

# AC: An Integrated Source Code Plagiarism Detection Environment

Manuel Freire, Manuel Cebrián and Emilio del Rosal

Escuela Politécnica Superior

Universidad Autónoma de Madrid

28049 Madrid, Spain

`{manuel.freire, manuel.cebrian, emilio.delrosal}@uam.es`

## Abstract

Plagiarism detection in educational programming assignments is still a problematic issue in terms of resource waste, ethical controversy, legal risks, and technical complexity. This paper presents AC, a modular open-source plagiarism detection program. The design is portable across platforms and assignment formats and provides easy extraction into the internal assignment representation. Multiple similarity measures have been incorporated, both existing and newly-developed. Statistical analysis and several graphical visualizations aid in the interpretation of analysis results. Finally, a case study is performed on two programming assignment data sets, one of which was automatically-generated.

# Contents

<b>Title</b>	<b>1</b>
<b>1 Introduction</b>	<b>3</b>
<b>2 State of the art and motivation</b>	<b>7</b>
<b>3 Design of AC</b>	<b>14</b>
3.1 Distance integration . . . . .	15
3.1.1 Compression similarity distance . . . . .	16
3.1.2 Token counting similarity distance . . . . .	18
3.1.3 Infrequent substrings similarity distance . . . . .	21
3.2 Threshold recommendation and outlier detection . . . . .	22
3.2.1 Discordancy tests . . . . .	24
3.2.2 The Discordancy test choice: the Hampel identifier . .	34
3.3 Submission filtering . . . . .	38
3.4 Visualization . . . . .	43
3.4.1 Graph representation . . . . .	44
3.4.2 Histogram visualization . . . . .	47
<b>4 Practical examples</b>	<b>49</b>
4.1 Artificial submissions . . . . .	50
4.2 Real submissions . . . . .	55
<b>5 Conclusions and further work</b>	<b>58</b>
<b>Bibliography</b>	<b>63</b>

# 1 Introduction

The development of the World Wide Web and the increasing standardization of electronic documents has led to a greater incidence of plagiarism in many aspects of life. According to a recent article in Nature [20], incidence of plagiarism has also reached the scientific community. However, it is much more widespread in the case of undergraduate students, where educators are generally ill-equipped to face the technological challenges posed by plagiarism. Realizing this, different national educational authorities (for instance, those of the United Kingdom [14, 8]) have begun funding projects dedicated to study the impact and growth of plagiarism, and to propose adequate measures.

Exact figures are unknown, since successful plagiarism is by definition not detected, but are believed to be high and growing [13, 27]. Alexander Aitken, one of the leading experts in operating plagiarism detection software, asserted in a personal communication prior to 2001 that for any (USA) student corpus, 10% of submissions are plagiarized [14, p. 4].

Two major types of documents are being targeted by undergraduate plagiarism: essays and computer submissions, although plagiarism cases in art degrees have also been reported [49, p. 4]. This paper focuses on source code plagiarism.

Plagiarism detection in programming courses is tedious and extremely time consuming for graders. Additionally, it is emotionally and legally risky for student and educator alike [24]. University experience also shows that even minor plagiarism levels can cause a mistrust for the work of students

which can lead to baroque examinations to prove the authenticity of each student's work, or to the relative weight of possibly-plagiarized practical submissions in the final grade to be far lower than the actual share of effort they truly required from the student. Ignoring the problem posed by plagiarism results in unfair grading, and may have an avalanche effect in plagiarism incidence levels.

There are two facets in the prevention of plagiarism in computer programs. The first is of ethical and normative nature, and involves fighting the deeper causes of plagiarism and selecting appropriate academic and legal deterrents. This facet is examined, for example, in [7] and [34]. The second facet is directly related with software engineering, and addresses the technical measures required to detect plagiarism within a set of submissions.

The present work is focused on this second facet, and seeks to assist educators in the task of plagiarism detection by following the approach of [29]: monitoring excesses of collaboration which can signal anomalous behaviors in students.

Among the most usual causes for plagiarism, we may find the following:

- Low ethical and/or technical preparation of the students.
- Ambiguity in assignments: poorly understood exercises are more likely to suffer plagiarism.
- Low clarity in the university's guidelines on student collaboration.
- Bad course planning causing an excessive work load.

In this light, a plagiarism detection tool can be considered as a sanity check,

to be used in the diagnosis of the relative health of the teaching environment. If this check is to be objective, we feel that plagiarism detection tools should be updated and augmented with modern technologies as those discussed in Sect. 2.

The main contribution of this research is the design and development of AC, a novel plagiarism detection tool available at

**<http://tangow.ii.uam.es/ac>**

Notable features of AC include the following:

- Intuitive display of results, providing visualizations such as graphs and different types of histograms, allows visual exploration of analysis results. From the visualizations, plagiarism suspects can be selected in order to perform detailed manual analysis.
- Extraction/filtering utility designed to ease the initial task of preparing assignment submissions for analysis. Submission preprocessing would otherwise be a repetitive and error-prone task; this utility seeks to automate the process without being specific to any particular institution.
- Implements a novel distance integration architecture, allowing similarity distance algorithms to be combined, compared and refined. Many distance algorithms found in the scientific literature on plagiarism detection are included, together with several others developed by our research group.
- Incorporates a novel threshold recommendation system based on statistical outlier detection. Distances lower than this threshold are sug-

gested for manual analysis, establishing an heuristic starting point to help grader's unveiling of plagiarism patterns.

- Open Source code, featuring a modular design that allows easy customization of both algorithms and visualizations. AC can therefore be analyzed, extended and adapted for any particular requirements.
- Stand-alone platform-independent program which can be executed in any computer with a JAVA runtime environment. There is no need to send the submissions to a server in a different institution, avoiding privacy concerns. Furthermore, a desktop application can deliver interactive visualizations which would be difficult to perform online.
- Intended to be a long-term supported tool by means of facilities such as version source-control system, bug-tracking record, forums and an updated website. Transition to Sourceforge or other Open Source collaborative software management portals is currently under study.

The remainder of this paper is organized as follows: Sect. 2 motivates the need for AC by means of a review of the state of the art in tools for plagiarism detection; Sect. 3 details several design aspects which make AC a novel tool: data extraction, graphical visualization, metric integration and a threshold recommendation system. Sect. 4 presents AC at work using two practical case studies: the first with artificially generated source code and the second with real submissions as coded by students. Finally, sect. 5 presents conclusions and outlines future work.

## 2 State of the art and motivation

Although the concept of plagiarism is somewhat vague, there is general agreement on the desirable detection capabilities that a copy-detection algorithm should have, in the spirit of schemata from [46] and [38]. The following list enumerates possible types of source code manipulations which may be performed to disguise the code origin of plagiarized code:

### 1. Text replacement, within

- (a) identifiers (variables, function names, labels, etc.)
- (b) textual strings within a program.
- (c) comments.

### 2. Code reordering

- (a) Internally in an instruction (e.g. `if (!A && B)` transforms into `if (B && !A)`.)
- (b) Internally within a method or function (change instructions order.)
- (c) Reordering method or functions within the same source file.
- (d) Interchanging methods or functions between different source files.

### 3. Code rewriting, within

- (a) an instruction (e.g `i++` transforms into `i=i+1`.)
- (b) a method/function: ‘if-then-else’ engineering, loop unrolling, etc.

### 4. Spurious code insertion/deletion

- (a) Error checking (eg.: `if (parameter == NULL)`.)
- (b) Extra debugging code (`printfs`, `try-catch`, `#ifdef`, etc.)

## 5. Source code mixing

- (a) Mixing fragments of plagiarized code with non-plagiarized, original code.
- (b) Mixing plagiarized code from several sources (multiple plagiarism.)

The above list is rather exhaustive, and should enough different techniques be combined in a single plagiarized submission, the result would probably escape detection even by experienced graders. However, the goal of cheating is usually to save time and effort, and the effort required to perform extensive manipulations on a plagiarized program can be greater than the effort required for full rewrite. Note that a full rewrite will generally not be registered as plagiarism, since the great majority of plagiarism detection tools only analyze similarities at a lexical level. Understanding and comparing the informal, higher-level semantics of programs would be considerably more complicated.

In general terms, available automatic plagiarism detection tools deal without any problems with manipulations of type 1 (textual substitution), by using a lexer for each programming language and comparing only tokens, probably ignoring comments entirely. Manipulations of type 2 (reordering) are more difficult to detect, but fortunately this degree of difficulty is directly proportional to the source code knowledge required to perform these manipulations.



Code rewriting (type 3) is one of the most affordable manipulations for students, since only working knowledge of the programming language is required. However, it is difficult to detect for algorithms that only seek similarities in token sequences. Manipulations of type 4 (extra code insertion) are also hard to detect, even for humans, because they can be considered both as a valuable improvements to the program's robustness, or as camouflage for a possible fragment of plagiarized code.

Finally, source code mixing (type 5) is similar to code reordering: a student would generally need to have a good knowledge of the program's structure to successfully mix code from different sources, independently of the origin of the second source. Should the second source be original, un-plagiarized code, this raises the question of how much is too much. Plagiarism detection tools can only point to probable cases of plagiarism, but the final decision is always in the hands of a human grader. The value of automated plagiarism detection is in narrowing down the search by highlighting the clearest cases. No grader can be reasonably expected to perform the 5050 possible pairwise comparisons required to check for plagiarism within a corpus of 100 submissions; a program, on the other hand, can perform such a check in seconds.

With similar ideas in mind, several plagiarism detection tools have been implemented since the 1960s. MOSS [3], SIM [22], YAP [33], JPLAG [42] and SID [11] are probably the most widespread within the academic community. The degree to which these tools withstand the manipulations listed above has only been studied in small-scale experiments; results can be found in [21, 11, 47, 6].

A qualitative study based on general properties and the distances they deploy to find similarity can be found in an excellent review by Lancaster and Culwin [32]. According to their conclusions, two systems seem to outperform the rest: MOSS and JPLAG. However, when comparing their performance on a single corpus [14], no tool significantly outperformed the other. This and other studies [41, 42] suggest that there is no “silver bullet” for plagiarism detection, and the degree of success achieved by a tool depends mainly on its ease of use for prospective graders.

Lancaster and Culwin enumerate a set of comparative properties which can aid a grader or institution in the selection of a plagiarism detection tool. An abridged version follows:

- i. *Locality* of the tool – the tool’s software may be downloaded and the processing can be performed locally, or it may be web-based, with processing taking place on a remote server (maybe in another country).
- ii. *Breadth* – number of programming languages and variants that the tool can process.
- iii. *Privacy* – the engine may only be available to its host institution, or it may be a public tool, available for widespread use.
- iv. *Documentation* – existence of technical information about the tool in the form of papers and reports, or even source code and associated program documentation, which facilitate the comprehension, validation and replication of the plagiarism detection algorithms used within the tool and the tool itself.

- v. *Algorithms* – quality and variety of the similarity distances incorporated into the tool.
- vi. *Visualization* – The existence or not of a Graphical User Interface (GUI) incorporated into the tool.
- vii. *Preprocessing* –The existence of a collection tool for submissions, or the difficulty of the process of assembling the corpus in some predetermined way prior to analysis.
- viii. *Support* – The availability of long-term technical support of the tool.

Our aim when designing and developing a new plagiarism detection tool is to fill some gaps pointed out in the literature. We enunciate those gaps referencing the above points.

The *locality of the tool* is gaining importance due to increased awareness of intellectual property issues. Data protection laws from several countries and universities current explicitly prohibit exportation of identifiable data to external organizations (see for instance [2, 1]). This would prevent graders from using external, web-based engines. Even though multiple local plagiarism detection tools have been developed, most of them are either out of availability (TEAMHANDIN [15], SAXON [45], COGGER [16]), heavily platform-dependent, and/or unsupported (YAP3, SIM, SHERLOCK [30], JONES [28, 29], DETECTACOPIAS [35], BIG BROTHER [27]). AC is completely stand-alone, and can be run in any computer with a suitable JAVA runtime environment.

Regarding *Breadth* and *Privacy*, many systems lack comprehensive documentation on their internal logic and on the algorithms they use. Some authors oppose public access to this information, since this could lead to circumvention of these algorithms by cheaters. This can be considered akin to “security by obscurity”. Although winds of change are blowing (for instance, MOSS is starting to detail its so far unclear functioning [46]), few systems have actually opened up their source code to external scrutiny.

However, the implementation details of plagiarism detection tools are as important as their underlying algorithms. The effectiveness of these algorithms is strongly dependent on “details” such as thresholds and parsing methods. Additionally, if students within an institution learn that certain types of cheating are likely to be caught, while others seem to be much safer, the types of plagiarism that the tool will encounter will vary very fast. This calls for an equally dynamic and tunable plagiarism detection tool, which can be adapted to stop leaks as soon as they are detected. A closed, centralized development model is unlikely to provide the necessary flexibility [25]. AC is therefore being developed as an open-source tool.

Similarity distance algorithms, point **v.** above, play a major role in plagiarism detection. If distances are reliable, they can be used to narrow down the search for plagiarized submissions; otherwise, the grader will be none the wiser. Many similarity distances have been proposed in the literature, some of them intuitive and others less so, each with different strengths and weaknesses. Even though all of them have been shown effective in locating plagiarism (although in a rather unstructured experiments) little correlation between them has been observed. The similarity rankings produced

by MOSS and JPLAG are clear examples of this [14, p. 14]. AC provides a framework that allows easy integration of different similarity distances, in an effort to complement the weaknesses of one with the strengths of another. Most current tools follow the opposite approach: one, or at most two similarity distances are used, frequently small variations of the same pattern-matching algorithm.

Regarding the visualization of results, the current tools tend to provide a very simplistic overview of the actual similarities. Detailed analysis of pairs of submissions is usually available, but support for the investigation and validation of clusters of similarities is generally absent. One of the scarce exceptions is the system by Ribler and Abrams [43], but the tool itself is no longer available on their web page, and their paper provides very few implementation details. AC seeks to change this situation by incorporating multiple visualizations, supporting high-level overviews, clique detection, and detailed analysis of individual submissions within the whole population, and one-to-one submission comparison. Further comparison of our visual facilities with the ones of the MOSSCLIQUES POST processing software [40] will be performed as soon as the latter is incorporated into MOSS.

The existing variability in submission delivery formats, generally imposed by the graders, frequently becomes a major usability bottleneck. Plagiarism detection tools expect submissions to be structured according to rigid specifications, while each university or subject will probably adhere to a different standard. Many current tools are therefore highly targeted to particular educational environment. AC provides an extraction/filtering utility that automates the conversion from the submission format to the in-

ternal submission structure, designed to work with almost all packaging and compression formats. This tool is essential to perform analysis of hundreds of submissions with tens of files in each in real-world settings. Without it, the manual intervention required to reach a common format would probably discourage many graders from performing automated plagiarism detection at all.

Regarding requirement **viii.**, many existing systems are supported through a web-site where documentation can be found, and provide support through individual queries by e-mail. AC's web page includes detailed system documentation and tutorials for new users, and doubles as a software project site, including internal program documentation, forums, issue-tracking support, and as an access point to the source code versioning system.

### 3 Design of AC

The design of AC relies on a general framework for similarity measurement. This framework is based on requiring all similarity distance algorithms (also referred to as *similarity tests*) to return, for each pair of submissions  $A$  and  $B$ , a single value to represent the similarity distance between both, according to that algorithm. Similarity distance are expected to be symmetrical and normalized between 0 (identical sources), and 1 (minimal similarity, no common information). Therefore, this distance is also referred to as *normalized similarity*. Normalization can be trivially performed by calculating all distances, and scaling all values so that they fall between 0 (corresponding to the lowest observed distance) and 1 (highest observed distance). This type

of normalization should only be used as a last recourse, since information on the actual range of observed distances is lost in the process.

Standardization is desirable because it simplifies comparisons between different test results, and makes the tests themselves interchangeable from the point of view of further processing steps: data sets can be examined using different similarity distance algorithms, simplifying the evaluation of the relative strengths and weaknesses of each. It is easy to add new tests, and to compare their results against those of existing ones. Additionally, results from multiple tests can be easily composed into higher-level tests.

### 3.1 Distance integration

The output of a test, or a composition of tests, is a *distance matrix*  $\mathbf{D}$  of  $N$  rows and columns. Each of the  $N$  columns and rows of this matrix represents a submission, so that a single cell  $\mathbf{D}_{AB} \in [0, 1]$  contains the similarity distance between the submissions  $A$  and  $B$ .

An ambitious goal of AC is to use multiple “orthogonal” similarity measures at once. Orthogonality is used informally in the sense of specialization on different sets of similarity dimensions (for instance, common token sequences, comment similarities, or matching token frequencies). Individual similarity distance tests can be composed into higher-level tests, allowing a modular approach to test construction: what one test does not uncover, another will highlight, and a well-designed composition can yield the best of both. Many metrics that are each accurate in measuring particular sets of similarity dimensions are expected to be more accurate than a single, complex test, which would be harder to tune or develop.

A composition approach also has advantages when exposing configuration controls to end-users: it is easier to configure a set of blocks, each with a well-defined goal and few parameters, than a single monolithic block with many goals and a larger set of possibly interdependent internal implementation parameters.

This modular design of graphical interface, metrics and test visualizations pitches AC as a powerful testbench for source-code plagiarism research.

### 3.1.1 Compression similarity distance

A natural (yet not discovered until the late 90s) measure of similarity is based on the observation that two strings  $p$  and  $q$  are similar if the basic blocks of  $p$  are in  $q$  and vice versa. String  $p$  can be described by referencing the blocks belonging to  $q$ . If both strings are indeed similar, the description of  $p$  using that of  $q$  will be very simple.

This is essentially how a file compressor would operate on the concatenated  $pq$  sequence: the compressor would attempt to eliminate redundant information from the merged sequence; if information from  $p$  is present in  $q$ , or vice-versa, the compressed size of  $pq$ ,  $C(pq)$ , will be smaller than  $C(p)+C(q)$ , since part of the information contained in  $p$  can be described by referencing suitable blocks of  $q$ . This similarity measure was formalized by Rudi Cilibrasi and Paul Vitányi in [12], giving rise to the concept of *normalized compression distance* (NCD), based on the use of compressors to provide a measures of the similarity between sequences. This distance is defined as



follows, and may then be used to cluster strings.

$$\text{NCD}(p, q) = \frac{C(pq) - \min\{C(p), C(q)\}}{\max\{C(p), C(q)\}} \quad (1)$$

where  $pq$  is the concatenation of strings  $p$  and  $q$ , and  $C(x)$  denotes the length of the text  $x$  compressed using some compression algorithm which asymptotically reaches the true entropy of  $x$  as the length of  $x$  tends to infinity.

This idea is very powerful, because it can be applied in the same way to any type of sequence, be it music, text, or genetic data. There is no need to use specific features of sequences. The distance from sequence  $p$  to sequence  $q$  is simply a measure of the difficulty of using  $p$  describe  $q$  and vice versa.

The first proposal to use this similarity distance for plagiarism detection can be found in [45]; however, it was not rigorously studied until the development of the SID engine (available at <http://genome.math.uwaterloo.ca/SID/>), experimentally demonstrating that the use of NCD as a text similarity distance within a plagiarism detection tool has advantages over other similarity distances [11].

NCD has been incorporated into AC in two flavors: the *Zip* test operates on tokenized source code (and is therefore dependent on the availability of a suitable parser) while the *Gross Zip* distance operates on raw source code. The latter is specially useful when analyzing code written in programming languages for which it is difficult to obtain a parser. In both cases, compression is done by means of the standard JAVA compression library using the classical Lempel-Ziv algorithm [50].

### 3.1.2 Token counting similarity distance

The token-counting test is an example of attribute-counting test [30], in the sense that it does not analyze a program’s structure, centering instead on the distribution of certain structure-independent attributes. The main idea behind the test is that, if two submissions share a substantial amount of code, then the frequency of each fundamental token in both submissions will be similar.

A vector is built for each submission  $A$ , such that  $\mathbf{v}_A = (t_0, \dots, t_n)$ , where  $t_i$  represents the number of occurrences of token  $i$  in  $A$ ’s source files. The euclidean distance between each pair of vectors,  $\mathbf{v}_{A_j}$  and  $\mathbf{v}_{A_k}$ , is then used as a similarity measure between  $A_j$  and  $A_k$ .

Since the general contract of a similarity test requires results to be normalized in the range  $[0, 1]$ , normalization is performed by normalizing both vectors before checking their distance. Therefore,

$$\text{dist\_token}(A, B) = \frac{\mathbf{v}_A}{|\mathbf{v}_A|} \cdot \frac{\mathbf{v}_B}{|\mathbf{v}_B|} = \cos(\widehat{\mathbf{v}_A \mathbf{v}_B}) \quad (2)$$

This is the vector-space model distance used in typical information retrieval applications, as described in [4]. However, it is not robust when used with programming submissions, since it is trivial to convert from certain tokens to other, semantically equivalent constructions in a way that completely thwarts this similarity comparison. It should therefore be used only in conjunction with stronger tests, as support.

However, it must be noted that this test is completely resistant to attempts to disguise plagiarism by any structural obfuscation attempts, since

it is only the relative frequency of tokens that is taken into account, and not their location within the source code. For instance, refactoring a program into different sets of functions would have a very small impact on token distribution, and would therefore fail to mislead the token-counting test.

The variance distance is a refinement on other tests, because it does not compute a similarity directly from the source code of submissions. Instead, it refines a distance matrix in an attempt to assign higher similarities to “left-most” outliers within individual distributions (see section 3.4.2).

The intuitive goal is, when refining the distance between submissions, to give lower distances to submissions that appear as a “left-most spike” in each other’s individual histograms. In a more formal characterization, let  $A$  be a submission and  $\mathbf{D}_A$  the set of distances from  $A$  to all other submissions  $X_i$ , as found in a distance matrix  $\mathbf{D}$ :

$$\mathbf{D}_A = \mathbf{D}_{AX_1}, \dots, \mathbf{D}_{AX_n} \tag{3}$$

That is,  $\mathbf{D}_A$  is a row from the distance matrix  $\mathbf{D}$  that is to be refined as a result of running this test.

Let  $B$ , another submission, have an individual distribution  $\mathbf{D}_B$ . If  $\mathbf{D}_{AB}$  is a left-most outlier (that is, is uncommonly low) in both  $\mathbf{D}_A$  and  $\mathbf{D}_B$ , then there is a strong case for  $A$  and  $B$  to have much in common, regardless of the actual value of  $\mathbf{D}_{AB}$  within the global distribution of distances (which may be moderate enough as to not currently arouse suspicion).

Since we are only interested in *left-most outliers* (uncommonly low distances within a distribution), outliers to the right should not be considered.

Therefore, our calculations substitute classical standard deviations of a distribution  $\mathbf{D}_X$  for *left standard deviations*, or  $\text{LSD}(\mathbf{D}_X)$ , defined as the standard deviation calculated exclusively using those distances  $\mathbf{D}_{XY}$  such that  $\mathbf{D}_{XY} < \widehat{\mathbf{D}}_X$ , that is, lower than the mean of  $\mathbf{D}_X$ . Distances greater than the mean do not contribute towards the LSD.

The following metric is used to determine the *outlierlyness* of  $A$  in  $B$ :

$$\text{out}(A, B) = \frac{\max(\widehat{\mathbf{D}}_B - \mathbf{D}_{AB}, 0)}{\text{LSD}(B)} \quad (4)$$

That is, the degree in which  $A$  is a left-most outlier of  $\mathbf{D}_B$  is 0 if  $A$ 's distance is greater than  $\mathbf{D}_B$ 's mean, and positive and proportional to the number of left-most standard deviations of the difference otherwise.

Using this definition, the degree of *outlier importance* that we are willing to attribute to this fact is expressed by

$$\text{OI}(A, B) = 2^{-\text{out}(A, B) \cdot \text{out}(B, A)} \quad (5)$$

This exhibits the desirable traits of being symmetric ( $\text{OI}(A, B) = \text{OI}(B, A)$ ), and bounded in the interval  $]0, 1]$ , since  $\text{OI}(A, B) = 1$  when  $\text{out}(A, B)$  or  $\text{out}(B, A)$  are zero, and  $\simeq 0$  when both are relatively high. The outlier importance can then be factored into a *corrected distance*  $\mathbf{D}'$  with a weighing constant  $k_{oi}$ :

$$\mathbf{D}'_{AB} = \text{OI}(A, B) \cdot k_{oi} + \mathbf{D}_{AB} \cdot (1 - k_{oi}) \quad (6)$$

A value of 0.4 for the outlier-importance constant  $k_{oi}$  has proven useful

in our experiments, and is currently the default value when the “previous test”  $D$  is a compression test. Once the compression test’s results are refined with this test, the end effect on a histogram display is that gaps are enlarged, causing a few distances to become lower, and therefore “uncovering” previously hidden similarities between submissions.

### 3.1.3 Infrequent substrings similarity distance

This test is based on the observation that plagiarized sources are likely to share substrings which are infrequent within the whole corpus. Therefore, shared, infrequent substrings can be considered a strong indicator of similarity. Code which has been provided by the course instructor is likely not to be infrequent, as are common idioms in the corresponding programming language. These would be ignored by the test. However, if substrings shared by a pair of submissions can not be found anywhere else in the corpus, it is reasonable to assume a common origin.

For instance, given submissions  $A$  and  $B$  within a corpus of submissions written in the C programming language,  $A$  and  $B$  will likely share the substrings `printf` or `for(int i=0; i<n; i++)`, since these are common C idioms. Likewise, if requested in assignment’s statement, the substring `int functionRequestedByStatement()` would also be common to most submissions found in the corpus. On the other hand, substrings such as like `// A very peculiar mis-spelt comment` or `fflush(a_file)` are much scarcer in C submissions, and when found only in  $A$  and  $B$ , can be considered as indications of a common origin.

Based on this idea, the *Infrequent Substrings* distance builds a data

structure (a sort of *Patricia* tree [36]) that stores every substring within a minimum and maximum length, as well as the sources where that substring appears. Later on, this data structure can be used to assign a similarity distance to each pair of submissions within the corpus, based on the relative frequency of their shared substrings. The current implementation is the following:

$$\text{dist\_substring}(A, B) = \frac{f_{AB}}{n} \times \frac{l_{AB}}{L_{\max}}$$

where  $f_{AB}$  stands for the frequency within the corpus of the most infrequent substring shared by  $A$  and  $B$ ,  $n$  is the number of submissions within this corpus,  $l_{AB}$  stands for the length of the most infrequent substring shared between  $A$  and  $B$ , and  $L_{\max}$  is the maximum substring length considered. In other words, those couples of submissions that share a very rare and long substring receive very small distance values.

The *Substring* distance is still under development and subject to constant improvements. Nevertheless, it has already been used with promising results in [10], demonstrating that it is a valuable alternative to other similarity distance tests. Moreover, this test’s results can be used to highlight areas of identical (and, more important, infrequent between two submissions), making it potentially valuable during visual side-by-side similarity inspection.

### 3.2 Threshold recommendation and outlier detection

The result of performing similarity calculations is a large (quadratic in the number of submission in the corpus) number of pairwise distances, a number

of which will be very low in case of plagiarism occurrence. The goal of a plagiarism detection tool is to flag these distances as suspicious. A human grader can then locate cases of plagiarism by visual inspection of only the (few) likely candidates, without requiring analysis of the remaining set of (many) non-plagiarized pairs.

When based only on either the bulk of distance values, the process of flagging likely candidates depends on the determination of a *similarity threshold* value. If a pair of submission has a similarity distances which falls below this threshold, it will be flagged for inspection. Otherwise, the pair will not be flagged. Of course, the use of a single threshold value incurs the risk of false positives and false negatives. Graders should be aware of these risks, and exercise caution when heeding its recommendations. However, a threshold value is a useful starting point to start a search, and determining a sane initial value can greatly increase the ease of use of a plagiarism detection tool.

However, the task of locating a good initial threshold has received no attention in the literature, and is therefore left completely to the personal decision of the grader. The following example illustrates the problems raised by this issue. A grader is asked to locate cases of plagiarism within a corpus of size 5. There will be  $\binom{5}{2} = 12$  pairwise distances, say, 0.1, 0.15, 0.3, 0.6, 0.6, 0.62, 0.62, 0.64, 0.64, 0.65. It seems intuitive to flag 0.1 and 0.15 as surprisingly small when compared to the rest; but how should 0.3 be considered? The example may be driven further: assume now that the distribution of the above distances is known to be uniform between 0 and 1: would the grader still flag 0.1 and 0.15 as surprising? probably not, because

those distances are fairly probable within a uniform distribution. With no prior assumption about the mechanism which generated the data, detection skills rely too much in the subjectivity of the grader.

Under certain conditions, it is possible to quantify the ‘amount of surprise’ presented by a given value within a sample. This problem is completely equivalent to the one of finding *outliers*, a classical problem in statistical data analysis. Quoting Barnet and Lewis [5, page 7]

We shall define an outlier in a set of data to be an observation (or subset of observations) which appears to be inconsistent with the remainder of that set of data.

The reader would have probably realized that the term ‘surprisingly small distances’ referred to throughout this work is equivalent to the term *lower outliers* as used in statistics. In other words, the set of distances which should be flagged for further visual inspection during plagiarism detection corresponds to what a reasonable statistical test will consider as lower outliers within the full set of distances. As far as the authors are aware, this is the very first application of outlier detection techniques to plagiarism detection despite its natural analogy.

### **3.2.1 Discordancy tests**

Having motivated the equivalence of both problems, the next step is to make use of *discordancy tests*, that is, statistical tests used to detect outliers in a sample.

More to the point, we are interested in designing discordancy test for



two scenarios:

- **Scenario A** We want to detect outliers in the *full ensemble of distances* i.e. in the set  $\{D(i, j)$  for all submissions  $i, j\}$ . The size of this set is  $\binom{n}{2}$ , where  $n$  is the number of submissions.
- **Scenario B** We want to detect outliers in the *individual ensembles of distances*. Fixing the submission  $k$ , we want to detect outliers in  $\{D(k, j)$  for all submissions  $j \neq k\}$ . If the submission corpus has  $n$  submissions, then there are  $n$  different ensembles of distances, one for each submission. The size of each individual ensemble is therefore  $n - 1$ .

Generally, the conditions for finding outliers in scenario A are stronger than the ones in scenario B: outliers in scenario A are outliers with respect to the full ensemble of distances, while outliers in scenario B are outliers only with respect to certain individual ensembles they belong to. The choice between one and the other depends on the level of deterrence desired by the professor.


Once an scenario is chosen, discordancy tests are designed as follows: the  $n - k$  biggest distances  $x_n \geq x_{n-1} \geq \dots, x_{n-k+1}$  are assumed to originate from non-plagiarized pairs of submission with an underlying probability model  $F$ . Meanwhile, the  $k$  smaller  $x_k \geq x_{k-1} \geq \dots \geq x_1$  are assumed to originate from a set of plagiarized pairs with underlying probability model  $G$ . Of course,  $k$  is unknown to the grader.

Two options are available for selecting a discordancy test: *non-parametric* or *parametric* discordancy tests.


## NON-PARAMETRIC DISCORDANCY TESTS

Working with different submissions, students, programming languages, and similarity distances makes it likely to result in many different probability distributions ( $F$ s and  $G$ s in our formulation) to be found. This favors the non-parametric approach, since non-parametric tests automatically generate a specific discordancy test from the data. Non-parametric tests can therefore be expected to have wider applicability.

Non-parametric discordancy test work in the following way: First, they adjust some flexible model (such as gaussian mixtures [44], neural networks [9], or k-means [37]) to the data, penalizing the complexity of the model to avoid overfitting. Once the non-parametric model is obtained, it is possible to calculate the probability of generating data points below a certain value, a probability which can then be used to design discordancy tests. An excellent review of this field can be found in [26].



(a) Algorithm Analysis (AA;  $N=58$  in C) (b) O-O Programming I (OOPI;  $N=108$  in C) (c) Computer Networks II (CNI;  $N=224$  in C)



(d) Operating Systems (OS,  $N=21$  in C) (e) Information Data Structures II (IDSII;  $N=57$  in C) (f) Methodology and Technology of Programming II (MTPPI,  $N=57$  in C)

Figure 1: Figueiredo-Jain kernels for the full ensemble (scenario A) of NCD pairwise distances of several ensembles. Name of dataset, number of submissions and programming language is listed for each.

**Scenario A** Experiments performed indicate that, due to a non-negligible percentage of submissions corresponding to plagiarism, the complexity punishment used when building the non-parametric model may be not powerful enough to leave outliers out of the model. When this observation holds, discordancy tests constructed on top of the contaminated model are condemned to fail to diagnose outliers. This problem is related to the *masking* and *swamping* phenomena in the statistical outliers literature.

To illustrate this, we use the Figueiredo-Jain Gaussian Mixture Model [18]<sup>1</sup>, considered to be one of the best unsupervised learning algorithms, to build non-parametric models for several submissions corpora.

As it is displayed in figure 1, this method wrongly places one of the gaussian kernels on the lower outliers region of several source code submission<sup>2</sup>, while the desired outcome would be to place the kernels only on non-plagiarized data. This is due to the very high incidence of plagiarism found in the corresponding corpora (typically around 5-10%, although incidence rates as high as 40% have been observed in the university where the authors are affiliated).

This unsupervised method considers it worthwhile to dedicate a kernel to model distances corresponding to plagiarized submission, eventually incorporating them into the model. These outliers then *mask* the others, making them undetectable. Preliminary experiments therefore show that non-parametric discordancy tests may not be robust when confronted with

---

<sup>1</sup>A Matlab implementation is available at <http://www.it.lut.fi/project/gmmbayes/downloads/src/gmmbayestb/>

<sup>2</sup>The corpora used throughout experiments were obtained from several graduate courses of the Computer Science and Telecommunications degrees at the Universidad Autónoma de Madrid.

high rates of plagiarism.

However, in cases where plagiarism incidence is very low (not shown here), this method is able to detect outliers correctly.

**Scenario B** Non-parametric model construction requires a moderately large sample size to correctly build a model which represents the probability model of the underlying generation mechanism. While in scenario A we have  $\binom{n}{2}$  such distances, we are restricted to  $n - 1$  in scenario B, a considerably smaller number which could make it dangerous to use non-parametric approaches, with a big probability of data overfitting.

We are aware of submission corpus whose size ranges from 13 to 224. It is clear that with sizes such as 224, the non-parametric model would serve, but not with sizes as small as 13. We are interested in implementing a robust outlier detection method which works for as many situations as possible, and therefore we discard non-parametric models also for scenario B.

#### PARAMETRIC DISCORDANCY TESTS

Due to the detected weakness of non-parametric outlier approaches, the more classical parametric approach has been considered. In this manner, different discordancy tests are specifically designed for each parametric family. Outliers would be therefore considered as outlying respect to such family.

The main issue is to analyze which probability distribution model the distance in absence of plagiarism for scenarios A and B. Once this probability family has been identified, then it is possible to design discordancy tests to

diagnose whether some distance (or set of distances) have been generated by that distribution. In other words, given a sample assumed to be drawn from a particular probability distribution, determine which distances of the sample deviate from the rest in a way which is not explained by the assumed probability model.

**scenario A** Experiments performed on several submissions corpora evidence that the normal distribution is a very good probability model for the distances in absence of plagiarism, i.e. is a good  $F$  for scenario A in our formalism. This is shown in figures 2, 3, 4, and 5, where it succeeds in modeling different corpora written in different languages, and using several similarity distances are used (we test the ones available in AC).

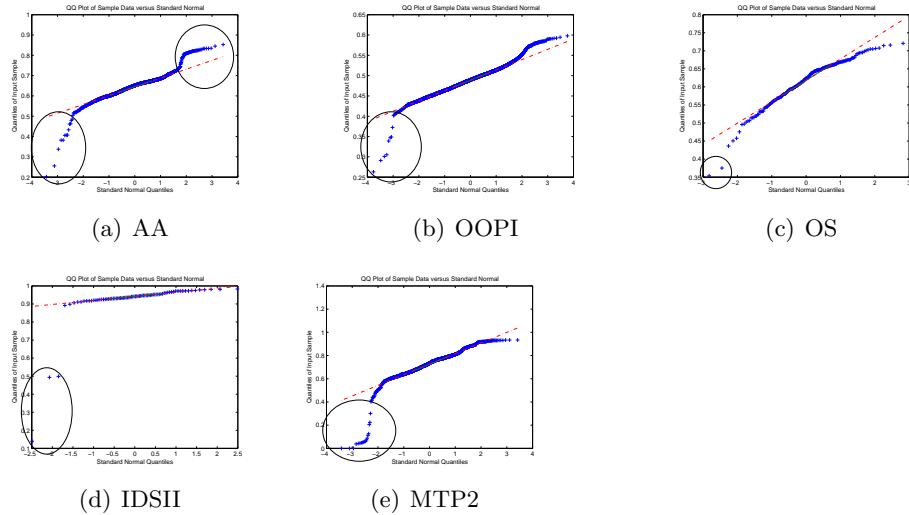


Figure 2: Q-Q plots for the full ensemble of NCD pairwise similarity distances (scenario A). Regions which deviate from normality are surrounded with an ellipse.

The plot used in the figures is known as a *Q-Q plot* (where ‘Q’ stands for




Figure 3: Q-Q plots for full ensemble of NCD-Variance pairwise similarity distances (scenario A).

*quantile*), a graphical tool for diagnosing differences in distributions from normality in a population from which a random sample has been taken. Quantiles of the comparison distribution are displayed on the horizontal axis, and order statistics of the sample are displayed on the the vertical axis. A Q-Q plot of a sample from the normal distribution is expected to approximate a straight line, especially near the center. In the case of substantial deviations from linearity, the statistician would reject the hypothesis of normality.

The overwhelming majority of the points approximate the expected straight line, providing clear support for normality. Only some small regions deviate from this behavior: they are surrounded with an ellipse on the plots.

These circles appear either in the left tail (lowest distances of the en-

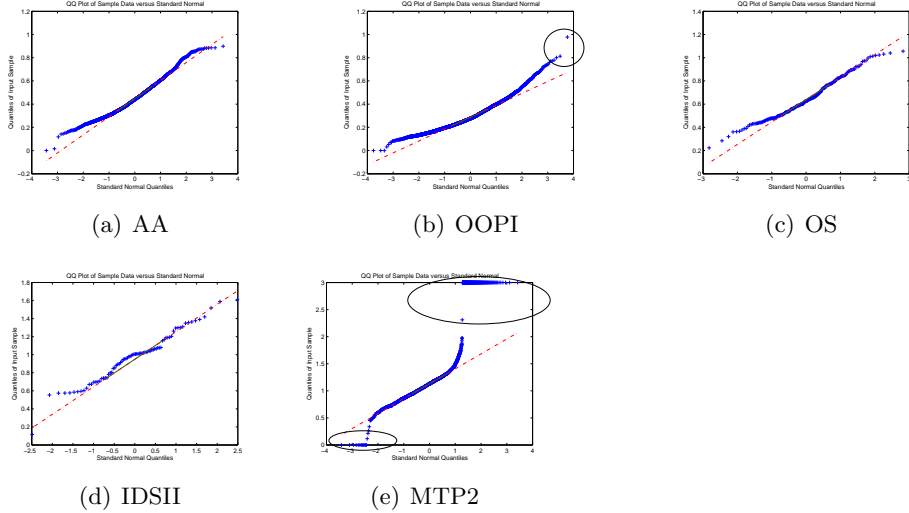


Figure 4: Q-Q plots for the full ensemble of Token counting pairwise similarity distances (scenario A).

semble) or in the right tail (largest ones). A detailed (human) examination of the corresponding submissions reveals the following:

- Left tail deviation is explained by pairwise distances both from submissions which were plagiarized (in this formalism, belong to distribution  $G$ ) and from duplicate submissions by the same authors (which, although not strictly plagiarism, also belong to  $G$ ).
- Right tail deviations are caused by large distances between ‘regular’ submission and ‘anomalous’ submission, or between anomalous submission themselves. This last category includes assignment submissions which contain strange (and often incorrect) answers to the original statement, submission which are incomplete or syntactically incorrect (do not compile) or submission which contain extraneous



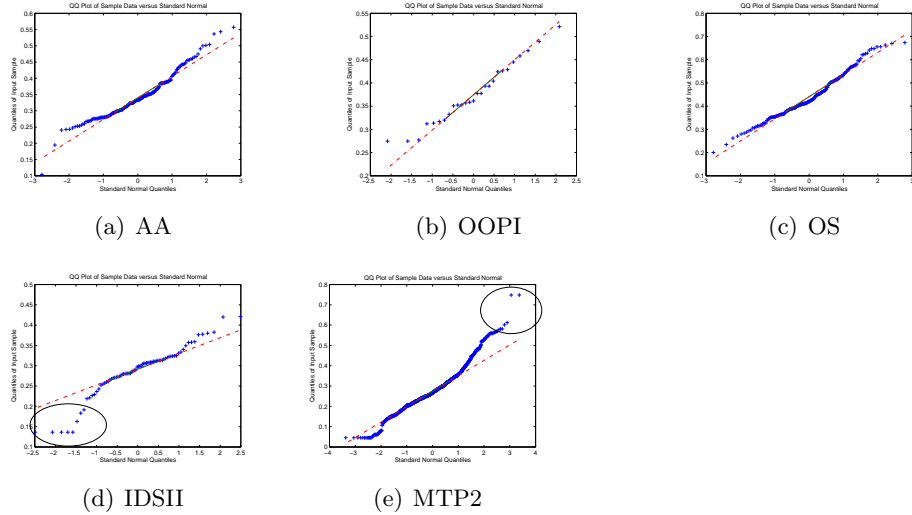


Figure 5: Q-Q plots for the full ensemble of Infrequent Substring pairwise similarity distances (scenario A).

files which were not included in any other. Distances between regular submissions and anomalous ones, and between anomalous submissions themselves tends to be near to 1, forming a right tail cluster.

Therefore we find that, in scenario A, distances are generated by three probability distribution models. The main model, which generates the distance in absence of plagiarism (the majority), is the normal distribution. Other (unknown) probability model generates low distances which lay on the left tail of the normal distribution, and are plagiarized distances. The remaining distribution generates large distances laying on the right tail of the normal distribution, and serves to model pairwise distances in which an anomalous submission is involved.

**scenario B** We analyze scenario B by computing all individual ensemble of distances within a corpus to determine the underlying probability distribution. Experimental results show that distances in this scenario exhibit a behavior analogous to scenario A.

This can be observed in figures 6, 7, 8 and 9, where we plot the distances between a selected submission of some corpus and the rest of the corpus. This behavior is exhibited by the overwhelming majority of submissions, although we display the Q-Q s plot of a only a selected submission of each corpus as a proof of concept.

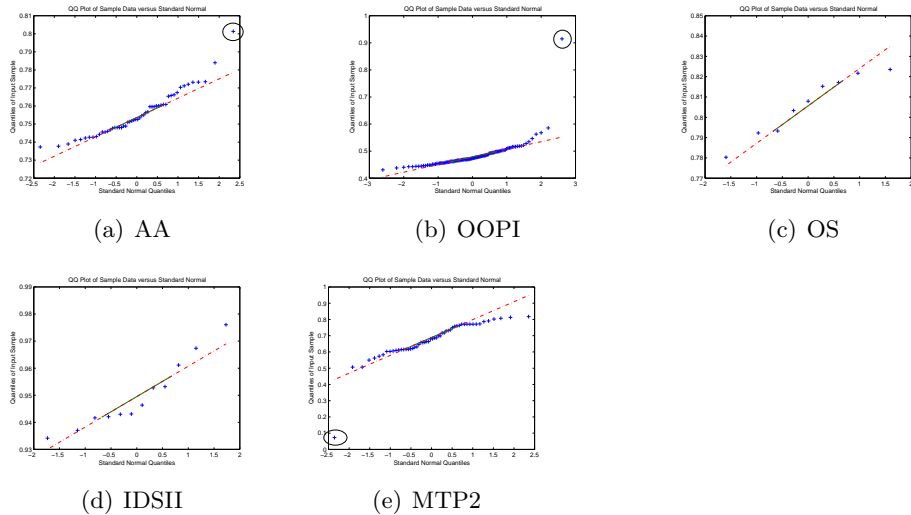


Figure 6: Q-Q plots for NCD similarity distances between a selected member of some corpus and the rest of the corpus (scenario B).

### 3.2.2 The Discordancy test choice: the Hampel identifier

There is strong statistical evidence that the probability distribution for modelling in absence of plagiarism is the normal distribution, and that plagia-

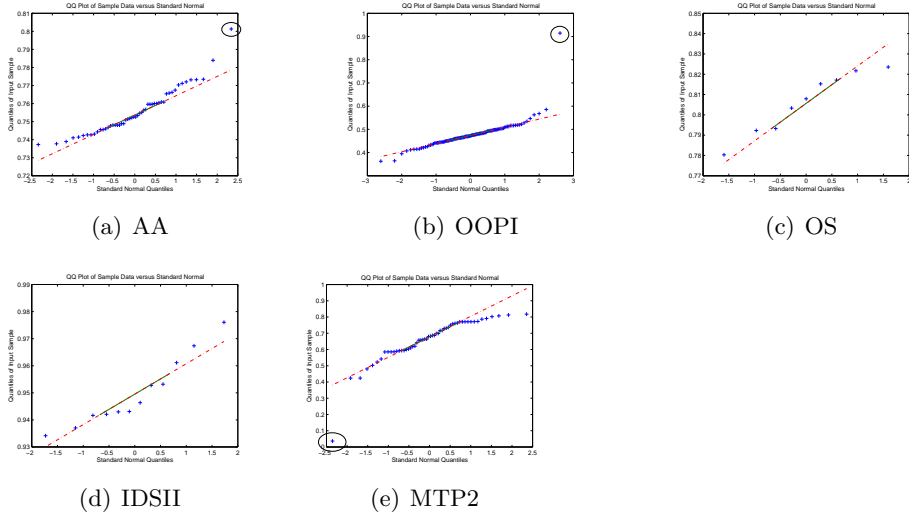


Figure 7: Q-Q plots for NCD-Variance similarity distances between a selected member of some corpus and the rest of the corpus (scenario B).

rized distances tend to be on the left of this distribution, but significantly separated. Plagiarism can then be considered as a generational mechanism which contaminates the normal distribution samples with (a few) low distances.

It is possible to use a discordancy test able to extract the main (normal) tendency of the sample and to generate a model to detect outliers. Distances that are deviations from normality are not a problem if an adequate discordancy test is chosen which is neither influenced by lower outliers (plagiarism) nor from upper outliers (anomalous distances). The normality discordancy test should not be affected by outliers of either sign.

There are several discordancy tests in the literature that fulfill these conditions [5, chap. 6]. The *Hampel identifier*, is maybe the most extensively studied (see [23]). Additionally, there is empirical data proving that

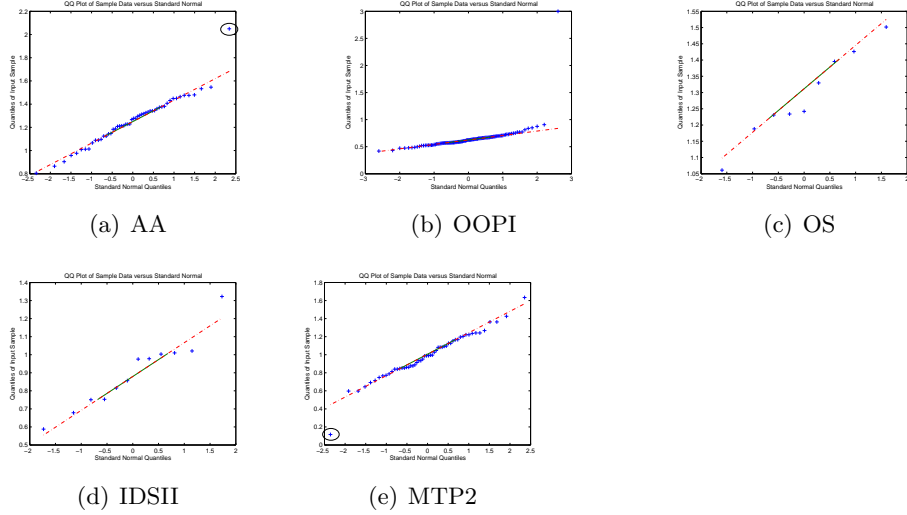


Figure 8: Q-Q plots for Token Counting similarity distances between a selected member of some corpus and the rest of the corpus (scenario B).

it outperforms other tests in many applications [48, 39].

The Hampel identifier works as follows. Let  $X_{(1)}, X_{(2)}, \dots, X_{(N)}$  be the ordered distances  $X_1, X_2, \dots, X_N$ . Let  $M$  be the *median* of the sample, and  $S$  be the *median absolute deviation from the median* of the sample.

The Hampel identifier, adapted to lower outliers, is a rule which identifies as plagiarisms all distances of the sample  $X$  satisfying  $(M - X)/S > g(N, \alpha)$ , where the function  $g(N, \alpha)$  serves for standardizing the identifier in the following way (see [17], p. 783, standardization (4)):

$$\text{Probability(no outliers in the sample)} = \text{Probability} \left( \frac{|X_{(N)} - M|}{S} < g(N, \alpha) \right) = 1 - \alpha. \quad (7)$$

The function  $g(N, \alpha)$  does not have an analytic form and is estimated in AC

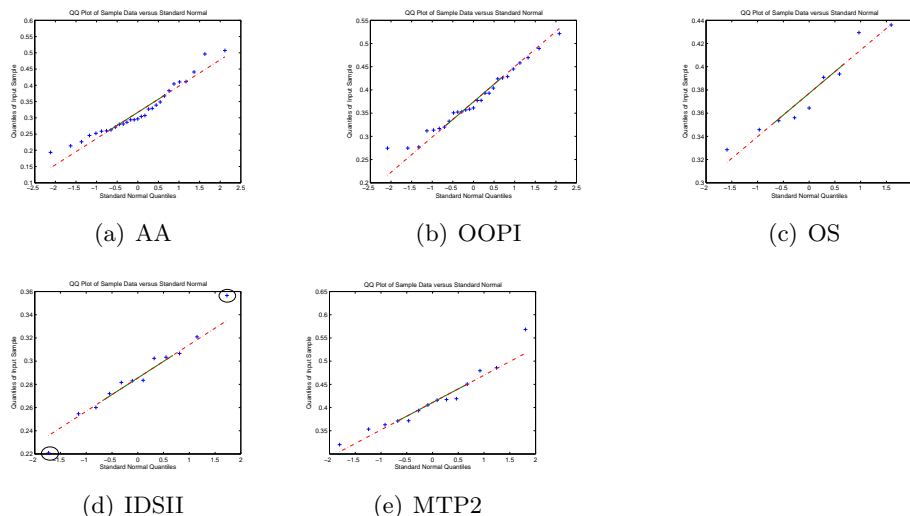


Figure 9: Q-Q plots for Most Infrequent Substring similarity distances between a selected member of some corpus and the rest of the corpus (scenario B).

by means of a Montecarlo simulation<sup>3</sup>.

This method can handle a large number of outliers, and is resistant to the problems appreciated in the non-parametric approach. We have incorporated it into both the and Graph+Population visualization (scenario A, see figure 10) and the Individual visualization Histograms (scenario B, see figure 11). We use the thresholds  $\alpha = 0.01$  and  $\alpha = 0.05$ , which are the most used in the literature.

More experimental work is needed to determine optimal thresholding values ( $\alpha$  values), and how this is affected by aspects such as the programming language, the size of source code or the type of submission in the corpus.

<sup>3</sup>A Matlab implementation to locate critical values of the Hampel identifier ( $g(N, \alpha)$ ) has been made available at <http://www.mathworks.com/matlabcentral/fileexchange/loadFile.do?objectId=14769>

Future lines in this research are detailed in section 5.

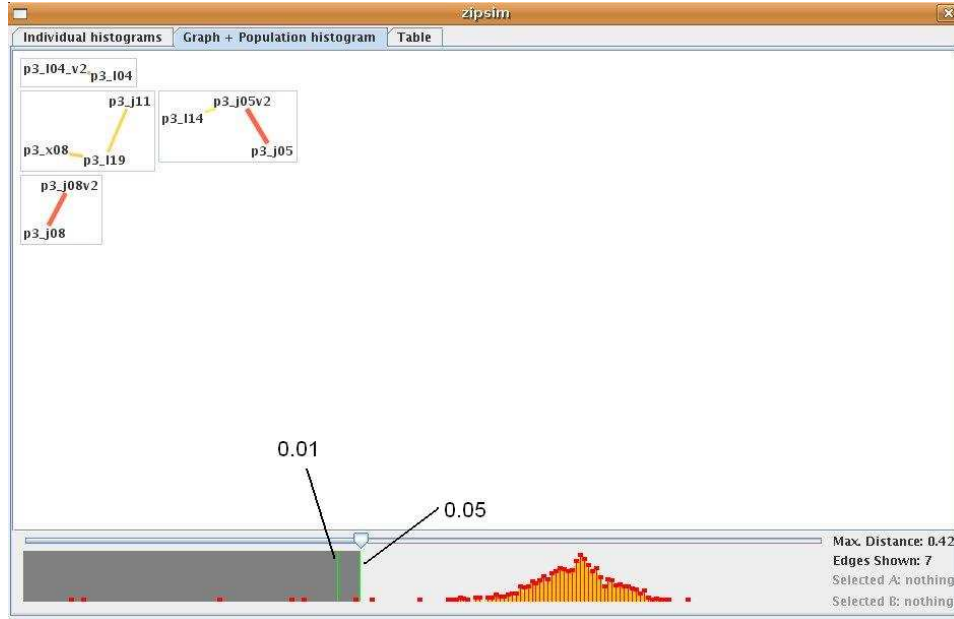


Figure 10: Use of Hampel identifier to recommend two thresholds ( $\alpha = 0.05$  and  $0.01$ ) for plagiarism (outlier) detection in the Graph+Histogram visualization (scenario A).

### 3.3 Submission filtering

Since the goal of plagiarism detection in an academic context is to assist a human grader, it is important to make this assistance as extensive as possible. The first step when performing any type of automated plagiarism detection is to prepare a set of submissions for analysis. All plagiarism detection systems rely on a stringent organization of sources in order to perform the analysis. This is understandable, since accurate results require comparing apples to apples, and the introduction of “noise” can only result in distorted results. The task of converting assignment submissions into

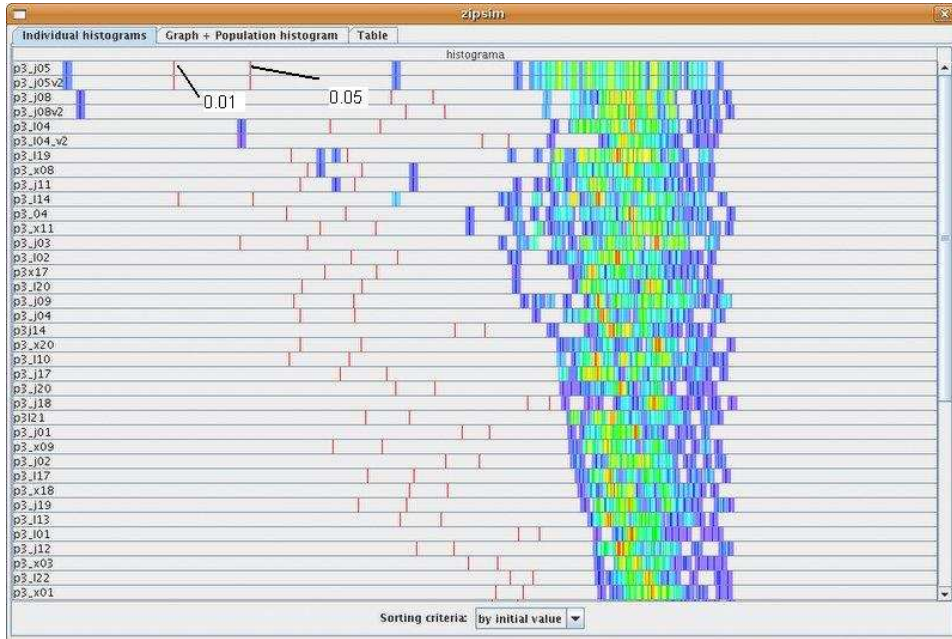


Figure 11: Use of Hampel identifier to recommend two thresholds ( $\alpha = 0.05$  and  $0.01$ ) for each Individual Histogram visualization (scenario B).

the internal, standardized input format expected by the plagiarism detection tool will be referred to as *submission filtering*. This step can be partially automated with suitable shell scripts, but such scripts are difficult to code without extensive experience, debugging them may be a time-consuming process, and they cannot be reused outside of the platform and institution for which they were originally developed. AC offers an easy to use front-end that automates the selection, extraction and sanitizing of submissions from their original turn-in formats into AC's internal submission input format.

The internal format used by AC to perform plagiarism detection is a one-level hierarchy, where folder names represent submission ids and each folder contains all the files of that submission that should be included in

the comparison. For instance, if submissions identified as SubA, SubB, and SubC are to be compared, AC would expect a folder structure similar to the following:

SubA/: `functions.c`, `main.c`

SubB/: `program.c`, `vector.c`, `matrix.c`

SubC/: `my_submission.c`

This section describes the submissions extraction utility included in AC to ease the submission filtering task, which can be divided into two main phases: assignment submission selection, and submission content selection.

### **Assignment submission selection**

During *submission selection*, the grader is expected to select submissions to be compared from their locations on the harddisk, networked drive, or any other type of storage. Currently, AC can access any filesystem visible to the underlying operating system. The user is expected to add one or more filesystem folders to the leftmost pane of fig. 12. Folders added in this way can now be browsed (using a familiar folder-tree metaphor). Additionally, files within folders can be browsed, and compressed files (termed *archives*) are automatically expanded (in-memory) to the folders that would result if they were to be extracted. AC currently handles *.zip*, *.rar* (optional, requires an external freeware program), *.tar* and *.tar.gz* archive formats. Note that after initial extraction, no further distinction is made between archives that have been expanded in-memory and normal folders: both behave exactly



alike. For instance, in fig. 12, the selection of `P1A14.zip` (an archive) and `P1A08` (a folder) in the left-hand tree result in undistinguishable file trees in the right-hand tree.

The goal of submission selection is to locate the “base folders” (or archives) of the submissions that should be compared. These base folders should then be passed on to the rightmost pane of Fig. 12, where they can be pruned of any extraneous content or subfolders, or otherwise sanitized before comparison. Two methods are available to perform this task: direct selection and boolean condition-based filtering. Direct selection of submissions involves highlighting the target submissions with the mouse (similar to other selecting operations in most GUI file dialogues), and clicking on the ‘>>’ button. However, if the folders are mixed with large quantities of non-relevant folders, this may be a lengthy manual operation.

### **Filters to automate selection**

In condition-based filtering, a boolean expression is used to determine what folders should be passed to the “submission contents” pane. Atomic conditions can match either a file or folder’s name, its complete path, or its contents (relevant only to files), and rely on Java’s built-in regular expression matching capabilities. Since Java’s regular expressions differ slightly from Perl’s or other programming languages, and teachers should not be required to learn a new regular expression language, helpful hints and examples are provided as tooltips where applicable. Composite expressions allow an expression hierarchy to be built, by aggregating several atomic or composite sub-expressions under a common heading. This can be ei-

ther ‘all conditions are met’ (logical *and*), ‘at least one condition is met’ (logical *or*), or ‘no conditions is met’ (logical *nor*). Any sub-expression’s effects may be tested at any time by clicking on the ‘test’ button (immediately highlighting the results), and sub-expressions may be edited, added and removed at any time. Once the user is confident in the expression, clicking on ‘Confirm’ will apply the top-level expression to the entire pane, and select all results for insertion in the *Assignment Contents* pane.

From the point of view of plagiarism detection, AC’s automated extraction tool is a key feature, and the authors consider its lack a serious drawback of other systems. From the point of view of user interface design, we believe that this is an intuitive and powerful graphical approach to nested boolean queries.

### **Submission content selection**

The purpose of the Submission Contents pane (leftmost pane of Fig. 12) is twofold: to flatten all subdirectories present in submissions (AC expects to work on a single-level hierarchy), and to remove all files that, although present in the submission, are of no interest for plagiarism detection. Subdirectory flattening is performed automatically (again, in-memory, without writing anything to disk). Uninteresting file removal can be automated via filtering, with the same semantics as described above: this is the purpose of the lower filtering block in the center of Fig. 12. File removal can also be performed manually, by selecting files and clicking on the ‘-’ icon.

Since file names may be misleading, AC’s extraction tool allows teachers

to double-click on any file to see its contents, and expressions in the submission content selection filter can select files based on their actual contents (instead of relying on their names or paths). Furthermore, the built-in tokenizer may choke on certain malformed expressions. When viewing a file, it is also possible to perform in-memory edits, a useful feature when minor syntactical errors are present.

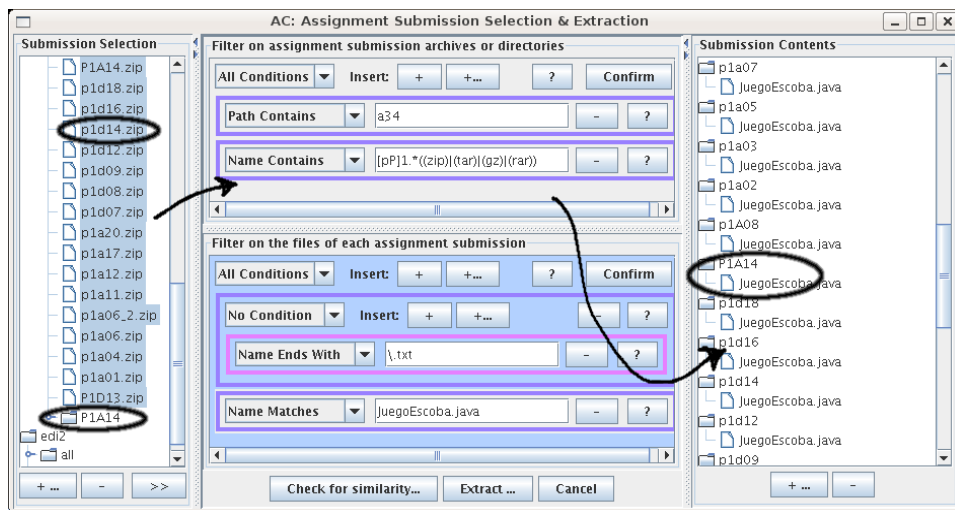


Figure 12: Screenshot of the filtering interface

### 3.4 Visualization

Once a test has been performed, users are expected to inspect the resulting distance matrix and take appropriate action if evidence of plagiarism is found. A simple approach would be to locate the lowest distances (ie.: greatest similarities) and manually inspect the relevant submissions. This approach can find the most flagrant cases, but identification of possible plagiarism patterns between submissions from a large matrix of numbers requires patient

and tedious analysis. Quoting R. W. Hamming, “the purpose of computing is insight – not numbers”. In this case, the desired insight is the location of cases of plagiarism.

Several graphical, interactive representations of test results (referred to as *visualizations*) have been developed to ease this task, and their effectiveness is one of the strongest points of AC. Visualization is also critical in allowing users to assess the meaningfulness of the outputs obtained from each test: plagiarism detection shares many characteristics with outlier detection, where outliers are only such within the context of a broader distribution.

Three distinct visualizations have been developed. The simplest, *distance list* (fig. 14) presents a list of all distances, with columns for “submission A”, “submission B” and “Distance”. The list supports sorting by any of its columns with a single click, and is intended for the common operations of finding the smallest distances and locating the distance values of a given submission. Double-clicking on any row presents a side-by-side comparison of the source code of the row’s submissions.

The input for any particular visualization is a simple distance matrix, and therefore individual visualizations can be used to analyze the results of any test, regardless of its actual implementation. These visualizations could actually be used to analyze any distance matrix, even if it were unrelated to programming submission plagiarism detection.

### 3.4.1 Graph representation

Graphs are useful to visually scan for “islands” (*connected components*): groups of submissions that are similar to each other, yet less similar to

others outside this group. It can be a valuable aid to locate and characterize groups of students that have collaborated with each other. In the case of a submission  $A$  that appears to have plagiarized from sources  $B$  and  $C$ , a graph display can quickly answer the question of whether  $B$  is also similar to  $C$  or not.

A screen capture of the graph visualization is presented in fig. 15. This is a graph where vertices represent submissions, and an edge is included for each similarity lower than a given threshold. Vertices without edges are omitted from the display, and for clarity, edges that are deemed redundant are also elided. Edge colors and widths are determined by the degree of similarity between their endpoints: red, thick edges denote high similarity, while green, thin edges are used to convey low similarity.

Clicking on any graph vertex results in the associated submission's source code being displayed. Clicking on any edge launches a side-by-side comparison of both of the submissions it is connected to.

The threshold below which graphs edges are not included is set with a slider, which ranges from 0 (only edges which correspond to "exact copies" are shown) to 1 (all edges are included). To aid the teacher in selecting a good threshold, the slider is placed above a histogram of the frequencies of each distance in the test's distance matrix. The shaded part of the histogram of fig. 15 represents the portion of edges that are currently being used in the graph.

Not all edges below the threshold are included, because this would result in an unnecessarily cluttered graph. Edges that do not belong to each connected component's *minimum spanning tree* (MST) can frequently be

removed without significant loss of information. This approach, suggested by Whale [47], results in a drastic simplification of each connected component, and speeds layout considerably. However, removing all edges not in the MST risks the loss of edges that have very high relevance (but result unnecessary in the construction of the MST), while other edges with lower relevance (but a role within the MST) are included. This can lead to a false impression of the component's structure. To characterize these edges that are important but do not belong to the MST, the following approach has been adopted: if a component has  $n$  vertices, the lowest-distance  $n$  edges are also retained, regardless of whether they belong to the MST or not. This additional heuristic ensures that structures such as triangles are not lost, with a minimal increase in graph complexity.

Graph generation and layout is delegated to CLOVER [19], a graph visualization framework developed by one of the authors. The framework provides fast automatic layout, even for large graphs (tested to several hundred vertices and tens of thousands of edges). CLOVER can also visualize clustered graphs (graphs with a superimposed hierarchical clustering, where clusters can be abridged and represented as single vertices). Beyond plagiarism detection, a clustered graph representation could prove very useful to analyze broad trends in similarity: do the main connected components reflect different approaches to the problem statement? What do these approaches differ in?. However, this research line has not yet been pursued.

### 3.4.2 Histogram visualization

A third type of visualization is represented in fig. 16. This view presents rows of “individual histograms”, generated for each of the submissions. While the general histogram used for distance threshold selection within the graph visualization displays the frequency of all distances found in the distance matrix, an individual histogram only displays distances for a particular row of the matrix. That is, the individual histogram for submission  $A$  would be generated from  $\mathbf{D}_A$ , the set of distances from  $A$  to all other submissions. In the display of fig. 16, each histogram contains a header which identifies the submission for which it was generated.

By default, individual histograms within the histogram visualization are displayed in a condensed format where color coding, instead of bar height, is used to convey distance frequency. We have termed this abbreviated histogram a *hue histogram*; it is reminiscent of the colored bars found in chemistry analysis for similar purposes [31]. Within hue histograms, blue is used to represent low frequency, and red for high frequency; intermediate hues (in traditional rainbow order) are used for the remaining values. The use of hue histograms allows large space savings to be achieved, and enables easy comparison of neighboring histograms.

#### Reading Histogram Visualizations

The histogram display is perhaps the most cryptic at first, but it is also the most powerful, as it allows an overview of the distribution of similarities within each individual which can be very enlightening for a trained eye.

A key observation is that similarity “spikes” to the left of the similarity distribution in a given row are likely to correspond to cases of plagiarism. For example, if a submission  $A$  is very similar to another submission  $B$ , and a large gap exists between  $\mathbf{D}_{AB}$  and the rest of the distances  $\{\mathbf{D}_{AX_i} : X_i \neq B\}$ , this will show up as a left-most spike in  $A$ ’s individual histogram. This distribution would be exactly what would be expected if  $A$  and  $B$  have common code which is not shared with any other submission. Imagine that  $A$  decides to obscure this fact, and obfuscates code in an effort to avoid detection. This will result in an increase of  $\mathbf{D}_{AB}$  – but, if the distance test used to compute the distances is robust, all other distances within  $\mathbf{D}_A$  will also experiment increases, and the general shape of the histogram will be retained<sup>4</sup>. Therefore, “damning spikes” (outliers to the left of a distribution) can be used as an indicator of plagiarism.

Conversely, a smooth individual histogram without gaps in the leftmost part of the distance distribution is an indicator that no plagiarism has taken place, since this is the expected distribution when similarity is only due to a large number of independent programming decisions.

### **Interaction with Histograms**

Interaction with the histogram is simple. Selecting or unselecting any set of rows toggles between bar histogram (selected) or hue histogram (unselected) display for the affected rows. Hovering the mouse over a histogram brings up a tooltip which displays the names of the submissions that are nearest to

---

<sup>4</sup>This observation has been used when developing the “Variance similarity” (see section ??) distance test, which attempts to correct for possible skew whenever left-most spikes are detected in a distribution.



the current similarity for the current row (since quantization is performed to generate histograms, several submissions may share the “same” similarity). Double-clicking on a row brings up a display of the source files of the submission that corresponds to that row.

The rows of the histogram visualization can be sorted by smallest distances, by gap or by submission ID. Smallest-distances is the default ordering, and represents the histograms which have the smallest distances on top. This ordering is also useful to locate large gaps and for visual inspection, since two rows for submissions  $A$  and  $B$  which share an identical “smallest distance”  $d_s$  probably do so as a result of  $\mathbf{D}_{AB} = d_s$ . If they are indeed very similar, then their histograms can be expected to match too; this is trivial to inspect by matching colors in hue histograms (and may be easier than matching hills and valleys in a bar histogram representation). Sorting by gap uses a heuristic to estimate leftmost-gap sizes, and the results are used to sort the histograms. Finally, sorting by submission ID is self-explanatory.

## 4 Practical examples

This section illustrates how a grader would analyze a programming submission corpus with AC. An on-line tutorial on AC can be found at

<http://tangow.eps.uam.es/ac/>

## 4.1 Artificial submissions

The first example corresponds to a synthetic dataset. In this dataset we have full knowledge of which submissions in the corpus have been plagiarized. In a real-life setting, such a corpus can only be obtained if students have been previously asked to follow a specific behavior (otherwise, carefully disguised plagiarized submissions may slip through). Creating such a set of data is an expensive and time-consuming task, and raises certain ethical concerns, since teaching students how to plagiarize is generally frowned upon. However, the technique presented in [10] is a reasonable substitute: artificial sets of submissions are automatically generated by means of evolutionary computing, and include cases of plagiarism planned in advance.

This benchmark is available online at

<http://www.ii.uam.es/~mcebrian/plagiarism-benchmark/>

It is composed of 44 programs in *APL* that seek to implement the function  $\cos(\log x)$ . Thirty of the 44 programs are *original* programs, and were generated by totally independent evolutionary runs with different random seeds. Six of them are ‘mutational plagiarized programs’, obtained by applying the mutation operator to one of the original programs. Finally, eight of them are ‘recombination plagiarized programs’, which have been generated by merging two original programs by means of crossover evolutionary operators. To facilitate the task of identifying the source codes, the different generated programs have been labeled as follows: P1, P2, ... , P30.

Mutational-plagiarized programs use labels in the form  $MPx$  (where  $Px$  is the original source), and recombination-plagiarized  $PxRGP_y$  or  $PxRFP_y$ , where  $Px$  and  $P_y$  represent the original programs identifiers. A graph representing the relations and structure of this corpus is presented in figure 13.

Given this set of submissions, we run an un-tokenized *Zip* similarity distance, which produces results which can be analyzed using any of the visualizations discussed in Section 3. The first visualization corresponds to the row data (see fig. 14). This output is not very helpful to find suspects of plagiarism, since single values have no meaning for our purposes, unless they are examined under the perspective of the whole population.

The second visualization is depicted in figure 15. The graph on the top of the figure represents relations between submissions; the histogram below the graph displays a shaded portion, which is controlled by a slider. The graph only displays distances that fall within this shaded portion, allowing a grader to concentrate only on the shortest distances between submissions.

The interaction between slider and graph is detailed in section 3.4. The initial value for the threshold is obtained from the Hampel identifier, but the slider can be freely adjusted by the user. In this example, the recommendation of the Hampel identifier (0.33) will be followed (with 95% significance), as displayed in figure 15. All mutational-plagiarized programs (MP5, MP10, MP15, MP20, MP25 and MP30), and recombination-plagiarized (P10RGP5, P14RFP7, P15RGP20, P15RFP20, P1RFP30, P20RGP15, P22RFP5 and P5RGP10) are displayed. The only original submission that appear, although they have clearly not been copied or plagiarized, are P13, P14 and

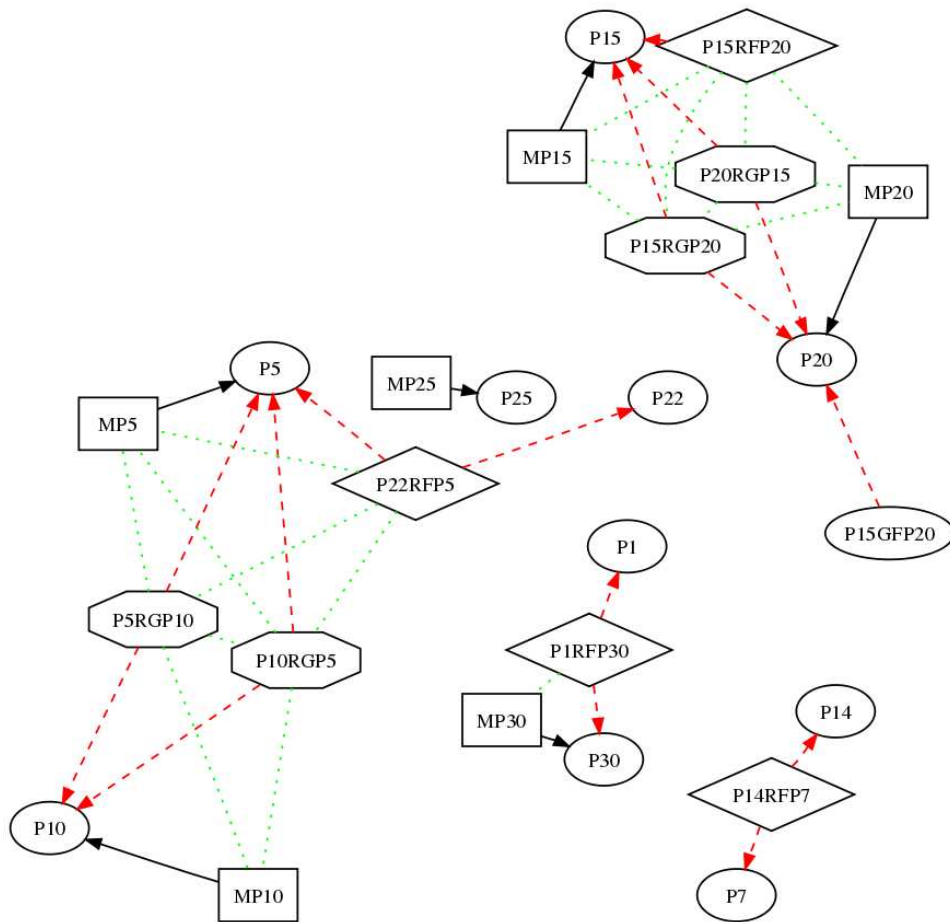


Figure 13: Plagiarism dependence of the artificial benchmark. Round vertices stand for original submissions, boxes for plagiarism using a single source, romboids and octagons for the two different types of plagiarism that use two sources. A black solid line between two submissions denotes that one is a plagiarism of the other, a red dashed lines denotes that one is one of the sources for plagiarism of the other and a green dotted line denotes that they are indirect copies, i.e. they share a common source of plagiarism.

Individual histograms	Graph + Population histogram	Table
Distance	One	The other
0.021857923	P30.APL	MP30.APL
0.062992126	P1RFP30.APL	P1.APL
0.11627907	P7.APL	P14RFP7.APL
0.15272728	P5.APL	MP5.APL
0.18978103	P15.APL	MP15.APL
0.1941392	P10.APL	MP10.APL
0.2037037	P20.APL	P15RFP20.APL
0.21311475	P30.APL	P18.APL
0.21851853	P10RGP5.APL	P10.APL
0.21857923	P18.APL	MP30.APL
0.227758	P15RGP20.APL	P15.APL
0.22797927	P4.APL	P13.APL
0.2287234	P13.APL	MP30.APL
0.23076923	P5RGP10.APL	P5.APL
0.23316061	P4.APL	P18.APL
0.23404256	P18.APL	P13.APL
0.23404256	P30.APL	P13.APL
0.24522293	P25.APL	MP25.APL
0.24870466	P4.APL	MP30.APL
0.24870466	P4.APL	P30.APL
0.24915825	P20RGP15.APL	P20.APL
0.2507965	P15RGP20.APL	MP15.APL

Figure 14: Test results visualized as a distance table.

P18. It must be noted that two submissions can be very similar by arbitrary reasons, specially if the assignment is small and simple as in this case with the function  $\cos(\log x)$ .

Although the graph visualization is already very useful to detect suspicious submissions, AC provides a third output, depicted in figure 16, which can improve our analysis. This output corresponds to the *hue histograms* described in section 3.4. The histogram for each row corresponds to distances between the row's submission and all others. This individual view provides a different perspective, from which new suspects can be found and others dismissed.

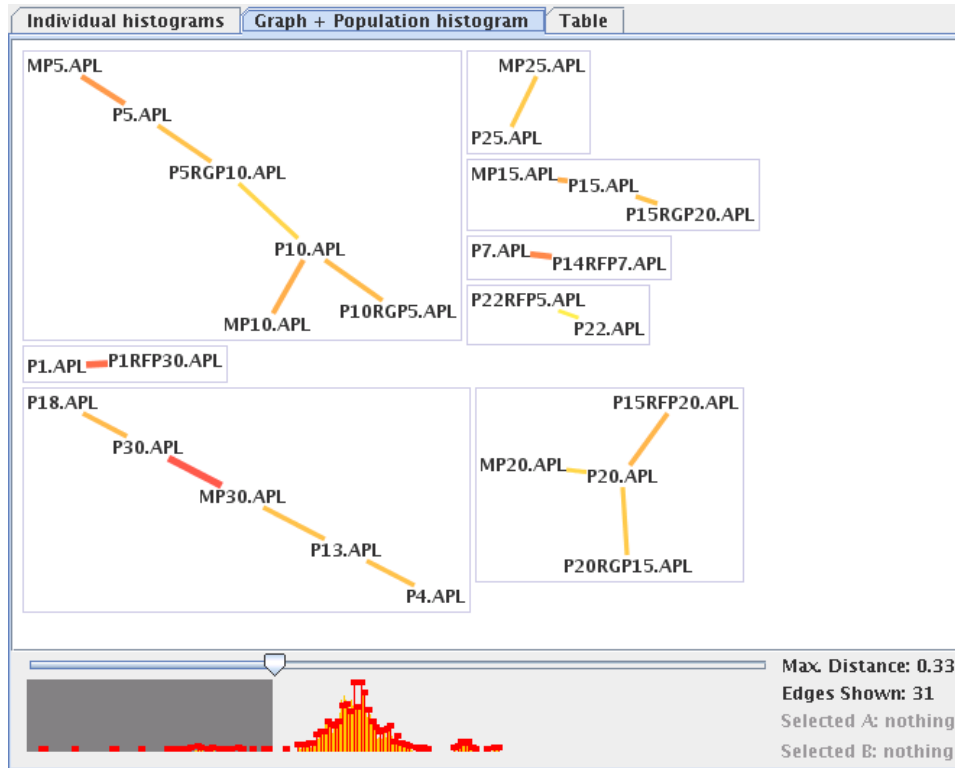


Figure 15: Test results visualized as a graph, with a histogram reflecting the frequency of each distance (ranging from 0, most similar, to 1) and a horizontal slider, used to select the maximum distance that is used for inclusion in the graph: only pairs of submissions with a distance lower than this threshold are included.

Certain conspicuous cases can be easily located. For instance, the lowest similarity distance value within the P25 histogram corresponds to its plagiarized sibling MP25; and is also clearly *outside* the distribution of other distances. The same applies to P15, whose plagiarized siblings MP15 and P15RGP20 also exhibit low values. Large gaps are found within the individual histogram for P25 between MP25 and all other distances. The same large gap can be found in P15's individual histogram between distances for

plagiarized siblings (MP15, P15RGP20) and those corresponding to all other programs. The meaning and relevance of these gaps is discussed in section 3.4.2.

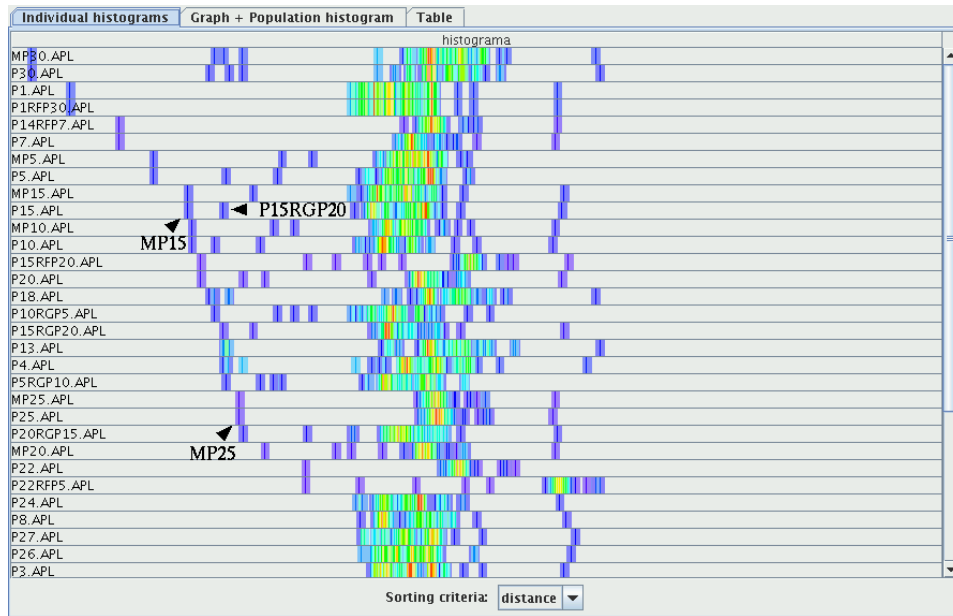


Figure 16: Test results visualized as individual histograms. Each row represents a color-coded histogram (blue is low, red is high) of the frequency with which other submissions have presented a given similarity to this one. Unexpected gaps in the leftmost side of the histogram suggest existence of plagiarism

## 4.2 Real submissions

This example follows the same steps as the previous one. However, the data to analyze corresponds to a real programming submission corpus: the Object Oriented Programming I corpus used in section, with 108 submissions written in JAVA. This corpus, as in many other real-world ones, has duplicate submissions from the same authors, due to last-minute bug-fixes. These can

be identified by the name of the submissions that share the same ID prefix, usually followed by a “v2” to denote a second version.

In the *Graph and Population histogram* tab (see fig. 17), the Hampel identifier (with 95% significance) used as a starting threshold can be seen to be highly accurate. All duplicate versions of the corpus are identified, together with two clearly plagiarized submissions. In the individual histograms view (fig. 18), further suspects of plagiarism can be identified.

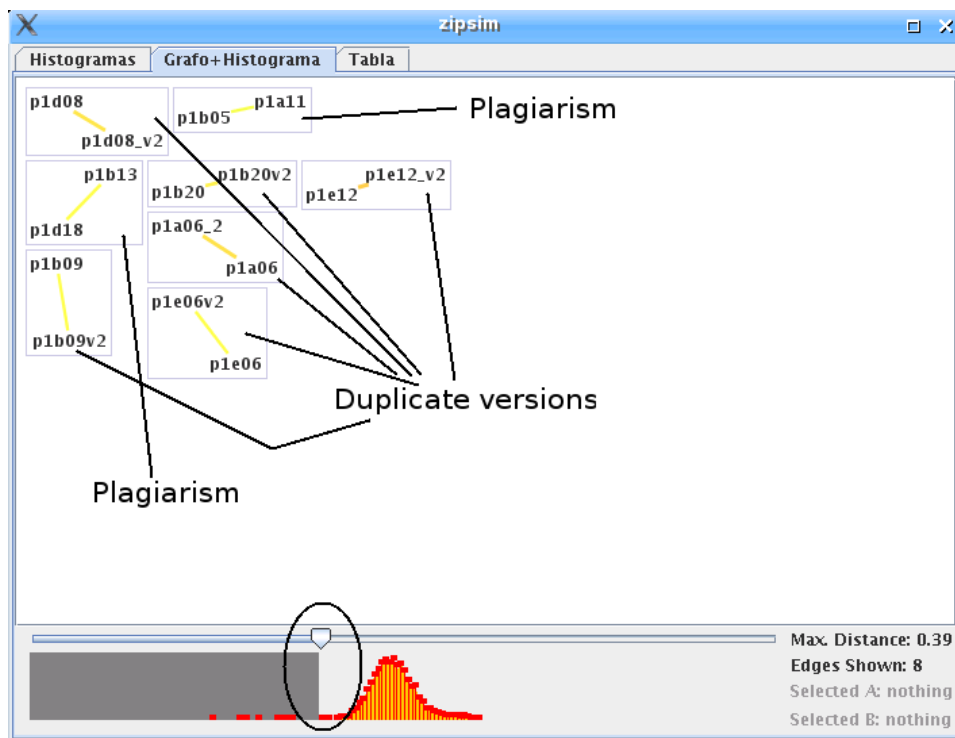


Figure 17: The population histogram and graph view for the Object Oriented Programming I corpus (108 submissions in Java).

The interpretation of individual histograms is again based on locating surprisingly low distance values when compared to an individual submis-



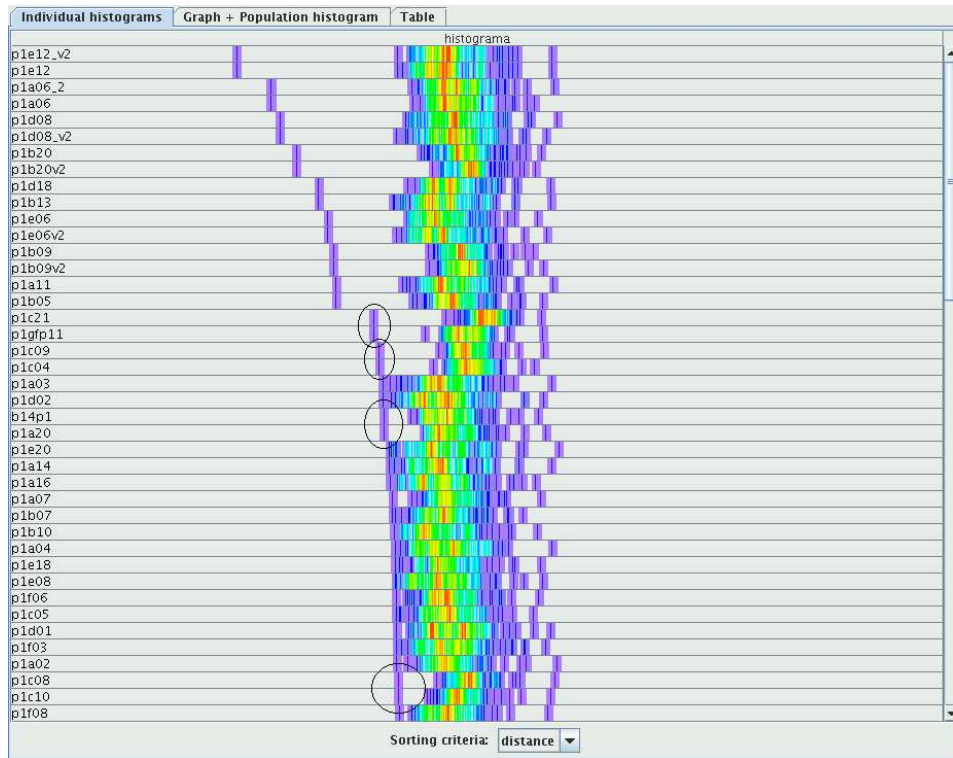


Figure 18: The individual histograms view for the Object Oriented Programming I.

submission’s distribution. For instance, distances from submission P1D18 to all other submissions are mostly in the range 0.4 to 0.6 (according to the *Zip* similarity distance algorithm). However, an isolated value falls clearly below this bracket, corresponding to P1B13. This indicates that, with high probability, similarities between P1D18 and P1B13 can be attributed to plagiarism. Visual inspection reveals that this is indeed the case.

Note that individual histograms are more valuable for outlier detection than the full histogram displayed in the graph visualization. Isolated “left” values found in individual histograms may be unusually low for the indi-

vidual histogram, but may be high enough not to arouse suspicion within the population histogram. Distances that are clearly outliers within individual distributions may therefore be average within the general distance population; and, as noted in section ??, can be seen as a natural result of obfuscation efforts during plagiarism in an effort to avoid detection. Although the individual histogram visualization is more powerful than the graph visualization, it is also more time-consuming. Each visualization has its own strengths and weaknesses.

In the example, there are 4 couples that present clear gaps in the left tails of their individual distributions, even though these distances were not identified as outliers within the full population histogram. These couples are B14P1-P1A20, P1C08-P1C10, P1C21-P1GFP11 and P1C09-P1C04. After a visual inspection of the similarities between their sources (an example of visual source comparison can be found in fig. 19), P1C21-P1GFP11 was found to be a clear case of plagiarism, P1C08-P1C10 was considered a false positive, and B14P1-P1A20 and P1C09-P1C04 were considered as probable partial plagiarism.

## 5 Conclusions and further work

The problem of plagiarism detection is a difficult one. The frontier between, on one side, random similarity or simple inspiration from others' work and, on the other side, blind cut+paste plagiarism is not clear-cut, and certain cases will always require a human grader to distinguish between what is acceptable and what is not. However, different algorithms and heuristics

```

// **
// * Devuelve los puntos de este jugador en la ronda actual. Se debe llamar
// * desde <code>JuegoEs</code> cuando se detecta final de ronda, para
// * actualizar los totales de cada jugador.
// * @param i n mero del jugador cuya puntuaci n se desea consultar
// * @return la puntuaci n de ese jugador en esta ronda
// **
private int puntosEnRonda(int i) {
    // A RELLENAR:
    // calcular aqui los puntos en ronda del jugador 'i'; se parte de 0, y se su
    // 1 punto para el el/los que mas sietes tenga(n)
    // 1 punto para el/los que mas oros tenga(n)
    // 1 punto para el/los que mas cartas tenga(n)
    // 1 punto por cada escoba conseguida
    // 1 punto para el que tenga el 7 de oros

    int NumSietes[] = new int[NumJugadores];
    int NumOros[] = new int[NumJugadores];
    int NumCartas[] = new int[NumJugadores];
    int Jug_max_sietes=0, Jug_max_oros=0, Jug_max_cartas=0, Jug_siete_oros=0;
    int max_sietes=0, max_oros=0, max_cartas=0;
    int contador = 0;
    int puntosJug=0;
    String carta_a_analizar;

    // Calculo para cada Jugador, los Oros, las Cartas, los Sietes, y quien tiene
    puntosJug = (int)datosJugador[i].getPuntos();
    for(int j=0; j< NumJugadores; j++)
    {
        NumCartas[j] = datosJugador[j].getEnBazas().size();
        for(int k=0; k < datosJugador[j].getEnBazas().size(); k++)
        {
            carta_a_analizar = datosJugador[j].getEnBazas().cartaEn(k).getId();

            if((carta_a_analizar.compareTo("10") == 0) || (carta_a_analizar.com
            (carta_a_analizar.compareTo("30") == 0) || (carta_a_analizar.com
            (carta_a_analizar.compareTo("50") == 0) || (carta_a_analizar.com
            (carta_a_analizar.compareTo("70") == 0) || (carta_a_analizar.com
            (carta_a_analizar.compareTo("00") == 0) || (carta_a_analizar.com

                NumOros[j]++;

            if((carta_a_analizar.compareTo("70") == 0) || (carta_a_analizar.com
            (carta_a_analizar.compareTo("78") == 0) || (carta_a_analizar.com
    }
}

```

```

// **
// * Devuelve los puntos de este jugador en la ronda actual. Se debe llamar
// * desde <code>JuegoEs</code> cuando se detecta final de ronda, para
// * actualizar los totales de cada jugador.
// * @param i n mero del jugador cuya puntuaci n se desea consultar
// * @return la puntuaci n de ese jugador en esta ronda
// **
private int puntosEnRonda(int i) {
    // A RELLENAR:
    // calcular aqui los puntos en ronda del jugador 'i'; se parte de 0, y se
    // 1 punto para el el/los que mas sietes tenga(n)
    // 1 punto para el/los que mas oros tenga(n)
    // 1 punto para el/los que mas cartas tenga(n)
    // 1 punto por cada escoba conseguida
    // 1 punto para el que tenga el 7 de oros

    int apuntos=0;
    int cont=0;
    int max7=0; /*numero maximo de 7's de un jugador*/
    int max0=0; /*numero maximo de oros de un jugador*/
    int maxC=0; /*numero maximo de cartas de un jugador*/
    String acarta;

    //Inicializamos las tablas*/
    int numCartasJugadores[] = new int[NumJugadores];
    int num7Jugadores[] = new int[NumJugadores];
    int numOrosJugadores[] = new int[NumJugadores];

    //puntos del jugador*/
    //tambien podemos sacar los puntos del jugador utilizando
    unicamente getPuntos(), ya que esta definida en esta clase**/
    apuntos = datosJugador[i].getPuntos();

    //numero de cartas de cada jugador*/
    for (int j=0; j < NumJugadores; j++) {
        cont = 0;
        /*numero de cartas que se ha llevado en bazas el jugador*/
        numCartasJugadores[j] = datosJugador[j].getEnBazas().size();
        for(int p=0; p < datosJugador[j].getEnBazas().size()
        /*tomamos cada una de las cartas del mazo */
        acarta = datosJugador[j].getEnBazas().carta
        /*traemos los 7 que tiene*/
        if ((ncarta.compareTo("70")==0) || (ncarta.
        || acarta.compareTo("78")==0) {
            num7Jugadores[j]++;
        }
    }
}

```

Figure 19: Visual comparison of two submissions

can be used to identify suspects of blatant plagiarism and flag the more complex cases, greatly simplifying the grader’s task.

This work has presented AC, a plagiarism detection tool which also doubles as a framework for research into source code plagiarism detection. AC offers many improvements over other tools described in current literature: the use of rich visualization greatly simplifies the task of analyzing the result of similarity tests; its stand-alone, cross-platform implementation does not raise privacy concerns found in web-based systems; and preparation of assignment submissions for automated plagiarism detection, overlooked by many systems, can be automated with a graphical user interface.

AC is released as open source software, and full implementation details are available to any interested parties. Lack of comparable information for

other plagiarism detection tools makes their results difficult to replicate or improve. Furthermore, the availability of AC under an open source license guarantees that other researchers can follow the project and participate in its development. Due to the modular design of the program, it is easy to integrate new similarity distance algorithms, allowing graders or researchers to compare or complement their performance. Many of these algorithms have been already incorporated into AC, including two novel approaches. Finally, AC includes the first application of outlier statistical methods to plagiarism detection, proving fast initial identification of suspicious submissions. The analysis of related experiments has revealed interesting aspects of the patterns found in typical submission corpora.

### **Further work**

The use of statistical methods has opened different lines for further research. Greater insight into the *Hampel identifier* threshold choice could be obtained from a student controlled experiment or by further work on an ongoing project where submission corpora are artificially developed [10]. A second line of research is concerned with the surprisingly high accuracy of the normal distribution in outlier identification, even when confronted with different corpora and different similarity distance algorithms.

Once two submissions have been deemed to be “very similar” to each other, a human grader is currently expected to visually compare both for evidence of plagiarism. In other systems, similar fragments from both submissions are highlighted for side-by-side inspection, but AC currently lacks this feature. The *Substring* similarity distance test is a good candidate to

identify such areas. Further refinement of the test itself and an extension to facilitate identification of *hot* areas during visual inspection is pending.

Although source code plagiarism detection is a relatively veteran field of research, few systems have undergone experimental validation. A typical approximation is to use a corpus of already-graded set of submissions (where cases of plagiarism were manually identified) and compare these cases to those found using the tool to be tested. However, both the original grader and the tool may fail to correctly identify all cases of plagiarism; and false positives are also possible. A better approach would require asking a group of students to write plagiarized versions of randomly selected submissions within a carefully selected corpus where no plagiarism had occurred. This would allow experiments to be performed with a fully annotated corpus. An alternative and less labor-intensive approach is to use automatic programming to generate artificial sets of submissions; initial steps using this method have been described in [10]. Whatever the approach, validation will continue to be one of the main lines of inquiry.

## Acknowledgments

This work was supported by grant TSI 2005-08255-C07-06 of the Spanish Ministry of Education and Science. We would like to firstly thank K. Koroutchev for his seminal linux script implementation of AC. We would also like to thank A. Suárez, G. Martínez, J.R. Dorronsoro, M. Alfonseca, and specially P. Paalanen and S. Roberts for his technical help in the statistic analysis design of AC. Additional thanks to P. Haya, L. Shafti and R. Moriyon for providing us with real submission data sets. Finally, we would like to thank the Computer Science Department of the Universidad Autónoma de Madrid for their feedback in using the tool.

## References

- [1] JISC Legal Information Service. <http://www.jisclegal.ac.uk/>. 11
- [2] Spanish Agency for Data Protection. <https://www.agpd.es>. 11
- [3] A. Aiken et al. Moss: A system for detecting software plagiarism. *University of California–Berkeley*. See [www.cs.berkeley.edu/aiken/moss.html](http://www.cs.berkeley.edu/aiken/moss.html), 2005. 9
- [4] Ricardo A. Baeza-Yates and Berthier A. Ribeiro-Neto. *Modern Information Retrieval*. ACM Press / Addison-Wesley, 1999. 18
- [5] V. Barnett and T. Lewis. *Outliers in statistical data*. Wiley New York, 1994. 24, 35
- [6] KW Bowyer and LO Hall. Experience using MOSS to detect cheating on programming assignments. *Frontiers in Education Conference, 1999. FIE'99. 29th Annual*, 3, 1999. 9
- [7] B.F. Braumoeller and B.J. Gaines. Actions Do Speak Louder than Words: Deterring Plagiarism with the Use of Plagiarism-Detection Software. *PS: Political Science and Politics*, 34(04):835–839, 2002. 4
- [8] J. Bull, C. Collins, E. Coughlin, S. Developer, D. Sharp, and P. Square. Technical review of plagiarism detection software report. *CAA University of Luton* [http://online.northumbria.ac.uk/faculties/art/information\\_studies/Imri/Jiscpas/docs/jisc/luton.pdf](http://online.northumbria.ac.uk/faculties/art/information_studies/Imri/Jiscpas/docs/jisc/luton.pdf). 3

- [9] T. Caudell and D. Newman. An adaptive resonance architecture to define normality and detect novelties in time series and databases. *Proc. IEEE World Congress on Neural Networks*, pages 166–176, 1993. [26](#)
- [10] M. Cebrián, M. Alfonseca, and A. Ortega. Automatic Generation of Benchmarks for Plagiarism Detection Tools using Grammatical Evolution. In *Proceedings of the 9th annual conference on Genetic and Evolutionary Computation*. ACM Press New York, NY, USA, 2007, arXiv:cs.NE/0703134. [22](#), [50](#), [60](#), [61](#)
- [11] X. Chen, B. Francia, M. Li, B. McKinnon, and A. Seker. Shared information and program plagiarism detection. *Information Theory, IEEE Transactions on*, 50(7):1545–1551, 2004. [9](#), [17](#)
- [12] R. Cilibrasi and PMB Vitani. Clustering by Compression. *Information Theory, IEEE Transactions on*, 51(4):1523–1545, 2005. [16](#)
- [13] J. Clare. Computer plagiarism threatens the value of degrees. *Daily Telegraph*, 3(7):2000, 2000. [3](#)
- [14] F. Culwin, A. MacLeod, and T. Lancaster. Source Code Plagiarism in UK HE Computing Schools, Issues, Attitudes and Tools. *South Bank University Technical Report SBUCISM-01-02*, 2001. [3](#), [10](#), [13](#)
- [15] F. Culwin and J. Naylor. Pragmatic Anti-Plagiarism. *Proceedings 3rd All Ireland*, 1995. [11](#)



- [16] P. Cunningham and A.N. Mikoyan. *Using CBR Techniques to Detect Plagiarism in Computing Assignments*. Trinity College, Dept. of Computer Science, 1993. [11](#)
- [17] L. Davies and U. Gather. The Identification of Multiple Outliers. *Journal of the American Statistical Association*, 88(423):782–792, 1993. [36](#)
- [18] MAF Figueiredo and AK Jain. Unsupervised learning of finite mixture models. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 24(3):381–396, 2002. [28](#)
- [19] Manuel Freire. *An Approach to the Visualization of Adaptive Hypermedia Structures and other Small-World Networks based on Hierarchically Clustered Graphs*. PhD thesis, September 2007. [46](#)
- [20] J. Giles. Preprint analysis quantifies scientific plagiarism. *Nature* 444, pages 524–525, 2006. [3](#)
- [21] D. Gitchell and N. Tran. Sim: a utility for detecting similarity in computer programs. *Technical Symposium on Computer Science Education*, pages 266–270, 1999. [9](#)
- [22] D. Grune and M.H. Vakgroep. Detecting copied submissions in computer science workshops. *Informatica Faculteit Wiskunde & Informatica, Vrije Universiteit*, 1989. [9](#)
- [23] F.R. Hampel. The Breakdown Points of the Mean Combined with Some Rejection Rules. *Technometrics*, 27(2):95–107, 1985. [35](#)

- [24] J.K. Harris. Plagiarism in computer science courses. *Proceedings of the conference on Ethics in the computer age*, pages 133–135, 1994. [3](#)
- [25] G. Hertel, S. Niedner, and S. Herrmann. Motivation of software developers in Open Source projects: an Internet-based survey of contributors to the Linux kernel. *Research Policy*, 32(7):1159–1177, 2003. [12](#)
- [26] V. Hodge and J. Austin. A Survey of Outlier Detection Methodologies. *Artificial Intelligence Review*, 22(2):85–126, 2004. [26](#)
- [27] R. Irving. Plagiarism Detection: Experiences and Issues. *JISC Fifth Information Strategies Conference, Focus on Access and Security*, 2000. [3](#), [11](#)
- [28] E.L. Jones. Metrics based plagiarism monitoring. *Proceedings of the Sixth Annual CCSC Northeastern Conference, Middlebury, Vermont*, pages 1–8, 2001. [11](#)
- [29] E.L. Jones. Plagiarism monitoring and detection-towards an open discussion. *Proceedings of the seventh annual consortium for computing in small colleges central plains conference on The journal of computing in small colleges*, pages 229–236, 2001. [4](#), [11](#)
- [30] M. Joy and M. Luck. Plagiarism in Programming Assignments. *IEEE TRANSACTIONS ON EDUCATION*, 42(2):129, 1999. [11](#), [18](#)
- [31] Robert Kincaid and Heidi Lam. Line graph explorer: scalable display of line graphs using focus+context. In Augusto Celentano, editor, *AVI*, pages 404–411. ACM Press, 2006. [47](#)

- [32] T. Lancaster and F. Culwin. A Comparison of Source Code Plagiarism Detection Engines. *Computer Science Education*, 14(2):101–112, 2004. [10](#)
- [33] Thomas Lancaster and Fintan Culwin. Towards an error free plagiarism detection process. In *ITiCSE '01: Proceedings of the 6th annual conference on Innovation and technology in computer science education*, pages 57–60, New York, NY, USA, 2001. ACM Press. [9](#)
- [34] B. Martin. Plagiarism: a misplaced emphasis. *Journal of Information Ethics*, 3(2):36–47, 1994. [4](#)
- [35] Rodrigo Meza. DetectaCopias. Available at <http://www.dcc.uchile.cl/~rmeza/proyectos/detectaCopias/index.html>. [11](#)
- [36] Donald R. Morrison. Patricia: Practical algorithm to retrieve information coded in alphanumeric. *J. ACM*, 15(4):514–534, 1968. [22](#)
- [37] A. NAIRAC. A System for the Analysis of Jet Engine Vibration Data. *Integrated Computer-Aided Engineering*, 6(1):53–66, 1999. [26](#)
- [38] A. Parker and JO Hamblen. Computer algorithms for plagiarism detection. *Education, IEEE Transactions on*, 32(2):94–99, 1989. [7](#)
- [39] RK Pearson. Outliers in process modeling and identification. *Control Systems Technology, IEEE Transactions on*, 10(1):55–63, 2002. [36](#)
- [40] J.L. Popyack, N. Herrmann, P. Zoski, B. Char, C. Cera, and R.N. Lass. Academic dishonesty in a high-tech environment. *ACM SIGCSE Bulletin*, 35(1):357–358, 2003. [13](#)

- [41] L. Prechelt. . Mahpohl, and M. Philippsen. Jplag: Finding Plagiarisms among a set of Programs. Technical report, Technical Report 2000-1, Universitat Karlsruhe, March 2000. [10](#)
- [42] L. Prechelt, G. Malpohl, and M. Philippsen. Finding plagiarisms among a set of programs with JPlag. *Journal of Universal Computer Science*, 8(11):1016–1038, 2002. [9](#), [10](#)
- [43] R.L. Ribler and M. Abrams. Using Visualization to Detect Plagiarism in Computer Science Classes. *coordinates*, 5:6. [13](#)
- [44] SJ Roberts. Novelty detection using extreme value statistics. *Vision, Image and Signal Processing, IEE Proceedings-*, 146(3):124–129, 1999. [26](#)
- [45] S. Saxon. Comparison of plagiarism detection techniques applied to student code, Part II. *Computer Science project, Trinity College, Cambridge*, 2000. [11](#), [17](#)
- [46] S. Schleimer, D.S. Wilkerson, and A. Aiken. Winnowing: local algorithms for document fingerprinting. *Proceedings of the 2003 ACM SIGMOD international conference on on Management of data*, pages 76–85, 2003. [7](#), [12](#)
- [47] G. Whale. Identification of Program Similarity in Large Populations. *The Computer Journal*, 33(2):140, 1990. [9](#), [46](#)
- [48] R.R. Wilcox. *Applying contemporary statistical techniques*. Academic Press, 2003. [36](#)

- [49] M.J. Wise. Detection of similarities in student programs: YAP'ing may be preferable to plague'ing. *ACM SIGCSE Bulletin*, 24(1):268–271, 1992. [3](#)
- [50] J. Ziv and A. Lempel. A universal algorithm for sequential data compression. *Information Theory, IEEE Transactions on*, 23(3):337–343, 1977. [17](#)