

Social network aided plagiarism detection

Aljaž Zrnec and Dejan Lavbič

Aljaž Zrnec and **Dejan Lavbič**. 2018. **Social network aided plagiarism detection**, British Journal of Educational Technology (**BJET**), 48, pp. 113 - 128.

Abstract

The prevalence of different kinds of electronic devices and the volume of content on the Web have increased the amount of plagiarism, which is considered an unethical act. If we want to be efficient in the detection and prevention of these acts, we have to improve today's methods of discovering plagiarism. The paper presents a research study where a framework for the improved detection of plagiarism is proposed. The framework focuses on the integration of social network information, information from the Web, and an advanced semantically enriched visualization of information about authors and documents that enables the exploration of obtained data by seeking of advanced patterns of plagiarism. To support the proposed framework, a special software tool was also developed. The statistical evaluation confirmed that the employment of social network analysis and advanced visualization techniques led to improvements in the confirmation and investigation stages of the plagiarism detection process, thereby enhancing the overall efficiency of the plagiarism detection process.

Keywords

Social Network Analysis, Plagiarism detection process, Plagiarism visualization, Assessment, Facilitation, Learning management systems, Student-centredness, Case study

1 Introduction

The widespread availability of computers and mobile devices and the volume of content on the Web have changed the approaches to both teaching and learning processes. Simultaneously, the amount of plagiarism has increased enormously over the last few years due to the aforementioned changes. The act of plagiarism is defined as the unethical action of copying someone else's work (Youmans, 2011), and is usually considered as an offense, therefore we also have to improve the current plagiarism detection processes to be able to cope with the increasing amount of plagiarism in a more efficient way.

Culwin and Lancaster defined the Four-Stage Plagiarism Detection Process (FSPDP) (Culwin and Lancaster, 2001) used to systematically search for plagiarisms in a given set of documents focusing not only on similarity detection. FSPDP consists of four stages: collection, analysis, confirmation and investigation. Usually, all stages were performed by a human investigator, but with the advent of different plagiarism detection methods supported by computers, the first two stages in this process can be fully automated, and the latter two can only be partly automated (Makuc, 2013). The effectiveness of detection depends not only on the similarity engine (Ali et al., 2011; Hage et al., 2010) used in the second stage, but also on the rate of automation of the latter two stages in the process. Any similarity in the second stage, which is considered as positive, is further reviewed in the investigation stage. For a submission to be judged to contain plagiarism, the confirmation stage (3rd stage) must be completed, where the submission is examined and verified by a human investigator. This stage can also be fully automated, but false positive and false negative results may occur.

Today, most of the approaches (Ali et al., 2011) to the detection of plagiarism are focused on the first two stages, namely collection and analysis, leaving the investigator to perform the latter stages manually. That was the motivation behind conducting research work to propose a novel approach focusing on social aspects

of potential plagiarists, by taking into account their social network connections, activities and information from the Web, to support investigator's work in the third and the fourth stages of FSPDP, thereby making the plagiarism detection process more efficient. We believe that the plagiarism detection process can be improved by reducing the number of manual examinations of potentially plagiarized work. This could be achieved by the employment of new visualization techniques that enable a semantically enriched view of the relationships between possible plagiarists.

The paper is organized as follows. In the second section, we present the related work and propose a solution. In the third section, we describe our Social Plagiarism Detection Framework and supportive software tool. In the fourth section, we present the evaluation method for assessing the approach and discuss the obtained results. Finally, in the last section, we conclude the paper and discuss the possibilities for future research.

2 Related work

2.1 Review of related approaches and tools

According to authors (Mozgovoy et al., 2010), there are five different types of plagiarism, varying from verbatim copying to advanced types of plagiarism (Witherspoon et al., 2012) such as the copying of ideas and plagiarism in the form of translated text. The increasing use of computers and Web 2.0 tools have mainly had a positive effect on learning, but they also increase the possibility of using different types of plagiarism (Underwood and Szabo, 2003).

With the expansion of various types of plagiarism, especially with the proliferation of digital documents on the Web and the advent of social networks, many innovative approaches to plagiarism detection have also emerged. Several successful studies have been applied to traditional approaches (Mozgovoy et al., 2010, 2005; Stein et al., 2011), focusing on program code or plain text.

Early approaches to plagiarism detection heavily relied on methods that were based on string matching. Advanced methods include document parsing to extract the structure of the sentence and using a synonym thesaurus. All these methods do not perform well when faced with complex types of plagiarism (@ Mozgovoy et al., 2010) such as stealing ideas or text translations. Modern approaches are based on methods for natural language processing (Oberreuter and Velásquez, 2013), but they are still in their infancy.

So far, several of the above-mentioned approaches to plagiarism detection have been implemented in various types of software tools varying from autonomous applications to web services. Typically, applications are run locally and scanned for plagiarism within a given corpus of documents. On the other hand, there are web services that allow us to check for plagiarism among local corpuses and several sources on the Internet.

The main role of the plagiarism detection software tool is detecting similarities in program code, text or both. Some of the most commonly used tools today for detecting plagiarisms in computer source code are **Sherlock** (Joy and Luck, 1999; Mozgovoy et al., 2005), **JPlag** (Prechelt et al., 2002) and **Moss** (Schleimer et al., 2003). Their basic functionality is very simple. Selected submissions are run through a similarity engine, which provides pairwise results with potential plagiarisms. Modern software for plagiarism detection in source code is based not only on methods for string matching but also includes methods for searching lexical and structural modifications in programming code (Alzaharani et al., 2012a,b; Đurić and Gašević, 2013; Hein et al., 2012; Joy and Luck, 1999; Vrhovec et al., 2015). On the other hand, **WCOPYFIND** (Balaguer, 2009), **Ephorus** (Den Ouden and Van Wijk, 2011) and **TurnItIn** (Buckley and Cowap, 2013; Marsh, 2004; Rolfe, 2011) are tools for detecting plagiarisms in free text. They are used to find the amount of text shared between two or more plain text documents on the basis of fingerprinting (Introna and Hayes, 2011; Mozgovoy et al., 2010).

The user-friendliness of all the above-mentioned applications varies considerably. While some web services for detecting plagiarisms provide an intuitive user interface (eg, Ephorus), the majority of tools require a skilled user to operate them (eg, Moss) and they are not suitable for use by an ordinary teacher with the average computer skills. However, there is another consideration to be taken into account when using these tools. The vast majority of them do not support the work of the investigator through all four stages of the

plagiarism detection process (FSPDP). Current solutions are focused on the first and second stages (mainly on the second stage) of the process, which means that the investigator only gets a pairwise analysis (2nd stage) while the last two stages must be performed manually.

Despite the many different approaches and tools available, researchers (Mozgovoy et al., 2010) have also pointed out that currently available detection systems have several drawbacks which can be divided into two main categories:

- issues concerning the user-friendliness of today's detection tools (implementation of the system) and
- issues about the limitations of the existing technologies for plagiarism detection.

We also believe that the major drawback of current solutions is the inability to support all of the four stages in the plagiarism detection process. In fact, they can only be utilized in the first, and primarily in the second stage of the FSPDP, while the latter two stages (confirmation and investigation) have to be done manually by the investigator (eg, teacher), thus extending the time needed for the confirmation of plagiarized work and reducing overall efficiency of the FSPDP.

Moreover, several studies and analyses of social networks were the motivation behind our merging of information from social networks into the plagiarism detection process to counter the drawbacks of current approaches to plagiarism detection. Authors (Junco, 2012) aimed to identify relationship between Facebook use and academic performance. The research confirmed a negative relationship between time spent on Facebook and overall grade point average (GPA) achieved, as well as the time spent preparing for class. There have been multiple research studies conducted (Hew, 2011; Roblyer et al., 2010) highlighting attitudes toward social networking sites and student and faculty use of social networks. The results confirmed that students are more likely to use social networking sites and are significantly more open to the possibility of using similar technologies. Conclusions also suggest that social networks have very little educational use, as they are being used mainly to keep in touch with known individuals and students tend to disclose more personal information about themselves on social networking sites, hence exposing themselves to potential privacy risks. Furthermore, authors (Šubelj et al., 2011) successfully employed social network information in detecting automobile insurance fraud. The results provided some evidence that connectivity of users on social networking sites can have a predictive value in determining fraudulent activities like the detection of plagiarism.

2.2 Problem and proposed solution

The review of related approaches pointed out that the research studies on plagiarism detection are focused on the first two initial stages - collection (1st stage) and mainly on analysis (2nd stage). As we have mentioned before, the majority of existing approaches conclude their user support by providing information on the pairwise content similarity of documents, and leaving the investigator to perform the confirmation and investigation stages manually. In contrast, the proposed approach puts the emphasis on an integrated solution where we try to focus on the social aspects of possible plagiarists, by taking into account their social network connections, activities, as well as information from the Web. This provides improved support for plagiarism detection in confirmation (3rd stage) and investigation (4th stage). We argue that our approach facilitates plagiarism detection by providing the investigator with better support in the latter two stages; therefore, the confirmation or rejection of plagiarism in the third stage can be more efficient, consequently making the process of plagiarism detection also more efficient as a whole. The result of our approach represents the reduced number of potentially plagiarized work that an investigator has to examine manually, and the provision of new visualization techniques that enable a semantically enriched view of relationships among possible plagiarists. We also provide a tool to support the process that can visualize more corpora with additional information. This enables the investigators to have an overview of an author's plagiarism in the context of their previous work and work related to their colleagues, and not only by content similarity. The tool is intended for use by teachers who want to exclude the possibility of cheating among students.

In contrast to majority of the existing tools for plagiarism detection, the provided tool has two important advantages:

- it is more user-friendly and as a consequence it can be used by any teacher with a common level of computer knowledge; and
- it provides integrated support for an investigator in all four stages of FSPDP.

The tool is primarily designed for teachers or professors who teach programming, enabling them to find plagiarism easily in source code documents related to a particular assignment, or in a particular teaching assignment when compared with all previous assignments.

3 Social plagiarism detection framework

3.1 Description of proposed system

With the introduction of the Social Plagiarism Detection Framework (SPDF), we focus on the latter stages of the plagiarism detection process, namely in confirmation (3rd stage) and investigation (4th stage) as depicted in Figure 1.

The main contributions of our approach are as follows:

- integration of social network information and information from the Web that facilitates the plagiarism detection process; and
- an advanced semantically enriched visualization (semantic graph, co-occurrence matrix) of information about authors and documents that enables the exploration of data in search of advanced patterns of plagiarism.

The additional steps of SPDF and advantages compared with existing approaches based only on content similarity are depicted in Figure 2. Steps 1 - 5 are generally performed in content similarity matching in the plagiarism detection process, while we introduce steps 6 - 10.

In the confirmation stage, the system evaluates the content similarity report (Figure 2, step 5) provided in the analysis stage (Figure 2, steps 3 - 4) and performs an additional evaluation of general search engine results (Figure 2, step 9) and connections between authors on social networks (Figure 2, step 7). In the case of ambiguity, the investigator is provided with an option to review the social network analysis results. Based on all given information in the context (content similarity and connections between users on social networks), the investigator can confirm or reject pairwise plagiarism. The main benefit of our approach is the improved ranking of potential pairwise plagiarisms where social information as well as information from the Web are taken into account, thereby minimizing the effort required by the investigator in the confirmation stage. We argue, and provide a comprehensive evaluation of our findings in the fourth section, that the impact of social information is statistically significant in the plagiarism detection process.

3.2 Plagiarism detection framework definition

We can define P and D as nonempty sets of people and documents respectively

$$P = \{p \mid p \text{ is a person}\} \tag{1}$$

$$D = \{d \mid d \text{ is a document, written by } p \in P\} \tag{2}$$

where W is a set of pairs $\langle p, d \rangle$ with p as an author of a document d

$$W = \{\langle p, d \rangle \mid \exists p \in P \wedge d \in D : p \text{ is an author of } d\} \tag{3}$$

The set W is the result (Figure 2, step 2) of the collection stage (Figure 2, step 1) and is the direct input to the analysis stage (Figure 2, step 3). For this purpose, we can further define content similarity between documents

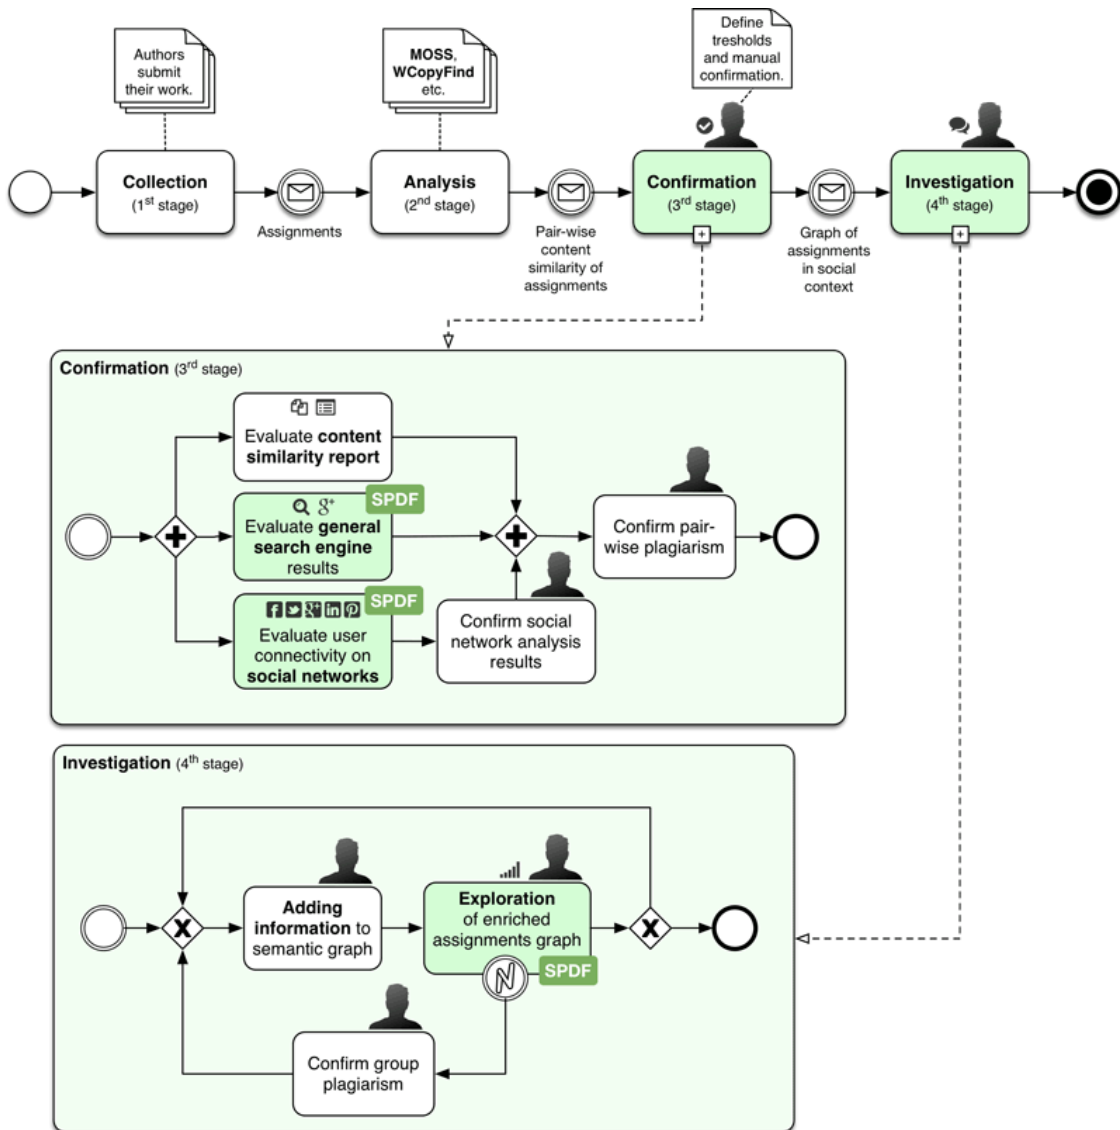


Figure 1: Social Plagiarism Detection Framework (SPDF)

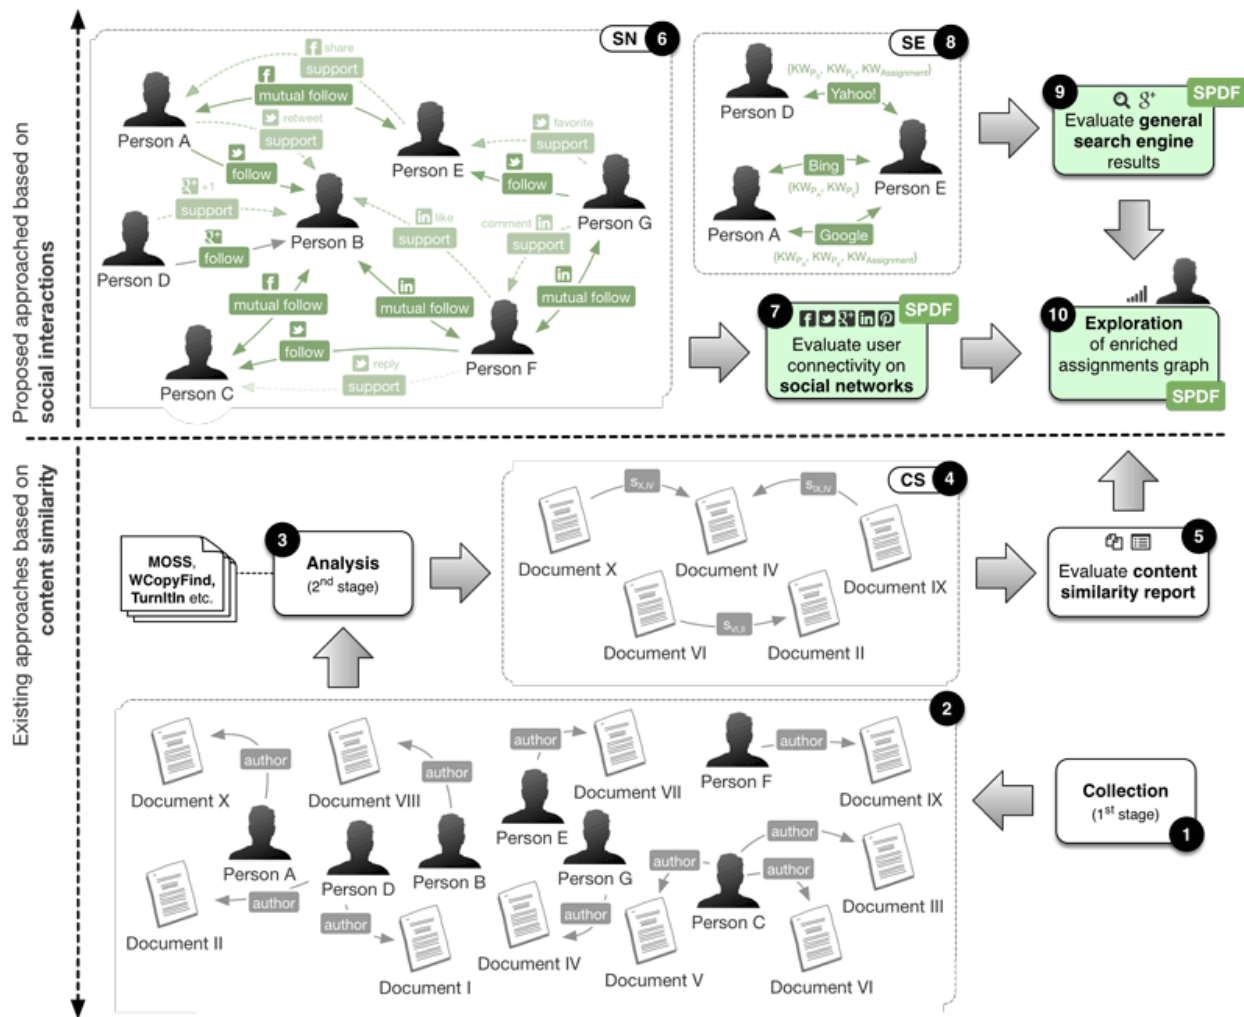


Figure 2: Contribution of SPDF compared with existing approaches

$$CS = \{\langle d_i, d_j, s_{ij} \rangle \mid \exists d_i, d_j \in D \wedge s_{ij} \in (0, 1] : s_{ij} \text{ is a share of content from } d_i \text{ in } d_j\} \quad (4)$$

as a set of pairs of documents d_i and d_j , with directed content similarity s_{ij} between documents (Figure 2, step 4).

By the integration of social network information in the plagiarism detection process, we introduce two measures: SN (Figure 2, step 6) and SE (Figure 2, step 8).

SN represents the connections between users on social networks considering pairwise actions between users. We define the following types of actions $T(A) = \{f^{\rightarrow}, f^{\leftrightarrow}, s^{\rightarrow}\}$ as (directed) follow f^{\rightarrow} , (undirected) mutual follow f^{\leftrightarrow} and (directed) support s^{\rightarrow} (eg, share, comment, reply, retweet, favorite, like, +1 etc).

When considering various social networks, we classify them into two distinct categories regarding user connections. The first group considers links between users as directed (unilateral following of users, eg, Twitter, Google+, etc), while the other group employs undirected links (users mutually confirm following, eg, Facebook, LinkedIn, etc). We can define the connections between users on social networks

$$SN = \{\langle p_i, p_j, A_{ij} \rangle \mid \exists p_i, p_j \in P \wedge \exists a_k^{(ij)} \in A_{ij} : a \text{ is an action between } p_i \text{ and } p_j\} \quad (5)$$

as a set of triples of people p_i and p_j , with set of actions A_{ij} between people p_i and p_j . Actions are further defined as set of triples $A_{ij} = \{a_k^{(ij)} = \langle nt, act, w \rangle\}$, where nt is the social network (eg, Facebook, Twitter, LinkedIn, Google+, etc), act is activity (eg, follow, share, comment, like, etc) and w is a user-defined weight of specific action.

When determining pairs $\langle p_i, p_j \rangle$ of connected people, a fuzzy search, implementing the Levenshtein distance algorithm, is performed that requires further action by the investigator in the case of ambiguity with multiple account and/or people matching.

We also introduce a set of related items from general search engine SE between people p_i and p_j (Figure 2, step 8) as

$$SE = \{\langle p_i, p_j, KW_{ij}, n \rangle \mid \exists p_i, p_j \in P, n \in \mathbb{N} : n \text{ is number of related items}\} \quad (6)$$

where n is the number of relevant search results involving people p_i and p_j , using a set of keywords KW_{ij} . This set of keywords KW_{ij} between pairs $\langle p_i, p_j \rangle$ of connected people, is user defined per assignment as follows:

$$KW_{ij} = KW_{p_i} \cup KW_{p_j} \cup KW_{assignment} \quad (7)$$

where KW_{p_i} and KW_{p_j} are keywords related to a person's information (eg, name, surname, etc.) and $KW_{assignment}$ is a set of assignment-related keywords to narrow down the result set.

In the process of plagiarism detection, the goal is to define the set of pairs of documents DP , where plagiarism has been confirmed

$$DP = \{\langle d_i, d_j \rangle \mid \exists d_i, d_j \in D : d_i \text{ is a plagiat of } d_j\} \quad (8)$$

When using existing approaches, the investigator, who performs the plagiarism detection process, tries to identify elements of DP , while considering CS and some tacit knowledge TK by investigation. The results of the confirmation and investigation stages can be defined as a function $check_{woSocio}$, performed by the investigator

$$check_{woSocio} : CS \times TK \rightarrow DP \quad (9)$$

1 st person	1 st document	2 nd person	2 nd document	CS	SN	SE	Total	Status
p_1	d_1	p_2	d_3	CS_{13}	SN_{12}	SE_{12}	t_{12}	confirmed
p_1	d_1	p_3	d_4	CS_{14}	SN_{13}	SE_{13}	t_{13}	confirmed
p_2	d_5	p_4	d_9	CS_{59}	SN_{24}	SE_{24}	t_{24}	rejected
p_7	d_6	p_9	d_8	CS_{68}	SN_{79}	SE_{79}	t_{79}	not-checked
...
p_{n-1}	d_{m-3}	p_n	d_m	$CS_{m-3,m}$	$SN_{n-1,n}$	$SE_{n-1,n}$	$t_{n-1,n}$	not-checked

Figure 3: Ranked table of pairwise possible plagiarism

The investigator evaluates the content similarity report consisting of pairwise documents and decides on classifying the event as confirmed, rejected or not-checked plagiarism.

When our proposed approach is utilized (Figure 3), the following function $check_{wSocio}$ is defined

$$check_{wSocio} : CS \times TK \times SN \times SE \rightarrow DP \quad (10)$$

that besides content similarity, also considers the social information of document authors (Figure 2, steps 6 - 9).

Figure 3 depicts a table of pairwise potential plagiarisms that the investigator has to examine in order to confirm or reject them. The investigator is provided with the full information about content similarity (CS), social network connectivity (SN) and general search engine matching (SE). The weighted average of the aforementioned factors between people p_i and p_j is defined as

$$Total_{ij} = w_{CS} \cdot cs_{d(p_i)d(p_j)} + w_{SN} \cdot SN_{ij} + w_{SE} \cdot se_{ij} \quad (11)$$

where $d(p_i)$ and $d(p_j)$ are documents authored by people p_i and p_j respectively.

The weights w_{CS} , w_{SN} and w_{SE} , where $\sum_i w_i = 1$, are user defined per assignment and allow investigators to define the importance of individual plagiarism detection factors (CS , SN and SE).

To enable investigators to have an overview of factor values and their distribution (Figure 3), we introduce an indication of clusters (**confirmed**, **rejected** and **not-checked** possible plagiarism) for individual factors. It is available as an interval with the minimum and maximum values for the selected factor, where colors (**red**, **orange** and **green** for **confirmed**, **not-checked** and **rejected** potential plagiarism respectively) are used to depict mean values for individual clusters.

Furthermore, we argue that the employment of $check_{wSocio}$ is more straightforward than $check_{woSocio}$ and enables the investigator to perform the confirmation stage more efficiently. This results in a reduction of the total number of documents suspected to contain plagiarism that the investigator has to manually review and confirm or reject the occurrence of plagiarism. For evaluation purposes, the supporting tool has been developed to test and compare the aforementioned scenarios.

3.3 Limitations

With the introduction of SPDF, the limitations of the proposed approach also have to be considered. The major limitation of the approach is the dependency on publicly available data from social networks. If we want the approach to be as much efficient as possible, the authors (users) should have publicly accessible profiles on several observable social networks. In the case of Facebook, it is interesting that the requirement is not so difficult to achieve, as studies about the user identity presentation (user profile) on Facebook show that users are willing to provide substantial amounts of personal data (Gross and Acquisti, 2005; Wilson et al., 2012). Although the awareness of privacy and security issues has increased over the past few years, several studies have revealed that many users still have publicly accessible profiles (Mazur, 2010; McKnight et al., 2011).

The second limitation is that the approach is suitable for smaller groups of authors (eg, in educational sectors or classes at University, High School, etc) because it requires gathering user data from social networks. Access to this kind of data is usually realized by using application programming interfaces (APIs) which are generally limited in terms of which and what data can be accessed and how often the data can be retrieved (Rieder, 2013). Facebook, for example, is very restrictive in terms of which data can be accessed, and it also determines the request frequency, making the possibility of analyzing large set of users very impractical. Facebook, Twitter, LinkedIn and others also reserve the right to close or modify these APIs, which represent an additional limitation to the proposed approach.

The approach also raises some ethical questions as it anticipates the automatic gathering of an author's data from social networks (Zimmer, 2010). The authors of analyzed documents are not active participants in the collection of their social data because they could provide inaccurate data about social connections and their activities on social networks. Therefore, they should be warned about this approach used for plagiarism detection that uses automatic inquiries about their profiles before their assignments are checked. However, it is also important to point out that all the information gathered about users is publicly available.

3.4 Plagiarism detection assistant tool

To support the proposed process, the plagiarism detection assistant (PDA) tool was developed in which the following functionalities are supported:

- creating and managing projects,
- integration of existing plagiarism detection tools,
- automatic acquisition of social network information and general search engine results,
- confirming/rejecting assignments, and
- advanced visualization.

The initial action in the process performed by an investigator is creating a project and collecting the submissions. Then the following steps include the preparation of data for the confirmation stage:

- performing content analysis by selecting the existing plagiarism detection tool, where pairwise content similarity report is retrieved; and
- acquisition of social network information and general search engine results for investigated authors.

After data are prepared, the investigator enters the confirmation stage as Figure 4 depicts (the names are pseudonyms). The goal of this step is to assign one of the following status to the pairwise assignments by two people who are being assessed for the possibility of plagiarism:

- **not checked** - similarity between assignment has not been considered yet,
- **rejected** - the assignments are not plagiarisms, and
- **confirmed** — the assignments contain plagiarized sections.

When making the decision, the PDA tool assists the investigator by providing an extensive report of matches found on assignments submitted by different authors as depicted in Figure 5. There is a history of all assignments and their corresponding content similarity enriched by the social network component. The colors used depict the severity of the warnings. In this way, the tool automatically ranks detected plagiarized

1 st person	2 nd person	Content similarity	Status	Actions			
WRIGHT, Richard	RUSSEL, Kevin	6 %	NOT CHECKED	View match	View result	Confirm match	Reject match
PERKINS, James	TAYLOR, Scott	7 %	REJECTED	View match	View result	Confirm match	Reject match
PARKER, Thomas	MOORE, Eric	8 %	CONFIRMED	View match	View result	Confirm match	Reject match
GREEN, Charles	MOORE, Eric	9 %	REJECTED	View match	View result	Confirm match	Reject match
GREEN, Charles	MILLER, George	9 %	REJECTED	View match	View result	Confirm match	Reject match

Figure 4: Pairwise assignment view by content similarity

sections depending on information obtained from social networks and activities on the Web, but it does not confirm suspicious assignments as plagiarism. We must emphasize that the proposed approach is not intended to be used as a replacement for conventional analyses of texts, but rather as a supporting tool for increasing the efficiency of the confirmation stage. When the investigator reviews all the provided information, they can make a decision and confirm or reject plagiarism. By performing these steps, the confirmation stage of the plagiarism detection process is concluded (see Figure 1) and the investigation stage can begin.

One of the views in the PDA tool within the investigation stage is depicted in Figure 6 where the support for advanced visualization is provided. The investigator can interactively explore the semantic graph and co-occurrence matrix equipped with information about the content similarity, connectivity on social networks and general search engine results. The data from social networks and the Web are collected by means of social network public APIs and Web scraping of publicly available data about authors. As we only do a pairwise analysis of data from a limited set of people, we do not have any problems with processing resources. By visualizing the context of the entire group under investigation (eg, class at University), the investigator can carry out plagiarism detection by exploring group plagiarism.

4 Evaluation

4.1 Method

Our approach was evaluated on a case study of 76 students taking one of the lectures in Computer Science at undergraduate level. Each student had to submit five programming homework assignments during the semester that were later checked for plagiarism. There were two experiments performed with two groups of evaluators that followed different approaches to the same dataset (76 students submitted 5 assignments, where 22 assignments were missing, so in total 358 assignments). Both groups of evaluators had the common goal - to identify plagiarism in the students' work. In the first approach, *check_{woSocio}* evaluators employed MOSS (Schleimer et al., 2003) and performed a manual investigation on pairwise content similarity, while in the second approach, *check_{wSocio}*, our method with additional social network analysis results was used. The information from social networks employed in the second approach was extracted from the public profiles of the students. In our case study, 54 students had publicly available information on Facebook and 43 students were active on Twitter. Students were informed about the use of all available public information in the process of plagiarism detection throughout the course.

The method used for the evaluation of the aforementioned approaches is a generalized linear model with logistic regression where the link function is defined as follows

$$g(Y) = \log_e \left(\frac{n}{1-n} \right) = \beta_0 + \sum_{j=1}^p \beta_j X_j \tag{12}$$

The logistic regression is applied to a situation where the response variable $Y = \textit{cheat}_{\textit{confirmed}}$ is dichotomous (0,1). The model assumes that Y follows a binomial distribution and it can be a fit to a linear model $g(Y)$.

GREEN, Charles



 1 confirmed

 0 matched

 1 confirmed

 0 matched

MOORE, Eric



 11 matched

 0 matched

 1 confirmed

 0 matched

Matches

There is **NO Google+ or LinkedIn connections** between Charles GREEN and Eric MOORE. ✕

There is **NO Facebook connections** between Charles GREEN and Eric MOORE, although you should define the correct profile for Eric MOORE. ✕

[Select Eric's MOORE Facebook profil](#)

There is a **Twitter connections**, where Charles GREEN follows Eric MOORE. ✕

General search engine results warning! ✕

There were **3 results** found using **Google search**, while averages are as follows:

(rejected) 0,4

3,0

(This result)

(non-checked) 0,3

0,9

(confirmed)

[View details](#)

Assignment	Content similarity	Status	Actions
Homework 1b	9 %	NOT CHECKED	View Confirm match Reject match
Homework 3a	25 %	CONFIRMED	View Confirm match Reject match

Figure 5: Extensive report on matching, including social network information

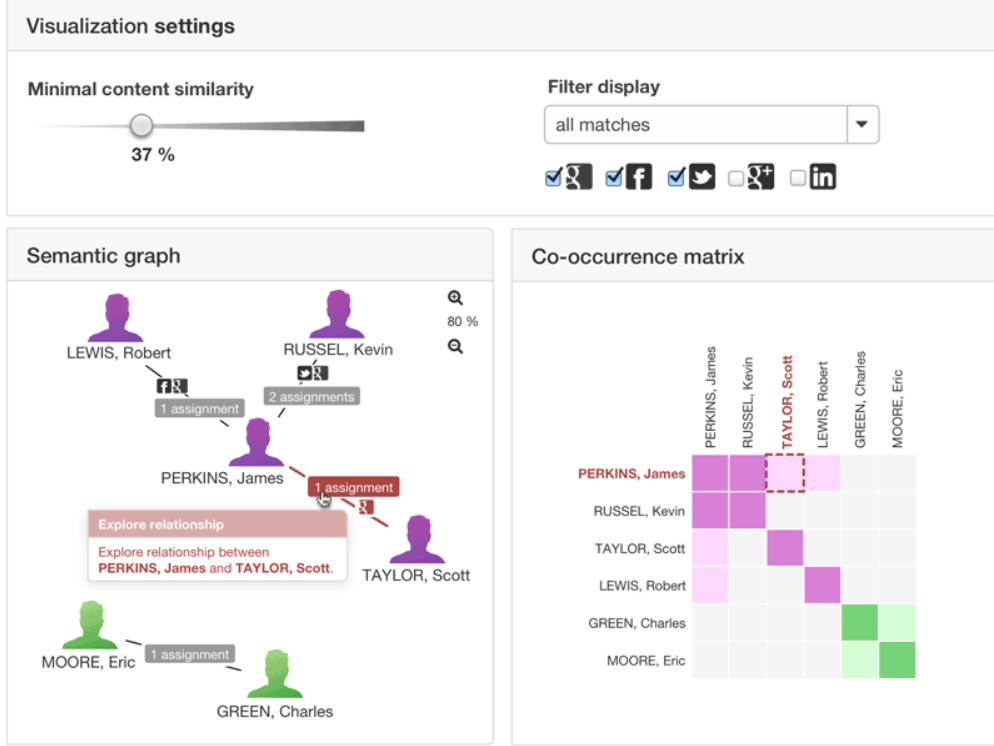


Figure 6: Analytical assistance provided by the PDA tool

The conditional mean of Y is the probability $\pi = \mu_Y$ that cheat is confirmed, given a set of X values. The odds that cheat is confirmed are $\frac{n}{1-n}$ and $\log\left(\frac{n}{1-n}\right)$ is the log odds or *logit*.

We have defined two models

$$check_{woSocio} : cheat_{confirmed} \sim match_{cs} \quad (13)$$

$$check_{wSocio} : cheat_{confirmed} \sim match_{cs} + match_{fb} + match_{tw} + se_{hits} \quad (14)$$

where $check_{woSocio}$ is a nested model within $check_{wSocio}$ with the same response variable $Y = cheat_{confirmed}$ and different predictors $X_{woSocio}$ and X_{wSocio} , where $X_{woSocio} \subseteq X_{wSocio}$. The predictor variables are as follows: $cheat_{confirmed}$ is $\{true, false\}$ factor with information about confirmed plagiarism from investigator; $match_{cs}$ is content similarity s_{ij} between documents d_i and d_j , where $\langle d_i, d_j, s_{ij} \rangle \in CS$; $match_{fb}$ is a $\{true, false\}$ factor, based on existence of $\langle p_i, p_j, \langle FB, follow, 1 \rangle \rangle \in SN$; $match_{tw}$ is a $\{true, false\}$ factor, based on existence of $\langle p_i, p_j, \langle TW, follow, 1 \rangle \rangle \in SN$ and se_{hits} is a number of search engine results n between people p_i and p_j , where $\langle p_i, p_j, KW_{ij}, n \rangle \in SE$ and KW_{ij} is a set of keywords including names and surnames of p_i and p_j and subject title.

The employment of models $check_{woSocio}$ and $check_{wSocio}$ is not intended to predict plagiarism, but rather for the ranking of potential pairwise plagiarized sections that the investigator can review and confirm in the latter stages of plagiarism detection.

To conclude our experiment, we performed a follow-through interviews with all of the students where they defended their work and evaluators determined if their submitted work was original. The evaluator's decisions were then used for $cheat_{confirmed}$ response variable to evaluate both models.

4.2 Results

When building a model *checkwoSocio*, the results show that the predictor variable *matchcs* is significant ($p \leq 9 \times 10^{-6}$) in predicting the response variable *cheatconfirmed*.

The next step was to build another model *checkwSocio* with integrated social network information, where the results of the second model show that all predictor variables *matchcs* ($p \leq 0,0025$), *matchfb* ($p \leq 0,0289$), *matchtw* ($p \leq 0,0904$) and *sehits* ($p \leq 0,0432$) are significant in predicting the response variable *cheatconfirmed*.

Then we were able to compare both models, which consist of variables that all have significant impact on the prediction. We performed ANOVA with likelihood ratios test (LRT) on both models. The measure used for comparison is deviance as a distance between two probabilistic models. Deviance can be regarded as a measure of lack of fit between model and data. Based on the results, we can conclude that the residual deviance in the 1st model *checkwoSocio* ($RD_1 = 48,932$) is significantly higher ($p \leq 6 \times 10^{-8}$) than in the 2nd model *checkwSocio* ($RD_2 = 12,479$). We can argue that the 1st model is a poorer fit to the data and that the 2nd model performs better.

To confirm that the results are meaningful, we have performed the test for overdispersion for both models that could lead to distort test standard errors and inaccurate test of significance. We performed fitting of the model twice - once with binomial family and second with quasibinomial family and the results confirmed that overdispersion is not a problem (the noncentral chi-squared test was not significant with $p_{woSocio} = 0,977$ and $p_{wSocio} = 0,990$). We also assessed the model adequacy by checking for unusually high values in the hat values, studentized residuals and Cook's D statistics. The results of these tests also confirmed that the models are adequate.

5 Discussion

As presented in the second section, the employment of advanced social network analysis approaches has been already successfully implemented in several business-oriented domains. In particular, our research was motivated from fraud detection in car insurance industry where some promising implementations can be found. The results of our approach demonstrate that there also exists a predictive value in utilizing social network information about authors of documents when detecting plagiarism in student's work.

The main goal of SPDF is to make the overall process of plagiarism detection (FSPDP) more efficient, which has several implications for education in general by impacting all participants of FSPDP—students and teachers. With more efficient support to plagiarism detection, teachers can focus more on high-performing students, while identification of ones that employ fraudulent actions is improved. Still, the time invested by teachers also incorporates some additional activities, such as disambiguation of students on social networks, when system cannot determine the unique person from a list of available ones (eg, people with the same name). But this invested time from teacher's perspective is rewarded with reduced set of potential pairs of documents to check for plagiarism. From the student's point of view, the employment of their social network information in plagiarism detection can have negative acceptance. This is a general concern, as people tend to protect their own online privacy. But nonetheless, SPDF employs only publicly available information from social networking sites and general search engines. If there is no information available online about students, then the proposed approach will not improve the plagiarism detection process and we will have to utilize only content similarity and perform all work in the 3rd and the 4th stages of FSPDP manually.

During the evaluation process of SPDF, there were several cases in our study when investigators (eg, teachers) of plagiarism commended the exploratory aspect of SPDF by supporting investigation and enabling them the traversal of information about documents and authors in integrated manner of a semantic graph and a co-occurrence matrix, thus making them more efficient in evaluating the knowledge of students. Time invested to plagiarism detection varied significantly between two groups of investigatorsthatwereinvolvedinourstudy.Thefirstgroupthatemployed only content similarity information had to review 52 pairs of assignments identified as possible plagiarism. The other group using SPDF approach identified only 17

pairs of possible plagiarism. The decreased number of possible plagiarism to review and confirm was due to availability of additional social network information and advanced semantically enhanced visualization of results that enabled investigators to identify clusters of students that collaborated and therefore handling them in a group. The social information also often provided investigators some intuition on who was the author (eg, general search engine results indicated student is collaborating in an open source project related to the subject of examination) and who was copying. The enriched view on possible plagiarism also presented a stronger evidence for investigators in the last stage of investigation where students were invited to an interview and to defend their work.

Besides confirming that SPDF performs significantly better than a content similarity approach in determining plagiarism of computer programming assignments at University, we applied the approach in other settings. One of the successful studies still in progress is in a High School environment where students in a class for their mother tongue language are required to write multiple essays about literary novels. As there are several assignments per semester, instructors are overloaded with work providing feedback to all students. The employment of SPDF was therefore a welcome upgrade from manual examination and content similarity evaluation process only. Instructors, with the help of SPDF, more quickly and efficiently identify plagiarism and focus on providing constructive feedback to students to support their progress.

6 Conclusion and future work

Plagiarism detection approaches mainly focus on the first two stages of FSPDP. However, that is not sufficient to discover authors who perform unethical acts relating to plagiarism, because we also have to deal with false positive and false negative results from the analysis stage.

With the proposed approach and PDA software tool, we are able to support the investigator's work effectively during the confirmation and investigation stages. In the confirmation stage, we can efficiently narrow the set of potential plagiarists from previous stages and in the investigation stage we can visualize the relationships among potential plagiarists with the additional semantic information. The evaluation of two different models, in the selected case study, demonstrates that the obtained results are significant. This provides evidence that the inclusion of social network information about the authors of texts assists the plagiarism detection process when compared with the approach where the information from the social networks and the Web is not employed during the manual decision-making process of confirmation and investigation of plagiarism as performed by a human investigator. In this place, it is also important to point out the restrictions of our approach. The major limitation is its dependency on publicly available data from social networks, inclination of users to provide public personal information and possibility to access these data. There also appear some ethical questions concerning automatic inquiries about user profiles but it has to be emphasized that all data from social networks and the Web that we utilized are publicly available.

Future research will focus on improving the framework by further analysis of social network connections between authors under investigation. The communication interactions will be analyzed by using advanced methods of text analysis. We will also try to find the correlation between the messages exchanged between authors and any plagiarism in their submitted documents.

References

- Ali, A. M. E. T., Abdulla, H. M. D., and Snásel, V. (2011). Overview and Comparison of Plagiarism Detection Tools. In *DATESO 2011: Annual International Workshop on DAtabases, TExts, Specifications and Objects*.
- Alzahrani, S., Palade, V., Salim, N., and Abraham, A. (2012a). Using structural information and citation evidence to detect significant plagiarism cases in scientific publications. *Journal of the American Society for Information Science and Technology*, 63(2):286–312.

- Alzahrani, S. M., Salim, N., and Abraham, A. (2012b). Understanding Plagiarism Linguistic Patterns, Textual Features, and Detection Methods. *IEEE Transactions on Systems, Man, and Cybernetics*, 42(2):133–149.
- Balaguer, E. V. (2009). Putting ourselves in SME’s shoes: Automatic detection of plagiarism by the WCopy-Find tool. In *SEPLN’09 Workshop on Uncovering Plagiarism, Authorship and Social Software Misuse*, pages 34–35.
- Buckley, E. and Cowap, L. (2013). An evaluation of the use of Turnitin for electronic submission and marking and as a formative feedback tool from an educator’s perspective. *British Journal of Educational Technology*, 44(4):562–570.
- Culwin, F. and Lancaster, T. (2001). Visualising intra-corporal plagiarism. In *5th International Conference on Information Visualisation*, pages 289–296. IEEE Computer Society Washington.
- Den Ouden, H. and Van Wijk, C. (2011). Plagiarism: Punish or prevent? Some experiences with academic copycatting in the Netherlands. *Business Communication Quarterly*, 74(2):196.
- Gross, R. and Acquisti, A. (2005). *Information revelation and privacy in online social networks*. ACM.
- Hage, J., Rademaker, P., and Vugt, N.-v. (2010). A comparison of plagiarism detection tools. Technical report, Department of Information and Computing Sciences, Utrecht University.
- Hein, J., Zobrist, R., Konrad, C., and Schuepfer, G. (2012). Scientific fraud in 20 falsified anesthesia papers: detection using financial auditing methods. *Anaesthetist*, 61(6):543–549.
- Hew, K. F. (2011). Students’ and teachers’ use of Facebook. *Computers in Human Behavior*, 27(2):662–676.
- Introna, L. D. and Hayes, N. (2011). On sociomaterial imbrications: What plagiarism detection systems reveal and why it matters. *Information and Organization*, 21(2):107–122.
- Joy, M. and Luck, M. (1999). Plagiarism in programming assignments. *IEEE Transactions on Education*, 42(2):129–133.
- Junco, R. (2012). Too much face and not enough books: The relationship between multiple indices of Facebook use and academic performance. *Computers in Human Behavior*, 28(1):187–198.
- Makuc, Z. (2013). *Methods to assist Plagiarism Detection*. BSc.
- Marsh, B. (2004). Turnitin.com and the scriptural enterprise of plagiarism detection. *Computers and Composition*, 21(4):427–438.
- Mazur, E. (2010). Collecting data from social networking Wb sites and blogs. In Gosling, S. D. and Johnson, J. A., editors, *Advanced methods for conducting online behavioral research*, pages 77–90. American Psychological Association, USA, Washington.
- McKnight, D. H., Lankton, N., and Tripp, J. (2011). *Social Networking Information Disclosure and Continuance Intention: A Disconnect*.
- Mozgovoy, M., Fredriksson, K., White, D., Joy, M., and Sutinen, E. (2005). Fast Plagiarism Detection System. In Consens, M. and Navarro, G., editors, *String Processing and Information Retrieval*, volume 3772 of *Lecture Notes in Computer Science*, pages 267–270. Springer Berlin Heidelberg.
- Mozgovoy, M., Kakkonen, T., and Cosma, G. (2010). Automatic student plagiarism detection: future perspectives. *Journal of Educational Computing Research*, 43(4):511–531.
- Oberreuter, G. and Velásquez, J. D. (2013). Text mining applied to plagiarism detection: The use of words for detecting deviations in the writing style. *Expert Systems with Applications*, 40(9):3756–3763.
- Prechelt, L., Malpohl, G., and Philippsen, M. (2002). Finding plagiarisms among a set of programs with JPlag. *Journal of Universal Computer Science*, 8(11):1016–1038.

- Rieder, B. (2013). *Studying Facebook via data extraction: the Netvizz application*. ACM.
- Roblyer, M. D., McDaniel, M., Webb, M., Herman, J., and Witty, J. V. (2010). Findings on Facebook in higher education: A comparison of college faculty and student uses and perceptions of social networking sites. *Internet and Higher Education*, 13(3):134–140.
- Rolfe, V. (2011). Can Turnitin be used to provide instant formative feedback? *British Journal of Educational Technology*, 42(4):701–710.
- Schleimer, S., Wilkerson, D. S., and Aiken, A. (2003). Winnowing: local algorithms for document fingerprinting. In *ACM SIGMOD International conference on Management of data*, pages 76–85. ACM.
- Stein, B., Lipka, N., and Prettenhofer, P. (2011). Intrinsic plagiarism analysis. *Language Resources and Evaluation*, 45(1):63–82.
- Underwood, J. and Szabo, A. (2003). Academic offences and e-learning: individual propensities in cheating. *British Journal of Educational Technology*, 34(4):467–477.
- Vrhovec, S. L. R., Trkman, M., Kumer, A., Krisper, M., and Vavpotic, D. (2015). Outsourcing as an Economic Development Tool in Transition Economies: Scattered Global Software Development. *Information Technology for Development*, 21(3):445–459.
- Wilson, R. E., Gosling, S. D., and Graham, L. T. (2012). A Review of Facebook Research in the Social Sciences. *Perspectives on Psychological Science*, 7(3):203–220.
- Witherspoon, M., Maldonado, N., and Lacey, C. H. (2012). Undergraduates and Academic Dishonesty. *International Journal of Business & Social Science*, 3(1):76–86.
- Youmans, R. J. (2011). Does the adoption of plagiarism-detection software in higher education reduce plagiarism? *Studies in Higher Education*, 36(7):749–761.
- Zimmer, M. (2010). “But the data is already public”: on the ethics of research in Facebook. *Ethics and Information Technology*, 12(4):313–325.
- Đurić, Z. and Gašević, D. (2013). A Source Code Similarity System for Plagiarism Detection. *The Computer Journal*, 56(1):70–86.
- Šubelj, L., Furlan, S., and Bajec, M. (2011). An expert system for detecting automobile insurance fraud using social network analysis. *Expert Systems with Applications*, 38(1):1039–1052.