MOSSAD by optimizing away its code insertions.

Unfortunately, this countermeasure suffers from several limitations. First, this approach requires that instructors manually map back matched assembly code fragments to the original source files, which is inconvenient at best. Second, and more seriously, generated assembly code—especially at high optimization levels—results in spurious matches due to compiler-generated artifacts like function prologues and epilogues. Third, it only works when optimizing compilers are available, and where it is possible to emit “clean” assembly code (e.g., without inlined calls to parts of a runtime library). It is not obvious if this is even possible currently with existing JIT compilers for languages like Python, Ruby, or JavaScript.

The results of this paper also suggest that more extensive code review, although costly in terms of human effort, could be an effective countermeasure to the MOSSAD line of attacks; empirical evaluation of this countermeasure is future work. Additionally, we hypothesize that integrating version control into the process of code assignment submission and subsequent grading could also mitigate the effects of MOSSAD; exploring this is also future work. We believe that both of these potential countermeasures have additional positive benefits beyond plagiarism detection: code review and version control are both useful skills for computer scientists and programmers, and are a way to give meaningful feedback to students.

8 RELATED WORK

This section describes other software plagiarism approaches not discussed earlier in the paper. We note that, in an effort to further evaluate MOSSAD’s effectiveness beyond Moss, JPlag, Sherlock, and Fett, we contacted the authors of these papers; all authors responded that their tools are not available, precluding a direct empirical comparison.

Static techniques: To combat students shuffling independent code segments to thwart plagiarism detection, Wise presents a string-similarity approach using Running-Karp-Rabin Greedy-String-Tiling, which can detect transposed subsequences in many source languages [Wise 1996]; this is the technique employed by Sherlock. Liu et al. and Chae et al. present graph-based approaches to detecting software plagiarism [Chae et al. 2013; Liu et al. 2006]. In the first case, plagiarism is detected by mining program dependence graphs. This method of plagiarism detection is effective because program dependence graphs are generally invariant in the face of plagiarism [Liu et al. 2006]. However, extending MOSSAD to adding unnecessary control flow to a plagiarized file could thwart this detection technique. Chae et al. also construct a control flow graph of a program; they then perform a random walk to compute the importance of each node, and finally generate a single score vector of the graph for comparison [Chae et al. 2013]. Son et al. uses the structural information of the program (obtained via parsing) to determine similarity [Son et al. 2006]. Luo et al. propose a binary-oriented program similarity method based on longest common subsequences of semantically equivalent basic blocks [Luo et al. 2014].

Dynamic techniques: Jhi et al. introduce value-based program characterization leveraging invariant values, in an effort to create a program plagiarism detection algorithm that is resilient to control flow and data obfuscation techniques [Jhi et al. 2011].

Code Clone Detection: Another potential method for identifying possible plagiarism attempts is via code clone detection. We evaluated SourcererCC [Sajjani et al. 2015], a recent code clone detector used to map code duplicates on Github [Lopes et al. 2017], in an effort to defeat MOSSAD. SourcererCC works at two granularities: file-level and block-level. We examined the output of running SourcererCC at the file-level granularity on a corpus of one of our C datasets and various MOSSAD variants. SourcererCC returned numerous near-miss clones, which the authors describe as clones “where minor to significant editing activities might have taken place in the copy/pasted fragments” [Sajjani et al. 2015]. However, SourcererCC returned nearly every file from the corpus as near-misses (these are almost all false positives). We do not believe that file-level operates at a sufficiently fine granularity for detecting code plagiarism that has been disguised using MOSSAD. SourcererCC’s block-level granularity only supports Python and Java; we were unable to test if this granularity is sufficiently fine to detect disguised plagiarism using MOSSAD with our projects, which are written in C.

9 CONCLUSION

This paper presents MOSSAD, a fully-automated program transformation system that defeats software plagiarism detection. We demonstrate its efficacy on Moss, the most effective and most widely-used such detector, as well as on JPlag. MOSSAD transforms source programs into one or even dozens of variants that all escape their detection algorithms while maintaining readability. Because MOSSAD directly strikes at the algorithmic underpinnings of Moss and similar systems, effectively coping with these will require innovation in software plagiarism detection.

ACKNOWLEDGMENTS

We thank Don Porter for providing our anonymized C code datasets. We also thank Neena Thota for providing time and space in her classroom for our case study, and her students for volunteering to be a part of the case study. Finally, we thank Alex Aiken for his thoughtful comments and for providing us with a research installation of Moss. This material is based upon work supported by the National Science Foundation under Grants No. CCF-1439008 and CCF-1617892.

REFERENCES

- Alex Aiken. Personal communication, Apr 2018.
- Alex Aiken. Plagiarism Detection. <http://theory.stanford.edu/~aiken/moss/>, April 2020.
- Ethan Baron. Stanford, UC Berkeley computer science cheating incidents part of widespread problem: report. *SiliconBeat*, May 2017. URL <http://www.siliconbeat.com/2017/05/30/stanford-uc-berkeley-computer-science-cheating-incidents-part-of-widespread-problem-report/>.
- Jess Bidgood and Jeremy B. Merrill. As computer coding classes swell, so does cheating. *The New York Times*, May 2017. URL <https://www.nytimes.com/2017/05/29/us/computer-science-cheating.html>.
- K. Bowyer and L. Hall. Experience using “MOSS” to detect cheating on programming assignments. *FIE’99 Frontiers in Education. 29th Annual Frontiers in Education Conference. Designing the Future of Science and Engineering Education. Conference Proceedings (IEEE Cat. No.99CH37011)*, 3:13B3/18–13B3/22 vol.3, 1999.
- Tracy Camp, W. Richards Adrion, Betsy Bizot, Susan Davidson, Mary Hall, Susanne Hambrusch, Ellen Walker, and Stuart Zweben. Generation CS: The growth of computer science. *ACM Inroads*, 8(2):44–50, May 2017. ISSN 2153-2184. doi: 10.1145/3084362. URL <http://doi.acm.org/10.1145/3084362>.
- Dong-Kyu Chae, Jiwoon Ha, Sang-Wook Kim, BooJoong Kang, and Eul Gyu Im. Software plagiarism detection: A graph-based approach. In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management, CIKM ’13*, pages 1577–1580, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-2263-8. URL <http://doi.acm.org/10.1145/2505515.2507848>.
- X. Chen, B. Francia, M. Li, B. Mckinnon, and A. Seker. Shared information and program plagiarism detection. *IEEE Transactions on Information Theory*, 50(7):1545–1551, Jul 2004. doi: 10.1109/tit.2004.830793.
- Joel Coehoorn. How do I ask and answer homework questions? <https://meta.stackoverflow.com/questions/334822/how-do-i-ask-and-answer-homework-questions>, 2019.

- Christian Collberg. the tigress c obfuscator. <https://tigress.wtf/>, 2020.
- EAB. Why cheating is on the rise in computer science classes. *EAB Daily Briefing*, Aug 2017. URL <https://www.eab.com/daily-briefing/2017/08/01/3-reasons-cheating-is-on-the-rise-in-comp-sci>.
- Forbes. Arjun Singh. <http://www.forbes.com/profile/arjun-singh/#15980dacc9c0>, 2018.
- Gradescope. Gradescope at SIGCSE 2019. <https://www.gradescope.com/sigcse2019>, Mar 2019.
- Yoon-Chan Jhi, Xinran Wang, Xiaoqi Jia, Sencun Zhu, Peng Liu, and Dinghao Wu. Value-based program characterization and its application to software plagiarism detection. In *Proceedings of the 33rd International Conference on Software Engineering*, ICSE '11, pages 756–765, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0445-0. URL <http://doi.acm.org/10.1145/1985793.1985899>.
- Mike Joy. Personal communication, Jul 2020.
- Mike Joy and Michael Luck. Plagiarism in programming assignments. *IEEE Transactions on education*, 42(2):129–133, 1999.
- Kylie Jue. Stanford CS department battles honor code violations. *The Stanford Daily*, Jun 2014. URL <https://www.stanforddaily.com/2014/06/03/stanford-cs-department-battles-honor-code-violations/>.
- Sudipta Kundu, Zachary Tatlock, and Sorin Lerner. Proving optimizations correct using parameterized program equivalence. In *Proceedings of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '09, pages 327–337, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-392-1. doi: 10.1145/1542476.1542513. URL <http://doi.acm.org/10.1145/1542476.1542513>.
- Claire Le Goues, ThanhVu Nguyen, Stephanie Forrest, and Westley Weimer. GenProg: A generic method for automatic software repair. *IEEE Transactions on Software Engineering*, 38(1):54–72, 2011.
- Chao Liu, Chen Chen, Jiawei Han, and Philip S. Yu. Gplag: Detection of software plagiarism by program dependence graph analysis. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '06, pages 872–881, New York, NY, USA, 2006. ACM. ISBN 1-59593-339-5. doi: 10.1145/1150402.1150522. URL <http://doi.acm.org/10.1145/1150402.1150522>.
- Cristina V. Lopes, Petr Maj, Pedro Martins, Vaibhav Saini, Di Yang, Jakub Zitny, Hitesh Sajjani, and Jan Vitek. Déjàvu: A map of code duplicates on GitHub. *Proc. ACM Program. Lang.*, 1(OOPSLA):84:1–84:28, October 2017. ISSN 2475-1421. doi: 10.1145/3133908. URL <http://doi.acm.org/10.1145/3133908>.
- Lannan Luo, Jiang Ming, Dinghao Wu, Peng Liu, and Sencun Zhu. Semantics-based obfuscation-resilient binary code similarity comparison with applications to software plagiarism detection. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*, FSE 2014, pages 389–400, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-3056-5. URL <http://doi.acm.org/10.1145/2635868.2635900>.
- Robert McMillan. University bans GitHub homework (then changes its mind). *Wired*, Feb 2015. URL <https://www.wired.com/2015/02/university-bans-github-homework-changes-mind/>.
- William Hugh Murray. Ubiquity: Cheating in computer science. *Ubiquity- An ACM Publication*, Oct 2010. URL <https://ubiquity.acm.org/article.cfm?id=1865908>.
- National Academies of Sciences, Engineering, and Medicine. *Assessing and responding to the growth of computer science undergraduate enrollments*. National Academies Press, 2018. doi: 10.17226/24926.
- Lawton Nichols, Kyle Dewey, Mehmet Emre, Sitao Chen, and Ben Hardekopf. Syntax-based improvements to plagiarism detectors and their evaluations. In *Proceedings of the 2019 ACM Conference on Innovation and Technology in Computer Science Education*, ITiCSE '19, page 555–561, New York, NY, USA, 2019. Association for Computing Machinery. ISBN 9781450368957. doi: 10.1145/3304221.3319789. URL <https://doi.org/10.1145/3304221.3319789>.
- Dominik Picheta. Code Obfuscation for the C/C++ Language. <https://arxiv.org/abs/2003.03449>, March 2020a.
- Dominik Picheta. Obfuscator. <https://picheta.me/obfuscator>, 2020b.
- Lutz Prechelt, Guido Malpohl, and Michael Philippsen. Finding plagiarisms among a set of programs with JPlag. *Journal of Universal Computer Science*, 8(11):1016–1038, 2002.
- B. K. Rosen. Program equivalence and context-free grammars. In *13th Annual Symposium on Switching and Automata Theory*, pages 7–18, October 1972. doi: 10.1109/SWAT.1972.18. URL doi.ieeecomputersociety.org/10.1109/SWAT.1972.18.
- Abhik Roychoudhury, K. Narayan Kumar, C. R. Ramakrishnan, I. V. Ramakrishnan, and Scott A. Smolka. Verification of parameterized systems using logic program transformations. In Susanne Graf and Michael I. Schwartzbach, editors, *Tools and Algorithms for Construction and Analysis of Systems, 6th International Conference, TACAS 2000, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS 2000, Berlin, Germany, March 25 - April 2, 2000, Proceedings*, volume 1785 of *Lecture Notes in Computer Science*, pages 172–187. Springer, 2000. ISBN 3-540-67282-6. URL https://doi.org/10.1007/3-540-46419-0_13.
- Hitesh Sajjani, Vaibhav Saini, Jeffrey Svajlenko, Chanchal K. Roy, and Cristina V. Lopes. SourcererCC: Scaling Code Clone Detection to Big Code. <https://arxiv.org/abs/1512.06448>, Dec 2015.
- 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*, pages 76–85. ACM, 2003.

- Arjun Singh, Sergey Karayev, Kevin Gutowski, and Pieter Abbeel. Gradescope: A fast, flexible, and fair system for scalable assessment of handwritten work. In *Proceedings of the Fourth (2017) ACM Conference on Learning Scale*, pages 81–88, April 2017. ISBN 978-1-4503-4450-0. doi: 10.1145/3051457.3051466.
- Temple F Smith, Michael S Waterman, et al. Identification of common molecular subsequences. *Journal of molecular biology*, 147(1):195–197, 1981.
- Jeong-Woo Son, Seong-Bae Park, and Se-Young Park. Program plagiarism detection using parse tree kernels. In *Lecture Notes in Computer Science*, pages 1000–1004. Springer Berlin Heidelberg, 2006. URL https://doi.org/10.1007/978-3-540-36668-3_122.
- stunnix.com. C/C++ Obfuscator. <http://stunnix.com/prod/cxxo/>, 2020.
- Michael J. Wise. Yap3: Improved detection of similarities in computer program and other texts. In *Proceedings of the Twenty-seventh SIGCSE Technical Symposium on Computer Science Education*, SIGCSE '96, pages 130–134, New York, NY, USA, 1996. ACM. ISBN 0-89791-757-X. URL <http://doi.acm.org/10.1145/236452.236525>.
- Lisa Yan, Nick McKeown, Mehran Sahami, and Chris Piech. TMOSS: Using intermediate assignment work to understand excessive collaboration in large classes. In *Proceedings of the 49th ACM Technical Symposium on Computer Science Education*, SIGCSE '18, pages 110–115, New York, NY, USA, 2018. Association for Computing Machinery. ISBN 9781450351034. URL <https://doi.org/10.1145/3159450.3159490>.