A Product- and Process-Oriented Tagset for Revisions in Writing

Written Communication
2021, Vol. 39 (1) 1–32
© 2021 SAGE Publications

Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/07410883211052104
journals.sagepub.com/home/wcx

**Rianne Conijn¹, Emily Dux Speltz²,
Menno van Zaanen³, Luuk Van Waes⁴,
and Evgeny Chukharev-Hudilainen²**

Abstract

The study of revision has been a topic of interest in writing research over the past decades. Numerous studies have, for instance, shown that learning-to-revise is one of the key competences in writing development. Moreover, several models of revision have been developed, and a variety of taxonomies have been used to measure revision in empirical studies. Current advances in data collection and analysis have made it possible to study revision in increasingly precise detail. The present study aimed to combine previous models and current advances by providing a comprehensive product- and process-oriented tagset of revision. The presented tagset includes properties of external revisions: trigger, orientation, evaluation, action, linguistic domain, spatial location, temporal location, duration, and sequencing. We identified how keystroke logging, screen replays, and eye tracking can be used to extract both manually and automatically extract features related to these properties. As a proof of concept, we demonstrate how this tagset can be used to annotate revisions made by higher education students in various academic tasks. To conclude, we discuss how this tagset forms a scalable basis for studying revision in writing in depth.

¹Eindhoven University of Technology, Eindhoven, Netherlands

²Iowa State University, Ames, IA, USA

³South African Centre for Digital Language Resources, Potchefstroom, South Africa

⁴University of Antwerp, Belgium

Corresponding Author:

Rianne Conijn, Human-Technology Interaction group, Eindhoven University of Technology,
PO Box 513 5600 MB Eindhoven, Netherlands.

Email: m.a.conijn@tue.nl

Keywords

writing analytics, writing process, revision processes, keystroke logging, eye tracking, annotation

Introduction

Revisions have commonly been argued to play an important role in the writing process (Fitzgerald, 1987; Flower & Hayes, 1981; Scardamalia & Bereiter, 1983). In one of the most cited definitions of “revision,” Fitzgerald (1987) describes the activity of revision in a very simple way: “Revision means making any changes at any point in the writing process” (p. 484). This broad definition also serves as a good starting point for this study. From this perspective, a revision does not necessarily need to correct an *error*; it can involve any change within the written text produced so far (external revision; Murray, 1978), as well as a change in the writer’s mind before text is written down (internal revision; Murray, 1978). Revision can have a wide range of effects. It may influence the outcome of the writing task, such as the writing quality of the written product, the writers’ development, such as the writers’ knowledge about the topic or about writing, and the writing process (Barkaoui, 2016; Fitzgerald, 1987). Given these effects, revision plays a central role in the learning-to-write literature (Allal et al., 2004).

To be able to study the effects of revision and distinguish between different types of revision, a variety of models of revision have been developed, sometimes embedded in more general writing process models (e.g., Flower et al., 1986; Scardamalia & Bereiter, 1983). In addition, several revision taxonomies have been developed to allow for categorization of different types of revisions. Revision taxonomies tend to be either product-oriented or process-oriented (Lindgren & Sullivan, 2006b). Product-oriented taxonomies primarily focus on the effect of the revision on the writing product, such as the orientation of the revision—for example, lower level surface versus meaning-level (semantic) revisions (Faigley & Witte, 1981). Process-oriented taxonomies typically focus on the process of making the revision, such as the time and place of the revision in the writing process—for example, episodes of revision at a single cursor location (Kollberg, 1996).

These variety of models and taxonomies of revision have proven useful for empirical studies, for example, when comparing revision behavior between skilled and less-skilled writers (e.g., Faigley & Witte, 1981) or for determining the effect of instruction on revision (e.g., Sengupta, 2000). However, previous approaches usually only take a few properties of revision into account. In addition, previous models and taxonomies frequently discuss

similar properties but sometimes use different definitions and terminology. Without a clear and uniform annotation guide, these models and approaches are hard to compare. Moreover, current advances in data collection and analysis, such as keystroke logging, eye tracking, and natural language processing, have made it possible to gain a more complete and in-depth analysis of revision compared with earlier studies. Yet, a complete overview of and approach to extracting a comprehensive set of properties is currently lacking.

Therefore, the present article aims to provide a comprehensive product- and process-oriented tagset of revisions, which can be used for analyzing writing products, such as final texts, and (typed) writing process data, such as keystroke logs or text-change logs. Because we aimed to create a comprehensive tagset, we considered all categories found in previous taxonomies of revisions and expanded on these categories with information that can be gained from keystroke logging, eye tracking, and natural language processing. This approach resulted in categorical as well as numerical features to describe revisions. We specifically selected the categories and features that can be used for any text and are not bound to a specific genre. As the current approach is not strictly used to categorize revisions but also to provide detailed descriptions of the revisions, we do not refer to this approach as a *taxonomy*. Rather, we refer to it as a *tagset* of revisions, which can be used to *tag* or annotate a wide range of features of revisions.

Our aim in developing the tagset was to allow for multiple categories per revision since revisions have properties on different levels, which are not necessarily mutually exclusive (Lindgren & Sullivan, 2006a, 2006b). The current tagset is restricted to features that can be automatically extracted using rule-based algorithms implemented in a programming language (e.g., JavaScript or R), and features that can be annotated manually using writing product or writing process data. This makes the tagset *measurable*, and hence relatively easy to implement in future studies. As a proof of concept, we show how the tagset can be used on an existing typing data set, consisting of keystroke data, screen replays, and eye-tracking data. We conclude by discussing how this tagset can be used in future studies of revision.

Models of Revision

One of the earliest writing process models distinguishes two subprocesses of revision (reviewing): *reading* and *editing* (Hayes & Flower, 1980). In this model, editing involves a change in the text which is spur-of-the-moment correction, while reviewing involves “systematic examination and improvement of the text” (Hayes & Flower, 1980, p. 18). The editing process can be

further divided into four modes: editing to adhere to writing conventions, editing to improve semantics, evaluating to improve readers' understanding, and evaluating to improve readers' acceptance (Hayes & Flower, 1980). In a later model, *evaluating* and *revising*, rather than reading and editing, were considered as the two subprocesses of reviewing (Flower & Hayes, 1981; Hayes et al., 1987). Here, evaluating refers to the process of identifying where changes need to be made in the text and revising to the actual change of the text.

Scardamalia and Bereiter (1983) expanded upon this reviewing process with the compare, diagnose, operate model. In this model, writers *compare* the mental representation of the text written so far with the mental representation of the intended text. If there is a perceived mismatch between these representations, they *diagnose* what needs to be changed and then *operate* this change on the text (Scardamalia & Bereiter, 1983). Flower et al. (1986) represented these three processes—detecting the problem, diagnosing the problem, and selecting a strategy—in a more detailed model of revision. In this model, writers start with the task definition, and then read their text to comprehend and to evaluate whether their goals are met. This results in a problem representation that can be ill-defined (merely a detection of the problem) or well-defined (a diagnosis of the problem). Based on the problem representation, the problem can be ignored or a strategy will be selected either to rewrite or revise (Flower et al., 1986).

Within the models of revision, different modes of processing can be distinguished. Flower and Hayes (1981) stated that evaluating can include different modes, such as evaluating the text written so far as well as evaluating the writers' planning or their unwritten thoughts and statements. Other researchers explicitly distinguished the process of revision in two modes: internal and external revisions, mostly measured using think aloud protocols (Murray, 1978). Internal revisions are defined as mental revisions that are made before transcription, while external revisions are revisions that are visible in the written production (Murray, 1978). In later studies, internal revisions were further subdivided into prelinguistic revisions, or changes made to nonlinguistic mental representations, and pretextual revisions, or revisions made to linguistic mental representations, both of which could affect conceptual content as well as formulation (Lindgren & Sullivan, 2006b; Stevenson et al., 2006).

In the current tagset, we solely focus on external revisions, as these are the only revisions that are directly visible in writing product and writing process data, such as keystroke logs or text-change logs.

Current Revision Tagset

In the following, we describe the different properties of revisions in the current revision tagset. These properties are described in relation to the previous models and taxonomies, with a specific focus on how these taxonomies could be used to create a comprehensive tagset that can be used to (manually or automatically) annotate revisions. In addition, we briefly describe how features related to these revisions can either be annotated manually or extracted automatically. A more extensive description of how to annotate the manual features can be found in the annotation guide in the online supplementary materials.

Revision Events

Before undertaking the task of annotating revisions, we first need a unit of analysis, for which we select a *revision event*, that is, a specific episode during writing where a single revision takes place. We define a revision event as an external revision, which starts from a deletion or insertion and ends with normal text production unrelated to the revision or a new external revision. Here, we consider text production as being related to the revision, as it is used to replace the deleted text.

For the purposes of the present tagset, a revision event begins when (1) the writer begins deleting characters in the text, or (2) the writer moves the cursor to a different location in the text and then begins producing new characters. The revision event is considered finished when (1) the writer initiates a new revision event as defined above, or (2) the writer continues with normal text production unrelated to the revision. The revision events can be, for a large part, automatically identified using rule-based algorithms based on the key-stroke data. However, it is difficult to automatically identify when a revision event is followed by an episode of new text production (second case for finishing a revision event). Therefore, the end of the revision is marked manually by data set annotators.

Revision Properties

After identifying the revision event, each revision event is described using a set of properties, based on previous literature (although the definitions, terminology, and usage of these properties sometimes differ). We distinguished nine different properties of external revisions. Four of the properties are product-oriented properties, focusing on the effect of the revision on the writing product: orientation, evaluation, action, and domain. The other five

properties are process-oriented, focusing on the process of making the revision: trigger, spatial location, temporal location, duration, and sequencing. For each property, we identified a wide variety of features related to the property. An overview of all features identified related to each of these properties can be found in Table 1.

A. Trigger. The trigger of the revision describes the cause of the revision. The different types of triggers identified in the literature include reading and evaluating the text written so far, or reading the writing task and evaluating the written text (Tillema et al., 2011). Other researchers have focused on revisions triggered by errors (Stevenson et al., 2006), also known as conventional revisions—revisions that are necessary to correct to fulfill linguistic requirements (Allal, 2000)—in contrast to optional or nonerror-triggered revisions. Error-triggered revisions are related to the evaluation or quality of the revision, that is, whether the linguistic requirements were already met at the starting point of the revision (Wobbrock & Myers, 2006). In the current tagset, error versus nonerror-triggered revisions are categorized under evaluation (C). As the trigger itself mostly happens internally, the trigger cannot be observed in the keystroke data. However, sometimes the result of revision is visible, which could in principle provide an indication of the trigger of the revision. For example, a typo revision (see orientation, B) might be triggered by evaluating the written text and discovering a typo. However, it can never be certain what has exactly triggered the revision. Therefore, we do not include internal triggers in the tagset.

With the advent of automatic writing evaluation systems, a revision during writing can also be triggered externally, that is, not by the writer itself. One common example is spelling or grammar checkers which trigger revisions in spelling or grammar (Figueredo & Varnhagen, 2006). As these external triggers can be identified in writing process data, they are included in the current tagset. Specifically, one automatic feature is included: type of feedback (i.e., no feedback, spelling, grammar, punctuation, and other).

B. Orientation. One of the most frequently used classifications of revision involves the orientation or the depth of the revision (Faigley & Witte, 1981). Faigley and Witte's (1981) taxonomy of revisions distinguishes surface and semantic revisions. Surface revisions are further subdivided into formal changes (i.e., spelling [tense, number, and modality], abbreviations, punctuation, and format) and meaning-preserving revisions that paraphrase but do not alter the meaning of concepts in the text. Semantic changes have a more fundamental impact on the actual text and can be divided into minor (microstructure) and major (macrostructure) changes which alter the summary of

Table 1. Features in the Revision Tagset per Property.

Property		Feature	Type
A. Trigger	1	Feedback type (no feedback, spelling, grammar, punctuation, other) ^a	Categorical
B. Orientation	1	Surface change	Binary
	1.1	Typography	Binary
	1.2	Capitalization	Binary
	1.3	Punctuation	Binary
	1.4	Spelling	Binary
	1.5	Grammar	Binary
	1.6	Cosmetics/presentation	Binary
	1.7	No change	Binary
	1.8	Wording/phrasing	Binary
	2	Semantic (deep) change	Binary
	2.1	Microstructure changes (supporting info, emphasis, understate, coherence, cohesiveness)	Categorical
	2.2	Macrostructure changes (overall aim, subtopic)	Categorical
C. Evaluation	1	Correct start	Binary
	2	Correct revision	Binary
D. Action	1	Action (insertion, deletion, substitution, reordering)	Categorical
E. Domain	1	Linguistic domain (subword, word, phrase, clause, sentence, paragraph)	Categorical
	2	Number of backspaces	Numeric
	3.1	Number of characters deleted	Numeric
	3.2	Number of characters inserted	Numeric
	4.1	Number of words deleted	Numeric
	4.2	Number of words inserted	Numeric
	5.1	List of part-of-speech tags words deleted	List
	5.2	List of part-of-speech tags words inserted	List
	6.1	List of chunk tags words deleted	List
	6.2	List of chunk tags words inserted	List
	7.1	Number of sentences deleted	Numeric
	7.2	Number of sentences inserted	Numeric
F. Spatial location	1	Word finished	Binary
	2	Intended word	String
	3.1	Word initial	Binary
	3.2	Clause initial	Binary
	3.3	Sentence initial	Binary
	4	Characters from leading edge	Numeric
	5	Words from leading edge	Numeric
	6	Precontextual (= 1 – contextual)	Binary
	7	Immediate (= 1 – distant)	Binary
	8.1	Characters from start sentence	Numeric
	8.2	Characters from start process	Numeric
	8.3	Characters from start product	Numeric

(continued)

Table 1. (continued)

Property		Feature	Type
G. Temporal location	1	Time from start process (min)	Numeric
H. Duration	1	Duration (seconds)	Numeric
	2	Pause before revision (seconds)	Numeric
I. Sequencing	1	Overrides previous revision	Binary
	2	Continues previous revision	Binary
	3.1	Repetitive (leading edge)	Binary
	3.2	Repetitive (immediate)	Binary
	4.1	Embedded revision (leading edge)	Binary
	4.2	Embedded revision (immediate)	Binary
	5.1	Sequence forwards (leading edge)	Binary
	5.2	Sequence forwards (immediate)	Binary
	6.1	Sequence backwards (leading edge)	Binary
	6.2	Sequence backwards (immediate)	Binary
	7	Time from previous revision (seconds)	Numeric
	8	Characters from previous revision	Numeric

^aThis feature was not extracted in the current study. The numbers indicate the hierarchy within the features. For example, B.I.1 indicates a surface change, and more specifically a typography revision.

the text. Meaning-preserving, microstructure, and macrostructure revisions are all further divided into additions, deletions, substitutions, permutations, distributions, and consolidations. In the current tagset, we treat these subcategories separately under revision action.

Other studies have used variations of Faigley and Witte's (1981) taxonomy on revision changes. For example, Lindgren and Sullivan (2006a) classified contextual revisions, revisions away from the leading edge of the text produced so far (see spatial location), into form and conceptual revisions. Form revisions are similar to Faigley and Witte's (1981) surface revisions but are further extended into typography (slip of the finger), spelling, grammar, punctuation, formatting, and meaning-preserving revisions. Several subcategories were further specified: spelling (i.e., revision, substitution, deletion, and homophone), grammar (i.e., verb tense, subject/verb agreement, preposition, conjunction, article, pronoun, genitive, adverb, and other), and format (i.e., punctuation, capitalization, paragraph, and other format). In addition, changes in language (L2 to L1 and L1 to L2) were added as meaning-preserving changes. The conceptual revisions are based on Faigley and Witte's (1981) semantic revisions and consist of content-based revisions (microstructure and macrostructure changes) and balance revisions, related to the topic

or audience (Lindgren & Sullivan, 2006a). Microstructure changes are changes that affect meaning but would not alter the gist of a text, whereas macrostructure changes are changes that would alter the summary of a text. Allal (2000) chose a slightly different approach and distinguished between semantics (lexical variations and changes of meaning), text organization (segmentation, connection, and cohesion), and spelling (lexical and grammatical aspects). Here, the first two could be considered semantic changes, while the latter constitutes surface changes. Sometimes, the orientation labels for the semantic changes are tailored to a specific genre; for example, for argumentative writing, F. Zhang, Hwa, et al. (2016) distinguished semantic changes in thesis/idea (claim), rebuttal, reasoning (warrant), evidence, and general content changes.

A few studies report on functions or purposes of revisions which we argue would be closely related to the orientation of the revision. For example, Monahan (1984) identified five purposes: cosmetic, mechanical, transitional, informational, and stylistic, where cosmetic, mechanical, and stylistic could be considered as surface changes, and transitional and informational as semantic changes. Similarly, Falvey (as cited in Min, 2006) distinguished grammatical, cosmetic (surface changes), texture (increasing cohesion and coherence), unnecessary expression, and explication (increasing explicitness) revisions.

For the current tagset, we primarily distinguish between surface and semantic revisions, and each of those categories encompasses multiple levels of subcategories. Surface changes are divided into formal changes (typography, spelling, grammar, capitalization, punctuation, cosmetics/presentation, and no change) and meaning-preserving changes (changes in wording/phrasing, which preserved the meaning of the sentence). All these categories were manually annotated, except for punctuation, capitalization, and no change. The revision was automatically annotated as punctuation when the change (difference between inserted and deleted text) only included punctuation characters; as capitalization if the change only included a change from lower case to upper case, or vice versa; and no change when the deleted text was the same as the inserted text. Typography and spelling were distinguished by applying the guidelines used in previous manual annotations (Stevenson et al., 2006; Wengelin, 2007). Additionally, we divided semantic changes (deeper changes that would alter the meaning) into microstructure and macrostructure changes. These changes were further divided in subcategories, based on Faigley and Witte (1981) and higher order and lower order concerns in writing center studies (Winder et al., 2016). Here we specifically only selected subcategories unrelated to genre. For microstructure changes, we included the following subcategories: adding or removing

supporting information, changing emphasis, understatement, coherence, and cohesiveness. Macrostructure changes consisted of altering the overall aim and adding or removing entire subtopics. If it was unclear which subcategory should be selected, for example, because a revision started in an unfinished word, all possible orientations needed to be selected (e.g., both wording and typographic error when only the first letter of a word was revised).

C. Evaluation. The evaluation of the result of a revision has typically been used for formal changes only: surface changes necessary to fulfill linguistic requirements (e.g., spelling, grammar, and punctuation). Possible evaluations that have been used include correct or incorrect (Allal, 2000); correct, incorrect, or neutral (Chanquoy, 1997); or successful or unsuccessful revisions (Stevenson et al., 2006). Moreover, some studies describe whether the starting point was correct: in other words, whether an error or nonerror was revised (e.g., from a correct nonerror into an error; Wobbrock & Myers, 2006). Crawford et al. (2008) evaluated the quality of all revisions, as opposed to formal changes only. As it is harder to determine whether semantic changes are “correct,” they renamed the labels to *increase*, *decrease*, and *neutral*.

In the current tagset, the evaluation of the revision was manually annotated. This was only done when the revision was oriented toward typography, spelling, or grammar, as these types of formal changes are necessary for satisfying linguistic requirements, and hence can be evaluated as correct or incorrect. In the current tagset, we describe whether the starting point was correct or incorrect and whether the result of the revision was correct or incorrect.

D. Action. Most of the actions or revision operations stem from work on the classification of single letter errors which consists of an extra letter (insertion), a missing letter (omission), a wrong letter (substitution), and a transposition of two adjacent letters (transposition; Damerau, 1964). These have been translated into four revision actions: deletion, insertion, substitution, and reordering (also referred to as transposition, reorganization, and permutations), respectively (e.g., Allal, 2000; Sommers, 1980). These actions could involve a single letter but also a larger linguistic unit (or *linguistic domain*, see below). Some researchers added one or more actions to these four actions. For example, Monahan (1984) added embedding, a specific form of insertion, where new information is embedded within the structure of already written text. Faigley and Witte (1981) included distribution, where a single unit is changed into more than one unit (e.g., one sentence is distributed into two), and its opposite, consolidation, where two or more units are combined into one unit.

In the current tagset, the revision action was automatically classified into insertion (addition), deletion, substitution, and reordering (transposition) of the given linguistic domain. For the revisions below word level, the revision action was determined by the edit distance of the deleted and inserted text, up to the manually annotated revision end, that is, based on the full revision event. The edit distance was calculated using the restricted Damerau-Levenshtein distance (Boytssov, 2011), which calculates the minimum number of insertions, deletions, substitutions, and transpositions to change a word (the deleted text) into another word (the inserted text). For the revisions at word level and above, the revision was coded as insertion if more words were inserted than deleted; as a deletion if more words were deleted than inserted; as a transposition if the same words were replaced, but in a different order; and as a substitution if the number of words inserted were equal to the number of words deleted, but not resulting in a transposition.

E. Linguistic domain. Several studies included the linguistic domain affected by the revision, also known as the level, scope, range, or size of the revision. In contrast to orientation, which is mostly focused on the semantics (meaning) of the revision, linguistic domain is focused on the linguistic properties or scope of the revision. Monahan (1984) distinguished subword (surface), word, phrase, clause, sentence, paragraph, and discourse level, while Stevenson et al. (2006) only distinguished below word, below clause, and clause and above. Some researchers have focused more on the smaller domains, for example, subgrapheme, grapheme, morpheme, word, and above-word (Lindgren et al., 2019). Others added domains such as theme (extended statement of one idea; see Sommers, 1980), punctuation (Crawford et al., 2008), or revisions made within the current functional component being realized (forward progressions) versus revisions made within the current component, but across component boundaries (communicative progression; Bowen & Van Waes, 2020). Xu (2018) took a lower level approach by merely counting the characters deleted or produced.

In the current tagset, the linguistic domain was manually classified into one of the following six categories: revision within a subword (part of a word; below word-level), word, phrase, clause, sentence, or paragraph. In addition to the manual features, several automatic features were extracted. These automatic features have been found to be related to the manual categories (e.g., Daxenberger & Gurevych, 2013), and hence might be used as a replacement of the manual annotation. The automatic features consisted of the number of backspaces, characters, words, and sentences inserted and deleted. The number of characters, words, and sentences inserted was calculated up to the manually annotated revision end. Note that the number of

backspaces and the number of characters deleted are related, but will differ, for example, when different revision techniques are used (e.g., selection of 10 characters and one backspace keystroke vs. 10 backspace keystrokes). Finally, if one or more words were deleted or inserted, we also extracted a list of part-of-speech tags of the words deleted or inserted, using the part-of-speech tagger in the Natural Language Toolkit in Python (Bird et al., 2009). Likewise, when more than one word was deleted or inserted, we extracted a list of chunk tags of the words deleted or inserted, using the chunk tagger in the Natural Language Toolkit in Python (Bird et al., 2009).

F. Spatial location. Revisions occur at different locations in the text produced so far. We refer to this as the spatial location in the writing process. We distinguished two slightly different operationalizations of the spatial location in the writing process. Lindgren and Sullivan (2006a, 2006b) distinguish pre-contextual and contextual revisions. Precontextual revisions are defined as revisions at the leading edge of the text produced so far. Contextual revisions are defined as revisions made when the writer moves away from the leading edge and makes a revision in a previously written and completed sentence. In this way, revisions away from the leading edge but in an uncompleted sentence are left uncategorized. Accordingly, Baaijen et al. (2012) used a simpler categorization and only distinguished between revisions at the leading edge and revisions elsewhere in the text. Contrary to precontextual and contextual revisions, Thorson (2000) distinguishes immediate (or intermediate) and distant revisions. Immediate revisions are revisions at the point of the cursor location, and distant revisions include revisions where the cursor moved away to start a revision elsewhere in the text. For distant revisions, the distance from the cursor location has been included, such as the number of lines below or above the cursor location (Van Waes & Schellens, 2003). Likewise, a distinction has been made between movements from the cursor location forward and backward in the text (Bowen & Van Waes, 2020). Xu (2018) added a third category to immediate and distant revisions: end revisions, which are revisions made after the completion of the whole text. In the current tagset, we treat this third category separately under temporal location in the writing process.

In the current tagset, we automatically identified several features related to spatial location. First, we distinguished precontextual and contextual revisions, where a precontextual revision is defined as a revision within zero characters from the leading edge. As Lindgren et al. (2019) recommended, we also classified a revision as precontextual if there were only invisible characters between the revision and the leading edge (e.g., a trailing white space). However, sometimes a writer first starts with a rough outline (e.g.,

intro-body-conclusion) and then starts to fill in the gaps. In this case, there will almost always be characters behind the cursor, and all revisions will be contextual. Therefore, we also automatically extracted immediate versus distant revisions, where immediate revisions are revisions at the cursor location. In addition, we included the number of words and characters from the leading edge (cf. Van Waes & Schellens, 2003). Last, we included the number of words from the start of the sentence, the start of the writing product (the number of words in the written product up to the revision), and the start of the writing process (the number of words typed so far).

In addition, we added several manual features of spatial location not previously discussed in revision taxonomy literature. Often, corrections are made when a word is not yet finished. This makes it difficult to manually annotate the revision, as it might be impossible to know what the writer was supposed to write, especially if only one or a few characters were typed (Lindgren et al., 2019). This is also a difficult task for current computer algorithms, as half-written words are especially hard to parse, for example, for part-of-speech tagging. Therefore, as a metric of uncertainty of the manual annotations, the annotators indicated whether the word in which the revision started was finished. If the word was not finished, the intended word was guessed (if possible). Last, the location up to where the characters are deleted, or where characters are inserted, could provide information on the orientation of the revision. For example, when characters are deleted up to the middle of a word, it would be less likely to be a semantic change compared with when characters are deleted up to the start of a sentence. Therefore, we also included whether the location of the revision was word initial, clause initial, and/or sentence initial.

G. Temporal location. Temporal location refers to the point in time when the revision is made. Temporal location has been defined in the literature both within and between drafts. Monahan (1984) described the temporal locations between drafts: prewriting stage, during the first draft, between drafts, during the second draft, and after the second draft. Others described the temporal location within a draft. Tillema et al. (2011) split the writing process into episodes of equal durations. With three episodes, for example, you could identify whether the revision took place in the begin, middle, or end of the writing process (Barkaoui, 2016; Van Waes & Schellens, 2003). Others only distinguished revisions made during text production and revisions made at the final writing stage after completing the whole text (Xu, 2018). Last, some researchers used a more computational approach and defined the temporal location as the time elapsed from the start of the writing process (M. Zhang, Hao, et al., 2016).

In the current tagset, only one automated feature was extracted for the temporal location of the revision: the time (in milliseconds) until the revision from the start of the writing process. Combined with the total time, the time until revision can be used to determine any episode in which the revision took place (e.g., begin, middle, and end). In addition, combined with the version number of the document, this can be used to determine any time within a writing session (e.g., end of first draft and beginning of final version).

H. Duration. One rather technical aspect of revisions that has been included in previous studies is the revision duration. The duration is usually measured in (milli)seconds from the start of the revision until the end of the revision (Xu, 2018). Duration has also been expressed as the length of an R-burst, where an R-burst is a sequence of text production involving a revision and bounded by long pauses (e.g., pauses >2 seconds; see Baaijen & Galbraith, 2018).

In the current tagset, two features related to the duration of the revision were automatically extracted. First, the duration of the revision was calculated by extracting the time from the first editing key press until the last editing key press. The last editing keystroke was extracted from the manually annotated revision end. Second, the pause time before the revision was calculated by extracting the time from the last key press before a revision until the first key press of the revision.

I. Sequencing. All previous properties of revision are related to a single revision event. Sequencing describes the relation of a revision to a preceding revision in time. As Lindgren and Sullivan (2006a) describe, some revisions might be single, independent revisions, while other revisions might be part of a series of revisions, also known as revision episodes. As examples, they describe three types of revision episodes, as suggested by Kollberg (1996). First, they discuss episodes of revision at a single cursor location. For example, a writer may start a sentence, delete it, start it again, delete it again, and then write the final full sentence. In this case, the start of the sentence (same cursor location) is revised twice. Second, Lindgren and Sullivan (2006a) describe episodes with embedded revisions, where a revision is made within a previous revision. Third, they note episodes with a sequence of revisions. For example, to maintain consistency, one might change the wording or spelling of a specific word and change that word throughout the text. Williamson and Pence (1989) also discussed a fourth category, where a revision “inspires” another revision without reviewing the text. For example, a revision of the last sentence of a paragraph might be followed by a revision in the beginning

of that paragraph. Another approach to examining the sequencing of revisions is the S-notation, which maps all deletions and insertions in the text to the spatial location in the writing process (Kollberg, 1996). With this S-notation, the nonlinearity of the writing process can be analyzed, and connected episodes of revision can be automatically identified (Kollberg, 1996; Severinson-Eklundh & Kollberg, 2001).

In the current tagset, four types of sequencing were automatically identified: repetitive revisions, embedded revisions, sequence forward revisions, and sequence backward revisions. Repetitive revisions are operationalized as revisions which start (in the writing product) within two characters from the start (in the writing product) of the previous revision. This could, for example, happen when a typo is incorrectly revised, resulting again in a typo, which is then correctly revised (e.g., poeple > peple > people). Embedded revisions are revisions where the range of character positions of the start and end of the revision (within the writing product) fall within the range of character positions of the previous revision. For example, this occurs when a sentence is inserted within the text, followed by a wording revision within that sentence. Sequence forward includes revisions where the subsequent revision was made further on in the text (forward from the last revision), while sequence backward includes revisions where the subsequent revision was made earlier in the text. All four types were calculated based on both the leading edge (precontextual/contextual) and cursor (immediate/distant) spatial locations. Last, we automatically extracted metrics related to the distance from the previous revisions: time and number of characters from the previous revision.

In addition to the automatic features, we added two manual features: overrides previous revision and continues on previous revision. A revision overrides the previous revision if it is repetitive, that is, changing the same linguistic domain at the cursor location. A common example of this is when a writer makes a typographical error and attempts to revise it, but then makes a typographic error again in that revision attempt. A revision continues on a previous revision if the previous revision caused the subsequent revision. For example, changing a noun from singular to plural might result in changing the verb to maintain subject-verb agreement.

Proof of Concept

As a proof of concept, we used our tagset to describe the revisions made by university students while conducting an academic writing task. The manual features were manually annotated by multiple raters, and the other features were automatically extracted.

Data Set

Our data set was created by sampling writing process data from a large data store containing anonymized writing process log files. These log files were recorded as part of various prior research studies conducted using a web-based word-processing tool that collects keystroke data and eye fixations during writing (Chukharev-Hudilainen, 2019; Chukharev-Hudilainen et al., 2019; Ranalli et al., 2018). In addition, this tool offers playback functionality to replay the screen recording combined with the gaze-point marker. To include a wide variety of revisions, a semirandom subset was collected from the data store, while ensuring a diversity of tasks and writer backgrounds.

The subset of data consisted of data from 20 graduate students and 20 undergraduate students (all speakers of English as their first language), and 25 learners of English as a second language (most likely undergraduate students based on the original study that contributed to this portion of the data set). The graduate students were prompted to write a 150- to 250-word abstract based on one out of four research articles (articles were counterbalanced). The undergraduate students wrote one out of two argumentative tasks (counterbalanced) adapted from the Test of English as a Foreign Language. The 65 participants had a total of 7,120 revision events ($M = 110$, $SD = 53$). For every revision event, all manual and automatic features listed in Table 1 were extracted.

Feature Extraction

For the manual annotations, a spreadsheet was created with (for every revision event; i.e., for every row) the revision event id, removed characters, and typed characters (columns). Next to this information, the annotator could further explore the revision event (in the context of the writing process) using the visual playback function in the interface of the word-processing tool, providing a keystroke-by-keystroke animated reconstruction of the text production process with an overlaid eye fixation marker. For all features that needed to be manually annotated, we created an extensive annotation guide. Within this guide, guidelines, explanations, and examples are provided for each label. This annotation guide was first created by the authors through several rounds of discussion, trial coding, and evaluation. Thereafter, four annotators were trained to manually annotate using the annotation guide. In each training round, we explained the guidelines to the annotator, the annotator coded a sample document, disagreements were discussed, and where necessary, we clarified the annotation guide with additional explanations or examples. The final annotation guide can be found in the online supplementary materials.

After each training round, the interrater reliability (IRR) was calculated to determine whether more training was needed. In addition, this showed which features specifically required the most attention, that is, which features received the lowest reliability. Krippendorff's alpha (Krippendorff, 2011) was used as IRR metric, as it allows for multiple raters, multiple measurement levels (e.g., numerical and categorical) and is reliable for coding sparse categories, as it focuses on the disagreement rather than the agreement. Krippendorff's α of 1 indicates perfect agreement, while $\alpha = 0$ indicates agreement as produced by chance. It has been argued that an $\alpha \geq 0.80$ is considered "good" and an $\alpha \geq 0.667$ "acceptable" (Krippendorff, 2004). We considered an annotator fully trained when the IRR of all features between author and annotator were similar to the IRR between two authors. In total, two to three rounds were needed to train each annotator, taking 3 hours on average.

After the training, the annotators independently annotated the documents. On average, it took the annotators 1 minute per revision to annotate. Of the 65 documents annotated, 15 (23%) documents were randomly selected to be annotated twice (by different pairs of annotators) to calculate the IRR of the fully annotated data set. The IRR of the manual annotations (after training) can be found in Table 3.

The automatic features were extracted using JavaScript and R as described in the Current Revision Tagset section. The annotated data set can be found at (Conijn et al., 2020).

Descriptive Statistics

The descriptive statistics of the numerical features are shown in Table 2. In the current data set no feedback was provided, so the trigger of feedback is always annotated as "no feedback." Regarding the orientation of the revision, most revisions were surface revisions (92%), and especially typographic revisions or wording/phrasing revisions. Only 14% of the revisions were deep revisions. An overview of all subgroups of surface and deep revisions can be found in Figure 1. Note that the total is more than 100% here, which indicates that several revisions are annotated both as surface and as deep revision (444 revisions, 6.2%). Thus, for some revisions, it was not visible in the keystroke data and screen replays with eye fixation whether the revision was a surface or deep revision.

Most of the spelling, grammar, punctuation, and capitalization revisions were correctly revised (85%). For some of the revisions, the previous form was already correct (5%). Most revisions were insertions (40%), one fourth of the revisions were deletions, and one fourth were substitutions. Only 3%

Table 2. Descriptive Statistics of the Numerical (Automated) Features in the Revision Tagset per Student.

Property		Feature	M (SD)
A. Trigger	1	No trigger	100% (0.0%)
B. Orientation	1	Surface	92.6% (6.2%)
	2	Semantic (deep)	14.1% (8.6%)
C. Evaluation	1	Correct start	4.7% (4.2%)
	2	Correct revision	85.2% (9.4%)
D. Action	1.1	Insertion	40.0% (15.5%)
	1.2	Deletion	25.2% (7.6%)
	1.3	Substitution	24.4% (10.3%)
	1.4	Reordering	3.0% (2.7%)
E. Domain	2	Number of backspaces	2.4 (0.5)
	3.1	Number of characters deleted	3.5 (1.4)
	3.2	Number of characters inserted	8.2 (8.7)
	4.1	Number of words deleted	1.1 (0.2)
	4.2	Number of words inserted	1.7 (1.5)
	7.1	Number of sentences deleted	0.02 (0.02)
	7.2	Number of sentences inserted	0.04 (0.09)
F. Spatial location	1	Word finished	51.0% (12.0%)
	3.1	Word initial	43.1% (10.9%)
	3.2	Clause initial	13.7% (6.7%)
	3.3	Sentence initial	10.2% (6.3%)
	4	Characters from leading edge	69.3 (91.4)
	5	Words from leading edge	11.6 (15.7)
	6	Precontextual (= 1 – contextual)	77.9% (17.1%)
	7	Immediate (= 1 – distant)	86.2% (10.1%)
	8.1	Characters from start sentence	69.2 (31.7)
	8.2	Characters from start process	814 (349)
	8.3	Characters from start product	817 (345)
G. Temporal location	1	Time from start process (minutes)	8.5 (4.0)
H. Duration	1	Duration (seconds)	3.1 (3.0)
	2	Pause before revision (seconds)	2.0 (1.1)
I. Sequencing	1	Overrides previous revision	13.8% (7.6%)
	2	Continues previous revision	14.6% (8.3%)
	3.1	Repetitive (leading edge)	23.9% (10.1%)
	3.2	Repetitive (immediate)	23.8% (10.4%)
	4.1	Embedded revision (leading edge)	0.2% (0.5%)
	4.2	Embedded revision (immediate)	0.2% (0.5%)
	5.1	Sequence forwards (leading edge)	8.1% (8.3%)
	5.2	Sequence forwards (immediate)	4.8% (5.2%)
	6.1	Sequence backwards (leading edge)	1.4% (2.1%)
	6.2	Sequence backwards (immediate)	1.2% (1.8%)
	7	Time from previous revision (seconds)	6.7 (3.4)
	8	Characters from previous revision	7.0 (11.1)

Figure 1. Descriptive statistics of the orientation subcategories.

of the revisions involved reordering of characters or words. The domain of the revisions was rather small. Two thirds of the revisions involved only a few characters, and 24% of the revisions involved a word. Only 8% of the revisions were larger than a word. This was also shown in the low number of characters inserted and deleted, on average. Half of the time, the revisions were made while the current word typed was not finished yet. In 43% of the revisions, the revision deleted up to a word-initial position or inserted at the start of a word.

Most revisions were made at the leading edge (precontextual revision) and at the point of inscription (immediate). The revisions were spread out over the full writing process, started after a pause of 2 seconds, and took 3 seconds, on average. A substantial amount of the revisions was found to follow-up on previous revisions, with 24% of the revisions repetitive, 8% sequencing forward based on the leading edge, and 5% sequencing forward based on the cursor position. Only 1% of the revisions was sequencing backward, and less than 1% consisted of embedded revisions.

Reliability

The IRRs (Krippendorff’s alpha) of the manual annotations are shown in Table 3. For comparison with previous studies, the percentage agreements are

Table 3. Interrater Reliability (IRR; Krippendorff's α) and Percentage Agreement of the Manually Annotated Features in the Revision Tagset, $N = 7,120$.

Property		Feature	IRR (α)	Percentage agreement
General	1	Revision start [Y/N]	0.96	99%
	2	Revision end	0.74	81%
B. Orientation	1	Surface	0.64	91%
	1.1	Typography	0.71	87%
	1.4	Spelling	0.74	95%
	1.5	Grammar	0.69	86%
	1.6	Cosmetics/presentation	0.83	99%
	1.8	Wording/phrasing	0.75	83%
	2	Semantic (deep)	0.59	86%
C. Evaluation	2	Deep (specify) ^a	0.22	84%
	1	Correct start	0.69	97%
	2	Correct revision	0.66	90%
E. Domain	1	Domain (specify) ^a	0.59	79%
F. Spatial location	1	Word finished	0.70	85%
	2	Intended word	0.71	80%
	3.1	Word initial	0.80	90%
	3.2	Clause initial	0.68	92%
	3.3	Sentence initial	0.82	97%
I. Sequencing	1	Overrides previous revision	0.55	90%
	2	Continues previous revision	0.27	82%

^aIRR is calculated once for the full category, as all labels are mutually exclusive.

included as well. The interrater reliabilities were acceptable to good for the annotation of the revision and the end of the revision, and for the features within the properties evaluation and spatial location ($\alpha = 0.66$ - 0.96).

For orientation, the features proved to be more difficult to annotate. The subcategories of surface revisions showed acceptable to good reliabilities ($\alpha = 0.71$ - 0.83), comparable to previous annotations of surface revisions (Daxenberger & Gurevych, 2012). However, the reliability for the overarching categories of surface revisions ($\alpha = 0.64$) and semantic revisions ($\alpha = 0.59$) and subcategories of semantic revisions ($\alpha = 0.22$) were low. The lower reliability for these features can be largely explained by the fact that these categories were fairly unbalanced. Most of the revisions were surface revisions, and the deep revisions were rather sparse in our data set. Krippendorff's alpha penalizes sparse categories more. This effect can also be seen in the comparatively high percentage agreement for these features

(84% to 91%). These percentage agreement scores are similar to previous studies that annotated orientation (e.g., Barkaoui, 2016; Stevenson et al., 2006).

To further explore the low reliability for the subcategories of semantic revisions, we examined the confusion matrix for these subcategories. First, it should be noted that there were only 91 revisions that were coded as a deep revision by at least two annotators. Hence, the IRR of the subcategories of a deep revision is based on only 91 revisions. For these revisions, the majority (55%) was coded as a supporting info revision. This indicates that the other types of deep revisions were very sparse in this specific data set, and hence, for these specific (short) writing assignments, detailing the subcategories of deep revisions might not be very valuable. Two common areas of confusion could be distinguished: emphasis was often confused with supporting info (6 out of 8 times) and cohesiveness was often confused with coherence (5 out of 9 times). This indicates that more attention needs to be provided to these categories within the training, especially as these categories may be barely present in the training data set (as was the case here). Therefore, we added additional examples of these categories in the annotation guidelines.

A relatively low IRR was found ($\alpha = 0.59$) for domain as well. Again, this might be due to some of the sparse categories in domain (e.g., revisions at the sentence or paragraph level), as the percentage agreement is higher: 79%. This percentage agreement is slightly lower compared with previous studies (e.g., Crawford et al., 2008; Min, 2006). The confusion matrix shows that the disagreements between the raters were mostly between adjacent categories (e.g., phrase vs. clause or subword vs. word). Thus, the annotations of domain still provide some insights into the size of the domain. In addition, the disagreements indicate that the training of the annotators needs to emphasize the differences between adjacent categories.

Last, sequencing proved to be difficult to annotate, and especially when it considered a “continues on previous revision” sequence ($\alpha = 0.27$). These categories have not been annotated in previous studies, so the reliabilities cannot be compared with previous work. Again, this may partly be because these revisions were sparse. In addition, the annotators indicated that it was sometimes hard to determine whether a revision continued on a previous revision, especially when it was interrupted by one small (e.g., typo) revision. Also, the intended word was not always known, making it difficult to identify whether the revision followed up on a previous revision. Fortunately, a wide variety of automated features are available, potentially making the manual annotation of sequencing unnecessary in many cases.

Discussion

In this study, we aimed to provide a comprehensive product- and process-oriented tagset of revisions. The tagset includes nine properties to describe external revisions: trigger, orientation, evaluation, action, linguistic domain, spatial location, temporal location, duration, and sequencing. As a proof of concept, we showed how keystroke logging, screen replays, and eye tracking could be used to both manually annotate as well as automatically extract features related to these properties.

Until now, revisions in keystroke logging studies have often been operationalized in a simplified way. More specifically, for practical reasons researchers limited the identification of revisions by counting the number of backspaces or deletions (e.g., Zhu et al., 2019). Although this approach can be used as a raw proxy of revisions—especially in the observation of less experienced writers—it ignores a wide range of revision events that originate from more complex key and mouse combinations. The current tagset makes it possible to study this wide range of revisions in depth. In addition, by providing a full overview of the different properties, including a detailed overview of the annotation guidelines, the tagset allows for informed choices and transparent descriptions of which features are included (and which are not). This enables comparisons among future studies of revision (as described in the Potential Uses of the Tagset for Future Work section).

First, the relatively low IRR for the orientation features, especially for the deep revisions, showed that regardless of the training, orientation is difficult to annotate. This stresses the need for elaborate annotation guidelines, which are currently not always available in annotation studies. In addition, the low IRR for the orientation features shows the difficulty for annotating revisions within the current sample of short texts with limited revisions (as is common in current writing studies). This might indicate that the subcategories of orientation, although they follow from theory, might not be useful in all writing tasks.

Second, the proof of concept provided insight into what types of information may be needed for the manual annotation of revision properties. The low reliability of the deep revision features might also indicate the difficulty to identify these features with screen replays and eye tracking alone. Triangulating the data with direct input from participants (e.g., through concurrent think-aloud or stimulated recall) might be necessary to get a better understanding of the writers' intentions behind these deep revisions. In general, the eye-tracking data added to replays of writing processes proved to be especially useful for making annotation decisions about some specific categories. The annotators indicated that they often used the replay with

eye-tracking data when they were unsure about certain aspects, such as the orientation or sequencing of the revision. For example, an eye fixation on a similar word previously written was a clear indication of a spelling revision (as opposed to a typographic revision). Likewise, an eye fixation on a previously written word could provide information on the intended word. For certain categories eye fixation was not necessary (e.g., evaluation, action, temporal location, and duration), indicating that, depending on which categories are needed, eye fixation does not always need to be extracted.

Third, the proof of concept showed the differences between the two approaches to categorizing spatial location. For spatial location, two approaches can be found in the literature: precontextual versus contextual revisions (Lindgren & Sullivan, 2006a, 2006b) and immediate versus distant revisions (Thorson, 2000). Although the approaches are quite similar, in 9% of the cases an immediate revision was a contextual revision, and in 7% of the cases a distant revision was a precontextual revision. This indicates that these approaches are still quite different and cannot be used interchangeably. We would argue that the “best” approach depends on the research question and the writing task at hand. For example, in tasks where students are required to make an outline up front, the immediate versus distant approach might work better, as otherwise many revisions would be considered contextual once students start to fill out the gaps in the outline. In addition, we included the number of characters away from the leading edge. This allows for a more detailed approach: the leading edge might be updated when the number of characters away from the leading edge stay the same for several revision events. In this way, not only invisible characters (cf. Lindgren et al., 2019) are ignored for leading edge calculations but also visible but perhaps currently “unused” characters, such as a trailing dot or the reference list.

Limitations

Although we did aim for a comprehensive tagset, the tagset does not consist of context-specific features. For example, F. Zhang and Litman (2015) included tags related to genre, such as changes in warrant, reasoning, or backing for argumentative writing. In addition, the tagset is focused on revision in text and does not consist of features specific to visual components, such as images or tables (cf. Leijten et al., 2014). However, we tried to keep the current tagset as context independent as possible. In this way, the tagset can be used in different contexts and, where necessary, be tailored to specific genres, languages, or tasks. For example, changes related to genre could be added to the orientation property, when needed.

A second limitation is that the revision tagset is built for and exemplified with typing (keystroke) data. Hence, it is not clear how the tagset may be applied to handwriting data (cf. Limpo et al., 2014). Conceptually, the revision tagset may still be used for handwriting data, as most of the features of revision hold for handwriting as well. However, some features may be irrelevant for handwriting, such as the number of keystrokes, and some annotation guidelines may need to be adapted, such as the inclusion of “revision for improving legibility” under the category cosmetics/presentation. Future work should determine to what extent all these features can be manually or automatically annotated for handwriting, using automatic tracking of handwriting (e.g., when using HandSpy; Alves et al., 2019).

Another limitation is that the proof-of-concept data set consisted of very few semantic (or deep) revisions. Hence, the IRR for these types of revisions were low, and we could not properly assess whether the annotation guidelines were sufficient in distinguishing between different types of semantic revisions. Even though the annotation guidelines were adapted from previous manual annotations of orientation (Stevenson et al., 2006; Wengelin, 2007), we feel future work should further test the manual annotation of semantic revisions. This could be done for example by applying the tagset to a data set with higher level revisions, such as expert writing, multisession writing, or intervention studies that require writers to respond to (high-level) feedback.

Last, the annotation for the sequencing feature is limited, as the manual and automated features only determined the sequencing in terms of the immediately preceding or following revision, rather than over multiple revisions. More complex analysis involving natural language processing and S-notation representations (Kollberg, 1996; Severinson-Eklundh & Kollberg, 2001) could be used to further investigate connected revision episodes over multiple revisions (Leijten et al., 2019). Alternatively, features from the other properties could be used to identify sequences of revisions. For example, the characters from the start of the writing product (spatial location) could be used to determine whether revisions are made within the same sentence or word.

Potential Uses of the Tagset for Future Work

To further detail how the tagset may be used to study revision in depth in future work, we provide some example research questions and specify how the tagset may be used to answer these. It is important to note here that it is not necessarily essential to include the full tagset in future work. Rather, the tagset is a basis for which the relevant properties, depending on the research

question at hand, should be selected. Therefore, in addition to describing some example research questions we also detail which features from the tagset would be useful for exploring those questions.

First, the tagset could be used to identify how the types of revisions differ across groups of writers (e.g., L1 vs. L2, novice vs. expert; cf. Stevenson et al., 2006); different types of tasks (e.g., various genres, timed vs. nontimed tasks, tasks differing in cognitive load; cf. Révész et al., 2017), or different phases in the writing product (e.g., begin, middle, end; cf. Barkaoui, 2016). One potential research question here could be the following: “Do L2 writers make more (repetitive) form revisions, such as grammar and spelling, compared with L1 writers?” Here, one would need to investigate features on orientation, related to form revisions (B.1.4, B.1.5), as well as features related to sequencing, such as the repetitiveness of the revision (I.3.1, I.3.2). Another research question could be the following: “Do writers make a less thorough and/or structured final review of their writing in a timed writing task, compared with a nontimed writing task?” For this, one would need the temporal location (G.1) to indicate the final part of the writing process, and one may use the orientation (B.1, B.2) as well as the variance in the spatial location (F.8.3, proportionally to the size of the writing product) to indicate the “thoroughness” of the final review and sequencing forward and backward (I.5, I.6) to indicate the “structuredness” of the final review. An example case study on using the current tagset to compare the revisions between L1 and L2 writers can be found in Conijn et al. (2020).

Second, the revision tagset could be used for more process-oriented analysis. Previous work has already indicated that cognitive writing processes (and hence also revision processes) differ over time during writing (Breetvelt et al., 1994). However, revisions are commonly mostly analyzed based on frequencies while the sequential nature is ignored. Temporal analysis, by including spatial (F.8.2, F.8.3) and temporal location (G.1) may provide more detailed information on the process of revision and its effect on the writing product and writing process. For example, they may be used to identify frequent patterns of revisions (cf. Zhu et al., 2019) or to find clusters of similar revisions. In this way, one could further investigate the sequencing of revisions, the causes and the effects of the revisions, and how this differs for different types of revisions. This can eventually be used for a better alignment of keystroke logging with writing processes, as different revisions “behave” differently (Galbraith & Baaijen, 2019).

Third, the current tagset also allows for classification of the manual features using machine learning. With machine learning, the manual features could be approximated automatically, resulting in a more efficient (less time-intensive) tagset. Some studies have already tried to build classification

models for some of the manually annotated features. For example, Xue and Hwa (2014) built a classifier to identify the sequencing of revisions: an indication of whether consecutive revisions belonged to the same mistake. For this, they included the following features: number of words between edits, change of tense, change of word order, change in same word set, edit distance of revised words, original word in dictionary, original and revised word same part-of-speech, and original and revised word are both prepositions. In a related study, Daxenberger and Gurevych (2013) tried to automatically classify the orientation and action features manually annotated in previous work (Daxenberger & Gurevych, 2012). Classifiers were trained to classify the labels for each of the 21 categories in the corpus, using character and word n -grams with $n = 1, 2$, and 3 . Likewise, F. Zhang and Litman (2015) also classified manually annotated orientation labels. As features, they included unigrams, spatial location of the revision (e.g., first sentence), textual properties, such as edit distance, named entity, and discourse marker, and language properties, such as part-of-speech, spelling, and grammar mistakes. All these studies showed that their machine learning algorithms beat the baseline classifier and hence that they could, at least to some extent, automatically classify the labels. However, the classification for fine-grained classes, such as distinguishing between a surface change within a word and a conventional change, were still found to be difficult (F. Zhang & Litman, 2015). The current tagset provides additional features that could be used to further improve on these automatic classification algorithms.

Implications for Educational Practice

The tagset can be used to provide a more detailed overview of revisions, which can be used by both teachers and students in educational practice, including class-based writing instruction as well as one-on-one writing tutoring sessions. This could be done by, for example, by visualizing the revision properties in the form of a dashboard. Dashboards are tools that provide an overview of students' tracked learning activities to promote awareness and reflection (Verbert et al., 2014). Accordingly, such a dashboard can be used by students to reflect on their revision process or by teachers to inform their feedback and instruction on the revision process. A small case study already showed that such visualizations on students' writing processes can be useful to discuss and reflect on the writing process with students (see, e.g., Vandermeulen et al., 2020). In addition, these visualizations may be used to adjust students' (implicit) conceptions about revision and revision processes, which are argued to differ between teacher and student, and novice and expert writers (Flower et al., 1986). Last, the dashboard can be implemented as an intervention targeted at improving revision strategies. The tagset can then in

turn be used for a more detailed evaluation of the impact of the intervention on the types and timing of the revision.

Conclusion

To conclude, in this article, we provided a comprehensive product- and process-oriented tagset of revisions, including nine properties of external revisions: trigger, orientation, evaluation, action, linguistic domain, spatial location, temporal location, duration, and sequencing. We have shown that many variables related to these properties can be automatically extracted using information from the keystroke log. In addition, replays of the keystroke data combined with eye tracking could be used for manual annotation of the features that cannot be extracted automatically. These manual features could eventually be classified using machine learning techniques, to provide a fully automated tool to provide detailed information on revisions. This tagset forms a scalable basis for researchers to study revision in depth. Moreover, the tagset can be used by students and teachers to reflect and improve upon students' revision processes.

Declaration of Conflicting Interests

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship and/or publication of this article: This work was supported, in part, by Iowa State University's college of Liberal Arts and Sciences (Signature Research Initiative) and EARLI Emerging Field Group (EarlyWritePro).

Supplemental Material

Supplemental material for this article is available online.

References

- Allal, L. (2000). Metacognitive regulation of writing in the classroom. In A. Camps & M. Milian (Eds.), *Metalinguistic activity in learning to write* (Vol. 6, pp. 145-166). Amsterdam University Press.
- Allal, L., Chanquoy, L., & Largy, P. (Eds.). (2004). *Revision cognitive and instructional processes* (Vol. 13). Springer.
- Alves, R. A., Leal, J. P., & Limpo, T. (2019). Using HandSpy to study writing in real time: A comparison between low-and high-quality texts in grade 2. In E. Lindgren & K. Sullivan (Eds.), *Observing writing: Insights from keystroke logging and handwriting* (pp. 50-70). Brill.

- Baaijen, V. M., & Galbraith, D. (2018). Discovery through writing: Relationships with writing processes and text quality. *Cognition and Instruction, 36*(3), 1-25. <https://doi.org/10.1080/07370008.2018.1456431>
- Baaijen, V. M., Galbraith, D., & de Glopper, K. (2012). Keystroke analysis: Reflections on procedures and measures. *Written Communication, 29*(3), 246-277. <https://doi.org/10.1177/0741088312451108>
- Barkaoui, K. (2016). What and when second-language learners revise when responding to timed writing tasks on the computer: The roles of task type, second language proficiency, and keyboarding skills. *Modern Language Journal, 100*(1), 320-340. <https://doi.org/10.1111/modl.12316>
- Bird, S., Loper, E., & Klein, E. (2009). *Natural language processing with Python*. O'Reilly Media Inc.
- Bowen, N., & Van Waes, L. (2020). Exploring revisions in academic text: Closing the gap between process and product approaches in digital writing. *Written Communication, 37*(3), 322-364. <https://doi.org/10.1177/0741088320916508>
- Boytsov, L. (2011). Indexing methods for approximate dictionary searching: Comparative analysis. *Journal of Experimental Algorithmics, 16*, 1-91. <https://doi.org/10.1145/1963190.1963191>
- Breetvelt, I., Van den Bergh, H., & Rijlaarsdam, G. (1994). Relations between writing processes and text quality: When and how? *Cognition and Instruction, 12*(2), 103-123. https://doi.org/10.1207/s1532690xci1202_2
- Chanquoy, L. (1997). Thinking skills and composing: Examples of text revision. In J. H. M. Hamers & M. Oortoom (Eds.), *Inventory of European programmes for teaching thinking* (pp. 179-185). Sardes.
- Chukharev-Hudilainen, E. (2019). Empowering automated writing evaluation with keystroke logging. In E. Lindgren & K. Sullivan (Eds.), *Studies in writing: Vol. 38. Observing writing* (pp. 125-142). Brill. https://doi.org/10.1163/9789004392526_007
- Chukharev-Hudilainen, E., Saricaoglu, A., Torrance, M., & Feng, H.-H. (2019). Combined deployable keystroke logging and eyetracking for investigating L2 writing fluency. *Studies in Second Language Acquisition, 41*(3), 583-604. <https://doi.org/10.1017/S027226311900007X>
- Conijn, R., Dux Speltz, E., van Zaanen, M., Van Waes, L., & Chukharev-Hudilainen, E. (2020). A process-oriented dataset of revisions during writing. In *Proceedings of the 12th Language Resources and Evaluation Conference* (pp. 363-368). European Language Resources Association. <https://aclanthology.org/2020.lrec-1.45>
- Crawford, L., Lloyd, S., & Knott, K. (2008). Analysis of student revisions on a state writing test. *Assessment for Effective Intervention, 33*(2), 108-119. <https://doi.org/10.1177/1534508407311403>
- Damerau, F. J. (1964). A technique for computer detection and correction of spelling errors. *Communications of the ACM, 7*(3), 171-176. <https://doi.org/10.1145/363958.363994>
- Daxenberger, J., & Gurevych, I. (2012). A corpus-based study of edit categories in featured and non-featured Wikipedia articles. In *Proceedings of COLING 2012*

- (pp. 711-726). COLING 2012 Organizing Committee. <https://aclanthology.org/C12-1044>
- Daxenberger, J., & Gurevych, I. (2013). Automatically classifying edit categories in Wikipedia revisions. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing* (pp. 578-589). Association for Computational Linguistics. <https://aclanthology.org/D13-1055.pdf>
- Faigley, L., & Witte, S. (1981). Analyzing revision. *College Composition and Communication*, 32(4), 400-414. <https://doi.org/10.2307/356602>
- Figueredo, L., & Varnhagen, C. K. (2006). Spelling and grammar checkers: Are they intrusive? *British Journal of Educational Technology*, 37(5), 721-732. <https://doi.org/10.1111/j.1467-8535.2006.00562.x>
- Fitzgerald, J. (1987). Research on revision in writing. *Review of Educational Research*, 57(4), 481-506. <https://doi.org/10.3102/00346543057004481>
- Flower, L., & Hayes, J. R. (1981). A cognitive process theory of writing. *College Composition and Communication*, 32(4), 365-387. <https://doi.org/10.2307/356600>
- Flower, L., Hayes, J. R., Carey, L., Schriver, K., & Stratman, J. (1986). Detection, diagnosis, and the strategies of revision. *College Composition and Communication*, 37(1), 16-55. <https://doi.org/10.2307/357381>
- Galbraith, D., & Baaijen, V. M. (2019). Aligning keystrokes with cognitive processes in writing. In E. Lindgren & K. Sullivan (Eds.), *Observing writing* (Vol. 38, pp. 306-325). Brill.
- Hayes, J. R., & Flower, L. (1980). Identifying the organization of writing processes. In L. W. Gregg & E. R. Steinberg (Eds.), *Cognitive processes in writing* (pp. 3-30). Lawrence Erlbaum.
- Hayes, J. R., Flower, L., Schriver, K., Stratman, J. F., & Carey, L. (1987). Cognitive processes in revision. In S. Rosenberg (Ed.), *Advances in applied psycholinguistics: Reading, writing, and language learning* (Vol. 2, pp. 176-240). Cambridge University Press.
- Kollberg, P. (1996, October 23-25). *S-notation as a tool for analysing the episodic structure of revisions* [Paper presentation]. European Writing Conferences. <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.13.4542>
- Krippendorff, K. (2004). *Content analysis: An introduction to its methodology* (2nd ed.). Sage.
- Krippendorff, K. (2011). *Computing Krippendorff's alpha-reliability* [Working paper]. https://repository.upenn.edu/asc_papers/43/
- Leijten, M., Van Horenbeeck, E., & Van Waes, L. (2019). Analysing keystroke logging data from a linguistic perspective. In E. Lindgren & K. Sullivan (Eds.), *Observing writing* (Vol. 38, pp. 71-95). Brill. https://doi.org/10.1163/9789004392526_005
- Leijten, M., Van Waes, L., Schriver, K., & Hayes, J. R. (2014). Writing in the workplace: Constructing documents using multiple digital sources. *Journal of Writing Research*, 5(3), 285-337. <https://doi.org/10.17239/jowr-2014.05.03.3>
- Limpo, T., Alves, R. A., & Fidalgo, R. (2014). Children's high-level writing skills: Development of planning and revising and their contribution to writing quality. *British Journal of Educational Psychology*, 84(2), 177-193. <https://doi.org/10.1111/bjep.12020>

- Lindgren, E., & Sullivan, K. P. (2006a). Analysing online revision. In K. P. Sullivan & E. Lindgren (Eds.), *Computer keystroke logging and writing: Methods and applications (studies in writing)* (pp. 157-188). Elsevier.
- Lindgren, E., & Sullivan, K. P. (2006b). Writing and the analysis of revision: An overview. In K. Sullivan & E. Lindgren (Eds.), *Computer keystroke logging and writing: Methods and applications (studies in writing)* (pp. 31-40). Elsevier.
- Lindgren, E., Westum, A., Outakoski, H., & Sullivan, K. P. H. (2019). Revising at the leading edge: Shaping ideas or clearing up noise. In E. Lindgren & K. Sullivan (Eds.), *Studies in writing: Vol. 38. Observing writing* (pp. 346-365). Brill. https://doi.org/10.1163/9789004392526_017
- Min, H.-T. (2006). The effects of trained peer review on EFL students' revision types and writing quality. *Journal of Second Language Writing, 15*(2), 118-141. <https://doi.org/10.1016/j.jslw.2006.01.003>
- Monahan, B. D. (1984). Revision strategies of basic and competent writers as they write for different audiences. *Research in the Teaching of English, 18*(3), 288-304.
- Murray, D. M. (1978). Internal revision: A process of discovery. In C. R. Cooper & L. Odell (Eds.), *Research on composing: Points of departure* (pp. 85-103). National Council of Teachers of English.
- Ranalli, J., Feng, H.-H., & Chukharev-Hudilainen, E. (2018). The affordances of process-tracing technologies for supporting L2 writing instruction. *Language Learning & Technology, 23*(2), 1-11. <https://doi.org/10.125/44678>
- Révész, A., Kourтали, N.-E., & Mazgutova, D. (2017). Effects of task complexity on L2 writing behaviors and linguistic complexity. *Language Learning, 67*(1), 208-241. <https://doi.org/10.1111/lang.12205>
- Scardamalia, M., & Bereiter, C. (1983). The development of evaluative, diagnostic and remedial capabilities in children's composing. In M. Martlew (Ed.), *The psychology of written language: Developmental and educational perspectives* (pp. 67-95). Wiley.
- Sengupta, S. (2000). An investigation into the effects of revision strategy instruction on L2 secondary school learners. *System, 28*(1), 97-113. [https://doi.org/10.1016/S0346-251X\(99\)00063-9](https://doi.org/10.1016/S0346-251X(99)00063-9)
- Severinson-Eklundh, K., & Kollberg, P. (2001). Studying writers' revision patterns with S-notation analysis. In T. Olive & C. M. Levy (Eds.), *Contemporary tools and techniques for studying writing* (Vol. 10, pp. 89-104). Springer. https://doi.org/10.1007/978-94-010-0468-8_5
- Sommers, N. (1980). Revision strategies of student writers and experienced adult writers. *College Composition and Communication, 31*(4), 378-388. <https://doi.org/10.2307/356588>
- Stevenson, M., Schoonen, R., & de Glopper, K. (2006). Revising in two languages: A multi-dimensional comparison of online writing revisions in L1 and FL. *Journal of Second Language Writing, 15*(3), 201-233. <https://doi.org/10.1016/j.jslw.2006.06.002>

- Thorson, H. (2000). Using the computer to compare foreign and native language writing processes: A statistical and case study approach. *Modern Language Journal*, 84(2), 155-170. <https://doi.org/10.1111/0026-7902.00059>
- Tillema, M., van den Bergh, H., Rijlaarsdam, G., & Sanders, T. (2011). Relating self reports of writing behaviour and online task execution using a temporal model. *Metacognition and Learning*, 6(3), 229-253. <https://doi.org/10.1007/s11409-011-9072-x>
- Van Waes, L., & Schellens, P. J. (2003). Writing profiles: The effect of the writing mode on pausing and revision patterns of experienced writers. *Journal of Pragmatics*, 35(6), 829-853. [https://doi.org/10.1016/S0378-2166\(02\)00121-2](https://doi.org/10.1016/S0378-2166(02)00121-2)
- Vandermeulen, N., Leijten, M., & Van Waes, L. (2020). Reporting writing process feedback in the classroom: Using keystroke logging data to reflect on writing processes. *Journal of Writing Research*, 12(1), 109-140. <https://doi.org/10.17239/jowr-2020.12.01.05>
- Verbert, K., Govaerts, S., Duval, E., Santos, J. L., Assche, F., Parra, G., & Klerkx, J. (2014). Learning dashboards: An overview and future research opportunities. *Personal and Ubiquitous Computing*, 18(6), 1499-1514. <https://doi.org/10.1007/s00779-013-0751-2>
- Wengelin, Å. (2007). The word-level focus in text production by adults with reading and writing difficulties. In M. Torrance, L. Van Waes, & D. Galbraith (Eds.), *Studies in writing: Vol. 20. Writing and cognition: Research and applications* (pp. 67-82). Brill. https://doi.org/10.1163/9781849508223_006
- Williamson, M. M., & Pence, P. (1989). Word processing and student writers. In B. K. Britton & S. M. Glynn (Eds.), *Computer writing environments: Theory, research, and design* (pp. 93-127). Lawrence Erlbaum.
- Winder, R., Kathpalia, S. S., & Koo, S. L. (2016). Writing centre tutoring sessions: Addressing students' concerns. *Educational Studies*, 42(4), 323-339. <https://doi.org/10.1080/03055698.2016.1193476>
- Wobbrock, J. O., & Myers, B. A. (2006). Analyzing the input stream for character-level errors in unconstrained text entry evaluations. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 13(4), 458-489. <https://doi.org/10.1145/1188816.1188819>
- Xu, C. (2018). Understanding online revisions in L2 writing: A computer keystroke-log perspective. *System*, 78(November), 104-114. <https://doi.org/10.1016/j.system.2018.08.007>
- Xue, H., & Hwa, R. (2014, June 23-25). Improved correction detection in revised ESL sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics* (Vol. 2, pp. 599-604). Association for Computational Linguistics. <https://doi.org/10.3115/v1/P14-2098>
- Zhang, F., Hwa, R., Litman, D., & Hashemi, H. B. (2016, June 12-17). ArgRewrite: A web-based revision assistant for argumentative writings. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations* (pp. 37-41). Association for Computational Linguistics. <https://doi.org/10.18653/v1/N16-3008>

- Zhang, F., & Litman, D. (2015). Annotation and classification of argumentative writing revisions. In *Proceedings of the Tenth Workshop on Innovative Use of NLP for Building Educational Applications* (pp. 133-143). Association for Computational Linguistics. <https://doi.org/10.3115/v1/W15-0616>
- Zhang, M., Hao, J., Li, C., & Deane, P. (2016). Classification of writing patterns using keystroke logs. In L. A. van der Ark, D. M. Bolt, W.-C. Wang, J. A. Douglas, & M. Wiberg (Eds.), *Quantitative psychology research: The 80th annual meeting of the psychometric society*, Beijing, 2015 (pp. 299-314). Springer. https://doi.org/10.1007/978-3-319-38759-8_23
- Zhu, M., Zhang, M., & Deane, P. (2019). Analysis of keystroke sequences in writing logs. *ETS Research Report Series*, 2019(1), 1-16. <https://doi.org/10.1002/ets2.12247>

Author Biographies

Rianne Conijn is an assistant professor in the Human-Technology Interaction group at Eindhoven University of Technology, Netherlands. Her research interests include the analysis and interpretation of (sequences of) online and learning and writing behavior to improve learning and teaching.

Emily Dux Speltz is a PhD student in the Applied Linguistics and Technology program at Iowa State University. Her research focuses on computer-assisted writing interventions and the investigation of writing processes.

Menno van Zaanen is a professor in Digital Humanities. He works at the South African Centre for Digital Language Resources (SADiLaR), which is hosted at North-West University in Potchefstroom, South Africa. His research focuses on the development and application of (mostly language independent) tools that deal with language processing.

Luuk Van Waes is a professor in Professional Communication at the University of Antwerp. He has been involved in different types of writing research, mainly focusing on using keystroke logging to observe (professional) writing processes.

Evgeny Chukharev-Hudilainen is an associate professor in the Applied Linguistics and Technology program at Iowa State University. His research program investigates cognitive processes underlying language acquisition and use, currently focusing on written language production.