DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU

HARVARD LIBRARY
Office for Scholarly Communication

Deploying AI Methods to Support Collaborative Writing: A Preliminary Investigation

The Harvard community has made this article openly available. Please share how this access benefits you. Your story matters

Citation	Gehrmann, Sebastian, Lauren Urke, Ofra Amir, and Barbara J. Grosz. 2015. "Deploying AI methods to support collaborative writing: a preliminary investigation." In Proceedings of the 33rd Annual ACM Conference Extended Abstracts on Human Factors in Computing Systems, Seoul, South Korea, April 18-23, 2015: 917-922.
Published Version	doi:10.1145/2702613.2732705
Citable link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:29399610
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP

DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD

Deploying AI Methods to Support Collaborative Writing: a Preliminary Investigation

The Harvard community has made this article openly available.
Please share how this access benefits you. Your story matters.

Citation	Sebastian Gehrmann, Lauren Urke, Ofra Amir, Barbara J. Grosz. Deploying AI Methods to Support Collaborative Writing: a Preliminary Investigation. CHI'15 work-in-progress, 2015 ()
Accessed	March 9, 2015 10:09:28 AM EDT
Citable Link	*Not Available*
Terms of Use	This article was downloaded from Harvard University's DASH repository. *WARNING: No applicable access license found.*

(Article begins on next page)

Deploying AI Methods to Support Collaborative Writing: a Preliminary Investigation

Sebastian Gehrmann
Nordakademie University
Elmshorn, Germany
gehrmann@seas.harvard.edu

Ofra Amir
Harvard SEAS
Cambridge, MA USA
oamir@seas.harvard.edu

Lauren Urke
Harvard SEAS
Cambridge, MA USA
laurenurke@college.harvard.edu

Barbara J. Grosz
Harvard SEAS
Cambridge, MA USA
grosz@eecs.harvard.edu

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author. Copyright is held by the owner/author(s).
CHI'15 Extended Abstracts, Apr 18-23, 2015, Seoul, Republic of Korea
ACM 978-1-4503-3146-3/15/04.
http://dx.doi.org/10.1145/2702613.2732705

Abstract
Many documents (e.g., academic papers, government reports) are typically written by multiple authors. While existing tools facilitate and support such collaborative efforts (e.g., Dropbox, Google Docs), these tools lack intelligent information sharing mechanisms. Capabilities such as "track changes" and "diff" visualize changes to authors, but do not distinguish between minor and major edits and do not consider the possible effects of edits on other parts of the document. Drawing collaborators' attention to specific edits and describing them remains the responsibility of authors. This paper presents our initial work toward the development of a collaborative system that supports multi-author writing. We describe methods for tracking paragraphs, identifying significant edits, and predicting parts of the paper that are likely to require changes as a result of previous edits. Preliminary evaluation of these methods shows promising results.

Introduction
Collaborative writing is a common activity for many people: scientists write papers and proposals together, lawyers draft contracts, legislators draft legislation. In recent years, collaborative writing of documents has become even more widespread with the support of frameworks that simplify distributed editing (e.g., Google Docs, Dropbox). While document processors provide

capabilities such as "track changes", "diff", and commenting, they lack intelligent information sharing mechanisms. They do not attempt to reason about the importance of changes to different authors and do not consider how edits might affect other parts of the document. As a result, significant coordination overhead remains: authors need to re-read the entire document frequently or rely on communication from their co-authors in order to keep track of the current state of the document and ensure that their edits are consistent with other edits.

This paper proposes methods for tracking paragraphs across revisions, identifying significant changes, and predicting paragraphs that are likely to be edited following a significant edit of a particular paragraph. The development of these methods is a first step toward the design of a collaborative system capable of (1) drawing authors' attention to edits that are important and relevant for them, and (2) pointing out other parts of the document that are likely to require further editing. Systems with such capabilities have the potential to improve coherence of documents and coordination among authors while reducing the amount of communication required between authors. An empirical evaluation of the proposed approach on a corpus of Wikipedia articles shows promising initial results.

Related Work

Prior work has studied coordination in collaborative writing [10, 8, interalia]), as well as the social aspects that arise as a result of edits and comments made by collaborators [1], and has developed tools for supporting such collaboration (e.g., Quilt [4]). Most closely related to our work are methods and tools for supporting improved *awareness* of collaborators about document changes. These include flexible diffing for reporting significant changes [9], methods for categorizing edits and presenting them to authors [5, 11, 12], methods for detecting and alerting authors about merge conflicts [6], and methods for detecting structure in text and helping co-authors create coherent documents [2]. Our approach differs from these prior works in that it leverages natural language processing (NLP) methods to automatically detect significant changes, without requiring authors to specify explicitly what changes are of interest to them. Furthermore, these capabilities go beyond prior approaches for supporting change awareness in that they can predict parts of the document that have not changed but that are likely to require edits as a result of other changes.

Approach

This section describes methods for: (1) tracking paragraphs; (2) identifying significant changes to paragraphs, and (3) predicting future changes to the document. Capability (1) is required for both monitoring changes and predicting them. Capability (2) can help identify important changes and decide whether to alert authors of a change. Capability (3) is required in order to draw authors' attention to other parts of the document that might require changes as a result of a recent edit. We describe our use of lightweight NLP techniques, which provide the foundation for these important capabilities.

Tracking paragraphs

As a document is being edited, paragraphs can be added, moved or deleted. Thus, we developed an algorithm for tracking paragraphs across revisions. The algorithm compares paragraphs based on their Levenshtein distance, that is, the number of changed characters required to move from one paragraph to another, including additions, deletions, and substitutions. We use the Levenshtein edit ratio (LR) measure, which is defined as follows:

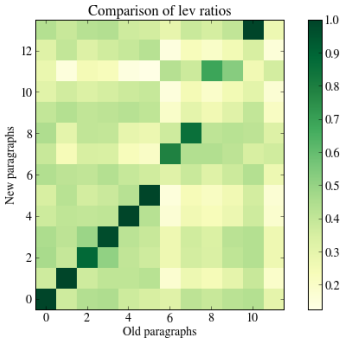

Figure 1: Levenshtein ratio between paragraphs in two versions of a document. A higher ratio is shown in darker color. Paragraphs with the highest ratio are matched.

$$LR(a,b) = 1 - \frac{Levenshtein Distance(a,b)}{max(|a|,|b|)}$$

A LR close to 1 indicates high similarity, while a ratio close to 0 indicates the opposite. For each revision of a document, the algorithm computes the LR between each paragraph in the current version and each paragraph in the previous version. Two paragraphs are mapped across the revision if they each have the highest LR with the other. If a paragraph in the new version has no matching paragraphs in the old version (i.e., none of its ratios are above a threshold of 0.4), we label the paragraph as an addition. If a paragraph in the old version is not matched with any paragraphs in the new version, we consider it deleted. Figure 1 illustrates this algorithm: some paragraphs found clear matches (e.g., 0, 1, 2), while paragraph 11 from the old version was deleted and 6, 9, 10, and 12 were added in the new version.

While Wikipedia provides a diff visualization using an algorithm that maps unique sentences between revisions[1], the use of LR with forward and backward mapping allows for more robust detection of paragraph movement even when there are significant changes to the content.

Identifying Significant Changes
We consider a change to be significant if it results in a noticeable change in the paragraph's topic and content. To detect such change, we compute the cosine similarity between word vectors that represent the two versions of a paragraph, using a Latent Semantic Indexing topic modeling approach [3]. If the cosine similarity between new and old versions of a paragraph is below a threshold of 0.8, we consider the edit significant. The threshold was determined empirically by manually evaluating differences between paragraphs with varied similarity values.

[1] http://en.wikipedia.org/wiki/User:Cacycle/diff

We chose a topic modeling approach for this task because, in contrast with the LR approach used for mapping paragraphs, topic modeling considers the content of the text. To illustrate, if two of a paragraph's sentences switch places, the LR would decrease despite no meaningful change in content. On the other hand, changing a couple of key words will slightly lower the LR, but could drastically affect the content.

Predicting Future Edits
We hypothesized that a paragraph that underwent a significant change would prompt edits in related paragraphs, and investigated which paragraphs are likely to change in future revisions as a result of such edits.

We considered three possible types of inter-paragraph relationships: (1) *proximity*, i.e., neighboring paragraphs; (2) *edit histories*, i.e., paragraphs that tended to be edited together in previous revisions, and (3) *topic similarity*, i.e., paragraphs with similar content. For (2), we labeled pairs of paragraphs that changed or remained unchanged together in at least 5 of the previous 10 revisions. We chose a window of 10 revisions as it allowed us to obtain a meaningful signal of correlation between edits, while not looking too far in the past where the content might be significantly different. For (3), we labeled pairs as related if their cosine similarity was above an empirically determined threshold of 0.4. (There are rarely paragraphs with a higher similarity than 0.4, while a lower similarity does not really capture a similarity in content.)

Empirical Evaluation

This section describes a preliminary evaluation of the proposed methods for tracking paragraphs, detecting significant edits, and predicting future changes using a corpus of Wikipedia articles and their revision histories.

Data

We used the complete revision histories of 41 different articles chosen from a diverse set of topics, ranging from famous people and places to mathematical algorithms to novels. We removed Wikipedia-specific tags that indicate formatting and other irrelevant data, and eliminated versions of articles under 150 characters, as they did not contain enough text. To focus on revisions that contained a substantial change, the versions with simple typo fixes (Levenshtein distance < 15) were eliminated.

Findings

This section describes findings from a preliminary empirical evaluation. We focused mostly on the prediction of future edits, but also manually evaluated paragraph mapping and significant change detection, which form the basis for edit prediction.

Detecting Significant Changes

Figure 2 shows the distribution of topic similarity between paragraphs in consecutive versions. As shown, most edits are minor. With the threshold of 0.8, we labeled about 15% of the edits as significant.

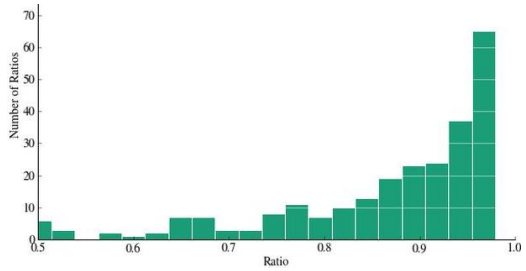

Figure 2: Topic similarity of paragraph pairs.

We manually evaluated a random sample of more than 100 mapped paragraphs to ensure that paragraphs were correctly mapped and to determine whether the method successfully classified edits as significant or insignificant. Figure 3 shows an example of a significant edit: while the bottom paragraph is clearly a revision of the same paragraph (recurring text shown in bold), the edit is significant because it adds new content and alters the tone of the text. Overall, we found that paragraphs were rarely mapped incorrectly (less than 5%), even when their content (as measured by topic similarity) changed significantly. Similarly, we found the classification of significant edits to be correct in most cases, though this is a more subjective measure which we plan to evaluate further in future work.

After complaints arose during the renovation in the late 1980s of the early residential colleges, a swing dormitory was built in 1998 to facilitate housing students during the ***massive overhaul*** *of buildings that had seen only intermediate* ***improvements in plumbing, heating, electrical and network wiring,*** *and general maintenance over their 35-to-80-year existence. The new Residence Hall is not a residential college, but houses students from other colleges during renovations.* ***It is commonly called "Swing Space" by the students; it's official name "Boyd Hall" is unused.***

In 1990, Yale launched a series of ***massive overhaul****s to the older residential buildings, whose decades of existence had seen only routine maintenance and incremental* ***improvements to plumbing, heating, and electrical and network wiring****. Calhoun College was the first to see renovation. Various unwieldy schemes were used to house displaced students during the yearlong projects, but complaints finally moved Yale to build a new residence hall between the gym and the power plant.* ***It is commonly called "Swing Space" by the students; its official name "Boyd Hall" is unused.***

Figure 3: A revised paragraph with a significant edit.

Predicting Future Edits

We confirmed our hypothesis that a significant edit to a paragraph often triggers edits in other paragraphs in the next revisions. This is demonstrated in Figure 4, where each line represents a paragraph that is edited significantly at time 0. After a period of relative inactivity, a significant edit triggers further adjustments in the near future (as represented by downward spikes). We evaluated

the three inter-paragraph relationships (proximity, edit history, and topic similarity) to determine which paragraphs are more likely to require further editing after a significant change occurs in the article.

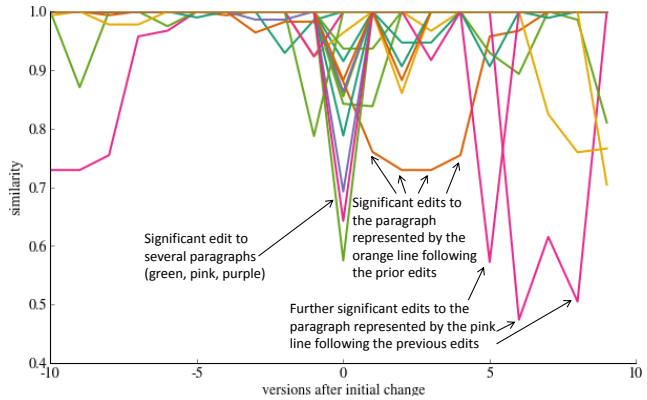

Figure 4: After a significant change at version 0, consequent edits are more frequent.

To evaluate the accuracy of the predictions of paragraphs that will require attention following a significant edit to another paragraph, we check whether the predicted paragraphs undergo a significant change in consequent revisions. Specifically, we iterate over all the revisions of each of the documents d, and for each revision d_t we use only information known at that time (t) to predict paragraphs that will change in revisions d_{t+1} to d_{t+10}. We look at two measures: (1) whether paragraphs related to a "triggering paragraph" that changed significantly in revision d_t underwent at least one significant change in revisions d_{t+1} to d_{t+10}, and (2) whether those paragraphs continue to be related by edit patterns to the triggering paragraph in revisions d_{t+1} to d_{t+10}, that is, whether they keep changing together (or remain unchanged together) more than half the time in the following 10 revisions. This second measure provides a stronger indication of a possible interdependency between the paragraphs.

The frequency of occurrence of each type of inter-paragraph relationship varied. On average, 1.6 pairs of paragraphs related by edit history were found per revision of the text. A pair related by topic similarity was only found once every ten versions. Proximity relationships always exist in each version, as each paragraph has at least one neighboring paragraph (and most have two).

With respect to the first measure, the relationship of edit history is much more predictive of a future significant edit than proximity, as shown in Table 1. We computed the proximity measure for the original paragraph as well as randomly sampled paragraphs and found a significantly lower likelihood than for paragraphs related by edit similarity.

Relationship	**Percentage**
Edit History	81%
Proximity	24%
Paragraph with original change	40%
Randomly sampled paragraph	15%

Table 1: The percentage of time a paragraph has been edited significantly in the 10 revisions following a significant change.

With respect to the second measure, we found that 71% of paragraph pairs related by edit history and 63% related by topic similarity continued being modified (or unmodified) together in the next 10 versions. (Proximity was not evaluated for this measure, as it did almost no better than the random control for the first, less strict, measure.) Further, we obtained a recall of 80% with predictions based on edit history and topic similarity. That is, only 20% of the paragraphs that should have been labeled as related were not included in the set of related paragraphs predicted by edit history and topic similarity.

Discussion and Future Work

Our long-term goal is to build a system that will improve the collaborative writing process. The methods we developed for tracking paragraphs throughout document revisions, detecting significant edits, and predicting future edits provide the basis for such a system. These capabilities can inform decisions about alerting authors to specific changes and places in the document that are likely to require revisions.

In future work, we will incorporate information about author identity to design personalized alerts for authors depending on the edits they have made. We also plan to develop additional algorithms for summarizing changes and to investigate the use of algorithms that measure text coherence (e.g.,Textiling [7]) in order to alert authors to parts of the text that could be improved. Finally, we will explore alternative interface designs for presenting the information chosen by the algorithms for authors to consider and ways to incorporate explanations for system recommendations (e.g., "we recommend reading the introduction because it is often edited following edits to the results section").

Acknowledgments. The work was funded in part by the Nuance Foundation. We thank Stuart Shieber for helpful discussions.

References

[1] Birnholtz, J., Steinhardt, S., and Pavese, A. Write here, write now!: an experimental study of group maintenance in collaborative writing. In *Proc. CHI 13'*, ACM (2013), 961–970.

[2] De Silva, N. H. *A narrative-based collaborative writing tool for constructing coherent technical documents*. PhD thesis, Citeseer, 2007.

[3] Deerwester, S. C., Dumais, S. T., Landauer, T. K., Furnas, G. W., and Harshman, R. A. Indexing by latent semantic analysis. *JASIS 41*, 6 (1990).

[4] Fish, R. S., Kraut, R. E., and Leland, M. D. Quilt: a collaborative tool for cooperative writing. In *ACM SIGOIS Bulletin*, vol. 9, ACM (1988), 30–37.

[5] Fong, P. K.-F., and Biuk-Aghai, R. P. What did they do? deriving high-level edit histories in wikis. In *Proc. OpenSym 10'*, ACM (2010), 2.

[6] Hainsworth, J., and Adviser-Cook, P. *Enabling truly collaborative writing on a computer*. Princeton University, 2006.

[7] Hearst, M. A. Multi-paragraph segmentation of expository text. In *Proc. ACL 94'*, ACL (1994), 9–16.

[8] Kittur, A., Suh, B., Pendleton, B. A., and Chi, E. H. He says, she says: conflict and coordination in wikipedia. In *Proc. CHI 07'*, ACM (2007), 453–462.

[9] Neuwirth, C. M., Chandhok, R., Kaufer, D. S., Erion, P., Morris, J., and Miller, D. Flexible diff-ing in a collaborative writing system. In *Prof. CSCW 92'*, ACM (1992), 147–154.

[10] Neuwirth, C. M., Kaufer, D. S., Chandhok, R., and Morris, J. H. Computer support for distributed collaborative writing: A coordination science perspective. *Coordination Theory and Collaboration Technology, ed. GM Olson, TW Malone, and JB Smith, Lawrence Erlbaum Associates, New Jersey* (2001).

[11] Papadopoulou, S., and Norrie, M. C. How a structured document model can support awareness in collaborative authoring. In *Proc. CollaborateCom 07*, IEEE (2007), 117–126.

[12] Tam, J., and Greenberg, S. A framework for asynchronous change awareness in collaborative documents and workspaces. *International Journal of Human-Computer Studies 64*, 7 (2006), 583–598.