# Detecting Interaction Patterns in Educational Collaborative Writing

NIELS SEIDEL, MARC BURCHART, and JOERG M. HAAKE, FernUniversität in Hagen, Germany
CLARA SCHUMACHER, Universität Potsdam, Germany
JAKUB KUZILEK, Humboldt Universität zu Berlin, Germany

Writing and collaboration are crucial skills in professional and academic settings. However, assessments of collaborative writing often focus only on the final text, overlooking individual contributions and the diverse strategies students employ during the writing process. To support teachers in interpreting and assessing student behaviour, we propose analysing writing data at a granular level, down to individual characters, and using user sessions as observation units to capture coherent interaction patterns. This paper introduces three methods for identifying interaction patterns in collaborative writing. The first method classifies session types by analysing features such as writing, reading, and communication behaviours, session length, and the number of group members collaborating synchronously. The second method identifies frequent sequences of session types by examining their order, enabling process analyses at an abstract yet manageable level compared to using log data. The third method focuses on text-level collaboration by evaluating the frequency of text passage modifications made by the original author or other group members. This approach quantifies individual collaboration and, at the group level, identifies isolated versus closely connected group members, shedding light on the mode of collaboration and degree of group cohesion.

We demonstrate these three methods in a case study involving two cohorts, $K_A = 294$ and $K_B = 242$ groups of up to 9 learners ($N_A = 1,848$, $N_B = 1,463$). The interaction patterns identified using these methods are intended to help teachers understand collaborative writing processes and identify situations where the participating learners require support.

CCS Concepts: • **Human-centered computing** → **Social network analysis**; • **Information systems** → *Clustering*; **Synchronous editors**.

Additional Key Words and Phrases: Collaborative Writing, Collaborative Learning, Learning Analytics

**ACM Reference Format:**
Niels Seidel, Marc Burchart, Joerg M. Haake, Clara Schumacher, and Jakub Kuzilek. 2025. Detecting Interaction Patterns in Educational Collaborative Writing. In *Proceedings of CSCW'25.* ACM, New York, NY, USA, Article CSCW 286, 25 pages. https://doi.org/10.1145/3757467

## 1 Introduction

In recent decades, synchronous and asynchronous collaborative writing has become a common practice for knowledge workers [7, 30, 40]. Higher education institutions, therefore, incorporate collaborative writing into their curricula (e.g., [49]), recognising it as essential preparation for academic and professional endeavours involving textual work. Collaborative writing assessment predominantly focuses on the final text [31, 56], overlooking individual contributions and the text

Authors' Contact Information: Niels Seidel, niels.seidel@fernuni-hagen.de; Marc Burchart, marc.burchart@fernuni-hagen.de; Joerg M. Haake, joerg.haake@fernuni-hagen.de, FernUniversität in Hagen, Hagen, Germany; Clara Schumacher, clara.schumacher@uni-potsdam.de, Universität Potsdam, Potsdam, Germany; Jakub Kuzilek, jakub.kuzilek@hu-berlin.de, Humboldt Universität zu Berlin, Berlin, Germany.


*CSCW'25, –, –,*

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM
https://doi.org/10.1145/3757467

creation process. However, students engage in text production with varying levels of involvement and employ diverse strategies [33, 43, 56].

Teachers must clearly understand student activities during collaborative writing to assess performance and identify support needs [18, 48, 56]. Ignoring individual contributions and the collaborative process limits their ability to detect and address problematic writing behaviours of individuals and within groups. This challenge increases in large classes with multiple groups or in distance education, where monitoring is particularly difficult. As a result, teacher support is often restricted to final assessments, missing opportunities for timely intervention during the writing process.

As a prerequisite for supporting teachers in interpreting and assessing student behaviours, such behaviours must be identified. In this paper, we, therefore, explore the central research question (RQ) of methods for identifying interaction patterns in collaborative writing. In CSCW, interaction patterns are recurring user behaviours when engaging with a system, including actions, reactions, and decisions in computer-mediated peer interactions [9, 25].

Research indicates that specific activity sequences recur during writing processes, enabling the identification of interaction patterns in collaborative writing processes [32]. Understanding these patterns requires analysing both individual and group-level writing processes. A suitable observation unit is crucial for capturing coherent interaction sequences [22]. We propose using individual user sessions, encompassing the actions performed within a session. This approach follows a three-step procedure:

(1) Classifying individual user sessions in collaborative writing: Individual user sessions are more common during complex tasks or extended collaboration with limited synchronous meetings. Classifying sessions into types based on recurring action patterns helps identify interaction trends. These session types (ST) offer a framework for analysing collaborative writing behaviour, including communication, coordination, and synchronous activities. This enables teachers to provide targeted support, fostering productive collaboration and mitigating sessions with less productive patterns. Thus, the first research question arises:

**RQ1:** How to classify individual user sessions in collaborative writing into session types?

(2) Identifying frequent sequences of session types: Collaboration over an extended period typically involves multiple sessions, which can be categorised as outlined in step (1). These sessions may follow various paths towards achieving the overall goal. Understanding the recurrent sequences of these sessions is crucial for gaining insights into the collaborative writing process. By analysing the sequence of user sessions, teachers can identify group interaction patterns in how learners of a group coordinate their efforts, distribute tasks, and progress through a sequence of user sessions of the same or different types. The second research question is therefore:

**RQ2:** What frequent sequences of user sessions exist in collaborative writing?

(3) Determine collaboration on text-level: Analysing individual text contributions and text-based group interactions is essential to understand how members support each other and work towards a common goal in text production. Examining these dynamics helps teachers to determine whether the process reflects true collaboration or consists of separate individual efforts.

Therefore, the third research question is formulated as follows:

**RQ3:** How can the extent to which learners genuinely collaborate in writing a text together be determined?

Answering these RQs is a prerequisite for supporting teachers in identifying problems in group writing processes, providing appropriate interventions (such as feedback and instructions), and measuring their interventions' success.

Existing approaches to analysing collaborative writing often overlook communication and coordination interactions [8, 33, 43] and insufficiently differentiate between synchronous and asynchronous collaboration [8] (RQ1, RQ2). Moreover, they rarely utilise fine-grained text-level interactions to identify meaningful interaction patterns (RQ3). The multidimensional nature of collaborative writing processes has not been effectively aggregated at the level of natural working units, such as user sessions (RQ1, RQ2). Additionally, prior attempts to identify interaction patterns have been constrained by limited empirical data [8, 33, 41, 43, 56] (RQ1–3). Consequently, these approaches failed to comprehensively capture and analyse interaction patterns in collaborative writing activities.

We propose the concept of interaction patterns to describe the interplay of group members' activities in creating a final text. Section 2 reviews related work to identify gaps and building blocks for analysing writing processes. Section 3 introduces new methods to measure and analyse individual and collaborative writing, enabling the identification of interaction patterns and problematic behaviours. Section 4 evaluates these methods through a case study of two large distance learning courses. Finally, we conclude with a discussion on the potential educational applications of our findings.

This paper presents new methods to identify interaction patterns in collaborative writing and their application in a case study of two large distance learning courses. CSCW/L researchers can use these methods to analyse group behaviour, while tool developers can leverage the measures and algorithms to create dashboards and awareness widgets. Teachers implementing collaborative writing projects may use these metrics to better support students needing additional assistance.

## 2 Related Works

### 2.1 Benefits and Challenges of Collaborative Writing in Education

For group learning, synchronous writing holds particular relevance. First, written text remains the dominant form of knowledge representation in teaching and learning [21]. Second, group members can work on documents simultaneously and coordinate their actions with other participants through the support of synchronous group awareness [13]. Third, collaboration in text production leads to improvements in content [31], expression [54], complete task fulfilment [44], fewer grammatical errors [44, 54], and shorter texts [44].

Collaborative writing groups in distributed learning settings should jointly produce a text matching the requirements of a teacher's task (e.g. [23, 31]). Teachers often assume that group writing functions smoothly and that the collaboration environment sufficiently supports the group's collaborative writing. However, teachers can only verify this occasionally for a small number of groups [23]. The extent to which there are problems with collaboration and a group is therefore not working and support is required cannot be verified by teachers with the help of common collaborative writing tools (e.g., Google Docs [2]). Referring to Hoppe et al. [19], inactivity is the most important indicator of inefficient group work. Strauß and Rummel [45] also mentions other problems that frequently occur in learning groups and are probably also encountered in collaborative writing. However, which information helps teachers interpret collaborative writing situations is still an open question [48].

In addition to the challenges mentioned above, which are due to a lack of monitoring, there are also difficulties in assessing individual performance. Collaborative writing assessment predominantly focuses on the final text [31, 56], overlooking individual contributions and the text creation process. Like other collaborative learning activities, students engage in text production with varying levels of involvement [19] and employ diverse strategies [32, 33, 55]. This oversight in assessment practices limits teachers in evaluating individual group members' participation and learning progress and from identifying issues arising from individual and group dynamics. These challenges increase with the number of groups and participants, making assessment of large numbers of collaborative writing groups more difficult.

### 2.2 Analysis of individual writing process

Writing process analytics aims to extract implicit knowledge about the temporal dynamics of the writing process from the sequences of user activities related to text revisions [43, 51, 56] and activities that accompany writing [19, 55]. Text revisions are either triggered by the user (e.g., [43, 51]) or triggered automatically by the system during the ongoing writing process. This results in differences in the granularity of the units considered in analysing a writing process. Revisions triggered by the user affect a variable number of text sections. Editing processes within a revision are not recorded, so the process analysis must be based on comparisons of versions of the text defined by the revisions (e.g., [22, 26, 43, 51]). The revisions triggered by the system are more continuous and concern short text sections of usually a few consecutive characters. A process analysis can thus be based on more continuous and fine-grained activity logs. However, these fine-grained logs are often inaccessible in frequently studied platforms like Google Docs [20, 43].

One simple methodological approach for analysing the writing process is to describe the activities for fixed periods using suitable metrics (e.g., number of text changes) and observe their changes over time [51]. Also, prior works that utilised information visualization techniques aggregate data to represent the process of collaborative learning [41, 43, 51], e.g., revision graphs or topic evolution graphs.

The temporal organization of user activities is reflected in sequences that may serve as conceptual units of analysis for understanding work routines in online knowledge collaboration. Analysing these activity sequences' distribution, evolution, and impact can enhance theories of online collective action, distributed work, and shared governance [22]. Sequence mining methods help to discover patterns across time or position in a given dataset of activity sequences. In sequence mining, the order of activities is essential. Consequently, possible permutations of the activities must be considered as potential candidates for frequently occurring sequences. To prevent less significant or noisy activities (e.g., long rows of scroll events) leading to less meaningful or biased patterns, an attempt is made to classify activities before the sequence mining. For instance, Hoppe et al. [19] grouped activities into coordination tasks, inactivity, monitoring, and minor and major text contributions.

For the identification of frequent sequences, methods like *PrefixSpan* [34] and *FPgrowth* [15] are utilised. However, these methods can only be used efficiently up to a specific sequence length and number of sequences.

In addition to searching for patterns, another approach is quantifying the similarity of pairs of sequences and identifying clusters of similar sequences (and thus similar user behaviours) using similarity matrices Hoppe et al. [19]. However, the methods used to determine the similarity of two sequences (e.g., Levenshtein distance [24]) are weak, as they only assume similarity if there is no temporal offset between two identical subsequences. Consequently, this method is only suitable for highly aggregated, coarse-grained activity sequences.

Another important line of research employs Markov chains to model transitions between states of user activities (e.g., [1]). In a first-order Markov model, the future state depends only on the current state, ignoring the sequence of preceding states. Higher-order Markov chains address this limitation by defining states as combinations of $n$ consecutive first-order Markov states [11]. Consequently, a future state is determined by a sequence of $n$ preceding activities, providing a more nuanced understanding of activity transitions. In contrast, the number of states exponentially increases, and some transitions are unlikely to be observed in smaller datasets.

Apart from these methodological limitations, another critical challenge is the lack of empirical studies that utilise large datasets for process analysis [33, 41, 43]. Identifying statistically reliable process patterns requires a substantially larger sample size, including more groups and participants.

Moreover, these patterns are heavily influenced by contextual factors such as the specific group task, the number of (active) group members, and the duration of the task or observation period [20]. Addressing these gaps is essential for advancing the robustness and generalizability of process analysis in collaborative settings.

This paper aims to extend existing approaches in analysing writing processes by identifying frequent sequences of user session types. We consider sessions as a natural and comprehensive unit of observation that covers all user activities supported by the collaborative writing environment, including communication and coordination tasks, which are often overlooked. This aggregation of data simplifies the interpretation of session type sequences. Additionally, we propose an efficient sequence mining method capable of identifying the most frequent session type sequences, even when dealing with large numbers of sequences from longer collaboration periods.

### 2.3 Analytics of collaborative writing processes

The analysis of collaboration processes has a long tradition in CSCL (e.g., [12, 17, 38]). The analysis of learning processes has been heavily influenced in recent years by research in the field of Learning Analytics, with collaborative learning playing only a minor role [52]. An example of joint research in CSCL and Learning Analytics is provided by Hoppe et al. [19], who analysed activity sequences in collaborative writing groups by aligning and clustering common activities. Another instance is the work of Perera et al. [36], who employed sequential pattern mining techniques to identify best practices among students collaborating on software development projects.

Graph representations for groups are commonly used in CSCL/CSCW and learning analytics. Suthers and Rosen [46], for instance, described the relationships of logged events in a collaboration space as contingency graphs. They assumed that relationships between events are contingent or incidental to the situation rather than causal or deterministic. They distinguished several contingency types, such as *media contingency*, where an event is connected to an object that another person previously manipulated. Similarly, topic-based collaboration networks have been proposed by Southavilay et al. [43]. Students represented as nodes in the network are connected if they contribute text to the same topic. These graph representations are commonly used for Social Network Analysis (SNA) [16, 50] and for pattern detection within these networks [14]. Zhang and Chen [56] applied SNA to collaborative writing by representing the relative number of written words through the size of the graph nodes. The links between the nodes indicate the number of revisions of words written by another group member. However, these text-level interactions have only been presented to students (N=26, groups of 4-6 students) in a group dashboard and have not been used as a research method to analyse collaborative writing.

Another strand of research examined collaborative writing processes to identify writing strategies. Onrubia and Engel [33] observed six small student groups performing a writing task and identified five writing strategies related to writing processes and four phases of collaborative knowledge construction. Zhang [55] analysed conversations and produced texts of 35 dyads of learners engaged in a collaborative writing task to determine individual comparative involvement in major aspects of the task. Through cluster analysis, five collaboration types were identified, each representing a distinct interactional pattern regarding pair member engagement. However, these strategies are derived from small samples and therefore have been analysed qualitatively. Olson et al. [32] analysed the data traces of collaborative writing from 96 documents produced in Google Docs by 32 groups of 4-5 students in three different cohorts of a project management class. By examining the documents, their revisions, and derived multi- and single-author sessions, they analysed various working styles, evidence of leadership, and the balance of participation. External graders scored the quality of the documents, revealing that document length, leader presence, and balance of participation were directly related to quality.

The studies discussed highlight various aspects of collaborative writing strategies and patterns. However, the reliability of these findings is limited due to the relatively small sample sizes and specific contexts. A significantly larger number of groups and participants is necessary to identify statistically reliable collaboration patterns. The current research presents a need for broader studies to enhance the generalizability of these findings across diverse writing tasks and educational settings. Moreover, prior research on text-level collaboration has relied on broad, asynchronous, and (user-led) revision-based analyses rather than examining changes at or near the character level that help to identify synchronous and asynchronous collaboration.

## 3 Classifying individual user sessions in collaborative writing

Collaborative writing environments, such as Google Docs or Etherpad Lite [47], enable users to create and edit shared documents collaboratively. These environments allow users to access content, position their cursors, and modify text through actions like adding, deleting, or changing properties. Many also support commenting on marked text. While asynchronous editing is possible, shared editors excel at synchronous collaboration by enabling concurrent access and providing group awareness features, such as user lists, remote cursor displays, and colour-coded text contributions, for coordination.

From the system's perspective, collaborative writing is a sequence of user actions on the shared document at specific times. The system maintains a detailed log of these actions, which can be replayed to trace document development. Each log entry typically includes a timestamp, user ID, action type, and action parameters, forming the foundation for logfile-based analytics.

For the presented method, the log data is grouped into user sessions. A user session represents a time-bound sequence of activities in the collaborative writing environment, described by the log data captured during that session. Sessions start after and end before a defined minimum period of inactivity (e.g., 30 minutes). Grouping user activities into sessions contextualises actions within specific timeframes, capturing bursts of focused effort and enabling a clearer understanding of user interactions. They naturally align with task-based activities, making it easier to evaluate collaboration efficiency and individual contributions.

Features regarding writing, reading, and communication behaviours can be extracted from each user session. Writing activity features include the frequency and amount of text added, removed, and formatted. Reading activity is measured by the frequency of scroll events in the text window and the number of text lines scrolled. Communication activities are represented by the frequency, overall text amount, average text amount of chat messages, comments anchored in the text, and comment replies. In addition, the session length is considered by features expressing both the text per time added to (or removed from) the document and the text written for communication purposes. Notably, the number of characters entered per unit of time provides insight into whether the text was copied from another location or composed within the editor. Sessions that exceeded or fell short of the mean session length plus or minus three times the standard deviation should be excluded as outliers.

We use clustering methods like k-means [27] to identify natural groupings and patterns in the session data. Clustering is an unsupervised learning method that does not require labelled data, making it particularly valuable when predefined session categories are unavailable [42, 53]. Manual labelling of sessions would not be feasible for large amounts of data. The goal is to uncover hidden patterns or groupings within the data. By analysing multiple features and describing each session, clustering can simultaneously consider these dimensions. This allows for identifying sessions with similar behavioural characteristics, even in complex feature relationships.

Grouping sessions into clusters simplifies the dataset by reducing complexity, enabling a more straightforward analysis of trends and behaviours. Instead of analysing individual session features, clustering considers the overlapping characteristics that many user sessions share. Thus, the clustering approach helps to understand the complex interaction of individual features while still allowing a comprehensive but understandable set of features to be considered. These clusters can highlight different types of user behaviour, such as engaged users, casual users, or explorers. The insights gained from such analysis can inform strategies for personalisation, design improvements, or targeted interventions.

Additionally, clustering can help detect rare but important behaviours, such as operating errors, unusual user journeys, or unique user segments. These insights can be critical for refining teaching instructions and enhancing user experiences.

The resulting types of user sessions offer insights into common interactions during user sessions. The frequency and distribution of these session types provide valuable information on how individual groups and individual group members collaborate on text production. It is possible to estimate the extent to which the identified session types are conducive and desirable to collaboration.

## 4 Identifying frequent sequences of user sessions in collaborative writing

In the previous section, user sessions were classified into different session types. We assume that typical users work over a longer period of time with phases of inactivity, leading to multiple sessions. We also assume that sequences of these sessions can be assigned to various session types and that there is a certain likelihood that one session type is followed by another. Consequently, specific sequences of session types are more likely than others. These sequences differ in their likelihood, length, and order of session types.

Especially longer sequences cannot be identified using Markov chains or state transition charts [28]. Instead, algorithms for frequent sequence mining (e.g. [15, 34]) are better suited for this purpose. However, known representatives of these algorithms, such as*PrefixSpan* [34] and *FPgrowth* [15], require a lot of resources and computing time. Given our specific application, we present a modified method (Algorithm 1) that delivers valuable results quickly and without extensive computational resources.

**Algorithm 1:** Identify frequent sequences of session types for given sequence length and minimum support

```
Input  : states, state_transitions, min_length, max_length, min_support
Output : Updated states

if max(str_length(states)) > max_length then
    return(states);
for i ← 1 to length(states) do
    if str_length(states [i]) ≥ min_length then
        last_item ← substr(states [i], str_length(states [i]), str_length(states [i]));
        targets ← subset(state_transitions, from == last_item and prob > min_support);
        for j ← 1 to length(targets) do
            if length(targets [j]$to) > 0 then
                states ← append(states, paste0(states [i], targets [j]$to));
return(unlist(states));
```

For this reason, we implemented an algorithm that determines candidate sequences of a fixed length, considering probable transitions between pairs of items.

In the first step, the probability of transitioning from one session type to another has to be computed. Transitions with a probability below a defined threshold (e.g., 0.05) should be omitted to improve clarity.

Second, based on the probabilities for the transition between session types, candidate sequences of a defined length (e.g., 2 to 10) whose successive transitions have a probability above a specified threshold (transition *support*) are determined. Thus, the number of theoretically possible sequences gets reduced to a manageable number of probable sequence candidates.

Third, to determine the most frequent sequences, the *support* measure is employed for each candidate sequence. The *support* is defined as the proportion between the number of sequences containing a specific sequence and the number of all sequences. It ranges from 0 to 1, where 0 means that the sequence did not occur, and 1 indicates that the sequence occurred in all activity sequences.

Finally, the algorithm outputs a list of frequent sequences within the given range of sequence length, sorted by the corresponding *support* in descending order.

In practice, frequent sequences can be computed for multiple groups (e.g., a cohort), individual groups, and individual group members. This approach allows for identifying frequent and, thus, dominant sequences of session types. In particular, it becomes possible to determine whether session types deemed less conducive to learning occur consecutively or are interspersed with desired session types. Additionally, frequent sequences can be analysed over time to evaluate the progression of collaboration.

## 5 Collaboration on Text-Level

The third method presented in this article aims to determine the extent to which learners genuinely collaborate in writing text together. This method is based on three key aspects. The first aspect considers the specific location in the text where a group member made a change. The second aspect focuses on the context of the text at that position, including which group member last inserted characters immediately to the left and right of the edit. The third aspect considers the complete history of editing operations, analysing how many edits each group member has made adjacent to characters written by others. These network analysis methods provide insights into collaborative text editing dynamics on individual and group levels.

*5.0.1 Preprocessing of changesets.* A detailed record of change operations is needed as a precondition for this method. Modern collaborative writing environments track individual or successive change operations in so-called changesets. A changeset is defined as the smallest set of subsequent text change operations stored by the collaborative writing environment to enable revisions and track individual contributions to the text. Change operations include adding and removing one or many characters. On average, a changeset affects two characters [5]. When text is pasted from the clipboard into the text, the changeset contains the entire pasted text. In addition, text formatting operations are also stored as a changeset.

Using changesets, it is possible to determine who contributed a character or a sequence of characters to the document at any given moment. To apply the method presented here (see Figure 1), for each changeset, the author of the adjacent left and right characters of the inserted or removed character in the current version of the document is identified. This involves counting how often someone has changed the document at a position where the neighbouring characters were previously contributed by the same or another group member. Simultaneously, it is counted whose text was supplemented or deleted by another group member. As described in Algorithm 2, these processing steps are repeated for every changeset, thus every change made to the text.

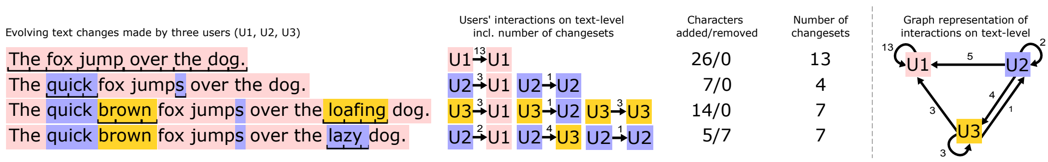


Fig. 1. Schematic example of how three users (U1, U2, U3) make changes to a text. Text changes are recorded as changsets and counted as interactions at text-level so that these interactions can be represented as a graph.

**Algorithm 2:** Process text changes and compute author relationships

```
input  : A data frame data_set with columns group_id, author_id, line_position, and
         text_changes
output : Author relations concerning text changes

groups ← distinct(data_set, group_id)
for j ← 1 to length(groups) do
    group_id ← groups[j]
    group ← filter(data_set, group_id)
    max_lines ← max(group[line_position])
    char_authors ← list(0, max_lines)
    // Iterate changesets of a group
    for i ← 1 to length(group) do
        changeset ← group[i]
        text_length ← changeset[text_changes]
        start_pos ← changeset[line_position]
        end_pos ← start_pos + text_length
        left_neighbor ← char_authors[start_pos]
        right_neighbor ← char_authors[start_pos + 1]
        author_relations ← append(group_id, author_id, left_neighbor, right_neighbor)
        // Update text additions and deletions
        if changeset[text_changes] > 0 then
            i ← start_pos
            while i < end_pos do
                append(char_authors, i, author_id)
                i ← i + 1
        if changeset[text_changes] < 0 then
            i ← start_pos
            while i < end_pos do
                remove(char_authors, i)
                i ← i − 1
```

*5.0.2 Graph Measures on Individual and Group Level.* Using the preprocessed changesets, it is possible to quantify students' text-related collaboration on individual and group levels. Therefore, the results of the preprocessed changesets are represented in an adjacent matrix, whose rows and

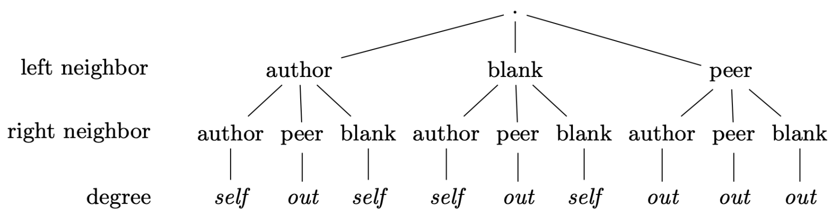


Fig. 2. Determination of self-degree and out-degree based on the authorship of characters to the left and right of the location where the text change occurred.

columns represent the group members. Each cell indicates the number of relationships between a changeset author (column) and the person whose text is adjacent to the position of the changeset (row). From this adjacent matrix, the pairwise relationships of group members are represented in a weighted graph $G = (S, L)$, where $S$ is a set of *nodes* representing students $s$ in the group, and $L$ is a set of paired and directed links. Each *link* $l$ $\{s_x, s_y\}$ represents an editing operation of student $s_x$ directly next to the text formerly added by $s_y$.

To characterise the graphs, three basic network analysis metrics are used at the individual level to describe the degree of collaboration: First, $selfdegree$ is the number of writing operations the author uses on the text written by this author. Secondly, the *outdegree* is the author's number of writing operations on text written by another author. Third, the *indegree* is the number of writing operations other authors have performed on the author's text. The distinction between $selfdegree$ and *outdegree* is shown in Fig. 2, e.g., if the character left to the position in the text was written by the author but the right neighbouring character by a peer, then the current text change is added to the outdegree. For these three measures ($selfdegree_length$, $outdegree_length$, $indegree_length$), the number and length of affected characters are also determined.

As a result, these graph metrics make it possible to characterise the collaboration between the group members when creating the text. For example, it is possible to identify group members who have been writing text without contributing to the text of the other group members ($selfdegree > 0$, $outdegree = 0$). It is also possible to identify those group member who have hardly written any longer texts of their own (low $selfdegree$) but who have extensively revised their peers' texts (high *outdegree*). It can also be determined whether the collaboration between pairs was based on reciprocity ($outdegree > 0$ and $indegree > 0$) or whether one person gave help rather than receiving it ($outdegree > 0$, $indegree = 0$).

At the group level, the number of active group members (*activeMembers*) and, thus, the group size is determined first. The *outdegree* and *indegree* of the group members are used to decide whether they received no, little, average, or a lot input on their texts (*noReceiver*, *lowReceiver*, *mediumReceiver*, *highReceiver*) or whether they contributed nothing, little, average, or a lot to the texts written by their team colleagues (*noHelper*, *lowHelper*, *mediumHelper*, *highHelper*). The assignment is based on the quartiles Q1–Q4: little (Q1), average (Q2+Q3), and a lot (Q4). In addition, the graph of a group is used to determine isolated nodes as *isolatedMember* for which the *outdegree* and *indegree* are equal to zero. In contrast, *closeCouples* have a particularly close relationship, where the mutual *outdegree* has a value in the 4th quantile range, indicating a strong reciprocal relationship in writing. *looseCouples* have a low-level relationship on both ties (1st quartile of the *outdegree*) and *averageCouples* have a medium relationship with an *outdegree* in the second or third quartile.

Table 1. Participants

| Cohort, Year | N | Gender (%) | | | Age (%) | | |
|---|---|---|---|---|---|---|---|
| | | Female | Male | n/a | <= 30 | > 30 | n/a |
| A, 2021/22 | 1,848 | 71.5 | 27.4 | 3.9 | 52.9 | 43.8 | 3.3 |
| B, 2022/23 | 1,463 | 67.3 | 31.4 | 1.3 | 52.7 | 46.8 | 0.5 |

These metrics on the group level enable comparisons across groups, in particular, to identify those groups that require attention because of unequal participation of group members. For example, groups can be identified in which one or more people write in isolation from the other group members ($isolatedMembers > 0$) or do receive no or little help from others ($noReceivers > 0$, $lowReceivers > 0$).

## 6 Case Study

The applicability of the three methods presented above is demonstrated below using a case study with two cohorts.

### 6.1 Methods

*6.1.1 Participants and Setting.* The data used for this study were collected during the same first-year course at a distance-learning university over two consecutive years. Over this period, the course and its structure remained unchanged. By including two distinct cohorts from the two years, we aim to replicate our results. The course was part of a compulsory module of the Bachelor's degree program in psychology at a large distance-learning university. Students could enrol in the course every semester of their choice. Consequently, it is not unusual for many students to enrol in the course but not complete it within the same semester.

Demographic information of the participants is shown in Table 1. Although the course language is German and a good command of this language is assumed, a diverse cultural background can be presumed [49].

Students were informed about data protection and the content and duration of data collection and provided written consent for participation according to the EU General Data Protection Regulations and research ethics guidelines by the American Psychological Association and the German Psychological Association. The local ethics officer approved the data collection.

Students were randomly distributed into groups of up to nine members, although not all later actively participated in group activities. The group's task was to summarise a scientific article. Collaborative activities were divided into two phases (T1 and T2), each lasting three weeks. In the first phase, the theory and methodology section of the article had to be summarised, while the second phase focused on summarising the results and discussion. The collaboration occurred entirely online, utilising tools provided by the Learning Management System (LMS) Moodle [29] and Etherpad Lite for joint learning and writing. Teaching was conducted entirely online.

*6.1.2 Collaborative Writing Environment and Data Collection.* Within the Moodle LMS, each group was allocated an individual collaborative writing space, referred to as a group instance. This space allowed members to contribute to the shared document, annotate text sections, and communicate via text chat. The group instance consists of an instance of *Etherpad lite* and a *group instance database*, separated per group as Docker containers for better maintainability and performance. Additionally, participants could utilise a forum within the Moodle course for broader discussions in the entire class.

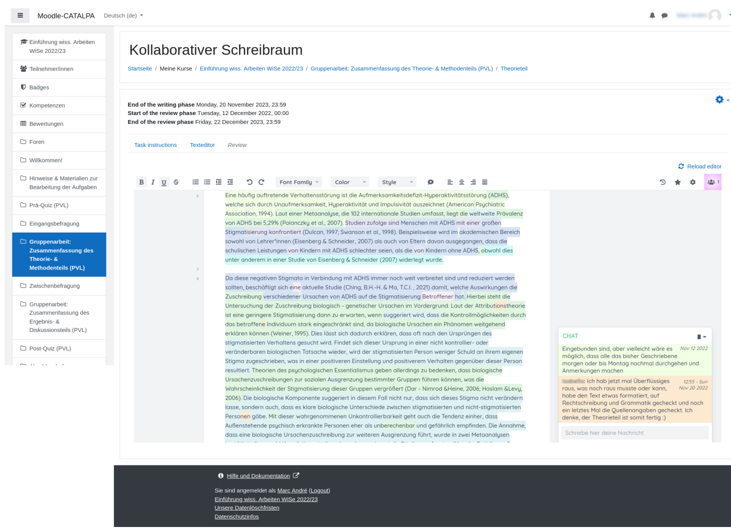


Fig. 3. Screenshot of the collaborative writing space provided in Moodle LMS using the collaborative writing (mod_cwr) plugin and *Etherpad lite*

Table 2. Collected data of both cohorts

| Cohort | Period | Participants | Groups | Changesets | Scroll Events | Comments | Chat |
|---|---|---|---|---|---|---|---|
| A | T1 | 1,848 | 294 | 2,010,044 | 685,446 | 1,960 | 3,056 |
| A | T2 | 1,842 | 294 | 1,466,667 | 596,578 | 1,670 | 3,892 |
| A | Total | | | 3,476,711 | 1,282,024 | 3,630 | 6,948 |
| B | T1 | 1,463 | 242 | 1,287,409 | 460,173 | 1,061 | 1,616 |
| B | T2 | 1,403 | 242 | 900,305 | 292,214 | 532 | 439 |
| B | Total | | | 2,187,714 | 752,387 | 1,593 | 2,055 |

Student interactions with the collaborative writing environment were collected in the LMS Moodle and the collaborative writing tool (*Etherpad lite*) in the form of event logs (e.g. 'event id', 'student id', 'group id', 'course id', 'type of learning action', 'text change set', 'timestamp'). Table 2 gives an overview of the number of events captured. The system architecture for large-scale collaborative writing and analytics is described in [5]. To implement further monitoring functionality, an Etherpad Lite Tracking Plug-In [5] was developed to log the following events for each user: (1) the student's connection times to the Etherpad Lite server, (2) the times in which the browser tab containing the Etherpad Lite editor is focused or not by the user, respectively, (3) the event when the student opens or closes the Etherpad Lite chat, (4) changes in the scrolling positions in the text document, and (5) the scrolling positions in the chat.

All data generated by the respective Etherpad Lite, including data from the additional *Tracking Plug-In* and *Annotation Plug-In* [5] is persisted in the *group instance database*, which was implemented using the document-oriented database management system (DBMS) *CouchDB*.

*6.1.3 Classification of User Sessions.* In total, 16 features have been extracted for each user session (Table 3). The features consider all available log data from Etherpad and Moodle related to the collaborative writing environment. They cover session length, concurrent users, text contributions, reading activity, and communication in terms of chat and discussion posts. In addition to the number of activities, their extent (e.g., scrolled lines, removed characters) and extent over time were considered.

We used k-means [27] as a clustering method to identify natural groupings and patterns in the session data. The number of clusters was determined by the distortion score and silhouette score [37]. The silhouette score measures how similar an object is to its cluster compared to others. It ranges from -1 to 1, with higher values indicating better-defined clusters. The distortion score is the mean sum of data points' squared distances to the cluster's centre.

Table 3. Features used for clustering user sessions

| Feature | Description | Value range |
|---|---|---|
| concurrent_users | Number of users simultaneously online | [1,N] |
| session_length | Session length in minutes | [0,N] |
| text_added | Number of characters added to the text | [0,N] |
| text_removed | Number of characters removed from the text | [0,N] |
| text_formatting | Number characters formatted in the text | [0,N] |
| text_added_pm | Number of characters added per minute | [0,1] |
| scroll_count | Number of scroll activities | [0,N] |
| scroll_length | Number of lines scrolled in all scroll activities | [0,N] |
| scroll_length_pm | Number of scrolled lines per minute | [0,1] |
| discussion_count | Number of discussions | [0,N] |
| discussion_total_length | Number of characters written in the discussion | [0,N] |
| discussion_total_length_pm | Number of characters written in the discussion per minute | [0,1] |
| chat_count | Number of chat messages | [0,N] |
| chat_total_length | Number of characters written in the chat messages | [0,N] |
| chat_total_length_pm | Number of characters written in the chat messages per minute | [0,1] |

*6.1.4 Frequent Sequences of Session Types.* To apply the method proposed in section 4, sequences of session types have been compiled for every user (e.g., "ST3, ST1, ST1, . . . , ST4"). To mark the boundaries of the sequence, an element has been added to the beginning and end of a sequence (e.g., "START, ST3, ST1, ST1, . . . , ST4, END"). Then, the transition probability for all pairs of session types, including boundary markers, was computed. In the first step, the probability of transitioning from one session type to another was computed for the whole population and terciles of group sizes (2-6, 7-8, and 8-10 persons). Transitions with a probability below 0.05 have been omitted to improve clarity.

In the second step, based on the probabilities for the transition between session types, we determined all candidate sequences of lengths 2 to 10 whose successive transitions have a probability above a defined threshold value of 0.01 (transition *support*). Third, to determine the most frequent sequences, the *support* measure has been employed for each candidate sequence. The *support* is defined as the proportion between the number of sequences containing a specific sequence and the number of all sequences. It ranges from 0 to 1, where 0 means that the sequence did not occur, and 1 indicates that the sequence occurred in all activity sequences.

*6.1.5 Collaboration on Text-Level.* Changesets in *Etherpad lite* are decoded with the *EasySync Protocol* [3]. As a first step, the changesets had to be decoded to obtain information on the type of change operation (add, remove, format), the position of the change (line number, line position), the affected characters, and their length. Then, the relationship of the group members could be computed using Algorithm 2. Graphs were constructed for each group based on these relationships. In these graphs, the nodes represent individual group members, while the directed edges signify interactions. Specifically, an edge points from one group member to another when they edit text adjacent to a character written by the latter. The network metrics shown in Table 4 have been calculated from each graph.

Table 4. Network metrics on individual and group level

| Metric | Description | Value range |
|---|---|---|
| *outdegree* | | [0,N] |
| *indegree* | | [0,N] |
| *noReceiver* | If a member's *indegree* is equal to zero. | true/false |
| *lowReceiver* | If a member's *indegree* is a value in the 1th quantile range. | true/false |
| *mediumReceiver* | If a member's *indegree* is a value in the 2nd or 3rd quantile range. | true/false |
| *highReceiver* | If a member's *indegree* is a value in the 4th quantile range. | true/false |
| *noHelper* | If a member's *outdegree* is equal to zero. | true/false |
| *lowHelper* | If a member's *outdegree* is a value in the 1th quantile range. | true/false |
| *mediumHelper* | If a member's *outdegree* is a value in the 2nd or 3rd quantile range. | true/false |
| *highHelper* | If a member's *outdegree* is a value in the 4th quantile range. | true/false |
| *activeMembers* | Number of individual members who have shown any activity. | [0,N] |
| *isolatedMember* | Number of members where *outdegree* and *indegree* are equal to zero | [1,N] |
| *looseCouple* | Number of members being part of a pair whose mutual *outdegree* is a value in the 1th quantile range. | [1,N] |
| *averageCouple* | Number of members being part of a pair whose mutual *outdegree* is a value in the 2nd or 3rd quantile range. | [1,N] |
| *closeCouples* | Number of members being part of a pair whose mutual *outdegree* is a value in the 4th quantile range. | [1,N] |

## 6.2 Results

The results were computed using partly parallelised R routines, executed on an Apple M1 MacBook Pro with 8 CPU cores, 8 GPU cores, and 16 GB RAM. Each cohort took about four hours to complete. The entire code used for the analysis and the resulting data are available as open source and publicly accessible under https://github.com/CATALPAresearch/cscw2025-collaborative-writing-patterns.

*6.2.1 Classification of User Sessions (RQ1).* For k ranging from 2 to 20, we used a silhouette analysis [37] to find the best k representative clusters for the supplied dataset, using measurements of the average distance inside the clusters and the average distance between the clusters. For $k = 6$, the best result was obtained ($silhouette\_score_A = 0.14$, $silhouette\_score_B = 0.25$, $distortion\_score_A = 31.0$, $distortion\_score_B = 16.6$). To keep the clusters consistent across multiple executions of the code, a seed value was used instead of a random initialisation of k-means.

Table 5. Session types and their features for cohort A

| Cluster | $ST1_A$ | $ST2_A$ | $ST3_A$ | $ST4_A$ | $ST5_A$ | $ST6_A$ |
|---|---|---|---|---|---|---|
| Sessions (%) | 16.79 | 1.87 | 3.00 | 67.45 | 3.24 | 7.65 |
| Groups (%) | 99.66 | 46.44 | 87.80 | 100.00 | 52.54 | 97.63 |
| User (%) | 83.89 | 19.60 | 33.99 | 99.41 | 32.02 | 62.48 |
| Concurrent users | 1.29 | 5.47 | 1.59 | 1.10 | 5.18 | 1.51 |
| Session length | 24.51 | 103.80 | 103.36 | 2.48 | 31.71 | 57.4 |
| Text added | 200.58 | 399.96 | 1614.90 | 13.96 | 71.84 | 671.87 |
| Text removed | 67.86 | 211.66 | 483.16 | 5.36 | 33.35 | 214.36 |
| text formatting | 1.52 | 3.69 | 10.93 | 0.13 | 0.64 | 4.52 |
| Text added/minute | 8.49 | 3.83 | 16.30 | 3.26 | 2.33 | 12.51 |
| Scroll count | 46.46 | 266.69 | 146.08 | 11.32 | 61.23 | 91.12 |
| Scroll length | 499.43 | 3026.34 | 1674.85 | 121.10 | 695.80 | 1001.96 |
| Scroll length/minute | 23.21 | 30.51 | 16.68 | 96.93 | 27.34 | 18.36 |
| Discussion count | 0.19 | 0.31 | 0.61 | 0.02 | 0.03 | 0.41 |
| Discussion length | 16.63 | 19.81 | 55.07 | 1.37 | 1.09 | 41.49 |
| Discussion length/minute | 0.77 | 0.19 | 0.57 | 0.28 | 0.04 | 0.79 |
| Chat count | 0.15 | 2.10 | 0.78 | 0.02 | 0.71 | 0.52 |
| Chat length | 18.49 | 104.73 | 69.03 | 2.33 | 37.44 | 43.80 |
| Chat length/minute | 0.85 | 0.94 | 0.66 | 0.53 | 0.99 | 0.76 |

As shown in Table 5 and Table 6, the size of the clusters is represented by the relative number of user sessions ranging from 3.00 % to 67.45 %. A significant portion of participants (9.45 % – 99.41 %) performed the session types identified. These session types occurred in at least a quarter of the groups.

In the following, we identify the characteristics of the 6 session types and assign them characterising names. It is noteworthy that these session types could be identified in both cohorts. The matching between the session types of both cohorts is based on similar feature ranges.

*Brief solo writing sessions* ($ST1_A$, $ST1_B$) are characterised by a low session length and moderate text changes. Here, individuals work alone (lowest no. of concurrent users), focusing on short, solo writing tasks. Minimal discussion and chat activity suggest focused, individual work without much collaboration.

*Intensive collaborative sessions* ($ST2_A$ ) are distinguished by the most concurrent users, indicating significant collaboration. Session length and text changes per user are high, showing substantial engagement with content. High scrolling and chat activity indicate dynamic interaction and communication among participants.

*Observing collaborative writing sessions*($ST2_B$) are similar to $ST2_A$ but show no text editing activities while a lot of communication is taking place.

*In-depth individual writing sessions* ($ST3_A$, $ST3_B$) feature the highest amount of text added and removed, suggesting intense writing and editing. Though the no. of concurrent users is low, indicating primarily individual work, the high text formatting and discussion length indicate detailed, thoughtful engagement with the task.

*Quick review sessions* ($ST4_A$, $ST4_B$) are characterised by the shortest sessions and minimal or no text changes, reflecting quick reviews or checks rather than active writing. Almost all users participate in such sessions, making it the most common but least engaged activity type.

Table 6. Session types and their features for cohort B

| Cluster | $ST1_B$ | $ST2_B$ | $ST3_B$ | $ST4_B$ | $ST5_B$ | $ST6_B$ |
|---|---|---|---|---|---|---|
| Sessions (%) | 18.12 | 1.48 | 3.19 | 65.2 | 3.2 | 8.82 |
| Groups (%) | 96.91 | 26.91 | 71.09 | 99.82 | 36 | 93.09 |
| User (%) | 70.21 | 9.49 | 21.68 | 96.97 | 19.72 | 48.98 |
| Concurrent users | 1.24 | 3.89 | 1.97 | 1.11 | 5.17 | 1.35 |
| Session length | 23.85 | 95.14 | 96.58 | 2.72 | 33.76 | 53.45 |
| Text added | 204.55 | 0 | 1375.58 | 27.96 | 75.51 | 499.91 |
| Text removed | 68.02 | 0 | 447.93 | 11.03 | 43.51 | 159.95 |
| Text formatting | 1.5 | 0 | 8.94 | 0.27 | 0.68 | 3.4 |
| Text added/minute | 8.87 | 0 | 15.29 | 17.11 | 3.33 | 10.04 |
| Scroll count | 23.92 | 290.98 | 25.45 | 6.73 | 33.47 | 41.52 |
| Scroll length | 251.01 | 3045.34 | 269.6 | 69.28 | 348.59 | 438.59 |
| Scroll length/minute | 11.55 | 34.21 | 2.69 | 48.3 | 12.62 | 8.36 |
| Discussion count | 0.07 | 0.44 | 0.08 | 0.01 | 0.05 | 0.13 |
| Discussion length | 7.15 | 32.9 | 11.43 | 0.67 | 3.74 | 13.7 |
| Discussion length/minute | 0.34 | 0.4 | 0.11 | 0.14 | 0.09 | 0.26 |
| Chat count | 0.05 | 0.88 | 0.09 | 0.01 | 0.19 | 0.12 |
| Chat length | 7.13 | 65.53 | 10.1 | 1.25 | 10.34 | 13.91 |
| Chat length/minute | 0.31 | 0.64 | 0.1 | 0.33 | 0.23 | 0.27 |

In *shallow collaboration sessions* ($ST5_A$, $ST5_B$) the number of concurrent users is similar to $ST2_A$/$ST2_B$, but the cluster is marked by a lower number of chat and discussion activities while having moderate reading and low to zero writing output. Despite the high number of concurrent users, these sessions seem less productive.

In *productive writing and discussion sessions* ($ST6_A$, $ST6_B$) writing is balanced with communication, evident from substantial text changes and moderate discussion/chat lengths. This looks like a blend of individual writing with moments of collaboration, indicative of sessions that alternate between focused writing and team discussions.

*6.2.2 Frequent Sequences of Session Types (RQ2).* Table 7 shows frequent sequences of session types that could be identified by at least one-quarter of the students. A huge portion of frequent sequences is identical for both cohorts. Only a few session types appear among the frequent sequences. ST4 (Brief solo writing sessions), as the most frequent session type, occurs in every sequence listed in Table 7. Most users also start their collaborative activities with ST4. *Solo writing sessions* (ST1) often followed ST4 or preceded them. In cohort *A*, ST6 (Productive writing and discussion sessions) are also followed by ST4.

*6.2.3 Collaboration on Text-Level (RQ3).* On the individual level, a strong correlation between the self-degree and out-degree could be determined ($r_A = .97, r_B = .96$, $p < 0.0001$). With a high self-degree, the out-degree increases less strongly. Group members who write extensively still collaborate significantly with their peers, but not to the same extent as they edit their text.

Due to the task's voluntary nature at the group level, the number of active group members varied between 2 and 9. Although contributing to the other's text, 85 (*A*) and 89 (*B*) group members received no help from their teammates at the text-level. In contrast, 55 (*A*) and 57 (*B*) group members did not help others despite benefiting from the edits made by their team members. A total of 29 (*A*) and 33 (*B*) group members wrote the text in isolation from the other group members (Fig. 5).

Table 7. Frequent session type sequences in cohort $A$ and $B$ with support (Sup) greater than or equal to 0.25 and a sequence length (L) between 3 and 6.

| Sequences in cohort $A$ | L | Sup |
|---|---|---|
| ST4>ST4>ST4 | 3 | 0.74 |
| ST4>ST4>ST1 | 3 | 0.61 |
| START>ST4>ST4 | 3 | 0.60 |
| ST4>ST4>ST4>ST4 | 4 | 0.57 |
| ST1>ST4>ST4 | 3 | 0.55 |
| ST4>ST1>ST4 | 3 | 0.54 |
| ST4>ST4>ST4>ST1 | 4 | 0.46 |
| ST4>ST4>ST1>ST4 | 4 | 0.43 |
| ST4>ST1>ST4>ST4 | 4 | 0.41 |
| ST4>ST4>ST4>ST4>ST4 | 5 | 0.41 |
| ST1>ST4>ST4>ST4 | 4 | 0.40 |
| START>ST4>ST4>ST4 | 4 | 0.39 |
| ST4>ST4>END | 3 | 0.37 |
| ST4>ST4>ST6 | 3 | 0.36 |
| ST4>ST6>ST4 | 3 | 0.34 |
| ST4>ST4>ST4>ST4>ST1 | 5 | 0.32 |
| ST6>ST4>ST4 | 3 | 0.31 |
| ST4>ST4>ST4>ST1>ST4 | 5 | 0.31 |
| ST4>ST4>ST4>ST4>ST4>ST4 | 6 | 0.30 |
| ST4>ST4>ST1>ST4>ST4 | 5 | 0.29 |
| ST4>ST1>ST4>ST4>ST4 | 5 | 0.29 |
| ST1>ST4>ST4>ST4>ST4 | 5 | 0.29 |
| ST4>ST1>ST1 | 3 | 0.26 |

| Sequences in cohort $B$ | L | Sup |
|---|---|---|
| ST4>ST4>ST4 | 3 | 0.54 |
| START>ST4>ST4 | 3 | 0.41 |
| ST4>ST4>ST1 | 3 | 0.40 |
| ST1>ST4>ST4 | 3 | 0.38 |
| ST4>ST1>ST4 | 3 | 0.37 |
| ST4>ST4>END | 3 | 0.36 |
| ST4>ST4>ST1>ST4 | 4 | 0.26 |
| START>ST4>ST4>ST4 | 4 | 0.26 |
| ST4>ST4>ST4>ST1 | 4 | 0.25 |
| ST1>ST4>ST4>ST4 | 4 | 0.25 |
| ST4>ST1>ST4>ST4 | 4 | 0.25 |

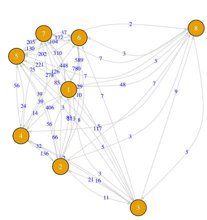

Fig. 4. Dense network of group members showing intensive mutual text editing

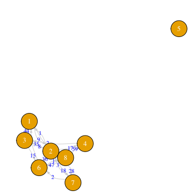

Fig. 5. An intensively collaborating group with one isolated group member

Couples are assumed when two group members edit each other's text. Concerning the maximum number of possible couples, only 45.7 % ($A$) and 45.1 % ($B$) of them could be identified in the groups. Thus, less than half of the possible relationships with another group member were reciprocal. Of the identified couples, a share of 21.2 % ($A$) and 20.9 % ($B$) showed low interaction with other group members (1. quartile). As depicted in Fig. 4, most couples worked very closely (4. quartile, close couples, $A$: 29.3 %, $B$: 29.3 %) or at least more closely (2. and 3. quartile, average couples, $A$: 49.5 %, $B$: 49.8 %) together on the text. Despite significant overall text production, there were no instances of pairs assisting each other in only 12 ($A$) and 10 ($B$) groups, respectively.

### 6.3 Discussion

This case study aimed at exploring learners' interaction patterns in collaborative writing. We employed clustering based on similar user sessions to answer this question by considering features describing the writing, reading, and communication activities in synchronous and asynchronous sessions of various durations. Based on these session types, we derived recurrent sequences of session types. Furthermore, we adapted network analysis to the text-level to identify collaborative writing patterns. Two extensive datasets about an educational writing task have been used to demonstrate the applicability of the methods presented in section 3–5. The results of the analyses based on these three methods and the two data sets are discussed in the following.

*6.3.1 Classification of User Sessions (RQ1).* We found 6 clusters of user sessions that vary in frequency depending on time, group, and individual. These 6 clusters are nearly consistent across the two cohorts and can be found in almost all groups and a large proportion of learners. Thus, we assume these session types are stable patterns regarding the particular type of task. The share of individual session clusters on the overall number of sessions may vary depending on particular instructions (e.g., about how and when to collaborate) and additional collaboration support provided by the collaborative writing environment (e.g., daily reports summarising the contributions group members made).

Concerning related works investigating collaborative writing patterns and strategies [8, 18, 20], we have been able to include asynchronous and synchronous writing, including text contributions, reading, time spent, and communication behaviour into the analysis.

The extent to which the provided communication facilities are used depends on the groups' needs to accomplish coordination tasks. For instance, it might be possible to discuss how to split work in a group of three using a text chat. Still, it becomes challenging in a group of six people without the possibility of audio or video chat.

The frequency of session types occurrences can serve as an indicator of the roles individual group members take. For instance, we can observe whether they actively participate in collaborative sessions, contribute significantly, engage in discussions or tend to be more reserved. Based on the relative frequency of the session types, the preferred behaviours of group members can be identified. These individual tendencies can also be used to assess the state of the group's collaboration and group progress. The current status and progress are observable on both individual and group levels.

One might argue that individual session features could be analysed separately. However, this approach would overlook the overlapping characteristics within the six clusters. For example, one might focus on the duration of synchronous collaboration or the amount of text produced, but not see whether synchronous collaboration coincides with high or low text production. Thus, the applied clustering approach helps to understand the complex interaction of individual features while still allowing a comprehensive but understandable set of features to be considered.

In practice, user sessions can be continuously assigned to one of the identified session types. The frequency of session type occurrences can be provided to individual group members to support reflection about group processes. In addition, teachers could be informed about emerging session types.

Future research is necessary to identify and compare session types that occur under different conditions (e.g., task type, group size, task duration, and rich media communication).

*6.3.2 Frequent Sequences of Session Types (RQ2). Brief solo writing sessions* (CS4) are both cohorts' most frequent session types. In addition to subsequent ST4 sessions, there are also *Solo writing* *sessions* (ST1) interwoven. While solo writing characterised session sequences in cohort *B*, cohort *A* with *Productive writing and discussion sessions* shows at least some frequent synchronous collaboration between solo writing sessions.

Although Etherpad Lite offers the option of synchronised collaboration, the texts were usually produced in a sequence of individual sessions. Collaboration is, therefore, predominantly asynchronous. Several studies have shown that asynchronous collaboration in writing can be effective and efficient [18, 19]. The network analysis results about the collaboration on the text-level confirm this assumption for most participants.

It becomes apparent that the writing process mainly involves short and successive individual work steps (ST4). This way of working could be associated with review and coordination tasks, especially in larger groups, that require checking progress and agreed-upon contributions. It raises the question of whether the group members could also obtain this information via other communication channels. However, it cannot be ruled out that these short sessions also express a certain hesitation in developing the text further and thus increase the intensity of text production. This could indicate deficits in the coordination of group tasks, such as a lack of task distribution and effective time management. However, the teacher could support group coordination tasks through specific instructions.

Regarding the further improvement of the collaborative writing environment, special attention should be paid to supporting asynchronous group awareness [35]. In addition to the existing supporting instruments (e.g., timeline, activity log), Schümmer and Lukosch [39] proposed periodic reports (e.g., daily/weekly emails about relevant progress information), change indicators (e.g., mark text sections that changed since the last session), and aliveness indicators (e.g., visualise group members' recent amount of text contributions) to support asynchronous awareness.

To promote synchronous collaboration, teachers could emphasise its benefits and suggest strategies for groups to achieve this (e.g., weekly jourfix, writing tandems).

*6.3.3 Collaboration on Text-Level (RQ3).* The findings suggest that the methods employed can be applied to (1) pinpoint groups facing problems writing together and (2) inform teachers or groups about the issues at hand, enabling interventions by either instructors, automated agents, or self-regulation by the groups themselves through improved group awareness support. Furthermore, these findings could facilitate a comprehensive analysis of the underlying causes of issues within challenging groups, for instance, by enhancing task instructions or refining the design of the collaboration environment to include more effective awareness cues.

The 85 (*A*) and 89 (*B*) group members who have not received help on the text-level appear to be many. The extent to which someone receives help with writing from other group members depends on their own text contributions. So if they write early and a lot, they have a higher chance of their texts being supplemented or updated by others. Besides this temporal effect, a low text contribution and conflicts among group members cannot be ruled out as a reason for receiving no or little help.

The possibility of helping others to write also requires that others have already written something. Other reasons for providing no or little help can be seen in the limited and very time-limited collaboration or a too strict division of tasks. Teachers could address the non-helping person directly to prevent this and encourage more effective collaboration. Addressing reciprocity, e.g., in a group dashboard [10], is another strategy. Group members are more motivated to help others when they become aware that others have already assisted them.

There may be several reasons why some learners wrote in isolation from their group members. On the one hand, there may have been temporal reasons, i.e., one-off text contributions were made very early or last. On the other hand, we could observe a strict division of tasks. The extent to which internal group conflicts contributed cannot be determined based on the data analysed here. Even if only a small proportion of learners wrote the text in isolation, it indicates a need for action by teachers, who should ensure that all group members work together on the text. A teacher dashboard could be useful to increase teachers' awareness of these learners [48]. We have already developed prototypes that visualise the network metrics for many groups and users. Future research could investigate the interventions teachers derive from a teacher dashboard indicating these issues and whether and how these interventions result in behavioural changes in the students' collaborative writing process.

The existence of *closeCouples* indicates well-functioning collaboration across many groups. Since *closeCouples* imply a symmetry of individual contribution, Burkhardt et al. [6] refers to it as a dimension to measure the quality of collaboration. However, close collaboration between more than two group members has been observed but not fully analysed. Also, the temporal aspects and modes of synchronous writing are worth analysing in more detail. Referring to Voltmer et al. [49], heterogeneous sociodemographic characteristics may have been conducive to establishing close collaborations.

We empirically refined the writing strategies proposed by Hoffmann et al. [18] (summative and integrative text construction) as collaborative writing patterns. We identified and visualised collaboration at the text-level using a similar approach to Chen et al. [8]. Additionally, we applied network measures to describe collaboration at the individual and group levels, enabling us to identify dysfunctional groups. Compared to Hsu [20], who (manually) identified communication patterns occurring during collaborative writing activities, our approach includes writing activities that were not part of the group members' communication. We, therefore, cover the entire writing process through an automatic analysis, but ignore communication so far.

The presented methods have been applied to analyse the collaboration during the processing of a learning task. There is no reason why these methods should not be used to analyse collaborative writing activities unrelated to learning.

*6.3.4 Limitations.* The methodology used in this paper is applicable for analyzing and describing collaborative writing activities. However, several limitations of this case study should be considered. These limitations can be categorized into three main areas: participant-related constraints, methodological constraints, and technical limitations.

*Participant-Related Constraints.* A key limitation concerns the experience and formation of the student groups. First, as the participants were first-semester students, they had little prior experience in collaborative writing and received no specific training in this regard. This may have influenced both their writing behaviour and their ability to effectively engage in the collaborative process. Additionally, group formation was random, increasing the likelihood of mismatches between members. Unlike naturally formed groups based on relationships or shared interests, some participants may not have been well aligned in terms of collaboration styles, potentially impacting group dynamics.

*Methodological Constraints.* Several methodological limitations stem from the nature of the study design and data collection approach. While well-defined in terms of expected results, the group task remained open-ended regarding the process to achieve those results. This variability in approach may have led to differences in collaboration strategies, making direct comparisons more challenging. Furthermore, the clustering of user sessions relied on each participant generating a significant number of interactions. This requirement, in turn, was dependent on the study setting, which needed to involve (i) a sufficiently extensive, (ii) freely coordinated, and (iii) a longer-lasting task to provide meaningful data. Another methodological limitation is that the study did not analyze external communication. Since students may have coordinated outside the collaboration platform—through email or messaging apps—important aspects of their teamwork might not have been captured. Regarding the interactions on text-level, it was assumed that all contributions were useful or positive. A more detailed analysis of whether rewriting, editing, or communication (within or outside of the collaborative writing environment) could have triggered conflicts, affecting group cohesion and the quality of collaboration, is part of future work.

*Technical Limitations.* Several constraints also stem from the technical capabilities of the collaboration environment and the data collection mechanisms. First, the platform did not provide notifications about activities in shared workspaces. This required students to manually check progress or coordinate externally. Moreover, detailed monitoring of writing activities was restricted. The system could not record keystroke-level changes; instead, text changesets were logged at a slightly higher level of granularity, averaging two characters per changeset ($M_A = 2.22, SD_A = 2.36$, $M_B = 2.16, SD_B = 2.36$). This limitation reduced the ability to analyze fine-grained writing behaviours. Additionally, AI-generated text was not detected, as no extensive copy-paste behaviour was observed. The employed methods focused solely on human-human interactions through shared text editing, commenting, and chat. Finally, the analysis did not account for the quality of the jointly produced text, such as linguistic sophistication or use of technical language. The text created by the groups was not graded by the teachers. Instead, admission to the written examination depended on whether students had sufficiently engaged with the collaborative writing task (determined by using user activity count and student peer review scores). This examination did not include any task related to the writing assignment. It was also impossible to consider information about prior performance since the participants were in their first semester and admission tests (e.g., GPA) do not exist in Germany. However, there is an indication of a positive correlation between the network density of the groups [4] and the average peer review ratings. Indicators of text quality and their relationship to interaction patterns go beyond the scope of this paper and require further investigation.

## 7 Conclusion

Teachers require a clear understanding of student activities during collaborative writing to assess performance and identify support needs. To support teachers in interpreting and assessing student behaviours, we proposed using fine-grained log data and individual user sessions as a suitable observation unit for capturing coherent interaction patterns within a session. Based on these sessions, we proposed three methods for identifying interaction patterns capturing the interleaving of individual and collaborative activities during collaborative writing: The first method (RQ1, section 3) determines session types by analysing key features, including writing, reading, communication behaviours, session length, and the number of group members collaborating synchronously. These session types help teachers understand students' different roles in the writing and collaboration process. The second method (RQ2, section 4) uncovers frequent sequences of session types by examining their order, allowing for process analyses at a higher level of abstraction while maintaining manageability compared to raw log data. Frequent session type sequences provide teachers insights into the collaboration processes concerning session types that are more conducive or less conducive to writing, collaborating, and, thus, learning. The third method (RQ3, section 5) investigates text-level collaboration by assessing the frequency of text passage modifications by either the original author or other group members. This approach quantifies individual collaboration efforts and differentiates between isolated and tightly connected members at the group level, offering teachers more profound insights into the collaboration mode and overall group cohesion.

To test the applicability of our approach, we presented a case study where we applied the methods to analyse two cohorts of the same first-year course at a distance-learning university over two consecutive years. Collaborative writing activities from 536 groups with 3,311 participants were analysed. The extent of the data included in the study made it possible to determine behavioural patterns of individual user sessions using the methods described in section 3 and 4 (RQ1/RQ2) and the collaboration at the text-level using the method described in 5 (RQ3). Concerning the identification of session types (RQ1), six characteristic clusters of recurring session patterns were identified based on the individual user sessions, which differ in duration, the number of synchronously collaborating group members, the extent of text changes, and the reading and communication behaviour. Regarding identifying frequent sequences of session types (RQ2), the most frequent sequences consisted of a less desirable session type characterised by its short duration and minimal or no text changes, reflecting quick reviews or checks rather than active writing. The investigation of text-level collaboration (RQ3) revealed interaction patterns regarding one-sided and mutual help among group members and isolated writing activities. The analysis of both cohorts yielded very similar results, which supports the robustness of our methods and confirms their applicability.

The proposed methods could help to enhance the monitoring and assistance provided to groups engaged in collaborative writing. CSCW/L researchers can use these methods to analyse group behaviour, while tool developers can leverage the metrics and algorithms to create dashboards and awareness widgets. Teachers implementing collaborative writing projects may use our metrics to better identify and support students needing additional assistance.

Future research on classifying session types could include collaboration metrics at the text-level, e.g., to consider group cohesion. Text-level collaboration could consider a context window beyond the adjacent neighbouring character. Finally, different collaboration tasks, group sizes, and writing support features should be analysed to search for generic writing patterns.

## Acknowledgments

This research was supported by the Center of Advanced Technology for Assisted Learning and Predictive Analytics (CATALPA) of the FernUniversität in Hagen, Germany. Thanks to Stefan Stürmer, Laura Fröhlich, Verena Walpurger, Björn Fiesseler, Jan-Bennet Voltmer, Natalia Reich-Stiebert, and Jennifer Raimann.

## References

[1] Olli Aksela, Joni Lämsä, and Sanna Järvelä. 2024. Secondary School Students' Enacted Self-Regulated Learning Strategies in a Computer-Based Writing Task–Insights from Digital Trace Data and Interviews. *Technology, Knowledge and Learning* (2024). https://doi.org/10.1007/s10758-024-09789-4

[2] Alphabet Inc. 2024. Google Docs. https://docs.google.com/. Accessed: 2024-12-08.

[3] AppJet Inc. and Etherpad Foundation. 2018. Etherpad and EasySycn Technical Manual. https://github.com/ether/etherpad-lite/blob/master/doc/public/easysync/easysync-full-description.pdf.

[4] Natahlie Bick, Laura Froehlich, Jan-Bennet Voltmer, Jeniffer Raimann, Natalia Reich-Stiebert, Niels Seidel, Marc Burchart, SE Martiny, Nikitin J, Stürmer Stefan, and Martin A. 2024. Virtually isolated: social identity threat predicts social approach motivation via sense of belonging in computer-supported collaborative learning. *Front. Psychol* 15, 1346503 (2024). https://doi.org/10.3389/fpsyg.2024.1346503

[5] Marc Burchart and Joerg M. Haake. 2024. Supporting Collaborative Writing Tasks in Large-Scale Distance Education. *IEEE Transactions on Learning Technologies* 17 (2024), 1051–1068. https://doi.org/10.1109/TLT.2024.3355791

[6] Jean-Marie Burkhardt, Françoise Détienne, Anne-Marie Hébert, Laurence Perron, Stéphane safin, and Pierre Leclercq. 2009. An approach to assess the quality of collaboration in technology-mediated design situations. In *European Conference on Cognitive Ergonomics: Designing beyond the Product — Understanding Activity and User Experience in Ubiquitous Environments* (Helsinki, Finland) *(ECCE '09)*. VTT Technical Research Centre of Finland, FI-02044 VTT, FIN, Article 30, 9 pages.

[7] Montserrat Castelló, Otto Kruse, Christian Rapp, and Mike Sharples. 2023. *Synchronous and Asynchronous Collaborative Writing*. Springer International Publishing, Cham, 121–139. https://doi.org/10.1007/978-3-031-36033-6_8

[8] Chih-Ming Chen, Ming-Chaun Li, and Cheng-Kai Liao. 2022. Developing a collaborative writing system with visualization interaction network analysis to facilitate online learning performance. *Interactive Learning Environments* 0, 0

(2022), 1–20. https://doi.org/10.1080/10494820.2022.2028851

[9] G D Chen, C Y Wang, and K L Ou. 2003. Using group communication to monitor web-based group learning. *Journal of Computer Assisted Learning* 19, 4 (dec 2003), 401–415. https://doi.org/10.1046/j.0266-4909.2003.00045.x

[10] Yuxin Chen, Cindy E. Hmelo-Silver, Susanne P. Lajoie, Juan Zheng, Lingyun Huang, and Stephen Bodnar. 2015. Using Teacher Dashboards to Assess Group Collaboration in Problem-based Learning. *The Interdisciplinary Journal of Problem-based Learning* 2, 21 (2015).

[11] Wai Ki Ching, Eric S Fung, and Michael K Ng. 2004. Higher-order Markov chain models for categorical data sequences. *Naval Research Logistics (NRL)* 51, 4 (2004), 557–574. https://doi.org/10.1002/nav.20017

[12] Dorian Doberstein, Tobias Hecking, and H Ulrich Hoppe. 2017. Sequence Patterns in Small Group Work Within a Large Online Course. In *Collaboration and Technology*, Carl Gutwin, Sergio F Ochoa, Julita Vassileva, and Tomoo Inoue (Eds.). Springer International Publishing, Cham, 104–117.

[13] Carl Gutwin and Saul Greenberg. 1996. Workspace Awareness for Groupware. In *Conference Companion on Human Factors in Computing Systems (CHI '96)*. ACM, New York, NY, USA, 208–209. https://doi.org/10.1145/257089.257284

[14] Feifei Han and Robert A Ellis. 2021. Patterns of student collaborative learning in blended course designs based on their learning orientations: a student approaches to learning perspective. *International Journal of Educational Technology in Higher Education* 18, 1 (2021), 66. https://doi.org/10.1186/s41239-021-00303-9

[15] Jiawei Han, Jian Pei, Yiwen Yin, and Runying Mao. 2004. Mining Frequent Patterns without Candidate Generation: A Frequent-Pattern Tree Approach. *Data Mining and Knowledge Discovery* 8, 1 (2004), 53–87. https://doi.org/10.1023/B:DAMI.0000005258.31418.83

[16] Derek L. Hansen and Marc A. Smith. 2014. *Social Network Analysis in HCI*. Springer New York, New York, NY, 421–447. https://doi.org/10.1007/978-1-4939-0378-8_17

[17] CindyE. Hmelo-Silver, Rebecca Jordan, Lei Liu, and Ellina Chernobilsky. 2011. Representational Tools for Understanding Complex Computer-Supported Collaborative Learning Environments. In *Analyzing Interactions in CSCL*, Sadhana Puntambekar, Gijsbert Erkens, and Cindy Hmelo-Silver (Eds.). Computer-Supported Collaborative Learning Series, Vol. 12. Springer US, 83–106. https://doi.org/10.1007/978-1-4419-7710-6_4

[18] Christian Hoffmann, Nadine Mandran, Cédric D'Ham, Sébastien Rebaudo, and Mohamed Anis Haddouche. 2022. Development of Actionable Insights for Regulating Students' Collaborative Writing of Scientific Texts. In *Educating for a New Future: Making Sense of Technology-Enhanced Learning Adoption*, Isabel Hilliger, Pedro J Muñoz-Merino, Tinne De Laet, Alejandro Ortega-Arranz, and Tracie Farrell (Eds.). Springer International Publishing, Cham, 534–541.

[19] H Ulrich Hoppe, Dorian Doberstein, and Tobias Hecking. 2021. Using Sequence Analysis to Determine the Well-Functioning of Small Groups in Large Online Courses. *International Journal of Artificial Intelligence in Education* 31, 4 (2021), 680–699. https://doi.org/10.1007/s40593-020-00229-9

[20] Hsiu-Chen Hsu. 2020. The impact of task complexity on patterns of interaction during web-based asynchronous collaborative writing tasks. *System* 93 (2020), 102328. https://doi.org/10.1016/j.system.2020.102328

[21] Ulla Karvonen, Liisa Tainio, and Sara Routarinne. 2018. Uncovering the pedagogical potential of texts: Curriculum materials in classroom interaction in first language and literature education. *Learning, Culture and Social Interaction* 17 (2018), 38–55. https://doi.org/10.1016/j.lcsi.2017.12.003

[22] Brian C Keegan, Shakked Lev, and Ofer Arazy. 2016. Analyzing Organizational Routines in Online Knowledge Collaborations: A Case for Sequence Analysis in CSCW. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW '16)*. Association for Computing Machinery, New York, NY, USA, 1065–1079. https://doi.org/10.1145/2818048.2819962

[23] Jenell Krishnan, Soobin Yim, Alissa Wolters, and Andrew Cusimano. 2019. Supporting Online Synchronous Collaborative Writing in the Secondary Classroom. *Journal of Adolescent & Adult Literacy* 63, 2 (jun 2019), 135–145. https://www.jstor.org/stable/48556227

[24] Vladimir I Levenshtein. 1966. Binary Codes Capable of Correcting Deletions, Insertions and Reversals. *Soviet Physics Doklady* 10 (February 1966), 707.

[25] Chen-Chung Liu and Chin-Chung Tsai. 2008. An analysis of peer interaction patterns as discoursed by on-line small group problem-solving activity. *Computers & Education* 50, 3 (2008), 627–639. https://doi.org/10.1016/j.compedu.2006.07.002

[26] Ming Liu, Rafael A. Calvo, and Abelardo Pardo. 2013. Tracer: A Tool to Measure and Visualize Student Engagement in Writing Activities. In *2013 IEEE 13th International Conference on Advanced Learning Technologies*. 421–425. https://doi.org/10.1109/ICALT.2013.129

[27] S. Lloyd. 1982. Least squares quantization in PCM. *IEEE Transactions on Information Theory* 28, 2 (1982), 129–137. https://doi.org/10.1109/TIT.1982.1056489

[28] Dennis Menze, Niels Seidel, and Regina Kasakowskij. 2022. Interaction of reading and assessment behavior. In *DELFI 2022 – Die 21. Fachtagung Bildungstechnologien der Gesellschaft für Informatik e.V.*, Peter A. Henning, Michael Striewe, and Matthias Wölfel (Eds.). Gesellschaft für Informatik, Bonn, 27–38. https://doi.org/10.18420/delfi2022-011

[29] Moodle [n. d.]. Moodle. https://moodle.org/. Accessed: 2024-12-08.

[30] Christine M Neuwirth, David S Kaufer, Ravinder Chandhok, and James H Morris. 1990. Issues in the design of computer support for co-authoring and commenting. In *Proceedings of the 1990 ACM Conference on Computer-Supported Cooperative Work (CSCW '90)*. Association for Computing Machinery, New York, NY, USA, 183–195. https://doi.org/10.1145/99332.99354

[31] Thao Thu Thi Nguyen and Hoang Yen Phuong. 2021. The Impacts of Collaborative Writing on EFL Students' Paragraph Writing Performance. *International Journal of Science and Management Studies* 4, 14 (2021), 177–190. https://doi.org/10.51386/25815946/ijsms-v4i4p117

[32] Judith S. Olson, Dakuo Wang, Gary M. Olson, and Jingwen Zhang. 2017. How People Write Together Now: Beginning the Investigation with Advanced Undergraduates in a Project Course. *ACM Trans. Comput.-Hum. Interact.* 24, 1, Article 4 (mar 2017), 40 pages. https://doi.org/10.1145/3038919

[33] Javier Onrubia and Anna Engel. 2009. Strategies for collaborative writing and phases of knowledge construction in CSCL environments. *COMPUT EDUC* 53, 4 (2009), 1256–1265. https://doi.org/10.1016/j.compedu.2009.06.008 Learning with ICT: New perspectives on help seeking and information searching.

[34] J. Pei, J. Han, B. Mortazavi-Asl, H. Pinto, Q. Chen, U. Dayal, and M. C. Hsu. 2001. PrefixSpan: Mining sequential patterns efficiently by prefix-projected pattern growth. *Proceedings - International Conference on Data Engineering* (2001), 215–224. https://doi.org/10.1109/icde.2001.914830

[35] Yu Peng, Yanyan Li, You Su, Kailiang Chen, and Shiyan Jiang. 2022. Effects of group awareness tools on students' engagement, performance, and perceptions in online collaborative writing: Intergroup information matters. *The Internet and Higher Education* 53 (2022), 100845. https://doi.org/10.1016/j.iheduc.2022.100845

[36] Dilhan Perera, Judy Kay, Irena Koprinska, Kalina Yacef, and Osmar R. Zaïane. 2009. Clustering and Sequential Pattern Mining of Online Collaborative Learning Data. *IEEE Transactions on Knowledge and Data Engineering* 21, 6 (2009), 759–772. https://doi.org/10.1109/TKDE.2008.138

[37] Peter J Rousseeuw. 1987. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* 20 (11 1987), 53–65. https://doi.org/10.1016/0377-0427(87)90125-7

[38] Daniel K Schneider, Barbara Class, Kalliopi Benetos, Julien Da Costa, and Valérie Follonier. 2014. Learning Process Analytics for a Self-study Class in a Semantic Mediawiki. *Proceedings of The International Symposium on Open Collaboration* (2014), 20:1–20:4. https://doi.org/10.1145/2641580.2641605

[39] Till Schümmer and Stephan Lukosch. 2007. *Patterns for computer-mediated interaction.* Wiley, Hoboken, NJ.

[40] M. Sharples. 1993. *Computer supported collaborative writing.* Springer Verlag.

[41] Antonette Shibani, Simon Knight, and Simon Buckingham Shum. 2018. Understanding Revisions in Student Writing Through Revision Graphs. In *Artificial Intelligence in Education*, Carolyn Penstein Rosé, Roberto Martínez-Maldonado, H Ulrich Hoppe, Rose Luckin, Manolis Mavrikis, Kaska Porayska-Pomsta, Bruce McLaren, and Benedict du Boulay (Eds.). Springer International Publishing, Cham, 332–336.

[42] Minseok Song, Christian W Günther, and Wil M.P. Van Der Aalst. 2009. Trace clustering in process mining. In *Lecture Notes in Business Information Processing*, Vol. 17 LNBIP. 109–120. https://doi.org/10.1007/978-3-642-00328-8_11

[43] Vilaythong Southavilay, Kalina Yacef, Peter Reimann, and Rafael A Calvo. 2013. Analysis of Collaborative Writing Processes Using Revision Maps and Probabilistic Topic Models. In *Proceedings of the Third International Conference on Learning Analytics and Knowledge (LAK '13)*. Association for Computing Machinery, New York, NY, USA, 38–47. https://doi.org/10.1145/2460296.2460307

[44] Neomy Storch. 2005. Collaborative writing: Product, process, and students' reflections. *Journal of Second Language Writing* 14, 3 (2005), 153–173. https://doi.org/10.1016/j.jslw.2005.05.002

[45] Sebastian Strauß and Nikol Rummel. 2021. Problematic interaction patterns during online-collaboration. A library and a survey. In *Proceedings ofthe 14th International Conference on Computer-Supported Collaborative Learning - CSCL 2021*, C. Hmelo-Silver, B. De Wever, and J. Oshima (Eds.). International Society of the Learning Sciences, 19–26. https://doi.org/10.22318/cscl2021.19

[46] Daniel Suthers and Devan Rosen. 2011. A unified framework for multi-level analysis of distributed learning. In *Proceedings of the 1st International Conference on Learning Analytics and Knowledge (LAK '11)*. Association for Computing Machinery, New York, NY, USA, 64–74. https://doi.org/10.1145/2090116.2090124

[47] The Etherpad Foundation. [n. d.]. Etherpad: A real-time collaborative editor for the web. https://github.com/ether/etherpad-lite. Accessed: 2024-12-08.

[48] Anouschka van Leeuwen, Nikol Rummel, and Tamara van Gog. 2019. What information should CSCL teacher dashboards provide to help teachers interpret CSCL situations? *International Journal of Computer-Supported Collaborative Learning* 14, 3 (2019), 261–289. https://doi.org/10.1007/s11412-019-09299-x

[49] Jan-Bennet Voltmer, Natalia Reich-Stiebert, Jennifer Raimann, and Stefan Stürmer. 2022. The role of multi-attributional student diversity in computer-supported collaborative learning. *The Internet and Higher Education* 55 (2022), 100868. https://doi.org/10.1016/j.iheduc.2022.100868

[50] Stanley Wasserman and Philippa Pattison. 1996. Logit models and logistic regressions for social networks: I. An introduction to Markov graphs andp. *Psychometrika* 61, 3 (1996), 401–425. https://doi.org/10.1007/BF02294547
[51] Michael Wininger. 2014. Measuring the evolution of a revised document. *Journal of Writing Research* 6, 1 (2014), 1–28.
[52] Alyssa Friend Wise, Simon Knight, and Simon Buckingham Shum. 2021. *Collaborative Learning Analytics*. Springer International Publishing, Cham, 425–443. https://doi.org/10.1007/978-3-030-65291-3_23
[53] Jiuyun Xu and Jie Liu. 2019. A Profile Clustering Based Event Logs Repairing Approach for Process Mining. *IEEE Access* 7 (2019), 17872–17881. https://doi.org/10.1109/ACCESS.2019.2894905
[54] Yu-Fen Yang. 2010. Students' reflection on online self-correction and peer review to improve writing. *COMPUT EDUC* 55, 3 (2010), 1202–1210. https://doi.org/10.1016/j.compedu.2010.05.017
[55] Meixiu Zhang. 2019. Towards a quantitative model of understanding the dynamics of collaboration in collaborative writing. *Journal of Second Language Writing* 45 (2019), 16–30. https://doi.org/10.1016/j.jslw.2019.04.001
[56] Meixiu Zhang and Wenting Chen. 2022. Assessing collaborative writing in the digital age: An exploratory study. *Journal of Second Language Writing* 57 (2022), 100868. https://doi.org/10.1016/j.jslw.2022.100868

Received 2 July 2024