Digital Commons @ George Fox University

Faculty Publications - Department of Electrical Engineering and Computer Science

Department of Electrical Engineering and Computer Science

10-2018

Plagiarism Detection Avoidance Methods and Countermeasures

Brian R. Snider
George Fox University, bsnider@georgefox.edu

Follow this and additional works at: https://digitalcommons.georgefox.edu/eecs_fac

Part of the Systems and Communications Commons

Recommended Citation

Snider, Brian R., "Plagiarism Detection Avoidance Methods and Countermeasures" (2018). *Faculty Publications - Department of Electrical Engineering and Computer Science*. 24.
https://digitalcommons.georgefox.edu/eecs_fac/24

This Article is brought to you for free and open access by the Department of Electrical Engineering and Computer Science at Digital Commons @ George Fox University. It has been accepted for inclusion in Faculty Publications - Department of Electrical Engineering and Computer Science by an authorized administrator of Digital Commons @ George Fox University. For more information, please contact arolfe@georgefox.edu.

PLAGIARISM DETECTION AVOIDANCE METHODS AND COUNTERMEASURES *

Daniel Rohwedder and Brian R. Snider
Department of Computer Science and Electrical Engineering
George Fox University
Newberg, Oregon 97132
drohwedder16@georgefox.edu, bsnider@georgefox.edu

ABSTRACT

Plagiarism is a major problem that educators face in the information age. Today's plagiarist has a near limitless supply of well-written articles via the internet. Due to the scale of the problem, detecting plagiarism has now become the domain of the computer scientist rather than the educator. With the use of computers, documents can be conveniently scanned into a plagiarism detection system that references public web pages, academic journals, and even previous students' papers, acting as an "all-seeing eye."

However, plagiarists can overcome these digital content detection systems with the use of clever masking and substitutions techniques. These systems cost universities tens of thousands of dollars, and also infringe upon intellectual property ownership rights without the informed consent of individual students. In this work, we examine the efficacy of commercial plagiarism detection systems when used against some selected masking techniques, and then present a simple countermeasure to combat the aforementioned detection avoidance technique.

INTRODUCTION

Masking plagiarism through digital means allows students to submit plagiarized documents that bypass a detection system, silently threatening academic integrity. In

Simkin and McLeod's ethical study "Why Do College Students Cheat?", they found staggering statistics from a survey of students taking a 300-level business course: "Of the 144 respondents, 87 (60%) reported that they had cheated an average of 6.1 times" during the course [14]. Although plagiarism is just one aspect of academic dishonesty, with the rise of easily accessible digital work it is only becoming more convenient and effective to plagiarize coursework.

While most instructors and institutions use commercial tools like Turnitin [15] (the premier detection system), the cost and legal challenges surrounding the integrity of student intellectual property in for-profit systems present a need for alternative detection methods. For example, given an estimation that each student produces 30 pages of writing per semester which must be tested for plagiarism, the cost for Plagium to check this work would be roughly $9,600 per semester per 1,000 students (that is, $0.08 per page 1 × 4 courses × 30 pages per course × 1,000 students = $9,600; although, in practice, a "deep scan" of a single-page document on Plagium costs anywhere from $0.08 to $0.33.) [11]. Turnitin's costs are hidden from public view, requiring a quote from a salesperson, and that quote is not usually shared. An anecdotal report of Turnitin's pricing as published by Times Higher Education quotes a University of Glasgow employee as stating the base cost as roughly £25,000 ($39,000) per year [10], a shockingly high cost for an automated system. The employee noted significant price hikes on two consecutive academic years, 31 and then 13%, for 2013-2014 and 2014-2015, respectively.

Turnitin's automated system stores and reuses previously-checked student papers to cross-reference new submissions, as a guard against student's copying previously submitted papers. This database-centered approach is common in commercial plagiarism detection, but Purdy identifies those systems as a possible abuse of students' intellectual property [12]. To effectively use those papers, Turnitin gains a license to submitted papers, even after students are finished using the service [16]. However, because this license may change at any time without prior notice to users, students' privacy may be at risk. Unless students or institutions request that specific documents be deleted, those documents will continue to stay in the system, used for any later-defined services, possibly including selling essays to the highest bidder, who may use them to infer personal information or to sell intellectual property of students.

Background

Current research revolves around lexical analysis rather than understanding and combating masking techniques. Established algorithms and systems include citation-based analysis, machine learning-based paraphrase recognition, and writing style analysis. Gipp's citation-based plagiarism detection approach tries to find "co-citations," other documents which cite the same documents as the subject of current analysis [3]. This approach allows for pattern analysis using citations, instead of the traditional lexical symbols and patterns, possibly detecting paraphrased documents. Chitra and Anupriya's machine-learning paraphrase recognizer also provides good accuracy (81% correct paraphrase detection) on realistic test corpora, including one built from articles written by the UK Press Association, and some derived newspaper articles [2].

Oberreuter and Velasquez's artificial intelligence-based text mining system brings a unique approach to detecting possible plagiarism without the traditional use of student-populated databases [9]. By examining changes in word choice, even classic books can be easily and automatically analyzed by an algorithm, visualizing changes in vocabulary throughout the work. One application of such a system was used to determine the authorship of classic writings, ranging from religious texts to plays, revealing significant changes in style without a database to store similar texts [8].

There is extensive research on optical character recognition (OCR) tools, which can recognize text in a fairly noisy environment, according to Nell [7]; however, those tools do not appear to have been employed yet in commercial systems like Turnitin. Without using OCR, it is impossible for a quick text-dump to catch all plagiarism masking techniques, leaving enterprising plagiarists an unguarded vector to exploit in established systems.

Project Goals

This exploratory work has two goals: 1) to identify detection avoidance methods which allow plagiarized content to escape detection in systems such as Turnitin or Plagium, and 2) to test the efficacy of countermeasures to those methods. An obvious first category of detection avoidance methods are masking methods, which may be as simple as one-to-one character replacements-white-colored non-whitespace characters replacing spaces, near-visually-identical ASCII-Unicode character replacements, white-colored filler text in the margins, or some combination of these techniques.

Masking may be more sophisticated: methods like Microsoft Word macros, Portable Document Format (PDF) script injections, or even exploiting Dynamic Data Exchange (DDE, recently used to execute malware according to ICT Monitor Worldwide [5]) within rich office suite documents still are yet to be explored in this work. Countermeasures to the basic techniques typically manifest as simple scripts, relying on external tools like ImageMagick and Google's OCR-based Tesseract to convert masked text back into plain-text representations of what the document appears to be [13, 14].

METHODOLOGY

We tested the effectiveness of three detection avoidance techniques with three different commercial services. First, we identified three widely-advertised and readily-available commercial services that offer plagiarism detection functionality to academic institutions: Turnitin, Plagium, and CopyLeaks. Next, we identified three straightforward-to-implement detection avoidance techniques based on a review of online student-facing resources and personal interactions with students connected with our institution. Then, we created a completely- and blatantly-plagiarized word processing document using LibreOffice by deliberately copying and pasting from Wikipedia's "Plagiarism" article [17]. We submitted this plagiarized document to each of the three commercial services and recorded the percent-similarity measure for each service.

After testing the basic efficacy of each service with unmodified text content, we then identified and implemented three detection avoidance techniques. For our first technique, we simply masked selected portions of the text by replacing whitespace characters with white-colored non-whitespace characters (i.e., with a white-colored letter "i"). We refer to this technique as our "whitespace replacement" technique.

For our second technique, we wrote a script that used ImageMagick's pdfimages command-line utility to capture a Portable Network Graphics (PNG) format image file of each page of the plagiarized document, and inserted the image in place of the page's text content via drag-and-drop. Crucially, we added white-colored procedurally-generated random-topic filler text using a 2-point-font to the document margins so that the detection systems would not reject the document as containing no textual content. We took care to ensure that the filler text was of correct spelling but did not guarantee correct English syntax or semantics. We refer to this technique as our "image foreground with margin filler" technique.

Finally, for our third technique, we replicated our second technique but used higher-quality screenshots of each page manually taken using *xfce4-screenshooter* and set the image as the background image of each page, and included the margin filler text here as well. We refer to this technique as our "image background with margin filler" technique. We then proceeded to submit each modified variant of the document to each of the three commercial services, and recorded the percent-similarity measure for each. To test a possible detection avoidance countermeasure, we ran the image foreground and background with filler variants through a script utilizing ImageMagick and Tesseract to extract the original text content out of the page image. We again submitted these preprocessed versions to each of the three systems and recorded the corresponding percent-similarity measures.

RESULTS

Table 1 lists percent-similarity measures for documents employing detection avoidance techniques in three existing commercial plagiarism detection systems. Note that image foreground and background with filler techniques easily avoid detection, with the commercial systems reporting 0% similarity to prior works. The simple whitespace replacement technique also completely avoids detection by two of the three tested systems.

Avoidance Technique	Similarity (%) by System		
	Turnitin	Plagium	CopyLeaks
Unmodified text	100%	100%	100%
Whitespace replacement	100%	0%	0%
Image foreground	0%	0%	0%
Image background with margin filler	0%	0%	0%

Table 1: Percent-similarity measures for detection avoidance techniques by system

However, after processing the documents through our preprocessing script, we observed much higher percent-similarity measures for these same documents, despite the OCR accuracy errors evident in the preprocessed output. We report the percent-similarity measures for documents employing an image background detection avoidance technique without and with preprocessing in Table 2. We found that all three commercial systems were able to correctly identify the document as plagiarized, despite such abysmal initial performance.

Image background with margin filler	Similarity (%) by System		
	Turnitin	Plagium	CopyLeaks
Without preprocessing	0%	0%	0%
With preprocessing	96%	83%	97%

Table 2: Percent-similarity measures for the image background with margin filler technique? without and with OCR-based preprocessing by system

CONCLUSIONS

We found it trivially easy to beat these expensive, established systems using straightforward and easily-reproducible detection avoidance techniques. With the image-background or image-insertion methods reaching 0% of the visual text detected as similar to prior work, the efficacy of these systems when met with masking techniques must be questioned*: is it really worth the money and intellectual property risks for institutions to use these systems?*

However, there are few alternatives, Turnitin seemingly dominating the market with more than 15,000 participating institutions worldwide [15], it is the clear option for universities, high-schools, and middle schools. Even worse, Turnitin offers a service called WriteCheck [18] for students to test their own papers against 3 their system, charging $6.00-8.00 per paper scanned. While this may be a helpful service to some students who wish to double-check their work, nefarious students may utilize this system as a pre-screen to Turnitin itself, allowing plagiarism to be accurately masked.

Ironically, a significant benefit of using Turnitin is its large database of original content created by students. If Turnitin did not already have millions of submissions, which are retained indefinitely, it would have far fewer detections on resubmissions of older papers, allowing an ill-intentioned bad actor to sell prior work to current students. Because of the success of Turnitin, it becomes more of a "must-have" for universities, further improving itself and distancing itself from competitors. While it is not officially a monopoly, Turnitin's success has snowballed, eliminating much-needed competition and allowing for the previously mentioned price hikes.

Future Work

Independent algorithms and methods of detecting plagiarism are valuable in providing an alternative to services like Turnitin, as proposed by Rife [13]. If a database-focused system like Plagium or Turnitin infringes on students' intellectual properties too much, an algorithm like Oberreuter and Velasquez's could still be used to operate only on a single paper. Such a tool could be extremely effective without infringing heavily on students' intellectual property rights, only caching the papers which are being checked at the time, and not persisting and growing a massive database of all submitted papers.

Butakov suggests moving plagiarism detection systems to become on-campus services to protect student's intellectual property, allowing for copyright configurations to be determined by the hosting institution [1]. However, to make this idea a reality, systems must be hardened against common masking techniques, otherwise university-based systems could be rendered ineffective to a simple "find and replace." If this concept were fully realized, several institutions could provide different services-one hosting a database detection scheme, another a machine learning-based solution, perhaps even an implementation of Oberreuter and Velasquez's system, all offering different levels of intellectual property rights protection for students. This competition could influence the industry to prioritize students' intellectual property, dissuade providing students tools to beat the system, and cut down costs to be more in line with what it might cost to get this service from another university. Instead of spending thousands for the dubious privilege of furthering the efficacy of a third-party closed system which can and does abuse student intellectual property rights, institutions should instead opt to further the development of an open, accessible, and effective plagiarism detection system created and maintained by those who actively seek to maintain the sanctity and integrity of academic pursuit, while reducing the cost burden of higher education on our students.

REFERENCES

[1] Butakov, S., Dyagilev, V., Tskhay, A., Protecting students' intellectual property in the web plagiarism detection process, *International Review of Research in Open and Distance Learning*, 13 (5), 1-19, 2012.

[2] Chitra A., Rajkumar, A., Plagiarism detection using machine learning-based paraphrase recognizer, *Journal of Intelligent Systems*, 25 (3), 351-359, http://search.proquest.com/docview/1861880867, 2016.

[3] Gipp, B., *Citation-based Plagiarism Detection Detecting Disguised and Cross-language Plagiarism using Citation Pattern Analysis*, Wiesbaden: Springer Fachmedien Wiesbaden, 2014.

[4] Google, Google Tesseract, https://opensource.google.com/projects/tesseract, 2018.

[5] ICT Monitor Worldwide, Attacks exploit Microsoft Dynamic Data exchange protocol, http://search.proquest.com/docview/1965410027, 2017.

[6] ImageMagick Studio LLC, Convert, edit, or compose bitmap images @ ImageMagick, http://www.imagemagick.org/script/index.php, 2018.

[7] Nell, H., Quantifying the noise tolerance of the OCR engine Tesseract using a simulated environment, Student Thesis, Blekinge Institute of Technology, 2014.

[8] Oakes, M. P., Literary detective work on the computer, *Natural Language Processing*, Amsterdam, Netherlands; Philadelphia, Pennsylvania: John Benjamins Publishing Company, 2014.

[9] Oberreuter G., Velasquez, J. D., Text mining applied to plagiarism detection: the use of words for detecting deviations in the writing style, *Expert Systems With Applications*, 40 (9), 2013.

[10] Parr, C., Some turn away from Turnitin over price and reliability, https://www.timeshighereducation.com/news/some-turn-away-from-turnitin-over-price-and-reliability/2014856.article, 2015.

[11] Plagium, Plagium: plagiarism checker & plagiarism detection, http://www.plagium.com, 2018.

[12] Purdy, J. P., Anxiety and the archive: understanding plagiarism detection services as digital archives, *Computers and Composition*, 26 (2), 2009.

[13] Rife, M. C., Copy(write): intellectual property in the writing classroom, *Perspectives on Writing*, Fort Collins, Colorado; Anderson, South Carolina: The WAC Clearinghouse: Parlor Press, 2011.

[14] Simkin, M. G., McLeod, A., Why do college students cheat?*, Journal of Business Ethics*, 94 (3), 441-453, https://doi.org/10.1007/s10551-009-0275-x, 2010.

[15] Turnitin, LLC, Turnitin: technology to improve student writing, http://turnitin.com, 2018.

[16] Turnitin, LLC, Turnitin end-user license agreement, https://turnitin.com/agreement.asp, 2018.

[17] Wikipedia, Plagiarism, https://en.wikipedia.org/wiki/Plagiarism, 2018.

[18] Turnitin, LLC, Plagiarism Checker: WriteCheck by Turnitin, http://en.writecheck.com, 2017.