

PasteTrace: A Single Source Plagiarism Detection Tool For Introductory Programming Courses

Jesse McDonald
Information and Computer Sciences
University of Hawai'i at Mānoa
Honolulu, Hawai'i
Email: jamcd@hawaii.edu

Scott Robertson
Information and Computer Sciences
University of Hawai'i at Mānoa
Honolulu, Hawai'i
Email: scottpr@hawaii.edu

Anthony Peruma
Information and Computer Sciences
University of Hawai'i at Mānoa
Honolulu, Hawai'i
Email: peruma@hawaii.edu

Abstract—Introductory Computer Science classes are important for laying the foundation for advanced programming courses. However, students without prior programming experience may find these courses challenging, leading to difficulties in understanding concepts and engaging in academic dishonesty such as plagiarism. While there exists plagiarism detection techniques and tools, not all of them are suitable for academic settings, especially in introductory programming courses. This paper introduces PasteTrace, a novel open-source plagiarism detection tool designed specifically for introductory programming courses. Unlike traditional methods, PasteTrace operates within an Integrated Development Environment that tracks the student's coding activities in real-time for evidence of plagiarism. Our evaluation of PasteTrace in two introductory programming courses demonstrates the tool's ability to provide insights into student behavior and detect various forms of plagiarism, outperforming an existing well-established tool.

A video demonstration of PasteTrace and its source code, and case study data are made available at <https://doi.org/10.6084/m9.figshare.27115852>

Index Terms—education, plagiarism, introductory programming, coding, IDE, plugin, students, human-subject research

I. INTRODUCTION

Introductory Computer Science classes serve as the initial exposure to programming for many students, playing a vital role in building the foundation for subsequent higher-level courses [1]. It is through these courses students learn about the fundamentals of programming, such as variables, data types, control structures, and functions, which they build upon in their subsequent higher-level courses [2], [3]. Moreover, the importance of these courses extends beyond students learning about programming language semantics and syntax. Through these courses, students also develop and practice problem-solving, creativity, and logical reasoning skills, which they can apply to their programming assignments [4].

However, these courses can be challenging for many students, especially those with no prior programming experience. These challenges range from syntactic difficulties to conceptual misunderstandings as well as deficits in problem-solving, code organization, and debugging [5]. Consequently, students may unintentionally or intentionally engage in academic dishonesty, such as plagiarism [6].

Plagiarism is particularly troublesome in introductory classes. Since an introductory class is often a student's first

experience of programming, if a student passes an introductory class by plagiarizing, then it is likely that the student does not understand part or all of the course. If students are not corrected for plagiarizing in their introductory programming classes, they may develop a false sense of their abilities. This lack of understanding due to plagiarism could leave them unprepared for more advanced courses and future career challenges.

A. Goal

In the field of computing education, several techniques and tools have been developed to identify different types of plagiarism [7], [8]. While these tools are valuable, not all of them are suitable for academic settings, especially in the context of introductory programming courses. Some of these tools are proprietary and closed-source, making it challenging for institutions to assess their accuracy and reliability. Additionally, certain plagiarism detection tools require the transfer of files to external services, raising concerns about potential violations of student privacy laws.

The goal of our research is to provide an effective and reliable open-source code plagiarism detection tool specifically tailored for academic environments, particularly in introductory programming courses. To this extent, we present *PasteTrace*, a single source plagiarism detection tool For introductory programming courses.

II. BACKGROUND

Plagiarism is generally defined as “The practice of taking someone else’s work or ideas and passing them off as one’s own.” [9] However, for this paper, we define plagiarism as “*copying someone else’s code to complete an assignment.*” We exclude copying ideas for several reasons, mainly because it is difficult to determine if an idea has been copied or just independently rediscovered; the general culture of computer science is to iterate on others ideas (often without attribution [10]); and because at the introductory level, all the required ideas have been thought of by so many people that determining an origin would be next to impossible. Additionally, we exclude the possibility a student could “cite” their copied code to avoid plagiarism because the introductory level involves teaching coding, not copying existing code.

86 A. Existing Techniques for Detecting Code Plagiarism

87 The general method of plagiarism can be summarized as
88 follows: Given a set A of submissions, find a set P such that
89 $\forall p \in A$, if $\exists q \in A$ such that $q \neq p$ and $D(p, q)$, then $p \in P$
90 and $q \in P$ for some plagiarism detection function D . However,
91 defining D effectively can be challenging [11].

92 Researchers at Stanford University developed a program
93 called MOSS (Measure of Software Similarity), which
94 is capable of giving each assignment pair a similarity
95 score [12]. This approach is robust against name changes,
96 code reorganization, and white space changes by analyzing
97 a semi-random subset of k-gram hashes based on each
98 assignment where k has been manually tuned for each
99 language.

100 Others have attempted to define D by ignoring syntax and
101 exclusively using code semantics [13] so that two pieces
102 of code match when they functionally do the same thing.
103 There are also many other efforts to define D using a
104 machine learning model [14]–[17]. Additionally, there are
105 also detection models that attempt to exclusively detect single
106 LLMs output [18]. However, these are not particularly reliable
107 at present[19] and have a high false positive rate[20]. AlSallal
108 et al. [21] used machine learning to detect stylistic variation
109 in essays. However, in our scenario, this approach would not
110 be ideal for introductory-level students who may not yet have
111 a defined “style.”

112 III. PASTETRACE

113 A. Single Source Plagiarism Detection

114 PasteTrace is an Integrated Development Environment
115 (IDE) based on the Processing IDE. Processing is an
116 open-source IDE for Java that is designed to be easy to learn
117 even for non-programmers [22]; as such, it is ideal for our
118 application.

119 When the modified IDE first launches, a persistent UUID
120 file is created on the machine that installed it. This UUID acts
121 as a machine or user ID, is referred to as the `InstallID`, and
122 is assumed to be constant throughout the class. Additionally,
123 whenever a project is created with the IDE, a second UUID
124 is created that is unique to that project. We refer to this as
125 the `ProjectID`. Both UUIDs are saved in a special hidden
126 metadata comment in the program source file. We will refer
127 to this as the `metaComment`.

128 Whenever a file is opened, the IDE checks the `InstallID`
129 of the file and machine. If they differ, the IDE notes the new
130 `InstallID` in the `metaComment` in an ordered list called
131 the `InfectionStack`. A similar stack is also used to track
132 paste events. As these stacks are effectively a single stack
133 with a way to discriminate between event types, we will use
134 `InfectionStack` to refer to both.

135 Additionally, when any part of the code is copied,
136 zero-width spaces (U+200B) are interspersed with the normal
137 letters to encode as much data as possible. The exact amount
138 of data varies based on the size of the copy, but it can
139 include several copies of the `InstallID`, `ProjectID`, and

`InfectionStack`. When pasted into a project, the encoded
140 data is compared to that of the project receiving the code, and
141 any mismatch UUIDs are added to the `InfectionStack`.
142 This encoding survives being sent over most messaging
143 software such as Discord, and is generally resilient to partial
144 pastes. Any paste without encoded data is logged as originating
145 outside the IDE. We should note that our testing revealed
146 that all email clients strip out the zero-width spaces, leading
147 to emailed code copies being identified as originating from
148 outside the IDE.
149

150 In addition to copy tracking, the IDE also logs the time,
151 location, and content of all edit events, including copy, paste,
152 cut, delete, and typing. Any mismatched UUIDs are included
153 in the log for the respective paste event. This data is also
154 included in the `metaComment`. Using this keylog, it is
155 possible to reconstruct the entire coding process and trace all
156 code pastes. Additionally, if the IDE loads a file without a
157 `metaComment`, it notes this as the first edit and logs the
158 initial state of the file.

159 Further, the UUIDs can be used to trace code authorship
160 through multiple student submissions. Additionally, analyzing
161 an entire class at the same time can be useful to identify false
162 positives, such as copies from previous assignments on the
163 student’s machine.

164 Even simple files such as Hello World would be detected, as
165 the detection does not rely on code comparison. Our method
166 is immune to traditional false positives as organically created
167 code is never considered plagiarized regardless of similarity
168 to another student.

169 B. Plagiarism Detection Scenarios

170 1) *Peer-to-peer (P2P) file sharing.*: In this situation, one
171 student writes the code and shares it with a peer, who then
172 submits the copy. With PasteTrace, there is no easy way to
173 submit someone else’s unmodified file without the UUIDs
174 being a clear red flag, and any attempt to edit the file will
175 mark it as infected. Additionally, introductory classes often
176 require students to add an “Author” comment to the top of
177 their files. This would need to be edited by anyone attempting
178 to cheat, and the original comment is preserved in the log.
179 Finally, any code shared via online messaging will be flagged
180 and, in general, traced. If a student copies from a previous
181 assignment, it will be traceable to which assignment and who
182 authored it.

183 2) *Collaboration.*: In this situation, two students work
184 together on an assignment that is supposed to be completed
185 individually. Collaborating students will have similar code
186 with overlapping timestamps at each edit. Additionally,
187 collaborating students tend to send code back and forth,
188 meaning both side’s `InfectionStack` will be populated
189 with each other’s UUIDs. This back and forth should make it
190 easy to distinguish from P2P.

191 3) *Theft.*: This is any form of P2P plagiarism where the
192 author does not know that it is happening. For example, a
193 student (the author) works on a shared computer and leaves
194 a file saved locally, and the plagiarizer finds the file when

195 attempting the same assignment. While it can be challenging
196 to discriminate it from P2P, theoretically, it is possible to
197 isolate the original file using the knowledge of the computer
198 setup (e.g., a computer lab with known `InstallIDs`), the
199 `InfectionStack`, and the timestamp log.

200 4) *Search.*: ‘This is when a student uses an internet search
201 engine to find an existing solution to an assignment. No matter
202 how a file is found, once it is loaded by the IDE, it will
203 be marked as an external file in the `metaComment`. If it
204 is submitted without being opened by the IDE, it will lack a
205 `metaComment`. If the source code is copied into an existing
206 project, it will be marked as external. The only viable source
207 that this does not detect is physically typing out the code
208 from a reference. However, this is theoretically detected by
209 analyzing the activity.

210 5) *Expert.*: ‘Similar to the Search scenario, the Expert
211 scenario involves obtaining a solution from paid sources (e.g.,
212 Chegg or Coursehero), or the solution is tailored by an expert
213 to the exact assignment given. Expert sources are detected in
214 the same way as Search.

215 C. Automated Detection

216 To facilitate a batch evaluation of student submissions,
217 we implemented a script that takes all student submissions
218 for a given assignment and analyzes them. The script is
219 designed to analyze a directory of student submissions by
220 building a graph using the `InfectionStack` to trace
221 shared files and identify shared machines. It categorizes paste
222 events and detects plagiarism based on various criteria, such
223 as the origin of pastes and the number of lines pasted.
224 Purely internal paste events, or ones that are too short to
225 track, are ignored. From there, pastes fall into 4 categories:
226 pastes originating from a project that was turned in by
227 another student, or pastes originating from a different students’
228 machine; pastes originating from the same machine, but a
229 different project; pastes that originate from an `InstallID`
230 that is not associated with any project turned in; and pastes of
231 external origin. In each case where plagiarism is detected, the
232 code pasted is also logged for human validation.

233 Pastes from another turned-in `ProjectID` or `InstallID`
234 that is associated with another student are immediately marked
235 as plagiarism. Pastes that originate on the same machine are
236 treated differently depending on the machine. If the machine
237 was not shared, the paste is assumed to originate from the
238 same student. However, there is the possibility they are trying
239 to use a loophole where the assignment is completed in another
240 project and then pasted to the final project to clean up the
241 keystroke log. To detect this, we decided to use a 50-line
242 threshold. This allows functions to be copied from previous
243 assignments but still detects entire projects being copied. If
244 the machine was shared, the paste is flagged as plagiarism.

245 Pastes that originate from a different machine that is not
246 associated with any student and a project not associated with
247 any student are assumed to be from the same student, similar
248 to pastes from the same machine; however, as this is less likely

to be the student and may be an expert source, we have set
the threshold to 20 lines.

Pastes that did not originate in the IDE are treated as
copying from online resources. While a certain amount of
online copying is permissible (like copying the name of a
function from a docs page), copying whole functions would
not be. We set the threshold for these events at three lines,
reasoning that 1 line is likely a line from a docs page, and
three lines are likely accidentally copying a new line before
and after the function within the docs page, but more than
three lines could be an entire function.

Additionally, we track the number of edits made after the
last plagiaristic paste event.

IV. PASTETRACE EVALUATION

To assess the effectiveness of PasteTrace, we conducted
a human-subject based study with students enrolled in
Introduction to Programming I and Introduction to
Programming II courses at the University of Hawai’i at
Mānoa. We collaborated with the instructors to provide each
class with a project that was challenging enough to encourage
plagiarism, but not so difficult that no students could complete
it on their own. One class was tasked with creating three
polymorphic-shaped objects that move around the screen and
bounce off the edges. The other class was given the task
of creating a function-defined fractal, such as a Mandelbrot
fractal. Students were instructed to use PasteTrace and were
expressly allowed to plagiarize the assignment. Extra credit
was offered to students who participated in this activity.
Further, we seeded two copies of the full solution on the
internet for each assignment.

The participating students were provided with a USB flash
drive containing PasteTrace, which was configured to save
project files within the flash drive. Participants received extra
credit when they turned in the flash drive. Further, students
were asked to self-report if they plagiarized to complete
the assignment. We should note that the data collection
was anonymous; the flash drive and its contents, along with
self-reported data, cannot be tied to any specific student.

As our evaluation was conducted by students enrolled in
two different courses, we report on our findings for each class
separately. As part of our evaluation, we also analyzed the
submitted code using a state-of-the-art tool, i.e., MOSS. The
files given to MOSS did not include our `metaComment`.

MOSS takes all code files submitted for an assignment as
input and then returns a list of file pairs. Each pair is given
a percentage score that represents what percentage of the
first file is also in the other file. Pairings are not necessarily
bijective, and there can be up to 2 items per pair of files,
each with a different percentage depending on which file is
listed first. The MOSS developers refer to this as the “Measure
of Software Similarity” and emphasize that similar is not
necessarily plagiarized, and manual review is required for
each pair. MOSS helpfully highlights the shared code. Only
file pairs with some similarity are returned. We manually
inspected all detected pairings for the ‘MOSS Output’ Column

304 of Tables I and II. As MOSS results require interpretation, we
305 categorized MOSS results in 3 ways.

- 306 1) Major Link: A large portion of the code is shared with the
307 listed file, enough that we would have counted the file as
308 plagiarized had this been a regular assignment. While this
309 is an arbitrary line, for these assignments "large" is "more
310 than 40% lines similar."
- 311 2) Minor Link: Moss reported that there is a link, but after
312 review, we determined this link was not significant and was
313 only composed of code that would naturally be similar given
314 the problem (e.g. a double nested for-loop using x and y
315 when iterating over an image). While again arbitrary, for
316 these assignments "less than 40%" was the dividing line.
- 317 3) No Plagiarism: Moss reported no links where this file was
318 the first in the pair.

319 Note that in these case studies, there is a clean divide at 40%
320 between Plagiarized, and not, but when using MOSS in the
321 real world, there is often not a clean division, and the pairs
322 require more thorough inspection to determine severity.

323 The code, raw and pre-processed data, automation reports,
324 and MOSS reports are available online¹.

325 A. Evaluation Results: Introduction to Programming I

326 We received six submissions from the first course. Each
327 submission has been arbitrarily assigned a single-letter
328 designation from A to F. The results are summarized in Table
329 I. Out of these six submissions, four successfully completed
330 the assignment, while two (A and D) did not.

331 First, examining the code submitted by students who did not
332 complete the assignment, we observe that Student A copied
333 and pasted from ChatGPT but prompted it wrong and did a
334 similar but different assignment. Plagiarism is easily detected
335 as there the `metaComment` is just a 75-line paste followed
336 by about 38 minor tweaks and formatting changes. Next,
337 Student D appears to have struggled: the code they copied is
338 a badly mangled AWT example and does not compile without
339 significant modification.

340 Examining the correct assignments, Student B completed
341 the assignment legitimately, and it is shown in the
342 `metaComments`. There are multiple well-organized files,
343 each organically coded with no large external pastes or
344 irregularly linear typing sections. There are copies between
345 files where code should be reused and modified, as well as
346 follow-up edits to make these modifications. Student E pasted
347 a large section of code from a different project with a matching
348 `InstallID` and then spent considerable time (around 8 hours
349 spread over a week or so) debugging and improving it. We
350 believe this to be a legitimate completion. Student F found
351 and submitted one of the seeded assignment files. It is logged
352 as a large foreign paste. Finally, we are unable to determine
353 if the code from Student C is legitimate.

354 MOSS only identified a small connection between Student
355 C and E. This connection is entirely composed of the expected
356 boilerplate code.

¹<https://figshare.com/s/43fbf89a815ba5c248cb>

B. Evaluation Results: Introduction to Programming II

357 We received 18 submissions from the second course. Each
358 submission has been given a single-letter designation from A
359 to R. The results are summarized in Table II. Out of these 18
360 submissions, 11 successfully completed the assignment, while
361 5 (C, G, L, M, and N) did not, and 2 (J and K) were empty
362 flash drives.
363

364 First, we examine the incorrect submissions. Student C
365 submitted a working Mandelbrot fractal written using Swing,
366 but it does not run in our IDE. Student G submitted copied
367 code for a Koch Snowflake written in Processing, which also
368 does not run. Student L initially copied the code of a Swing
369 program that makes a 3D rendering of a rain of rotating toroids
370 and then converted it to run in Processing. However, this is
371 not a solution to the actual assignment given. Student M used
372 the cross-copy-paste bypass method mentioned in ?? but failed
373 to explain their assignment. They claimed an IDE bug forced
374 them to use the exploit, and the solution provided did not meet
375 the assignment's requirements. Student N has created a 3D
376 solar system simulation using rotation and translation matrices.
377 However, it does not meet the assignment requirements.

378 We are unsure how C and G managed to find broken
379 solutions before finding working solutions. We tried searching
380 for the original source code, but were unable to locate it.
381 During the search we did find many solutions that would have
382 worked, but apparently the students did not find these.

383 Next, we examine the correct assignments. From the 11
384 submissions, 5 (E, F, H, I, and P) were flagged by MOSS as
385 having quite similar solutions with a similarity score of about
386 50%, so we shall discuss those first.

387 Student E's solution contains no large external pastes.
388 However when the edits are animated, the result is visibly
389 linear coding as if typing. The student claims they watched
390 a YouTube video of someone coding a Mandelbrot fractal
391 and followed along. While we consider E's solution to be
392 plagiarized, it is not part of the larger group unless the
393 group copied a source that itself is linked to the Video that
394 E copied. Student F and P independently copied a stock
395 example from the Processing website². Student H followed
396 an online video. After which, the final code was sent to I. The
397 `metaComment` for I is corrupted after this paste, but the final
398 solution of both students was significantly modified after from
399 this paste event, and the paste has little resemblance to either
400 final submission. The `InfectionStack` successfully tracks
401 this exchange.

402 Student A copied a Penrose Snowflake example using
403 multiple copies with no meaningful modifications. Likewise,
404 student B copied a Mandelbrot Fractal and made no
405 meaningful modifications. Students D and Q both asked
406 ChatGPT to generate a Mandelbrot and then copied the result
407 without modification. The results are nevertheless distinct from
408 each other. Student O copied a tree fractal example and then
409 modified it to be slightly different. Student R claims to have
410 used a video to help them make a Mendelbrot fractal complete

²<https://processing.org/examples/mandelbrot.html>

TABLE I
SUMMARY OF EVALUATION RESULTS FOR INTRODUCTION TO PROGRAMMING I

Student	Method	Visible in metaComment	Automated Check output	Moss Output
A	Plagiarized	Records large paste	Plagiarism Detected, 35 Edits	No Plagiarism
B	Legitimate	No Warning Signs	No Plagiarism Detected	No Plagiarism
C	Unclear	Large internal paste	Likely Plagiarized, 31 Edits	Minor link to E
D	Plagiarized	Records large paste	Plagiarism Detected, 31 Edits	No Plagiarism
E	Legitimate	Large Paste followed by many edits	Plagiarism Detected, 767 Edits	Minor link to C
F	Plagiarized	Large external paste	Plagiarism Detected, 0 Edits	No Plagiarism

TABLE II
SUMMARY OF EVALUATION RESULTS FOR INTRODUCTION TO PROGRAMMING II

Student	Method	Visible in metaComment	Automated Check output	Moss Output
A	Plagiarized	Records multiple large pastes	Plagiarism Detected, 1 Edits	No Plagiarism
B	Plagiarized	Records multiple large pastes	Plagiarism Detected, 0 Edits	Minor link to R
C	Plagiarized	Records multiple large pastes	Plagiarism Detected, 4 Edits	Minor link to R
D	Plagiarized	Records large paste	Plagiarism Detected, 0 Edits	Minor link to Q
E	Plagiarized	Visible Linear Coding	No Plagiarism Detected	Major Link to F, I, H, and P
F	Plagiarized	Records large paste	Plagiarism Detected, 105 Edits	Major Link to E, I, H, and P
G	Plagiarized	Records large paste	Plagiarism Detected, 78 Edits	No Plagiarism
H	Plagiarized	Records large paste	Plagiarism Detected, 78 Edits	Major Link to E, F, I, P
I	Plagiarized	Infected by H InstallID	Copied Directly from H	Major Link to E, F, H, P
J	Blank	Blank Submission		
K	Blank	Blank Submission		
L	Plagiarized	Records multiple large pastes	Plagiarism Detected, 11 Edits	No Plagiarism
M	Unclear	Records Large Internal Paste	Plagiarism Detected, 0 Edits	No Plagiarism
N	Legitimate	No Warning Signs	No Plagiarism Detected	No Plagiarism
O	Plagiarized	Records multiple large pastes	Plagiarism Detected, 303 Edits	No Plagiarism
P	Plagiarized	Records large paste	Plagiarism Detected, 38 Edits	Major Link to E, F, I, H
Q	Plagiarized	Records large paste	Plagiarism Detected, 3 Edits	Minor link to D and R
R	Legitimate	No Warning Sign	No Plagiarism Detected	Minor Link to Q

411 with zoom. However, there is no indication they copied from
412 the video, and we believe this to be a legitimate completion
413 of the assignment.

414 MOSS revealed weak connections between B and R, and
415 C and R, which were deemed to be incidental. Additionally,
416 a moderate connection was identified between D and Q, as
417 well as R and Q, despite the absence of a direct connection
418 between D and R. Notably, only two of these connections
419 were generated by the ChatGPT platform. Furthermore, the
420 identification of a strong connection between H and I, and
421 between F and P, was indicative of plagiarism within those
422 respective pairs. However, MOSS also flagged weaker yet still
423 notable connections among all four pairs and additionally with
424 E. It may not be accurate to include E in this group, as E
425 used a different source. It can be argued that MOSS correctly
426 identified the plagiarism despite the changes. However, the
427 similarities between the two groups are the same as the
428 similarities between each and E, and the code is quite common
429 to Mandelbrot fractals. As such, we believe this to be incorrect
430 with regard to plagiarism.

431 C. PasteTrace Vs. MOSS

432 Comparing directly to MOSS objectively is difficult as
433 MOSS requires a highly subjective analysis. However, if each
434 reported pairing from MOSS is considered correct, then MOSS
435 successfully detected 9 of 16 plagiarised assignments and had
436 two false positives. However, if our subjective analysis is used,
437 and only Major links are considered, it falls to 5 of 16 with 0

false positives. In comparison, our tool detected 15 of 16 and
438 had one false positive. 439

440 Due to its inherent design, MOSS is incapable of detecting
441 plagiarism unless two students have submitted similar code
442 files. Since our tool relies on the method used to create the
443 code, not just the end result, it can detect such events. 444

445 V. CONCLUSION

446 This paper introduces PasteTrace, a novel open-source
447 plagiarism detection tool specifically designed for introductory
448 programming courses. Our evaluation of PasteTrace
449 demonstrates the tool's capability to detect various forms
450 of plagiarism, including those from single sources, which
451 are often overlooked by comparative tools like MOSS.
452 Additionally, the tool's ability to track copy-paste events,
453 code origins, and the overall coding process provides
454 valuable insights into student behavior and potential academic
455 dishonesty. Our future work for PasteTrace includes refining
456 the detection algorithms, expanding compatibility with
457 various IDEs, and conducting large-scale studies across
458 diverse programming courses.

459 REFERENCES

- 460 [1] U. Omer, M. S. Farooq, and A. Abid, "Introductory programming
461 course: review and future implications," *PeerJ Computer Science*, vol. 7,
462 p. e647, 2021. 1
463 [2] P. Piwek and S. Savage, "Challenges with learning to program and
464 problem solve: An analysis of student online discussions," in
465 *Proceedings of the 51st ACM Technical Symposium on Computer
466 Science Education, SIGCSE '20*, (New York, NY, USA), p. 494–499,
467 Association for Computing Machinery, 2020. 1

- 467 [3] U. Nikula, J. Sajaniemi, M. Tedre, and S. Wray, "Python and roles
468 of variables in introductory programming: experiences from three
469 educational institutions," *Journal of Information Technology Education: Research*, vol. 6, no. 1, pp. 199–214, 2007. 1
- 471 [4] R. P. Medeiros, G. L. Ramalho, and T. P. Falcão, "A systematic literature
472 review on teaching and learning introductory programming in higher
473 education," *IEEE Transactions on Education*, vol. 62, no. 2, pp. 77–90,
474 2019. 1
- 475 [5] Y. Qian and J. Lehman, "Students' misconceptions and other difficulties
476 in introductory programming: A literature review," *ACM Trans. Comput. Educ.*, vol. 18, oct 2017. 1
- 477 [6] R. Fraser, "Collaboration, collusion and plagiarism in computer
478 science coursework," *Informatics in Education-An International Journal*,
479 vol. 13, no. 2, pp. 179–195, 2014. 1
- 481 [7] V. T. Martins, D. Fonte, P. R. Henriques, and D. da Cruz,
482 "Plagiarism detection: A tool survey and comparison," in *3rd Symposium on Languages, Applications and Technologies (2014)*,
483 Schloss-Dagstuhl-Leibniz Zentrum für Informatik, 2014. 1
- 484 [8] M. Jiffriya, M. Jahan, and R. Ragel, "Plagiarism detection tools and
485 techniques: A comprehensive survey," *Journal of Science-FAS-SEUSL*,
486 vol. 2, no. 02, pp. 47–64, 2021. 1
- 487 [9] O. U. press, ed., *Oxford English Dictionary*. 1
- 488 [10] J. P. Gibson, "Software reuse and plagiarism: a code of practice," *ACM SIGCSE Bulletin*, vol. 41, p. 55–59, July 2009. 1
- 489 [11] P. Walcott, "Attitudes of second year computer science undergraduates
490 toward plagiarism," *Caribbean Teaching Scholar*, vol. 41, p. 63–80, Dec.
491 2016. 2
- 492 [12] S. Schleimer, D. S. Wilkerson, and A. Aiken, "Winnowing: local
493 algorithms for document fingerprinting," in *Proceedings of the 2003 ACM SIGMOD international conference on Management of data*, ACM,
494 June 2003. 2
- 495 [13] M. Konecki, T. Orehovacki, and A. Lovrencic, "Detecting computer
496 code plagiarism in higher education," in *Proceedings of the ITI 2009 31st International Conference on Information Technology Interfaces*, IEEE,
497 June 2009. 2
- 498 [14] D. Heres, *Source Code Plagiarism Detection using Machine Learning*.
499 PhD thesis, Utrecht University, Aug. 2017. 2
- 500 [15] J. Y. B. Katta, "Machine learning for source-code plagiarism detection,"
501 2018. 2
- 502 [16] F. Ullah, J. Wang, M. Farhan, M. Habib, and S. Khalid, "Software
503 plagiarism detection in multiprogramming languages using machine
504 learning approach," *Concurrency and Computation: Practice and Experience*, vol. 33, Oct 2018. 2
- 505 [17] N. Awale, M. Pandey, A. Dulal, and B. Timsina, "Plagiarism detection
506 in programming assignments using machine learning," *September 2020*,
507 vol. 2, p. 177–184, Jul 2020. 2
- 508 [18] R. Koike, M. Kaneko, and N. Okazaki, "Outfox: Llm-generated essay
509 detection through in-context learning with adversarially generated
510 examples," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 38, Mar. 2024. 2
- 511 [19] J. H. Kirchner, L. Ahmad, S. Aaronson, and J. Leike, "New ai classifier
512 for indicating ai-written text," Jan 2023. 2
- 513 [20] A. M. Elkhayat, K. Elsaid, and S. Almeer, "Evaluating the efficacy
514 of ai content detection tools in differentiating between human and
515 ai-generated text," *International Journal for Educational Integrity*,
516 vol. 19, Sep 2023. 2
- 517 [21] M. AlSallal, R. Iqbal, V. Palade, S. Amin, and V. Chang, "An integrated
518 approach for intrinsic plagiarism detection," *Future Generation Computer Systems*, vol. 96, p. 700–712, Jul 2019. 2
- 519 [22] *The Processing Environment*, pp. 143–170. Berkeley, CA: Apress, 2007.
520
521
522
523
524
525
526
527