

# Reader-scale Text Mining and Visualization in the Secondary ELA Classroom

Nichole Nomura, University of Wyoming

Georgii Korotkov, Stanford University

Sarah Levine, Stanford University

Christine Wynn Bywater, Stanford University

Victor R. Lee, Stanford University

## Abstract

We report on our work collaborating with two Bay-Area middle-school (grades 6, 7, and 8) districts to design, implement, and research the use of text visualizations in supporting English Language Arts and data literacy learning outcomes. Through co-design, classroom research, and participatory professional development with English Language Arts teachers, the project developed and tested text analytics tools, visualizations, and pedagogical routines. This article overviews the constraints we encountered, including technical constraints related to devices, access, and use, and pedagogical constraints including teacher and classroom time, and highlights a few strategies developed in response to those constraints. Drawing on classroom video, co-design sessions with teachers, and qualitative interviews, we find that “small” data visualizations that visualize text mining data created on single texts or small excerpts may support student data literacy and self-efficacy as well as the development of a more expansive linguistic-experiential reservoir. We present a routine, “What did the computer miss?”, used by our collaborating teachers to support thinking with and about computing technologies that works especially well with single-/small- text data visualizations in the middle-school context, and conclude with suggestions and future directions for research and collaboration.

**Keywords:** middle school; ELA; co-design; scale; text mining; data visualization.

# Introduction

In this article we report on our experiences co-designing middle-school text visualizations to simultaneously support both ELA (English Language Arts) and data literacy learning outcomes by drawing on digital humanities methods and interpretive frameworks, especially those from computational literary studies. Digital humanities practice in secondary education and in community partnerships is hard, understudied (Georgopoulou et al. 2025), and frequently under-resourced (Arteaga 2014; Carnes and Smith 2024), motivating minimal approaches, as it is, despite these challenges, frequently valuable and rewarding for teachers, students, and researchers (Gil et al. 2018). Like our colleagues before us, we can report that working with school-managed computers is especially challenging, that even websites can be complicated on school-managed devices, and that forms of digital humanities that do not require devices (working with printed visualizations, for example) are still resource challenges to be solved in the middle-school context. Drawing on classroom video, co-design sessions with teachers, and qualitative interviews, we find that “small” data visualizations that visualize text mining data created on single texts or small excerpts may support student data literacy and self-efficacy as well as ELA learning outcomes via the development of a more expansive linguistic-experiential reservoir. We present a routine (“What did the computer miss?”) used by our collaborating teachers to support thinking *with* and *about* computing technologies when working with the single-/small- text data visualizations, and conclude with suggestions and future directions for research and collaboration.

We hypothesized that digital humanities methods and frameworks might serve to productively support students’ learning in both ELA and data literacy, while acknowledging that there is a meaningful difference between pedagogy for the development of future digital humanities researchers and the integration of digital humanities methods and tools into the disciplines as they are enacted in K12 spaces. Our project therefore emphasized the development of ELA and data literacy learning outcomes over digital humanities ones. Our project model was to iteratively co-design, test, and research text analytics tools, visualizations, and pedagogical routines to develop pedagogical content knowledge and to better understand the interactions between data literacy and ELA learning outcomes at the site of the text-based data visualization in secondary contexts. Our research team included experts in ELA, Data Science Education, Natural-Language Processing, Teacher Training and Professional Learning, and the Digital Humanities.

Our co-designers were a cohort of middle-school ELA teachers from two different public school districts in the Bay Area, teaching across grades six, seven, and eight. Teachers participated in an initial, day-long professional learning session, onboarding them to the project, introducing them to critical data-literacy concepts, and soliciting their needs and input for the year’s design projects. After that, we met once a month with the cohort from each district, combining new professional learning (such as the introduction of new kinds of visualizations or new routines) with mutual sharing (individual teachers debriefing a lesson with the group) and co-design (collaborating with the research team on pedagogical moves, most frequently in the form of prototype routines). Teachers had two individual qualitative interviews with a member of the research team—one halfway through the study, and one at the end of the academic year. Further, several teachers opted into additional research in the form of recorded classroom observations. Upon completion of the study, teachers had the opportunity to receive Continuing Education Units and received a stipend for their contributions. This article draws on recordings from the professional learning sessions, teacher interviews, and classroom observations.

# Constraint in Middle School Contexts

The categories and specifics of constraints we identify below will be recognizable, and potentially obvious, to educators who regularly work in middle school settings. For readers whose expertise is drawn from other educational contexts: these constraints are generalized conditions that can be found, with some variation, in secondary public education across the United States, especially in middle schools, and which we present based on our own expertise, the literature, and in conversation with our collaborating teachers. We highlight, therefore, the value of inter- and multi-disciplinary teams for this kind of work, including researchers, district partners, and practitioners.

## Pedagogical constraint in middle school contexts

The biggest pedagogical constraint we encounter in middle school contexts is time. Teacher-training/preparation time and instructional time are valuable resources, subject to ever-increasing demands. Teacher time is always at a premium, a theme that our teachers, veteran and new alike, emphasized heavily during our final interviews with them. Classroom time is likewise at a premium, and while a few of our teachers structured entire lesson sessions around a visualization, many found themselves working with visualizations in short, 5–15 minute sessions, especially when they used them as an introduction or recap to a text.

Tool failure was not a cost in teacher time that we were willing to pay. Experimentation with new methods and tools comes at a great time cost, even when those new methods and tools work well. Digital humanities methods and tools, to be frank, have a well-earned reputation for not working easily or without some tweaking—and that is before we factor in the demands of pedagogical experimentation. The professional digital humanities in the higher education space is working on an attitude that celebrates, or at least acknowledges, failure in both research (Spiro 2012; Dombrowski 2019) and pedagogy (Croxall and Warnick 2020). That culture does not necessarily hold true in the same ways, with the same incentives and rewards, when working with public school secondary teachers covering an already-full curriculum. It was, of course, possible that any given lesson plan would flop—these are the risks of teaching. But we did not want to add to those existing risks, and by shifting the majority of the risk that code would fail—what we might call the “incidental complexity” (Giannetti 2017)—to the research and development team, we freed up teacher time and effort to experiment with the pedagogical elements of the project. The research team, therefore, did all the coding and visualization creation work.

As teachers face a scarcity of time, they value flexibility and adaptability in the resources they use in their teaching. With that in mind, rather than offering fully developed, detailed lesson plans, we worked with teachers during our co-design sessions to develop routines—reworkable structures like short conversation starters or mini-activities to do with their students—aimed at developing pedagogical content knowledge rather than specific pedagogical content itself. These routines drew on the team’s own interdisciplinary pedagogical content knowledge, and were prototyped, tested, and refined in conversations with teachers. For example, “Big, Medium, Small” is a routine we developed for recognizing variability in a data visualization based on size or level. Asking students to check what is big, what is medium, and what is small in a visualization is a pedagogical structure that can be expanded and enacted in a number of ways. It supports teacher pedagogical content knowledge, as prior research shows students tend to focus on individual data points rather than the larger collection—often focusing on

more extreme values (Ben-Zvi and Arcavi 2001; Konold et al. 2015; Lee 2019). More deliberate, sustained attention and support is typically needed for students to see the whole range and aggregate. “Big, Medium, Small” scaffolds that sustained attention and support for teachers and students in a lightweight and flexible way.

Routines like “Big, Medium, Small” were designed to be memorable and adaptable. The routine support documents are typically one or two pages and are organized around an illustrative cartoon of a possible classroom interaction using the routine. This approach was appreciated by the teachers—one veteran ELA teacher, who was new to the data visualization project, said:

Not only did you all make the visualizations, but you also gave us the tools to use to implement those visualizations. That pairing is really powerful. Because it's like I didn't even have to spend any time thinking about it, and I don't have the time to spend to think about it, so I probably wouldn't have used it if I had been required to do that. But since it was just like pick a visualization, pick a routine, go. I think that was a really nice pairing.

While the comment above might invite comparisons to scripted curriculum, routines are *not* complete scripts. They're lightweight models for scaffolding conversations, with suggested uses and examples. This teacher (and the others in our study) did spend a lot of time thinking about routines—just not in their classroom preparation time. Time for thinking, testing, brainstorming, and considering examples was instead a regular part of the monthly professional development sessions.

## Technical constraint in middle school contexts

Technical constraints are often, in middle school contexts, related to financial ones. Chromebooks have significant technical limitations, especially in terms of software, but many schools use them because they are less expensive than other laptops. Digitally delivered content needs to accommodate the range of devices and operating systems found in middle schools, whether they work on a one-to-one computing model or have to run on the teacher's device. Other technical constraints are related to the nature of on-campus web and software access. Student browsing and ability to download or install software are not only constrained by the hardware and operating system of their devices, but also by school-site and district-level restrictions on the web pages and domains they can access. Under these constraints, we turned to a website for the delivery of our visualizations (we note that this move is not a perfect solution: a subpage of our website was accidentally blocked by the district, forcing a delay in a teacher's lesson planning). Provided the school's internet access was working, this was easier than software in terms of student-sided use.

We want to highlight here that creating the conditions for minimalist access can require shifting the expense elsewhere—in this case, to a team with significant web development and coding expertise, and to the hosting supported by a large institution. This shift in expense extends beyond initial development to ongoing maintenance, where technical debt becomes a persistent challenge that can threaten the

sustainability of educational technology projects. Minimalist digital humanities, and minimalist digital humanities pedagogy, must negotiate technical debt as a constraint.

Technical debt, a concept borrowed from software engineering, refers to the implied cost of future reworking required when expedient solutions are chosen over more robust ones.<sup>1</sup> In the context of our project, technical debt manifested most clearly in our dependency management—the complex web of software libraries and packages that enable our visualizations to function. Our project relies on approximately thirty different dependencies, each serving specific functions from text processing to visual rendering. While this modular approach allows us to leverage existing tools rather than building everything from scratch (a necessary pragmatism for the conditions of our project), it also creates a fragile ecosystem where updates to any single component can cascade into system-wide failures.

This reality underscores a fundamental tension in minimalist digital humanities, and one that refracts in pedagogical contexts: the tools that make computational methods accessible to teachers and students may require sophisticated technical infrastructure that educational institutions are often not equipped to maintain, and the choices we might be used to making as researchers may not be available. The apparent simplicity of a web-based visualization masks the complex technical ecosystem required to sustain it, creating long-term dependencies that can outlast both funding cycles and the technical expertise of their creators (Dombrowski 2022). In answering Risam and Gil’s “what are we willing to give up?” (2022), we sacrificed sustainability for our ability to support this particular cohort of teachers, in this particular year, with their particular needs. The website will live for as long as it lives, and then it won’t anymore. We have shared PDFs of the visualizations with the teachers, sacrificing the interactivity and accessibility facilitated by the website for durability in our long-term planning.

## Scale constraints in middle school contexts

Scholars like McLean Davies et al. have called for corpus-level digital humanities and database literacy in K12 spaces in order to build “postdigital literary literacy” (McLean Davies et al. 2020). Work in this vein could take similar approaches to college and professional digital humanities practice, at similar scales, and bring those approaches, corpora, and frameworks into the K12 classroom. We, however, found ourselves working at a different scale: with single and small texts. Our project did not plan to be a single-/small-text project from *a-priori* principles; rather, the co-design structure of the project, in which teachers brainstormed visualization ideas for the texts they would be teaching this year, resulted in our near-exclusive focus on single and small texts. While this is characteristic of middle school ELA, in which students are frequently working with the unit of a single text, reading individual examples while strengthening the interpretive skills that support synthesis in later years, this is different from the corpus-level approach we frequently find in literary text-mining.

Our 2024–2025 collaborating teachers almost exclusively requested visualization support on single-texts—only a single nonfiction collection was requested. A few requests could be worked with as individual texts or as a very small corpus—a set of four different personal narratives from a single collection, or the different acts within a play—and several of the novels were segmented into chapters. Our emphasis in this paper is not on the idea of a cohesive, single text, but rather on the smaller scale of these visualizations—a scale that is consistent with the existing, single-text reading tasks characteristic of middle school ELA.

Tools found in secondary pedagogical use (e.g. Wordle) have long supported single-/small-text use cases. Tom Lynch's website *Plotting Plots*, designed for secondary classroom use, presents visualizations almost exclusively of single texts (the only exception is Shakespeare's sonnets, and each sonnet is treated as if a chapter or scene, ordered sequentially in the plot and presented as a cohesive whole) ("Plots | Plotting Plots" 2021). The website as a whole does offer a few forays into corpus-level analysis—e.g. considering the entirety of Shakespeare's tragedy corpus as a contrast to Macbeth—but the blog posts and available plots by and large emphasize the single text. Since the word cloud software Wordle's launch in 2008, The National Council of Teachers of English (NCTE) has published articles aimed at secondary teachers that suggest use cases for word clouds on the scale of the single or small text: two paragraphs of Martin Luther King's "I have a dream" speech (Hagood 2012), student poll results for "choose two words that best describe Scrooge in the early stages of the story" (Turner and Hicks 2022), the highlighted words from a short passage think-aloud (Newman and Rosas 2016), "important passages in plays and novels" (Atkins 2011), and documents like "the Declaration of Independence, Roosevelt's Pearl Harbor speech, a slave narrative, or a soldier's letter home from a war" (Sims 2010). In these examples, a single- or small-text visualization has found a productive synergy with existing ELA practices.

Conversations with teachers, both during the design stage and in our year-end debrief, repeatedly approached the question of scale: requests for ways to focus (via reduction) the conversation, to filter, to make smaller, to make more readable. One teacher told us that she would probably never use a whole-novel word cloud, because it would be "overwhelming"—blending two kinds of scale problems, first, length, and second, variety:

I'm thinking of [the novel] *Summer of the Mariposas* ... That's 22 chapters, right? And then ... those chapters are so varied. ... it's like if you took one big word cloud ... it would be too much to digest with that type of novel, right? Because ... it's the hero's journey, so there's ... so many different variations ... that, having, like one big word cloud would be daunting.

Other teachers asked for increased abilities to filter the visualizations along a variety of information channels—either reducing the number of words, removing stop words, or filtering out different parts of speech.

The choice of how to make word clouds or other visualizations smaller is not a choice the visualization development team needs to make—and in fact, is probably not a choice the visualization development team *should* make. Teachers stressed that they wanted to make these choices. The challenge for the development team is facilitating making those choices possible, making the graphic user interface easy-to-use, and remembering that they need to work at the smaller-/single- text scale.

## **What Did the Computer Miss?: Routines for Teaching With Small Visualizations**

We prototyped, tested, and developed several routines this year to scaffold teacher pedagogical content

knowledge while teaching with data visualizations (Lee et al. 2025). Routines like “Big, Medium, Small” draw explicitly on data literacy concepts, while other routines like “Notice, Mean, Wonder” are interdisciplinary. Our most digital-humanities-aligned routine, one that supports the goals of thinking *with* and *about* computing technologies (Ringler 2024) in humanities contexts, is “What did the computer miss?”

“What did the computer miss?” is a pedagogical routine that supports students in comparing the computer’s representation of a text to their own interpretation or to the text itself. It supports any visualization where there is a calculated value, and variations of it can be used to support more general thinking *with* and *about* computing.

Follow-up questions include:

- Is the computer measuring the same thing we want to measure? (e.g. Does the computer’s “sentiment” capture the same thing as a human reader’s “sentiment”?)
- What are the possible errors a computer might make while trying to represent X?
- Let’s challenge the computer—can you find a place in the text that contradicts what we see on the graph?
- What do we know about how a computer produces this?

In our co-design sessions with teachers about this routine, including a session where we watched classroom footage of the routine in use and then analyzed it together, teachers emphasized a) the value of being able to prompt students to go back into the text for evidence and b) that it was important for them that students develop authority and skill in challenging computers and data in general. It should be emphasized that the teachers in this project were generally pro-data, or at least pro-data literacy—they see data as an essential or unavoidable part of their lives and student lives—and that their desire for students to feel able to challenge data is not a desire for students to blankly reject it.

Professional digital humanists in computational literary study vary in their epistemic commitments and their baseline level of belief that computers a) can and b) do “miss” things. In this vein, Mark Algee-Hewitt identifies three primary approaches to computational literary study (2024). Some tend to use computational modeling processes as a mechanism for producing empirical results, drawing on methods and standards for validity from the social sciences, data science, and statistics. Others use computational modeling’s results—whether they feel right or wrong to the trained eye of the humanist—to direct or redirect the scholar’s attention to the text. Others still argue for a hybrid approach that negotiates between the computational and the literary-critical, using operationalizing and modeling as the critical, ever-reflexive move of literary study itself. In all three approaches, a scholar’s ability to verify whether a computer missed anything—whether via gold-standard data or via their own literary-critical expertise—is an essential first, and potentially, repeated step. The degree to which computational literary studies asserts or needs to assert if a computer or algorithm “missed” (or is capable of missing) something is deeply related to questions of scale, whether that scale is framed in terms of statistically valid modeling or questions of adequate and responsible representativeness. Scale, both large and small, in its various forms, is one of axes computational literary studies thinks through as it considers the questions of

minimal computing—what we have, what we need, what we want, and what we are willing to give up (Risam and Gil 2022)—because what we “need” is professional digital humanities is often a convincing and valid (both internally and externally) argument.

Being able to assess whether a computational process misses things is also a necessary skill in data science education. That can be in terms of the principles of data feminism (D’Ignazio and Klein 2020, Lee et al. 2022) or in terms of statistical skepticism (Pfannkuch and Wild 2000). Such a question is also fundamental to definitions of algorithmic literacy (Oeldorf-Hirsch and Neubaum 2025), and especially important to understanding algorithmic systems as an interplay between humans and computers (Seaver 2019).

The question for our teachers and research team was how and why “What did the computer miss?” might be useful in service of ELA learning outcomes in a middle school context.

## **The Extreme Efferent and Human Memory: Affordances of Reader-scale Text Mining and Visualization**

We believe this routine works so well in the ELA classroom because of the small scale of the visualization. Small visualizations are especially useful in supporting what Lynch identifies as Rosenblatt’s *extreme efferent*—and we believe they can likewise support increased data literacy and self-efficacy as well as the development of a more expansive linguistic-experiential reservoir. That linguistic-experiential reservoir is a key component of many of the strategies used in the ELA classroom, from close-reading to creative writing.

Lynch’s “Electrical Evocations,” published in the practitioner-focused journal *English Education*, documents Lynch’s personal journey into literary text mining (his own encounter with a visualization of *Mrs. Dalloway*) and proposes a theoretical framework that incorporates the products and processes of literary text mining into Rosenblatt’s reader-response theory (Lynch 2019). The significant interventions are twofold: first, computational processing is a form of the “extreme efferent,” and therefore, part of, rather than distinct from, our existing literary-critical meaning making practices, which include efferent reading. Second, by conceptualizing the products of computation as computational associations available for referential meaning making, just as a dictionary or thesaurus is available for referential meaning making, Lynch pushes us to consider text-mining data as one resource, among many, for meaning-making rather than a challenge to it—one that we might teach students to use just as we teach them to use a dictionary.

The use of “What did the computer miss?” asserts the value of students’ linguistic-experiential reservoirs and is largely tractable because students have memory on par with that of the computer for the task. Teacher suggestions for using the routine focused on ways to equalize that memory even more—by looking at only one data point at a time or by requiring students to go back into the text to look for evidence to refute or confirm the visualization. Visualizations of text where computer RAM and human memory are closer to equal afford different pedagogical moves than those where computer RAM is orders of magnitude larger than a human’s memory, as we see in corpus-scale text mining visualizations.

Human memory, in this context, does not only refer to working memory or the ability to recall the text—it is augmented by the ability to flip through pages and access memory supports in the form of words on

the page or in notes. It blurs the lines between what is held in students' brains and what they know how to find easily enough to discuss in the moment—and our emphasis is on the workability of that memory in a class session. The single text visualization is at the appropriate scale for these readers' expertise.

In student responses to “What did the computer miss?”, we see negotiation of both the computer’s ability to be “right” and the leveraging of specific evidence, from a specific moment in the text, alongside the students’ own linguistic experiential reservoir.

For example, in a conversation about a dramatization of Charles Dickens’s “A Christmas Carol,” one participating teacher showed students a [sentiment analysis plot](#) of the script and asked whether “the computer got it right.” Many students challenged the computer. One student made the broad point that computers cannot feel or think like humans. Others pointed to a specific data point and contested it, as in this comment:

Well, the majority of the time, the computer got it right, except for in Act II, Scene 4, cause in Act II, Scene 4, Scrooge literally saw his own grave—I wouldn’t be happy, I don’t think he’d be happy if he saw his own grave.

In this response, the student makes several interesting moves that leverage their own linguistic-experiential reservoir and evidence from the text—contesting the computer’s judgment on the grounds of “literally,” before analogizing their feelings (I wouldn’t be happy) to Scrooge’s (I don’t think he’d be happy). It is on the grounds of “literally” that we see the student negotiating the relationship between data and text—between a computed value, and the words on the page: “literally” suggests that the words on the page may be worth more to this student, and they draw on the referential elements of their linguistic-experiential reservoir. To understand and contest that computed value, though, they must also dip into the aesthetic elements of their linguistic-experiential reservoir—their lived experiences of words about happiness, and the accrued meaning over time, with repeated lived exposure, of symbols and scenes like the one in which Scrooge sees his own grave. The framing of the question, teachers commented, supported the legitimacy of students’ experience and interpretations in contesting the visualization. The small scale of the visualization allows students to leverage the referential in a way that nearly-matches the computer: going back into the text for evidence (in other classroom videos and teacher reports, we observe students producing their own counts to contest or verify the visualization) or drawing on their own memory.

Professional literary critics perform the move of “What did the computer miss” on corpus-level visualizations all the time: what text-mining project has not encountered, at one point or another, a question that begins “But what about—?”. Many times, the researchers ask this question of themselves. It’s an essential move, but it is one that requires breadth in reading and fluency in the corpus sufficient to the task. Literary critics may acquire that kind of fluency after decades of professional reading in a subfield. Scaffolding a similar move for students who are not professional literary critics—or even for those who are—may require a shift in scale. Here, we have found a scale that supports middle grade readers.

There are risks to this routine—the question does not encapsulate all data and digital humanities literacy in itself. Computers, generally speaking, do not really get things “wrong” or “miss” things, but in this routine, they stand in, as synecdoche, for the algorithms, processes, and quite human decisions that do get things wrong and miss things. The suggested follow-up questions and the professional learning sessions we did with teachers both offered ways to complicate the question, which will ideally support later, more-nuanced distinctions between computers and the processes they run, and the assumptions of the programmers who write those processes.

There is also a risk inherent in the framing—in discussing how the routine played out in this class session, teachers highlighted the problem of students only needing to come up with a single response—in terms of both ELA and data literacy learning outcomes. It is much easier to demonstrate that the computer missed something than to assert that the computer missed nothing by checking all of its work. The classroom conversation about challenging the sentiment analysis for “A Christmas Carol” starts by identifying a moment in the text students thought was a strong counterexample, but, with teacher guidance, eventually builds to a discussion of something that resembles averages—while students wanted to focus on this one important point, the computer took into account all the words in the segment. This routine is therefore the starting place for discussion, not the end goal, and our review of classroom practice using this routine suggests that it facilitates mutual support between data literacy learning outcomes and ELA learning outcomes.

## Conclusion

Working within the constraints of middle school digital humanities pedagogy, we found an opportunity—we think visualizations of single and small texts may support students leveraging their linguistic-experiential reservoirs and further develop those reservoirs to include the computational extreme efferent. These visualizations are minimalist in contrast to research practice, and they have different aims. The teachers in our study talked about “going back into the text,” “engaging students,” and “learning to challenge data” as some of the primary effects of their visualization use. These are not the goals of visualization use in research—many digital humanities scholars have been scolded for using a visualization solely for the purposes of engagement. But in a middle school ELA classroom, the very real goal of getting a silent student to talk means that engagement matters differently; a lesson that may carry over, with modification, into post-secondary education. Visualizations that abstract only enough, while supporting returning to evidence, probably do not abstract sufficiently for most distant-reading claims, and a redundant research visualization is likely to be cut by an editor—but a redundant *teaching* visualization can productively constrain interpretation (supporting students in avoiding misconceptions), provide complementary information, and support deeper understanding by requiring comparison and transfer across the visualizations (Ainsworth 2006). The approach to middle school text visualization we took in this project—a visualization with a lightweight explanation of what it shows and how it was made, paired with general data and digital humanities literacy routines—will require new approaches and experimentation beyond our current research- and argument-centered practices.

That is not to say that digital humanities learning outcomes do not have a place in K12 education—we find the argument of McLean Davies et al. about the possibilities of working with large-scale digital objects like literary databases in supporting “postdigital literary literacy” very compelling (McLean Davies et al. 2020). We believe that our work with single-text and small-text visualizations ultimately

support this kind of learning outcome, scaffolding data literacy in a way that is productively foundational for later, large-data work. Small-text and single-text data visualizations support the development of students' linguistic-experiential reservoir, expanding them to include extreme efferent data like word frequency, but in a way that is still deeply contextualized within their own reading experiences and practices, which are learning outcomes of high importance to ELA teachers. These small visualizations may be a necessary scaffold for the work McLean Davies and collaborators have called upon us to begin; that is, corpus-level digital humanities and database literacy in K12 spaces, as they strengthen students' abilities to question operationalizing and visualization decisions by scaling the computer's memory for a task to something more comparable to their own.

We conclude with two practical notes for those interested in working in this space:

First, we need a range of visualization options. Not all our teachers are interested in making their own visualizations—a point that was made repeatedly in our debrief interviews. While a few teachers were comfortable with word cloud generators, many still expressed doubts about their ability to prepare different kinds of text for those generators. They offered that they would continue to use the visualizations we had prepared for them this year but were unlikely to seek out new ones if they had to make them. While that presents opportunities for tools like Voyant and other Graphical User Interfaces (GUIs) that facilitate visualization creation without code, it also suggests to us that Lynch's model in the website *Plotting Plots*—of creating the plots for teachers—will remain necessary in facilitating the adoption of text visualizations in secondary contexts, and will require a more generic approach to visualization that prioritizes pedagogical possibility and customization. For different student populations and circumstances, the pedagogical utility may differ, and we may need to modulate how argumentative, let alone how argumentatively sound, we expect a visualization to be. We therefore need further study on how teachers and students make meaning out of these visualizations in order to support the development of these specifically pedagogical visualization practices.

Second, there is a meaningful divide between pedagogical content knowledge designed to support the development of future digital humanities researchers, and the integration of digital humanities methods and tools into the pedagogical content knowledge for the disciplines as they are enacted in K12 spaces. We believe this divide is bridgeable but building that bridge requires collaboration. How do we structure our collaborations so that we can be deliberate about the relationship between digital humanities and other fields, and the relationship between secondary and postsecondary? We believe that responsible co-design and qualitative methods like interviewing, along with healthy interdisciplinary research relationships with colleagues in education, data science, and teacher training, are necessary for this kind of work. This approach can require a lot of work and resources in terms of time, money, and training. If minimalist K12 digital pedagogy is a balancing act between constraints, needs, and the things we are willing to give up, we stress that this is not the thing to give up. Give up corpora, give up interactive visualizations hosted on institutional websites, give up your time, give up open-source commitments, but do not give up building good collaborations with the teachers you work with.

## Acknowledgements

The authors thank our participating teachers, school district partners, and their students for their involvement in this project. This work was supported in part by funding from the National Science Foundation under Grant No. 2241483. The opinions expressed herein are those of the authors and do not

necessarily reflect those of the National Science Foundation. Thanks also to ELAlytics team members Dorottya Demszky, Liz Harris, Lena Phalen, Daniela Gamboa, and Deepak Varuvel Dennison.

## Notes

1. See Wiese et al, 2025's identification of technical debt challenges in terms of management, and Avgeriou et al 2016's consolidated "16162" definition of technical debt: "In software-intensive systems, technical debt is a collection of design or implementation constructs that are expedient in the short term, but set up a technical context that can make future changes more costly or impossible. Technical debt presents an actual or contingent liability whose impact is limited to internal system qualities, primarily maintainability and evolvability." [↑](#)

## References

Ainsworth, Shaaron. 2006. "DeFT: A Conceptual Framework for Considering Learning with Multiple Representations." *Learning and Instruction* 16 (3): 183–98.  
<https://doi.org/10.1016/j.learninstruc.2006.03.001>

Algee-Hewitt, Mark. 2024. "Computing Criticism: Humanities Concepts and Digital Methods." In *Computational Humanities. Debates in the Digital Humanities* 11. University of Minnesota Press.

Arteaga, Rachel. 2014. "Spar: Digital Humanities, Access, and Uptake in Rural Southwest Washington State." *New American Notes Online*. July 2014.

Atkins, Janet. 2011. "From the Secondary Section: Reading and Writing with Purpose: In and out of School." *English Journal*. <https://doi.org/10.58680/ej201118227>

Avgeriou, Paris, Philippe Kruchten, Ipek Ozkaya, and Carolyn Seaman. 2016. "Managing Technical Debt in Software Engineering (Dagstuhl Seminar 16162)." In *Dagstuhl Reports (DagRep)*, 6 (4). Schloss Dagstuhl—Leibniz-Zentrum für Informatik. <https://doi.org/10.4230/DAGREP.6.4.110>.

Ben-Zvi, Dani, and Abraham Arcavi. 2001. "Junior High School Students' Construction of Global Views of Data and Data Representations." *Educational Studies in Mathematics* 45 (1): 35–65.  
<https://doi.org/10.1023/A:1013809201228>.

Carnes, Jeremy, and Margaret K. Smith. 2024. "Bridging the Divide: Improving Digital Humanities Pedagogy by Networking Higher Education and Secondary Education Faculty in St. Louis." *IDEAH* 4 (2). <https://doi.org/10.21428/f1f23564.5127bfe6>.

Croxall, Brian, and Quinn Warnick. 2020. "Failure." In *Digital Pedagogy in the Humanities: Concepts, Models, and Experiments*, edited by Rebecca Frost Davis, Matthew K. Gold, Katherine D. Harris, and Jentery Sayers. Modern Language Association.  
<https://digitalpedagogy.hcommons.org/keyword/Failure>.

D'Ignazio, Catherine, and Lauren F. Klein. 2020. *Data Feminism. Strong Ideas*. The MIT Press.

Dombrowski, Quinn. 2022. "Minimizing Computing Maximizes Labor." *Digital Humanities Quarterly* 16 (2).

Dombrowski, Quinn. 2019. "Towards a Taxonomy of Failure." *quinndombrowski.com*, January 30. <https://quinndombrowski.com/blog/2019/01/30/towards-taxonomy-failure/>.

Georgopoulou, Maria Sofia, Christos Troussas, Evangelia Triperina, and Cleo Sgouropoulou. 2025. "Approaches to Digital Humanities Pedagogy: A Systematic Literature Review within Educational Practice." *Digital Scholarship in the Humanities* 40 (1): 121–37. <https://doi.org/10.1093/llc/fqae054>.

Giannetti, Francesca. 2017. "Against the Grain: Reading for the Challenges of Collaborative Digital Humanities Pedagogy." *College & Undergraduate Libraries* 24 (2–4): 257–69. <https://doi.org/10.1080/10691316.2017.1340217>.

Gil, Alexander, Roopika Risam, Stan Golanka, Nina Rosenblatt, David Thomas, Matt Applegate, James Cohen, Eric Rettberg, and Schuyler Esprit. 2018. "Digital Humanities in Middle and High School: Case Studies and Pedagogical Approaches—DH 2018." Mexico City: ADHO.

Hagood, Margaret C. 2012. "Risks, Rewards, and Responsibilities of Using New Literacies in Middle Grades." *Voices from the Middle* 19 (4): 10–16. <https://doi.org/10.58680/vm201219347>.

Konold, Clifford, Traci Higgins, Susan Jo Russell, and Khalimahtul Khalil. 2015. "Data Seen through Different Lenses." *Educational Studies in Mathematics* 88 (3): 305–25. <https://doi.org/10.1007/s10649-013-9529-8>.

Lee, Victor R. 2019. "Supporting Complex Multimodal Expression Around Representations of Data: Experience Matters." In *Critical, Transdisciplinary and Embodied Approaches in STEM Education*, edited by Pratim Sengupta, Marie-Claire Shanahan, and Beaumie Kim, 217–31. Advances in STEM Education. Springer International Publishing. [https://doi.org/10.1007/978-3-030-29489-2\\_12](https://doi.org/10.1007/978-3-030-29489-2_12).

Lee, Victor R., Daniel R. Pimentel, Rahul Bhargava, and Catherine D'Ignazio. 2022. "Taking Data Feminism to School: A Synthesis and Review of Pre-collegiate Data Science Education Projects." *British Journal of Educational Technology* 53 (5): 1096–113. <https://doi.org/10.1111/bjet.13251>.

Lee, Victor R, Elizabeth Finlayson Harris, Christine Bywater, Sarah Levine, and Dorottya Demszky. 2025. "Depicting Data Conversations for ELA Teachers through Routine Supports." Paper presented at Data Science Education K-12, San Antonio, Texas, United States. Data Science Education K-12 Proceedings, February.

Liam Lynch, Tom. 2019. "Electrical Evocations: Computer Science, the Teaching of Literature, and the Future of English Education." *English Education* 52 (1): 15–37. <https://doi.org/10.58680/ee201930312>.

McLean Davies, Larissa, Katherine Bode, Susan K. Martin, and Wayne Sawyer. 2020. "Reading in the (Post)Digital Age: Large Databases and the Future of Literature in Secondary English Classrooms." *English in Education* 54 (3): 299–315. <https://doi.org/10.1080/04250494.2020.1790991>.

Newman, Beatrice Mendez, and Penny Rosas. 2016. "Opening the Door for Cross-Disciplinary Literacy: Doing History and Writing in a High School to University Collaboration." *English Journal* 106 (2): 54–

61. <https://doi.org/10.58680/ej201628827>.

Oeldorf-Hirsch, Anne, and German Neubaum. 2025. “What Do We Know about Algorithmic Literacy? The Status Quo and a Research Agenda for a Growing Field.” *New Media & Society* 27 (2): 681–701. <https://doi.org/10.1177/14614448231182662>.

“Plots | Plotting Plots.” 2021. April 11, 2021. <https://plottingplots.com/plots/>.

Pfannkuch, Maxine, and Chris J. Wild. 2000. “Statistical Thinking and Statistical Practice: Themes Gleaned from Professional Statisticians.” *Statistical Science* 15 (2). <https://doi.org/10.1214/ss/1009212754>.

Ringler, Hannah. 2024. “Computation and Hermeneutics: Why We Still Need Interpretation to Be by (Computational) Humanists.” In *Computational Humanities*, edited by Lauren Tilton, David Mimno, and Jessica Marie Johnson. Debates in the Digital Humanities. University of Minnesota Press.

Risam, Roopika, and Alex Gil. 2022. “Introduction: The Questions of Minimal Computing.” *Digital Humanities Quarterly* 16 (2). <https://www.digitalhumanities.org/dhq/vol/16/2/000646/000646.html>.

Seaver, Nick. 2019. “Knowing Algorithms.” In *DigitalSTS: A Field Guide for Science & Technology Studies*, edited by Janet Vertesi, David Ribes, Carl DiSalvo, et al. Princeton University Press. <https://doi.org/10.1515/9780691190600>.

Sims, Porsche L. 2010. “Success with Ells: Using Your State’s Travel Websites to Promote Academic Vocabulary.” *English Journal*. <https://doi.org/10.58680/ej201011706>.

Spiro, Lisa. 2012. “‘This Is Why We Fight’: Defining the Values of the Digital Humanities.” In *Debates in the Digital Humanities*, edited by Matthew K. Gold.

Turner, Kristen Hawley, and Troy Hicks. 2022. “Digital Literacy (Still) Can’t Wait: Renewing and Reframing the Conversation.” *English Journal* 112 (1): 86–93. <https://doi.org/10.58680/ej202232072>.

Wiese, Marion, Kamila Serwa, Anastasia Besier, Ariane S. Marion-Jetten, and Eva Bittner. 2025. “Establishing Technical Debt Management: A Five-Step Workshop Approach and an Action Research Study.” arXiv:2508.15570. Preprint, arXiv, August 21. <https://doi.org/10.48550/arXiv.2508.15570>.

## About the Authors

**Nichole Nomura** is an Assistant Professor of Public Humanities in the Department of English at the University of Wyoming. She researches digital humanities pedagogy and how literature teaches/is taught, using methods from the digital humanities, literary criticism, and the educational social sciences. She holds a doctorate in English with a certificate in Digital Humanities from Stanford University, and a MA in Education from the Stanford Graduate School of Education.

**Georgii Korotkov** is a PhD Candidate in the Department of Slavic Languages and Literatures at Stanford University, with a minor in Computer Science. His most recent digital project is <https://vsesvit.vercel.app>. His latest publication, co-authored with Elena Ostrovskaya, Elena Zemskova, and Evgeniia Belskaia, “International Literature: A Multi-Language Soviet Journal as a Model of ‘World Literature’ of the Mid-1930s USSR,” appears in a collection that received the Best Edited Multi-Author Scholarly Volume Award from AATSEEL.

**Sarah Levine** is a teacher educator at Stanford University’s Graduate School of Education. She helps teachers use artificial intelligence (AI) to support students’ reading and writing skills. Before pursuing an academic career, Sarah taught secondary English at a Chicago public school for ten years. While there, she founded and ran a youth radio program that used digital audio production as a tool to help make writing and analysis relevant and real-world for students, and to build bridges between school and the world outside.

**Christine Wynn Bywater** is an Associate Director of the Center to Support Excellence in Teaching at Stanford’s Graduate School of Education where she co-leads the center’s mission and vision while spearheading initiatives in Digital Pedagogy. Current work focuses on developing co-designed curriculum resources to enhance AI literacy in teachers and students. Prior to CSET, Christine was a high school social studies teacher and a K-8 technology instructional coach. Christine holds a MA in Educational Leadership from New York University and a MS in Educational Technology & Media Leadership from CSU Long Beach. Christine also holds a Bachelor of Science in Secondary Social Studies Education with a minor in Psychology from New York University’s Steinhardt School of Education.

**Victor R. Lee** is an Associate Professor at the Stanford Graduate School of Education. He conducts research on learning, teaching, and technology particularly for K-12 settings. Current projects emphasize K-12 data literacy and data science education and the development of AI literacy in schools through partnership research with teachers and teacher educators. Lee obtained his doctorate in Learning Sciences at Northwestern University.



This entry is licensed under a Creative Commons [Attribution-NonCommercial-ShareAlike 4.0 International](#) license.

ARTICLES