Effort is All You Need: The Possibilities of Writing Analytics

Raymond Oenbring, *University of the Bahamas*

Scopus Abstract

This article describes the browser-based word processor and learning management system the Forge, a faculty-led open source project hosted by LibreTexts, an open education nonprofit. A guiding principle of the Forge system is that writing analytics, particularly those focused on student time and effort when composing, offer a potential solution to the challenges to academic integrity in the generative AI era. In addition to incentivizing students to honor the writing process by engaging in effort during the drafting and revision stages, the Forge provides instructors with previously unavailable insights into how students write assignments in real contexts, both individually and for classes as a whole. The Forge includes numerous different visualizations of student writing processes, including EssayMaps that present the editing history of the document; HeatMaps, which show the hours and days the student spent on the assignment; and WordChange graphs, which show the change in word count of the document every minute while the student is writing, highlighting spikes in the word count, and writing composed quickly. The system also provides advisory *effort* and *revision* scores to advise the instructor whether a deeper investigation of the assignment's analytics may be warranted. The Forge can be used for a variety of research studies looking at how students write assignments in real contexts, including developing profiles of student writing processes, assessing the effects of various pedagogical strategies and interventions, and understanding the ways students use generative AI writing tools.

Structured Abstract

- **Identification of Innovation:** This article describes the browser-based word processor and learning management system the Forge, a faculty-led open source project hosted by LibreTexts, an open education nonprofit. A guiding principle of the Forge system is that writing analytics, particularly those focused on student time and effort when composing, offer a potential solution to the challenges to academic integrity in the generative AI era. In addition to incentivizing students to honor the writing process by engaging in effort during the drafting and revision stages, the Forge provides instructors with previously unavailable insights into how students write assignments in real contexts, both individually and for classes as a whole.

- **Exposition of Innovation:** The Forge includes numerous different visualizations of student writing processes, including EssayMaps that present the editing history of the document; HeatMaps, which show the hours and days the student spent on the assignment; and WordChange graphs, which show the change in word count of the document every minute while the student is writing, highlighting spikes in the word count, and writing composed quickly. The system also provides advisory *effort* and *revision* scores to advise the instructor whether a deeper investigation of the assignment's analytics may be warranted.
- **Application of Innovation:** In its present form, the Forge can be used for a variety of research studies looking at how students write assignments in real contexts, including developing profiles of student writing processes, understanding the role that time-coded phenomena such as draft deadlines play in the composition process, understanding the differing strategies and processes students use to compose different types of assignments, assessing the effects of various pedagogical strategies and interventions, and understanding the ways students use generative AI writing tools.
- **Directions for Further Research:** The possibilities of writing analytics for writing instruction are still largely untapped by technologists and scholars. When applied across writing programs, tools like the Forge could provide unprecedented amounts of rich data about students' writing and writing processes for instructional, institutional, and research purposes. The development team is continuing to expand the platform's capabilities, including potentially and eventually adding optional functionalities that use limited AI to provide students with dynamic "teacherly" feedback on their writing.

Keywords*: academic integrity, Forge open source project, writing analytics, writing process*

1.0 Identification of Innovation

While compositionists have not traditionally been at the forefront of calls for basing student and institutional assessments on rich empirical data, a small but significant number of composition scholars are recognizing the importance of leveraging instructional and institutional data to help instructors and writing program administrators have a voice in institutional discussions (see, for example, Lang & Baehr, 2012; Moxley, 2013; Miller & Licastro, 2021; Warner, 2022). Arguing for data-intensive empirical studies of writing programs, Lang and Baehr (2012), for example, have contended that "writing program administrators, faced with increasing demands for accountability and assessment, as well as widely varying student populations, need to have ways of understanding the interactions of students, faculty, and administrators in their present program, both in the short term and longitudinally" (p. 173). Putting the divide between data-focused university administrators and data-hesitant compositionists in sharper relief, Miller and Licasto (2021) have suggested pithily that "if we do not use data, we may well be used *by* data" (p. 4). Typical studies under this writing analytics paradigm have relied upon common or bespoke learning management systems (e.g., Moxley, 2013; Kaufer & Ishizaki, 2023) and/or used corpus linguistic software packages to identify patterns in submitted completed/static student writing (Aull, 2015).

Further, despite a limited tradition of "writing teachers writing software" (see, for example, Miller, 2021, p. 159; LeBlanc, 1993), compositionists have traditionally not been at the forefront of developing writing technologies, leaving those developments to technologists whose priorities align more with business models than best practices in student learning—or even academic integrity. The most obvious manifestation of technologists pursuing their own priorities at the potential expense of student learning is, of course, generative AI writing tools, a broad heading including both systems that can generate custom writing (such as ChatGPT) and paraphrasing tools (e.g., QuillBot). The release of ChatGPT in late 2022 drew various responses, ranging from assertions that "the college essay is dead" (Marche, 2022) to others insisting upon ignoring the hype bubble. Regardless the effects of generative AI on student learning, the current social narrative regarding writing is important. Indeed, if Western society decides in the coming months that there's no point in having students write anything once they have mastered the alphabet and pencil-and-paper writing because AI can do the rest, that can't be good for students—or university writing instructors, right?

One obvious approach to the challenges to academic integrity presented by generative AI writing tools is simply to have students write all their work with pen and paper in class—and this is an approach that I, along with a number of other compositionists, have experimented with. Nonetheless, this can seem like a 19th-century solution to a 21st-century problem. In fact, one of my colleagues refers to this as the "*Little House on the Prairie"* approach. The "all in-class writing" approach can feel, however, like a betrayal of process philosophy—even though one can still emulate some elements of a process approach by having students write rough and final drafts in subsequent classes. Simply put, in a digital age, instructors need tools to: A) verify that the digital texts students present as their own work represent the skills students have actually acquired in class, and B) understand how students actually write in digital contexts. (Indeed, I can say personally as a composition instructor that my hesitance toward AI-powered writing tools is not that they "assist" students but that they allow students to avoid the effort that is key to intellectual growth.)

Most previous technologies developed specifically to respond to AI writing in instructional contexts have taken a "detection" approach, looking for the fingerprints of AI-produced syntax in writing; that is, these systems, like the static analytics approach, have only looked at submitted student writing. The problems with applications using this approach are many, the first being that these algorithms can easily be gamed by students engaging in a very modest amount of manual paraphrasing or by using paraphrasing applications (such as QuillBot). What's more, AI detection algorithms are prone to both false negatives (missing AI writing) and, more troublingly, false positives—labeling genuine student writing as AI writing (Dalalah & Dalalah, 2023). Perhaps more importantly, the algorithm fingerprint approach is unable to provide evidence other than black-boxed algorithmic hunches to support claims that the student is sidestepping the writing process.

Beyond the detection approach, process tracking, as advocated most notably by Mills (2024), offers an alternative. Here, the use of Google Docs (or similar applications) with revision history, with or without browser plugins to extend functionality, presents a way to track student activity while writing. While a step in the right direction, this approach and the technologies used for it are fundamentally limited in that they: A) do not take full advantage of the previously unavailable fountain of data provided about students' writing processes; and B) are still somewhat burdensome for instructors. Indeed, who wants to have to cycle through the revision history of every submitted student essay?

This article argues that writing analytics, particularly those focused on student time and effort when composing, offer a potential solution to challenges to the integrity of student writing processes in the generative AI era. While basing assessments around related notions like student effort, participation, and/or labor (e.g., Inoue, 2019) is lauded as more equitable than a purely product-quality focus, such notions have traditionally been vague and difficult to measure objectively (see, for example, Omizo, 2019). I contend, however, that detailed writing analytics provide a potential strategy to systematically measure and assess effort, and thereby make accurately determining genuine writing more feasible. (This focus on the importance of effort during the student writing process, while widely accepted by compositionists, is not, of course, universal; many people, including many faculty in other academic fields, view writing as a chore, a mere epiphenomenon secondary to thought, and thus are more amenable to outsourcing the writing process to AI.) More than just upholding academic integrity, writing analytics platforms would, when applied across classes and programs, provide large amounts of previously unavailable data, which could be used by instructors, writing researchers, and academic programs for insights into student learning and engagement. Indeed, in this article, I argue for the potential of a comprehensive writing analytics platform—bringing together student composition, instructor feedback and assessment, and learning analytics in the form of an open source project—to support genuine student effort and learning in the generative AI era and to provide previously unavailable insights into how students write in real contexts. Such a platform could be understood as a manifestation of the growing "open pedagogy" movement (see, Clinton-Lissel, 2021), a movement that, among other concerns, emphasizes the importance of transparency in the learning process.

2.0 Exposition of Innovation

The Forge is an open source project that has been iteratively developed by a research team consisting of a developer and academics at three US institutions and one Caribbean institution. Since December 2024, the Forge has been available for use through LibreTexts.org, one of the world's largest open educational resource websites, a faculty-led nonprofit supported by grants from the state of California and other major grantors. Built around Collabora, an online open source, browser-based word processor (i.e., the cloud version of the LibreOffice office suite), the Forge system is available as an easily installable Docker image. The entire Forge/Collabora system can be self-hosted by institutions in the interest of protecting student data and controlling costs. While the Forge development team initially, and understandably, focused on word processor technologies, the Forge now can provide similar time metrics for spreadsheet and presentation software, as well, allowing students to complete assignments using the Collabora/LibreOffice Calc and Impress applications. In the coming pages, I describe the major features of the Forge, particularly on the instructor side, including the Rewind feature, the EssayMap feature, and the WordChange graph.

Figure 1

Student Writing Interface

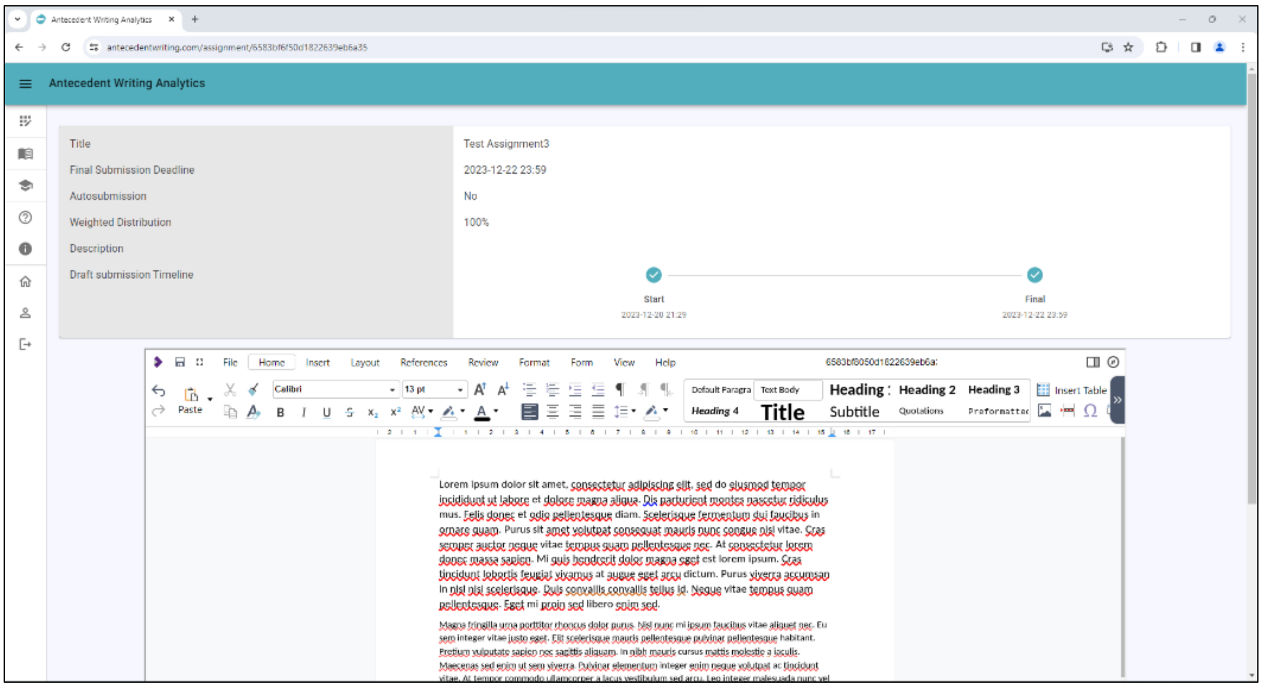

At its core, the Forge is a shared writing space between instructors and students that provides instructors with insights intoon students' writing processes as well as the ability to comment on and guide the writing process. Instructors create assignments, drafts, and deadlines. Students compose their work in a browser-based word processor. Figure 1 demonstrates the student writing interface. The student writing interface is built around the open -source Collabora word processor, having all of the major functionalities of commercial browser-based word processors, including: commenting, multiple concurrent writers, and track-changes functionality. Instructors can leave marginal feedback on drafts both as marginal (in the Collabora word processor) and share and as global comments.

Figure 2 shows the instructor dashboard. The submission timeline at the top shows time-coded information about the various drafts and student submissions. The top rail indicates the various stages and drafts of the assignment. As time passes, the lavender running icon, which when clicked brings up the current *ongoing* version of the text, moves along the *drafts rail*. The middle rail indicates the various submissions that the student makes for the assignment. When the instructor clicks on the arrows indicating the submission, the system brings up that submitted version of the document in read-only mode. The bottom rail is the c*omments rail*, which shows the various commented-upon versions of the document, from the instructor and from peer review.

Figure 2

Instructor Dashboard

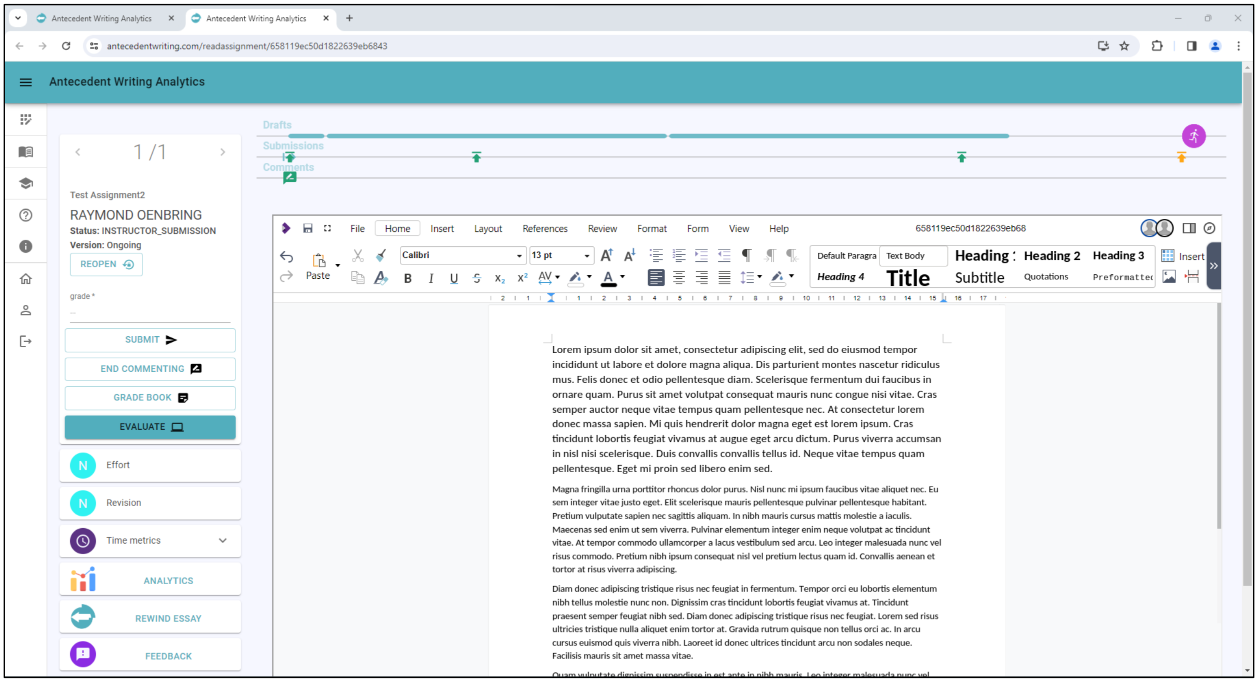

The Rewind feature shows an animation of the entire editing history of the document (see Figure 3), providing snapshots of the document state every 30 seconds while the student is writing. New language is highlighted in green, edited language is highlighted in orange, and deleted language is highlighted in red. Spikes or dips in the line graph above the progress line indicate significant sections of pasted-in or deleted text. Further, instructors can cue up the Rewind to investigate what happened at that moment.

To indicate to the instructor if a deeper dive into the document's editing history may be warranted, the system uses machine learning algorithms to provide *effort* and *revision* scores for each assignment on an A to F scale. (A more detailed description of the *effort* and *revision* algorithms is provided in a different manuscript, now under review at a different journal, and outside the scope of this piece, but broadly, we can say that the effort and revision scores consider all of the following: time on task; the types of edits [various tags include *typo correction, deletion,* and *in-text citation*]; and "fragmentation," —the extent to which the edits are nonlinear.) While these scores are meant to be advisory to the instructor and should not be seen as binding, the instructor may wish to include *effort* and *revision*, broadly understood, as weighted factors in the student's overall grade for the assignment to incentivize students to take full advantage of the writing and learning process. (While the Forge development team has not systematically or empirically investigated instructors' feelings about how using the Forge as an instructional tool affects instructor workloads, I can say personally that I have been using the Forge for multiple semesters, and I do not believe that it adds considerably to my workload. The more advanced analytics I outline in the coming pages should be understood as analytics available for instructors to investigate and navigate only if they desire.)

Figure 3

Rewind Feature

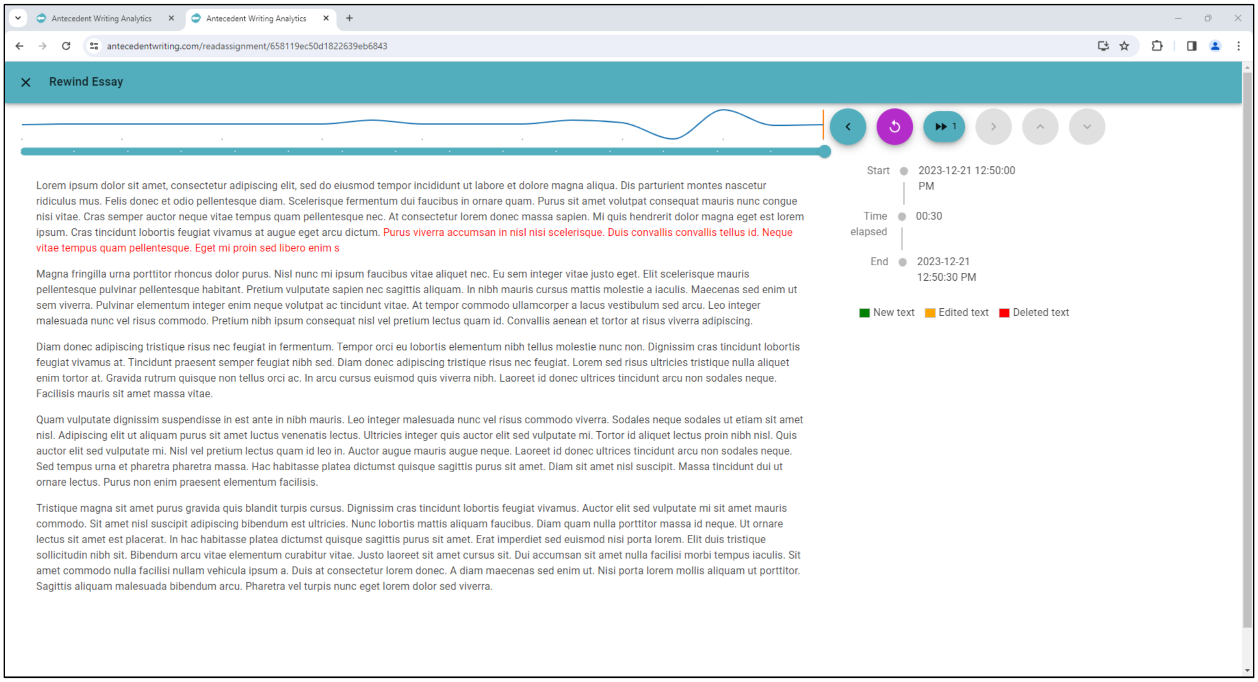

Figure 4 represents how the EssayMap feature appears on the instructor analytics dashboard. In EssayMaps, time is represented on the x-axis, and the physical location in the text of the edit is represented on the y-axis, with the y-axis flipped from a conventional graph so that the origin is at the top and further on in the text is at the bottom, like a written text. The left graph in Figure 4 represents the writing process of the document showing only the changes to the text, and the right graph includes time considerations, but does not include long gaps. Green circles represent paste- ins of text and red circles represent deletions of text, the larger the circle the larger the amount of text (see Figure 5). Figures 4 and 5 show relatively high-effort, high-revision assignments. Conversely, Figure 6 demonstrates an example EssayMap where for which a significant amount of the language in the text was pasted in rather than composed in the Forge word processor. If a student merely types in language that was previously composed, such as language from an AI-website, or the student writes the assignment with minimal effort, the EssayMap will appear as a mostly continuous diagonal line from the upper -left of the graph to the bottom right. (Note that the instructors have the option to make individual assignment-level visualizations—such as the EssayMaps, HeatMaps, and WordChange graphs—visible to the students in order close the loop on process transparency.)

Figure 4

Instructor EssayMap Feature, Showing a High-Effort Assignment

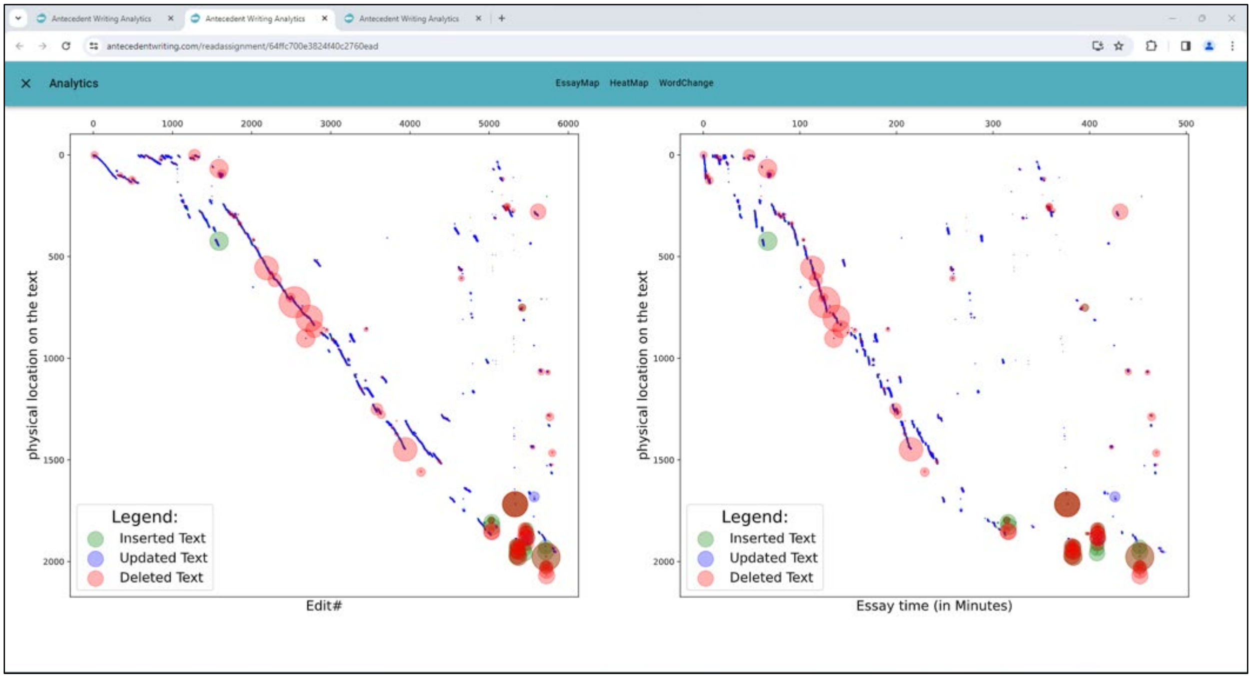

Figure 5

Explanation of EssayMap

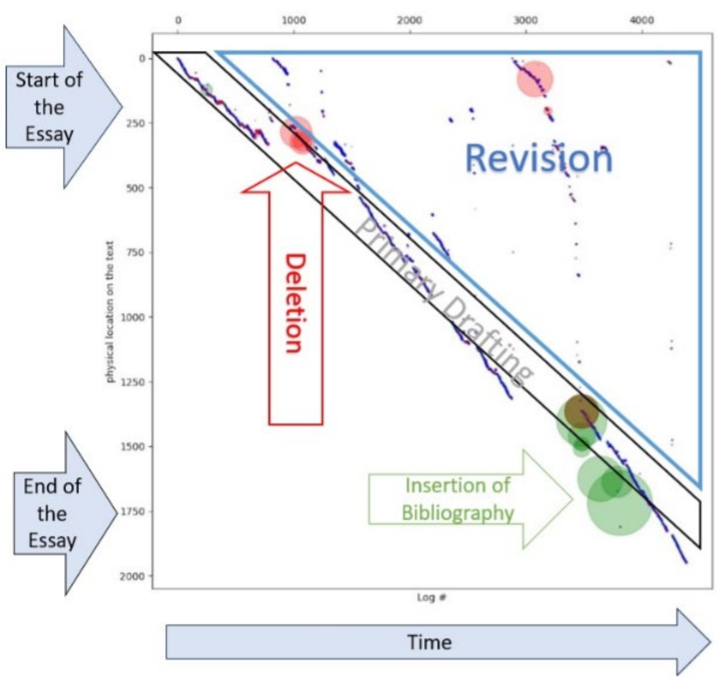

Figure 6

EssayMap of a Low-Effort Essay with Large Sections of Pasted-in Text

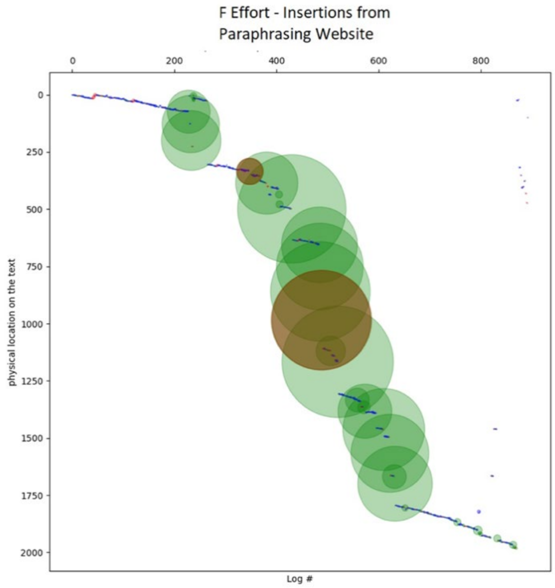

For the assignment represented by the HeatMap in Figure 7, the student spent multiple hours over multiple days on the assignment, which is a sign that the document was composed with some effort and care. Figure 7 shows the time periods when the student worked on the assignment. The columns represent the hours of the day, and the rows represent the days.

Figure 7

HeatMap Visualization of Student Activity While Writing the Assignment

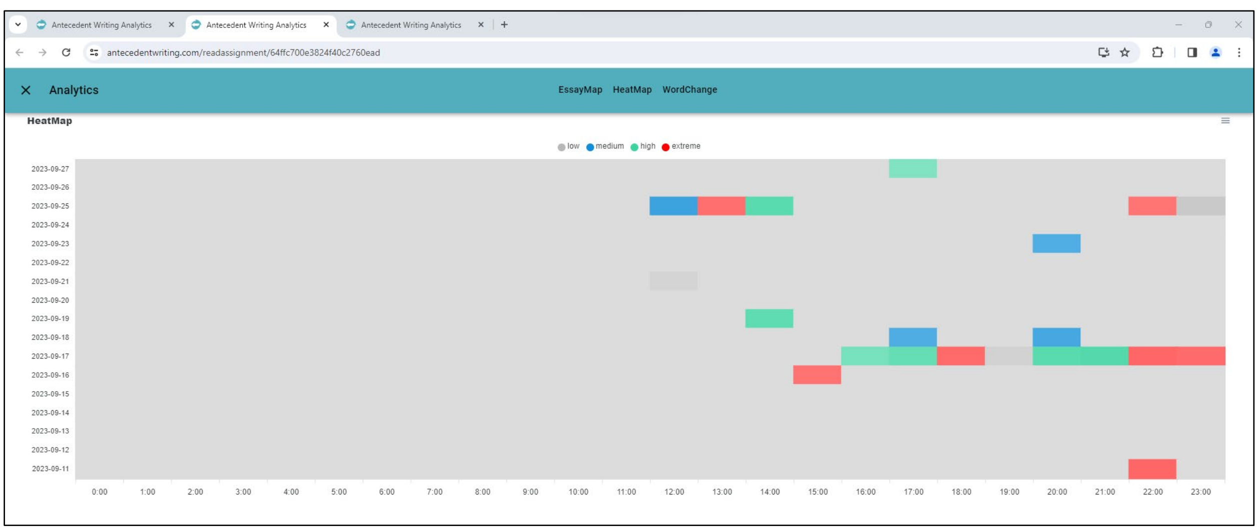

Figures 8 and 9 demonstrate the WordChange graph, which shows the change in word count for the document every minute the student was actively writing. The WordChange graph presented in Figure 8 is for a high-effort assignment, with the graph remaining mostly in the slowly composed yellow zone, and the student spending over eight hours in total writing the assignment. Dips into the negative/white zone indicate deletions of text. When dips immediately precede or follow similarly sized spikes as are seen on the right side of Figure 8, those indicate copy-paste-and-deletes or cut-and-pastes within the document. Conversely, Figure 9 is for a low-effort assignment, including sustained periods in the pink, more rapidly composed, range.

Figure 8

WordChange Graph for a High-Effort Assignment

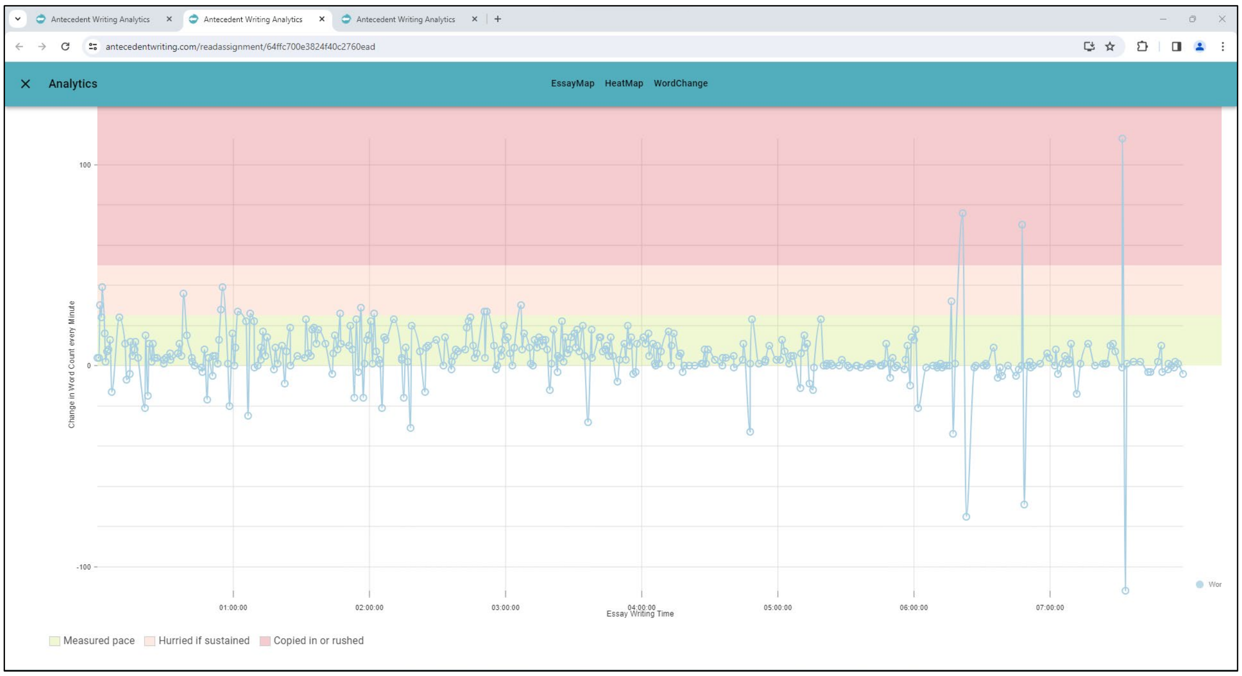

Figure 9

WordChange Graph for a Low-Effort Assignment

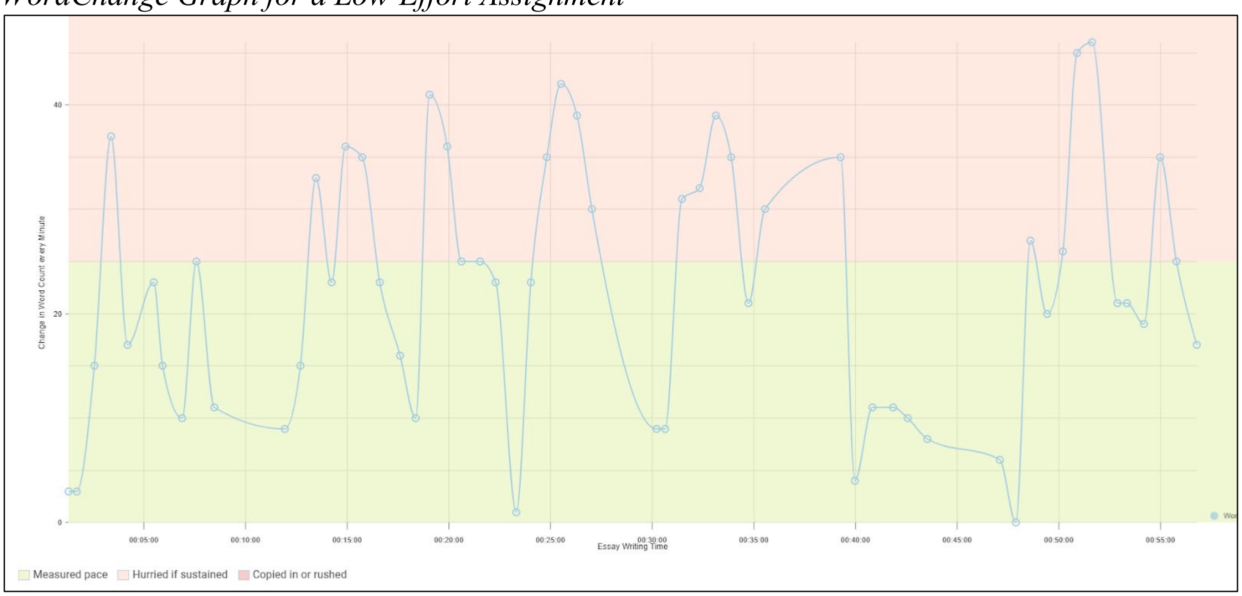

Figure 10 presents a class-level visualization of all student activity on various assignments in the class. The individual points on the graph represent words changed by the entire class over the course of eight-hour periods, with different color lines representing different assignments. This particular visualization, especially in combination with hour-level visualizations, is effective in demonstrating time-coded student activity by all students throughout the semester. For instance, spikes in activity commonly are associated with assignment deadlines. Broadly, instructors can use visualizations like this to understand the degree to which students extend their writing activity over multiple days and writing sessions. While all assignments will likely have spikes associated with their deadlines, instructors can use this metric to assess effectiveness of the strategies they use to encourage students to work on their texts in a more sustained manner. That is, flatter graphs are, broadly, better.

Figure 10

Class-Level Visualization of Number of Words Changed in Assignments Every Eight-Hour Period

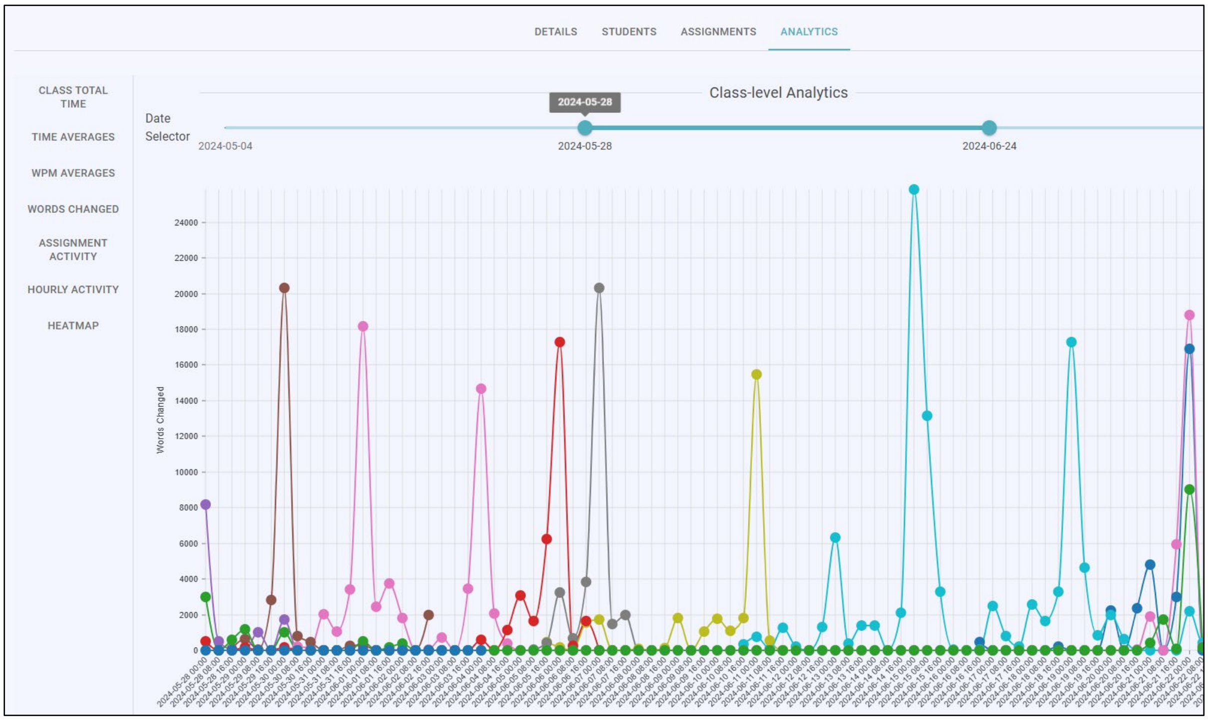

Figure 11 presents a class-level words-per-minute average graph, using boxplots for each assignment in the class. This particular visualization shows the distribution of student mean speeds of writing, allowing comparison across different assignments. The bottom whisker of each boxplot represents the 10th percentile in the class, and the top whisker of each boxplot represents the 90th percentile. The top and bottom of each box represent the 25th and 75th percentile, respectively. The centerline of the boxplot represents the median words per minute for the assignment. Broadly, assignments with a lower words-per-minute count may be composed with more attention and care.

Figure 11

Boxplot Diagrams Showing the Distribution of Words per Minute for Each Assignment

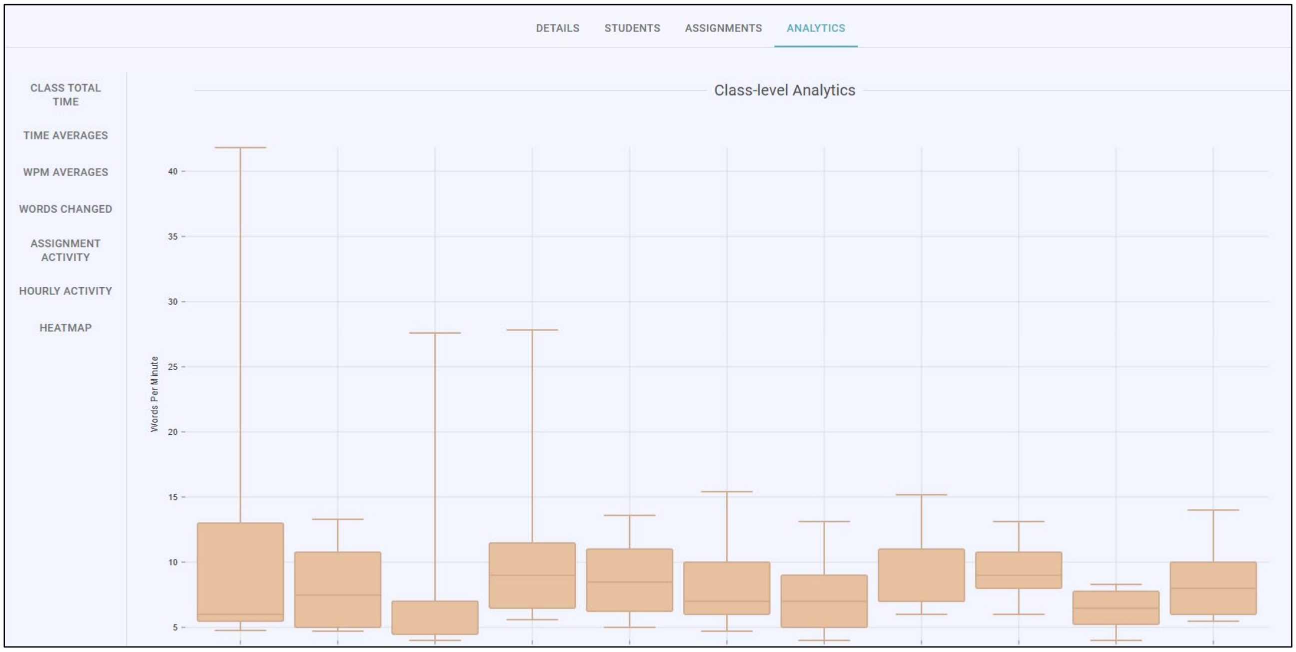

Figure 12 presents the Assignment Time feature, indicating students' total writing times throughout the class. Different assignments are represented by different color bars in the stacked bar chart. The chart is listed by total writing time, so the student with the highest writing time is presented at the top and the student with the lowest at the bottom. While the system does not currently have an analytics dashboard designed for an administrator or writing program administrator, time metrics such as those presented in Figure 11 seem the most promising candidates for inclusion in such a dashboard. (Further, to those concerned about the Forge's potential to be turned against instructors by administrators who wish to use it to surveil instructor activity, I would like to emphasize that the Forge is a faculty-driven project, hosted by LibreTexts, one of the world's largest open education websites founded and guided by a UC-Davis chemistry professor. As such, the Forge, like all LibreTexts technologies, will always aim to represent the best interests of faculty.)

Figure 12 presents the Assignment Time feature, indicating students' total writing times throughout the class. Different assignments are represented by different color bars in the stacked bar chart. The chart is listed by total writing time, so the student with the highest writing time is presented at the top and the student with the lowest at the bottom. While the system does not currently have an analytics dashboard designed for an administrator or writing program administrator, time metrics such as those presented in Figure 11 seem the most promising candidates for inclusion in such a dashboard. (Further, to those concerned about the Forge's potential to be turned against instructors by administrators who wish to use it to surveil instructor activity, I would like to emphasize that the Forge is a faculty-driven project, hosted by LibreTexts, one of the world's largest open education websites founded and guided by a UC-Davis chemistry professor. As such, the Forge, like all LibreTexts technologies, will always aim to represent the best interests of faculty.)

Figure 12

Stacked Bar Graph Showing Student Writing Time

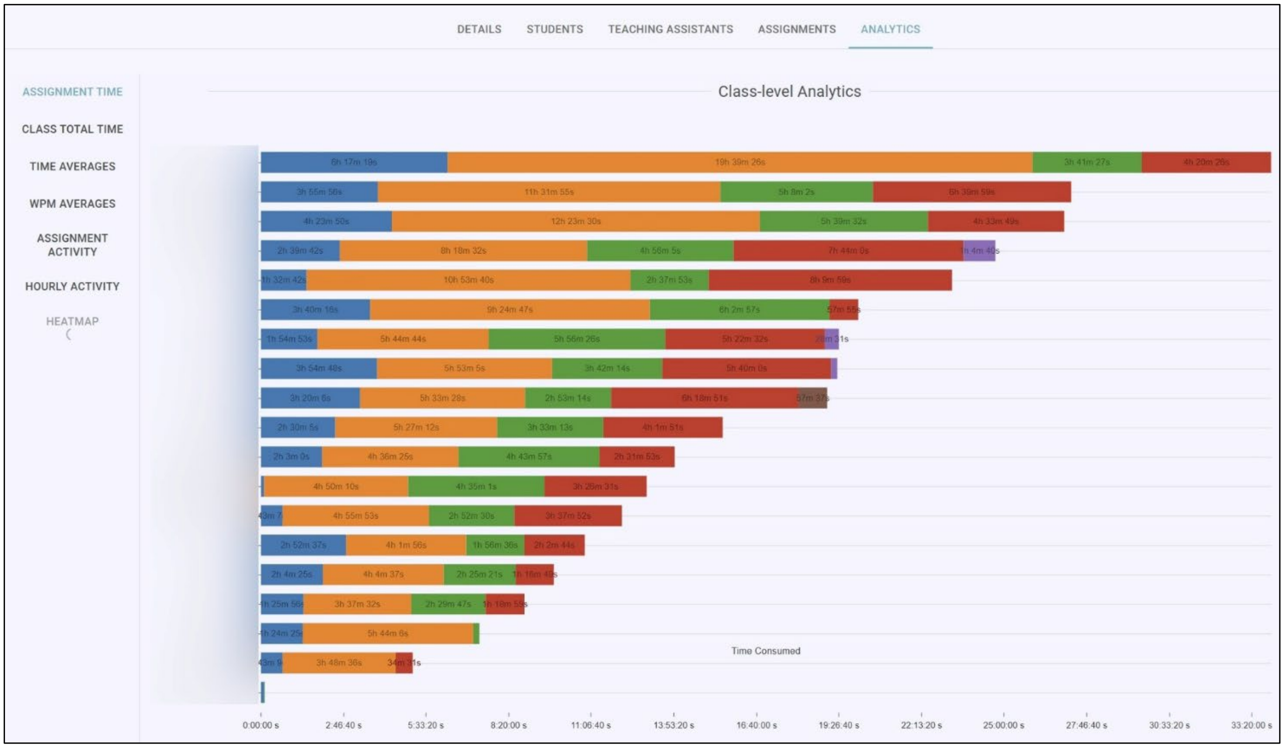

3.0 Application of Innovation

Academic assessment discourse increasingly focuses on measurable student learning outcomes (SLOs). However, objective, measurable SLOs have traditionally been difficult to achieve and monitor in writing education, and a writing analytics approach offers a potential advancement in this area. As previously suggested, most prior studies under the writing analytics paradigm have been ad hoc studies using submitted student work, using technologies like corpus linguistic tools (e.g., AntConc) rather than considering the dynamics of student writing processes. The challenges provided by the rise of generative AI suggest a pressing need for a single, unified platform that brings together student writing, instructor feedback, and writing analytics, both to support academic integrity and to provide rich insights into student writing processes to instructors, writing programs, and the research community. Indeed, while data analytics have become pervasive in our contemporary society, no comparable platform providing instructors and school administrators rich data analytics about students' writing processes exists.

While the team behind the Forge are expanding the system's analytics capabilities, the current features allow a wide variety of previously unavailable insights into student writing processes for instructional, administrative, and research purposes. At present, the Forge can be used for a variety of research tasks:

- developing profiles of student writing processes
- investigating the role that time-coded phenomena, including draft deadlines, play in the composition process

- providing insight into differing ways students write various types of genres and assignments
- assessing how student writing activity changes over the course of a class
- assessing how effective various pedagogical strategies are for encouraging student engagement in revision
- deepening our understanding of how students use generative AI writing and paraphrasing tools in their assignments, both to produce writing and to resist attempts by instructors to encourage original writing

As previously noted, the Forge system can be self-hosted by individual institutions or research teams in the interest of protecting student data and facilitating institutional review board compliance. (Indeed, if you are interested in hosting a deployment of the Forge at your institution, please contact us for technical guidance.)

4.0 Directions for Further Research

As Lang (2023) intimates, the writing analytics paradigm is still in its early stages, and possible applications for writing analytics tools using machine learning, both to support teaching and learning and to support research, are seemingly limitless. At present, generative AI writing tools have largely been designed to speed up the writing process, with the AI completing as much work for the writer as possible. The Forge team, meanwhile, wishes to use writer/student-facing AI for a different set of goals. A current team priority is to build out functionality that makes the system more active and “teacherly” during the student writing process. (Of course, instructors will have the option to turn this feature off.) Since the system is currently designed around iterative feedback from the instructor, the only dynamic feedback the system provides students while writing is traditional spell-checking and some limited grammar checking. Instructive feedback and/or controlled AI assistance (read: not doing all the work for the student) could potentially reduce the feeling that writing analytics systems like the Forge are largely about surveillance. Further, if the writing analytics paradigm takes hold, individuals will undoubtedly attempt to develop sophisticated typing bots that mimic more genuine student activity. Accordingly, having the system provide dynamic feedback to the student and then assess the student’s response to that feedback in their writing could further support students’ authentic engagement in the writing process.

Once the dynamic AI and analytics functionalities of the Forge are built out, the Forge could eventually support even more ambitious studies of student writing in university settings. Eventually, the system could facilitate research on students’ development of writing skills (e.g., use of argumentation strategies, documentation systems, and specific syntactic structures) throughout and across courses, as well as studies investigating the transfer of writing skills from one course to another and across a student’s program of study. At present, however, the Forge’s analytics functionalities are largely limited analytics within single sections.

As I have argued previously (see Oenbring, 2022), if compositionists wish to have a say in the future of writing technologies (and even the future of substantive student writing itself), we need to be more actively engaged in the development of writing technologies. I hope to have made the case in this piece for the Forge both as a pedagogical tool and as a research tool. (Of course, certain instructors may be more interested in the Forge primarily as a tool to induce compliance, and other instructors may be more interested in the Forge as a research tool.)

Nonetheless, as the events of the past few years have shown, major tech companies are more interested in the benefits of generative AI in the marketplace than the possible damage to academic integrity or student learning these tools can cause. Indeed, we might reformulate our previous assertion that *we must use data or we may well be used by data* to one focused on writing technologies: *we must develop writing technologies or be used by writing technologies.* Accordingly, the Forge team is delighted to be bringing our technology to the faculty-led, open education nonprofit LibreTexts.

Author Biography

Raymond Oenbring is Professor of English and Linguistics at the University of the Bahamas. His program of research includes rhetoric and composition, linguistics, and Bahamian English. He is managing editor of the *International Journal of Bahamian Studies*.

References

Aull, L. (2015). *First-year university writing: A corpus-based study with implications for pedagogy*. Palgrave Macmillan.

Clinton-Lisell, V. (2021). Open pedagogy: A systematic review of empirical findings. *Journal of Learning for Development*, *8*(2), 255-268.

Dalalah, D., & Dalalah, O. (2023). The false positives and false negatives of generative AI detection tools in education and academic research: The case of ChatGPT. *The International Journal of Management Education, 21*(2).

Inoue, A. (2019). *Labor-based grading contracts: Building equity and inclusion in the compassionate writing classroom,* 2nd ed. The WAC Clearinghouse; University Press of Colorado. https://doi.org/10.37514/PER-B.2022.1824

Kaufer, D., & Ishizaki, S. (2023). The DocuScope project: History, theory and future directions. In *Corpora and Rhetorically Informed Text Analysis: The diverse applications of DocuScope* (pp. 2-24). John Benjamins Publishing Company.

Lang, S. (2023). Editor's note: A new phase for *The Journal of Writing Analytics. Journal of Writing Analytics, 7.* https://doi.org/10.37514/JWA-J.2024.7.1.00

Lang, S., & Baehr, C. (2012). Data mining: A hybrid methodology for complex and dynamic research. *College Composition and Communication, 64*, 172-94. http://www.jstor.org/stable/23264925

LeBlanc, P. J. (1993). *Writing teachers writing software: Creating our place in the electronic age.* National Council of Teachers of English. https://eric.ed.gov/?id=ED357369

Marche, S. (2022). The college essay is dead: Nobody is prepared for how AI will transform academia. *The Atlantic*. https://www.theatlantic.com/technology/archive/2022/12/chatgpt-ai-writing-college-student-essays/672371/

Miller, B. (2021). The pleasurable difficulty of programming. In C. VanKooten & V. Del Hierro (Eds.), *Methods and methodologies for research in digital writing and rhetoric*. The WAC Clearinghouse. https://doi.org/10.37514/PRA-B.2022.1541

Miller, B., & Licastro, A. (2021). *Composition and big data*. University of Pittsburgh Press.

Mills, A. (2024). Writing process tracking is coming to Grammarly—and students and teachers need it. *Substack*. https://annamills.substack.com/p/writing-process-tracking-is-coming

Moxley, J. (2013). Big data, learning analytics, and social assessment. *Journal of Writing Assessment*, *6* (1).

Oenbring, R. (2022). ThesisWorkshopper: An automated thesis statement evaluator. *Journal of Writing Analytics*, *6*. https://doi.org/10.37514/JWA-J.2022.6.1.02

Omizo, R. (2019). Participation and the problem of measurement. In P. V. Banaji, L. Blankenship, K. DeLuca, L. Obermark, & R. Omizo (Eds.), *The rhetoric of participation: Interrogating commonplaces in and beyond the classroom*. Computers and Composition Digital Press/Utah State University Press. http://ccdigitalpress.org/rhetoric-of-participation

Warner, M. (2022). *A Zoomable assessment: Navigating the ecologies of writing program assessment.* [Doctoral dissertation, North Dakota State University]. https://hdl.handle.net/10365/33509