

A Taxonomy for Writing Analytics

Susan Lang, *The Ohio State University*

Laura Aull, *University of Michigan*

William Marcellino, *The RAND Corporation*



Abstract

This article proposes a functional taxonomy for the growing research specialization¹ of writing analytics (WA). Building on prior efforts to taxonomize research areas in WA and learning analytics, this taxonomy aims to scaffold a coherent and relevant WA research agenda, including a commitment to reflection, evidence-based propositions, and multidisciplinary as the research specialization evolves. To this end, the article offers a conceptual and practical overview of WA in the following sections: history, theorization, implementation paradigms, data, digital environments, analytic processes and uses, assessment, ethical considerations, and ongoing challenges. This overview highlights current limitations and needed WA research as well as valuable opportunities for the future of WA.

Keywords: assessment, data, ethics, programs of research, research principles, writing analytics

1.0 Introduction

In this article, we propose a functional taxonomy of writing analytics (WA), building on prior foundational work taxonomizing the nascent research lines in writing analytics. In “Writing Analytics: Conceptualization of a Multidisciplinary Field,” Moxley et al. (2017) categorized four potential programs of research in writing analytics:

¹ After much consideration, the authors chose to use the term “research specialization” to describe writing analytics. Our goal in using such a term is to avoid the finite label of *field* and thereby acknowledge the multiple points of entry for researchers across disciplinary traditions.

- **Educational Measurements.** Advances in educational measurement and the availability of large-scale assessment data have opened up new possibilities in measuring student competencies and skills in writing.
- **Massive Data Analysis.** Ever-increasing sources and volumes of data and the emergence of increasingly inexpensive and powerful computing power and storage mean that truly scalable data analysis is possible for most research communities.
- **Digital Learning Ecologies.** The emergence of specialized digital ecologies presents new opportunities to analyze a range of writing processes.
- **Ethical Philosophy.** Writing analytics as an enterprise presents a huge range of new questions about fair and just use of data in writing analytics, including concerns about bias in analytics, misuse of data, privacy concerns, and ownership of student data/student rights.

While there has been significant research at the intersections of these fields written under the label of *writing analytics*, prior efforts (e.g., Moxley & Walkup, 2016) to functionally organize the intersections into a more cohesive area of research have focused on identifying and tentatively categorizing fundamental components of the research specialization rather than creating a purposeful mapping that identifies both established and nascent research questions. The use of taxonomies and typologies has been the subject of some criticism, primarily from epistemological privileging of natural sciences and quantitative methods, and ideologies of scientism (Collier, LaPorte, & Seawright, 2012; Marradi, 1990). However, these criticisms fail to account for the analytic power of classification schemes, their potential rigor, and the potential new insights they can offer (Collier, LaPorte, & Seawright, 2012). They also fail to appreciate and articulate the interplay between quantitative and qualitative interpretive methods to make sense of and refine qualitative structures. We see great value in building a taxonomy, even an exploratory one, for WA practice areas and functions.

We chose a taxonomic approach for this effort over typology-building because our aim is clarity around different classes of real-world practices associated with writing analytics. Typologies are conceptual in nature, based on mental models of ideal types (Smith, 2002). Thus, a typological approach might impose preconceived structures on WA rather than take an open look at existing practices in WA. Typologies can serve important heuristic functions around conceptual clarity, but this is less applicable to the nuts and bolts concerns we have over practice areas such as data storage and choosing analytic software. Taxonomies are also marked by successive application of dimensions (Marradi, 1990). So for example, we do not simply divide horizontally between *data* as a practice area and the various *implementation paradigms* in WA. Rather, we look at fundamental divisions vertically within those practice areas.

Our contribution to this effort is to offer a functional taxonomy that, by articulating the intersections, maps out current and potential WA practice areas and considerations. Our taxonomic sketch adds to the programs of research noted in the inaugural issue of *The Journal of Writing Analytics* by outlining the practice areas of both applying WA (e.g., gathering data,

assessment) and accounting for context (privacy and data safety, implementation paradigms). From this taxonomy of practice areas, we then outline known challenges for the research specialization.

This effort is in some ways inspired by efforts to map out practice areas in learning analytics (LA), specifically the taxonomy proposed by Peña-Ayala (2018), which contains research profiles in LA, applications of LA, and underlying factors (context) for LA. Assisting us in our task to map the terrain is that, even in its short history, the research specialization now known as WA has assumed as its subject matter writing produced in a Vygotskian sense, that is, writing as a cognitive act influenced by intrapersonal, interpersonal, and neurological factors. Additionally, the rise of tools enabling corpus analysis of large data sets has enabled WA projects to flourish.

The following figure provides a working map of WA, illustrating the relationship between the four programs of research articulated in Moxley et al. (2017), current foci in WA publications and presentations², and ongoing challenges for future WA work. The latter areas of current and future work in WA roughly correspond to the sections of our discussion below. As the figure shows, some areas of WA can be mapped more fully and neatly than others, but our hope is that continued efforts like this article will help bolster coherence across the many areas and possibilities of WA.

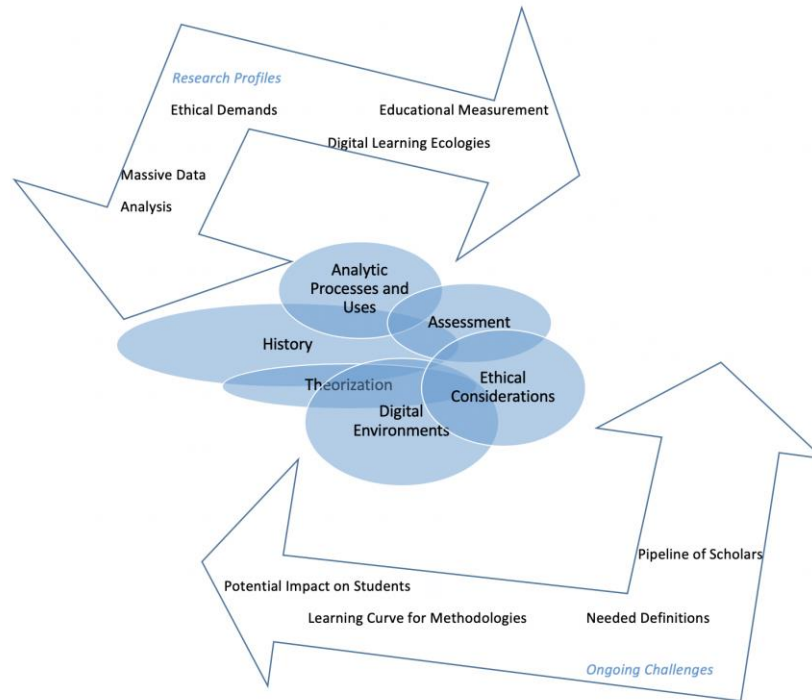


Figure 1. Writing analytics taxonomy.

² In constructing this taxonomy, we examined proceedings from such conferences as *International Conference on Writing Analytics*, *ACM*, and *LAK*; additionally, we ran broad searches for articles using keywords related to WA (many of which are listed above).

Writing analytics is a nascent entity with great promise but lacking the structures and interconnections of more mature fields. One of the challenges of creating this functional taxonomy for WA is in determining what untapped areas of research can and should exist alongside more frequently exercised research subjects. Much of the current research that labels itself as belonging to WA has come about as a chain of fortunate events—a few individuals, scholars, researchers, or administrators have been able to create ways to collect data concerning the phases of writing instruction and have either bought or built tools to analyze this data. Other scholars who would perhaps like to be part of this emerging specialization are seeking ways to become involved. The following pages provide a map of sorts for doing so, and a number of areas we've included in the taxonomy fall more into those for future research.

Overall, we hope that by offering a clear description of current applications of WA and the near-term challenges in the research specialization, this article will function as a scaffold for a coherent and relevant research agenda moving forward.

2.0 History

Writing analytics, as a discrete, named research specialization of inquiry, has a brief history, though its predecessors might be traced back to mid-20th-century computational linguistics and early corpus-based efforts that were both manual and computer-aided (see, e.g., Nelson, 1965; Nelson & Kučera, 1967; Wright, 1974). We could also view much of the activity of the research specialization as part of a response to Rich Haswell's (2005) call for more research in writing studies that is replicable, aggregable, and data-driven. WA research also provides a platform for enacting Driscoll's (2009) "skeptical view of empirical research," one that

does not change the act of research itself, but rather changes how one views research and what one assumes about it. A skeptical view of research includes the following: self-reflection and refutation, adherence to evidence-based propositions, proof not truth, and finally the acceptance of multiple methods of data. (p. 200)

In fact, much of the research done in the name of WA adheres to three of the four points of Driscoll's framework:

- Skeptical researchers should be skeptical of everything, including their own findings.
- Empirical research never claims to prove, but rather provides evidence.
- Empirical research does not build itself upon that which has been assumed, but rather that which has evidence.
- Empirical researchers are interested in gathering evidence from as many sources as possible and in drawing on the sources and data collection methods most valid and ethical for a given query.

The last point is perhaps most complex, as WA research is built on the idea of a textual corpus—but the methods that researchers use to study those corpora vary. And while discerning best

practices for valid and ethical inquiry remains a work in progress for every study, those working in WA will strive to produce research they see as ethical, valid, and reliable.

So while writing analytics can be seen as closely allied to and/or emerging from such areas of study as computers and writing, corpus linguistics, digital humanities, writing program administration, learning analytics, educational data mining, and technical communication, the specialization's own history as a distinct grouping begins in approximately 2015 or 2016. Shum et al. (2016) provide an initial definition of the term:

Broadly defined, writing analytics involves the measurement and analysis of written texts for the purpose of understanding writing processes and products, in their educational contexts. Writing analytics are ultimately aimed at improving the educational contexts in which writing is most prominent.” (p. 481)

One might consider that the specialization, from the outset, equally invokes methodological processes and the theory and content of writing instruction. In that Shum's definition risks privileging academic discourse, it begs important questions: What about the many important kinds of writing that occur outside of the educational context—in the workplace, in social media, and in other contexts—both on their own terms and vis-à-vis their impact on writing in educational contexts? Clearly, analytics has an application beyond the educational institution, although the research published thus far overwhelmingly resides in educational settings. As WA continues to develop, it will benefit from investigations of a range of writing contexts and genres, and it can be a source of information about the value and systematicity of many kinds of written discourse.

3.0 Theorization

To this point, no work has been published in which researchers have attempted to articulate an overarching theory of writing analytics. What binds the research specialization together are common methodologies, common exigencies, and common assumptions about the value of empirical research in humanistic areas where such work has not been historically common. Like fields such as technical communication and digital humanities, writing analytics emerged from the application of methodologies from fields such as corpus linguistics and computer science to fields such as composition studies, writing program administration, and assessment. A study identified as one in writing analytics generally requires 1) a corpus of texts of sufficient quantity to enable generalization inferences relevant to a given study, 2) one or more exigencies informed by programs of research (e.g., to add to or respond to prior research, or to answer a locally developed question), and 3) a particular set of research questions designed to make specific use of empirical techniques allowing inferences about situated language use—that is, inferences attentive to the interplay among individual cognitive processes, social practices, and larger linguistic and cultural patterns (Mislevy, 2018). However, one could be using theoretical lenses as diverse as post-colonial, deconstructionist, reader-response, process, or post-process to guide said sociocognitive and sociocultural examination of the data. A unifying thread regardless of

WA methods, inferences, and lenses includes thoughtful consideration of key premises discussed in our analytic processes and uses section below: that writing makes meaning in patterns and outliers across contexts and/or texts, the aggregated analysis of which facilitates unique insights.

Indeed, theorization efforts pose questions about what it means to develop a theory of a field or research specialization fundamentally concerned with application and analysis. Perhaps Carolyn Rude's (2009) "Mapping the Research Questions in Technical Communication" provides us with some guidance, as she explains how "[a]greement about research questions can strengthen disciplinary identity and give direction to a field that is still maturing" (p. 174). Rude's series of Venn diagrams that in turn examine research questions, books, and topics from tc.eserver.org over two decades assist in illuminating potential research directions for the field—not a theory of technical communication, but a research articulation. While the value of analysis and measurement is often contingent on situational context, determining a common set of core principles for conducting such research and labeling it part of writing analytics is within reach. To this end, we devote the following sections to considerations related to data, digital environments, analysis, assessment, and ethics.

4.0 Implementation Paradigms

Institutional adoption of writing analytics will require an "implementation paradigm"—a model that guides deployment of analytic efforts (Colvin, Dawson, Wade, & Gašević, 2017). Two important considerations for implementing a WA effort are the scale of the efforts and the technical capacity available for implementation. Smaller and larger efforts will make different choices about implementation, as programs/departments that include or work with computer scientists will make different choices than ones lacking organic programming and data analytics skills. We can thus visualize four broad general implementation paradigms, along two axes: technical capacity and scale of effort. In the figure below, the X-axis scales from less to more capacity in software coding and data science, while the Y-axis charts increasingly larger writing analytics efforts (e.g., a small effort within a department up to an enterprise effort for a university):

Writing Analytic Paradigms by Scale and Technical Capacity

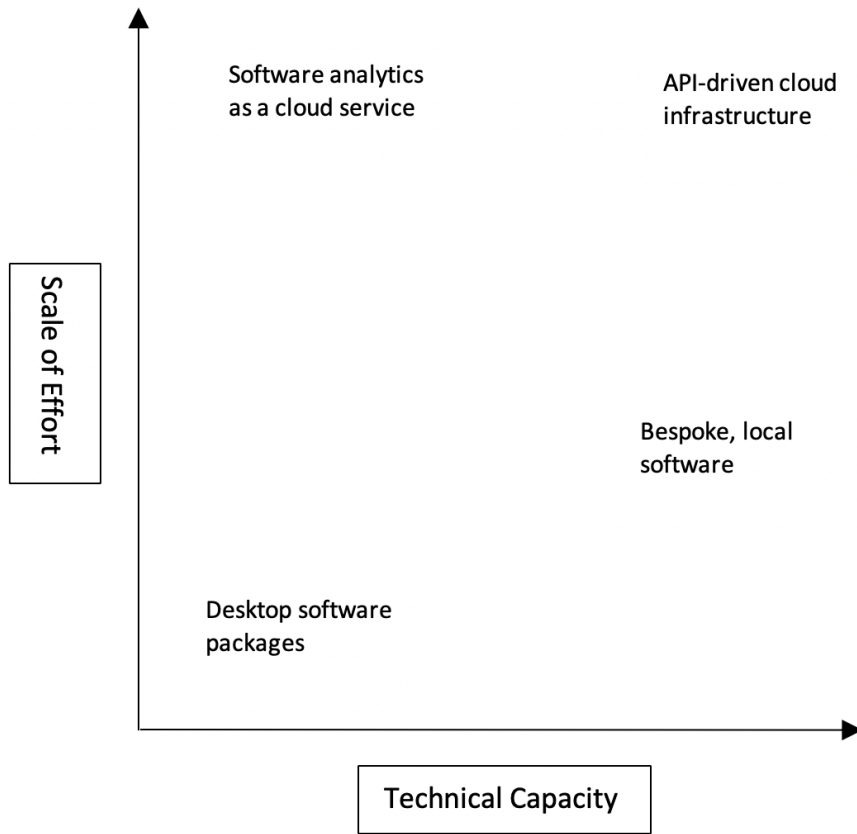


Figure 2. Writing analytics paradigms by scale and technical capacity.

4.1 Desktop Software Packages

We imagine a smaller effort in a traditional English department that offers a 101 general education writing course. In our scenario, the WA program aims to measure the genre performance gain between students’ diagnostic and final assessed argument paper, and the faculty in the program are firmly rooted in the humanities. Because they may not be equipped to write computer code or use automated programming interfaces (APIs), they will need to use existing traditional software with graphical user interfaces (GUIs). And because of the relatively small dataset, desktop PCs or laptops will have adequate computing power and storage for their WA effort. This paradigm is a scaffolded, relatively low-cost entry into WA. Software such as Tool for the Automatic Analysis of Lexical Sophistication (TAALES; Crossley & Kyle, 2018), is an example of a powerful word-level analytic program that could be implemented in a traditional

humanities-focused writing program with a low level of effort and without the need for specialized technical expertise.³

4.2 Software Analytics as a Cloud Service

An ambitious, large-scale effort requires significant computing power and data storage, both because of larger datasets and more resource-intensive capabilities (such as machine learning). We can imagine an enterprise-level effort across a large university to analyze outcomes for writing across the curriculum. The effort will be administered by both the English department and university writing center, and so they will still need user-friendly software with a GUI to conduct their analyses. However, local desktop/laptop computing using traditional software will not be able to handle this volume of data, and part of the effort might include piloting machine learning to support assessment. In such a case, a cloud-based software suite would offer the requisite power, storage, speed, and capabilities at a cost-effective price. RAND-Lex (<https://www.textinsight.org>) is an example of a scalable, user-friendly, cloud-based analytic suite powered by Amazon Web Services (AWS).

4.3 Bespoke, Local Software

At the other end of the technical capacity scale, we can imagine writing programs and English departments that include faculty with computer science and analytics skills who can create applications and tools that are scalable across one or more programs. The rise of digital humanities and data science broadly means more opportunities for building bespoke analytic systems, and the public availability of *libraries* (wide varieties of pre-written, community-shared code modules) for the Python and R programming languages can further leverage these efforts. Cross-department collaboration is also a possibility, where domain experts in writing instruction and computer science might work together to craft specific WA solutions for local use. This increased technical capacity may allow for more scale than desktop systems, for example by using distributed computing.⁴ DocuSope Classroom (Helberg et al., 2018) is an example of a very powerful style-level analytic tool that was first developed locally within the Carnegie Mellon English Department by faculty with rhetorical and coding expertise and is now poised to move to cloud deployment.

4.4 API-Driven Cloud Infrastructure

At the far ends of both scale and technical capacity would be WA efforts that borrow current business intelligence/business analytics (BI/BA) paradigms. Like enterprise-level BI/BA, an enterprise-level WA at a large university would require powerful computing and storage infrastructure, as well as analytics and machine learning to leverage the potential insights from

³ For reviews of several recent corpus-based programs, see the Tools and Tech Forum in *Assessing Writing*, 38 (October 2018).

⁴ Distributed computing frameworks such as Apache Spark or Hadoop allow many networked computers to act as one very large, efficient computer.

very large datasets. In terms of scale, this paradigm is like the analytics as a cloud service one, but because of much greater technical capacity, this effort will not need user-friendly software. Instead, researchers will write code that lets them directly hook onto the powerful APIs offered by many cloud computing vendors (APIs are software intermediaries that allow coders to connect to and automate software functions). A university-wide program that, for example, used all student writing to understand writing across the curriculum, combined with wider LA datasets, would benefit from this kind of enterprise-level infrastructure. Vendors such as Amazon and Microsoft have built highly efficient and scalable on-demand storage and computing services that can run software analytics at this level, using APIs. This is a different paradigm than using traditional software run through a graphical interface, potentially very powerful, but requiring coding skills.

5.0 Data

Core principles for establishing a writing analytics effort must include considerations for acquiring, preparing, storing, and accessing data. We address these in this section with some examples associated with data compilation and use.

5.1 Data Selection and Sources

Data acquisition starts with choices about what data sources are most relevant to a given analytics program, but also with choices about how that data is meaningfully partitioned into analytic categories. For example, an effort to analyze student writing for genre features might seem like a straightforward comparison of a corpus of Composition 101-level student papers against a gold-standard genre corpus for feature match. However, there are a number of choices possible in this example:

- If the program has a large number of non-native English speakers, should the student corpus be split demographically to look for L1 and L2 learner specifics? Are there other potential demographic splits that might help produce more granular insights into student writing performance?
- Should the student corpus be split in other, non-demographic ways? What insights might a comparison of first semester vs. second semester student writing produce?
- For the example genre performance analysis, what constitutes the appropriate benchmark? One choice might be a curated corpus of “A” graded papers within the program, a corpus that likely emphasizes a priori knowledge from instructors, embedded in grading rubrics. Such a corpus likely reflects explicit genre knowledge from instructors, but potentially de-emphasizes tacit, performative genre knowledge. By contrast, a curated, external corpus of genre or disciplinary examples might provide a better contrastive background for less-explicit/poetic genre features.

These examples highlight that what data to analyze, how that data is partitioned, and what data is chosen as comparative or baseline writing are all high-stakes choices as research commences. One additional, closely related task of equal importance is the ability to acquire institutional buy-in for these choices. All of these choices should accordingly be treated with the same consideration and transparency of data selection choices.

In academic settings, core sources for data likely include student writing, instructor feedback, and peer feedback. Depending on program size and variation within the program, other potential corpora might include syllabi and assignments/writing tasks from across a program. For comparative or baseline corpora, either ad hoc curation or existing public corpora may be used. On the one hand, it may be optimal to design a very specific kind of corpus, for example a gold standard corpus of disciplinary examples curated from professional journals. There are, in addition, publicly available existing corpora that may be useful. Two such examples for U.S. writing analytics efforts include the Michigan Corpus of Upper-Level Student Papers (MICUSP; <http://micusp.elicorpora.info>) and the Corpus and Repository of Writing (CROW; <https://writecrow.org/>). MICUSP contains senior undergraduate and graduate-level student papers, and can be sorted and filtered by grade-level, English nativeness, textual features (e.g., abstracts, literature reviews), genre, and discipline. CROW is a web-based archive that includes a first-year student writing corpus along with pedagogical materials and related research conducted by researchers at Purdue University, the University of Arizona, and Michigan State University. Finally, the University of Michigan Press has recently published a collection entitled *Developing Writers* and has made available the writing, interview, and student demographic data used in the studies featured in the collection (<https://www.developingwritersbook.org/pages/about/about-the-data/>).

5.2 Data Preparation

Data preparation will vary by the software/platform used for analytics, but it may include considerations related to file format, data schema, chunking, and segmenting.

- File format. Many popular text analysis tools, such AntConc or Wordsmith, require or do best with plain text files (*.txt files) in UTF-8 format. Other text formats, for example Microsoft Word, have text encodings that look good on screen (e.g., smart quotes) but come out garbled during analysis. Additionally, some text analysis suites can ingest PDFs, but likely use OCR technology that handles only English or English+Romance language texts. Finally, some software can use CSV files to upload large datasets that include some kind of metadata and data schema to enable machine learning or support ways to slice/analyze data.
- Data schema. Some applications for writing analytics will require a data schema, for example a *.csv file or Excel spreadsheet with column headings, such as ROW_ID (a number for each row/document), TEXT (the text/document to be analyzed), CATEGORY (demographic or source variable such as grade level,

gender, or section), and CODE (a label for data to train a machine learning algorithm).

- **Chunking.** Some analytic tests work best using whole documents, while other tests might require documents to be chopped up into uniform-length chunks. An example might be exploratory factor analysis to detect latencies in writing, such as argumentative strategies or genre moves. Thus, practitioners may need either software that includes chunking as a feature or programming support for prep.
- **Segmenting.** Similar to chunking, it may be useful to break apart functionally distinct segments of documents. If, for example, a given writing analytics project requires segmenting argumentative essays into introductory, main body, and concluding segments, that will need to be planned for as students submit documents, or post-hoc applied to documents as a preprocessing step.

5.3 Data Storage

The major choice in data storage is local vs. remote (i.e., cloud storage). Both offer different affordances and constraints, dependent on the size of the effort, implementation paradigm, and concerns around data security.

- **Local storage** likely makes more sense for smaller efforts, using locally installed desktop computer software. For example, a writing center director or writing program faculty might install one or more freeware and commercial analysis software packages on their computer and store data locally on their machine. This kind of storage has potential benefits: more control over the data because it is in hand, controlling access (which may align with Institutional Review Board or university data safeguarding requirements), and making the data custodian explicitly responsible for data backups.
- **Cloud storage** will be a requirement for any program that uses cloud-based analysis services. In place of a traditional software model where a person or institution buys one or more copies of a software product and installs them locally, software may be offered as a service, hosted remotely by large providers such as Microsoft or Amazon, accessed through a browser. To use a cloud-based analytic service means all of the effort for data protection and backup is included, and for enterprise-level efforts, the sheer data volume may make cloud-based storage relatively cheap and attractive.

6.0 Digital Environments

Writing analytics thrives in and depends upon digital environments, but researchers are only now beginning to address the range of issues emerging from research in these contexts. Digital writing environments have contributed clearly to writing analytics research as primary sites of much writing instruction, but until recently, most had not been designed with data extraction for research in mind. Early word processing applications, as well as most early digital writing

applications, such as HyperCard, StorySpace, and the Daedalus Integrated Writing Environment, lacked any features that we would recognize as useful for large-scale data collection and analysis—instead, their focus was on enabling a variety of types of writing, broadly defined, but that writing was generally kept on the author’s computer or shared in a very limited fashion. Even current Learning Management Systems, used by an increasing number of post-secondary writing programs, such as D2L, BlackBoard, and Canvas, do not emphasize the ability of these applications to assist in research about writing instruction at the level of extracted student text, instructor comments, or other key data points. And the application with undoubtedly the largest dataset of writing, Turnitin, focuses its attention on “preventing multiple forms of plagiarism, grading assignments, and safeguarding your institution’s reputation.” Research into other aspects of the instructional process does not seem a priority, even as this digital environment hosts huge corpora of student writing.

Over the last two decades, dedicated writing environments (Learning or Program Management Systems) developed by specific institutions have generated opportunities for more research, since they became repositories of student-generated texts, instructor feedback, and other associated data points. A sampling of these applications includes the Writers Exchange (Ohio State), MyReviewers (University of South Florida), Raider Writer (Texas Tech University), and ELI Review (Michigan State University).

Sustainability, however, remains a challenge for such projects, as ELI Review is the only one of the aforementioned applications still available for use; Writers Exchange and Raider Writer were retired without replacement, and USF launched a new platform, USF Writes, in Fall 2019, rendering MyReviewers a legacy platform in terms of departmental use. At this point, developing and sustaining useful data repositories in the context of a writing program remains a challenge to writing analytics researchers. Given this challenge, though, it is advisable to consider what such a data repository might contain in order to facilitate research endeavors. Larger-scale efforts for dedicated writing environments require cheap, fast, reliable long-term data storage. This may mean choosing between university computing infrastructure and commercial enterprise solutions such as Amazon Web Services or Microsoft Azure. These choices will likely entail the best balance between local needs and available data infrastructure and personnel (e.g., populations in need of support alongside affordable data extraction methods), but thinking through priorities before designing digital environments can ensure these choices are informed and proactive.

7.0 Analytic Processes and Uses

This section addresses analytic processes that aid discovering, understanding, and sharing meaningful patterns in written data, including how writing analytics findings are used for descriptive, predictive, or prescriptive ends.

7.1 Premises Underpinning Writing Analytics Processes

Before describing processes and uses, some premises underpinning writing analytics processes and uses bear mention upfront. Two premises of any analytic process in writing analytics include one, that writing makes meaning across contexts and/or texts, and two, that aggregated analysis of related patterns and outliers facilitates unique insights. These insights need not be at the expense of analytic processes like close reading that analyze one word, sentence, or text at a time, but they provide a different way of viewing texts, according to those patterns or distinctions that occur across them. Three additional, overlapping premises relate to both analysis and assessment.

The third premise relates to generalization inferences. Inferences that go beyond analyzed writing require clear evidence, and the more ambitious the inferences, the more evidence is required (Kane, 2013). Even as writing analytic tools make it possible to analyze increasingly-large corpora, the scope of generalization inferences must attend to the scope of the writing analyzed. Many interpretations, especially evaluations of error, involve generalizations over some conditions of observation, such as across tasks, occasions, raters, and contexts—any interpretation of a student’s writing proficiency based on a specific assessment task, for example, generalizes across these. In addition, many uses of writing analytic findings involve claims that go beyond generalizations (e.g., about relationships to non-test performance, such as that a student needs more or less writing support), and these claims also require other kinds of evidence (e.g., correlations with criterion measures). The validity of such inferences cannot be taken for granted but rather must be rooted in relevant, credible evidence with appropriate inferential force (Kane, 2013, p. 7; see also Cronbach & Meehl, 1955).

A fourth premise relates to probabilistic reasoning, or reasoning tasks based on evidence. Uncertainty plays a predominant role in any reasoning based on evidence; as Schum writes, “in any inference task, our evidence is always incomplete” and “rarely conclusive” (1994, p. xiii). Probabilistic assessment of evidence therefore requires an argument structured in terms of relevance and credibility and evidence characterized by an appropriate inferential force. In writing analytics cases of large masses of data, this premise places added value on knowledge of basic inference structures that support contextualized evidence interpretation.

The fifth premise is that analytic and assessment processes are epistemic choices: regardless of whether our goals are descriptive, predictive, or prescriptive, what we analyze influences what we conclude; it makes certain kinds of inferences possible and not others. Those inferences, in turn, influence what we notice and understand about collective and individual processes and practices related to language, communication, and learning. As writing analytics research grows, study design choices and attendant inferences therefore have a crucial role to play. Student writing, focused on the measurement of text features to better understand writing within education contexts, is the most common target of writing analytics and is a particularly important example: student writing is understood according to how it is assessed and analyzed, and so we need as much information as possible as we design and choose the writing analytic approaches that form our interpretative frameworks. Thus the “uses” described below could be labeled

“conceptualizations” for the ways that they take up writing analytics inferences in ways that shape how we understand those inferences.

7.2 Uses of Writing Analytics Inferences

Within writing analytics, the three categories of descriptive, predictive, and prescriptive uses are useful for illustrating how writing analytics inferences can be used and the epistemological implications for such uses. Importantly, we say “uses” here rather than processes, as these are not separate categories: arguably, any writing analytics data could be used in all three ways.

- *Descriptive use:* Writing analytics inferences are used to describe salient or characteristic features of types of writing, such as differences between multiple corpora representing writing fields, genres, or levels.
- *Predictive use:* Writing analytics inferences are used to predict performance beyond a given task, such as the use of a writing placement task to predict performance in a writing course.
- *Prescriptive use:* Writing analytics inferences are understood and used as prescriptions, or rules, for language use, such as for designing rubrics for writing tasks or for designing curriculum and instruction intended to help learners emulate and/or change features discovered through writing analytics research.

Examples from recent articles published in *The Journal of Writing Analytics* will help illustrate a range of writing analytics approaches and how they can be used for descriptive, predictive, or prescriptive ends. These articles focus on disciplinary differences, genre-based differences, and level-based differences.

For instance, Crossley, Russell, Kyle, and Römer (2017) studied the linguistic differences between science and engineering macro-disciplines and micro-disciplines and found significant distinctions in both lexical and cohesive strategies used by upper-level student writers in the Michigan Corpus of Upper-Level Student Papers (MICUSP). Using the Tool for the Automatic Assessment of Lexical Sophistication (TAALES; Kyle & Crossley, 2015), they showed, for instance, that engineering writing samples contain more frequent words that occur in a greater number of text categories than science texts, while simultaneously containing a greater number of academic words and academic formulas. These findings resonant with inferences drawn from multidimensional analysis by Hardy and Römer (2013), a writing analytics approach initiated by Biber (1988) that analyzes patterns of co-occurring lexico-grammatical features that show systematic differences according to spoken and written registers, fields, and genres. Multidimensional analysis is especially used for descriptive ends; as Hardy and Römer write,

It is not our intention to prescribe how students should write in these disciplines, because we recognise that there is a thin line between description and prescription in this context. One of the benefits of [mutidimensional analysis], however, is that it deals with continua and related dichotomies. (2013, p.184)

As previous writing analytic studies of disciplinary differences have focused on lexical differences, Crossley, Kyle, and McNamara's (2016) study extends them by investigating cohesive strategies using the Tool for the Automatic Analysis of Cohesion (TAACO). They found, for instance, that engineering texts contain more positive connectives and pronouns than science texts, while science texts contain more contrastive connectives and demonstratives. These inferences are presented in a *descriptive* way, in that they show comparative distinctions across fields without aiming to evaluate or prescribe said uses. Noah Arthurs' (2018) study in this journal as well, which instead focuses on expressions of stance and syntactic complexity (though the latter diverges from Biber's and Biber & Gray's definition), likewise concludes that disciplinary background as well as the disciplines in which students write influences student writing choices, and he recommends that the analysis of stance and syntactic features can provide "insights when looking at student writing in aggregate" (p. 172). In Crossley et al. (2017), the first pedagogical application recommended by the authors in the conclusion is likewise descriptive: they suggest that students and instructors can use a "contrastive rhetoric approach based on the reported findings" (p. 74) so that students would have the opportunity to understand or even examine differences between macro- and micro-discipline. This is echoed in Aull's (2017) study of genre-based differences in student writing in the same journal issue, which she suggests can be put to descriptive use, such that "identifying patterned discourse in academic writing can help make writing expectations and ontological orientations more transparent" (p. 9).

Crossley, Russell, Kyle, and Römer's (2017) second recommendation based on their writing analytics findings illustrates a more prescriptive use, our second of the three categories. Crossley et al. write that "teachers could provide students with discipline-specific writing guidelines for producing text samples that fit discipline expectations" (p. 74). This use positions the discipline-specific, lexical, and cohesive writing patterns they identify in MICUSP as target features for instruction, and presumably by extension, for assessment.

Finally, let us extend this example once more to our third category of predictive use. If we were to use Crossley et al.'s inferences in a study of high- and low-graded discipline-specific texts in order to determine whether the use of discipline-specific cohesive features correlated to higher grades in respective courses, we would be examining the predictive validity of these writing analytics inferences; then using students' use or non-use of highly-predictive features to place students into certain groups would apply this predictive use. Indeed, related studies have shown that the TAALES and TAACO tools have been able to predict expert judgments of essay quality and text coherence (see, e.g., Crossley et al., 2016; Kumar, Fraser, & Boulanger, 2017) and even student math performance (Crossley, Liu, & McNamara, 2017).

In a second example, Thomas Peele's (2018) study of first-year writing at City College of New York (CCNY) focuses not on discipline but on student level. Peele uses corpus linguistic analysis to analyze linguistic resources for counterarguments used in CCNY first-year writers' "engagement with opposing points of view," in comparison with Zak Lancaster's (2016) study of the same linguistic resources used by students at the University of Michigan and Wake Forest

University and in published academic writing in the Corpus of Contemporary American English. Peele finds, for instance, that the CCNY students are about equally likely to introduce objections to their arguments, but they are less likely to concede to those objections. Like many similar corpus studies, Peele frames the findings in terms of how to move the developing CCNY writers toward the writing of the reference corpora; for instance, he writes, “CCNY students demonstrate an unwillingness to use, or a lack of understanding of the necessary rhetorical moves, or a lack of know-how, or techné, of how to use them” (p. 89) and that these findings can be used in faculty development seminars. This is a prescriptive use of these findings, in that it suggests that the reference corpora offer prescriptions that students should learn.

Finally, some overall points. Descriptive use of language findings is as old as language and lexicography, but it is not the most common use in educational contexts. Prescriptive approaches, in which student language use is examined according to how it matches or diverges from conventional expectations, are far more common. Thus, as writing analytic approaches continue to grow, it will be valuable to stay vigilant about how we describe patterns conventionally valued in standardized written academic language without devaluing other patterns and how we ensure that multiple constructs of writing inform both writing analytics inferences and uses thereof. One powerful, descriptive use noted by Peele can be engaging students themselves in writing analytics research, so that they engage with “the enormous volume of work that has been generated” rather than writing “as if they were the first students on our campus to write” (p. 90). Thus, students’ own writing and that of their peers become objects of inquiry rather than only top-down assessment, and resulting descriptive knowledge may offer additional reference for students’ peer review and writing development.

This final point brings us back to premises regarding valid inferences noted at the start of the section: reasoning based on evidence always entails the “network of inferences and assumptions inherent in the proposed interpretation and use” (Kane, 2013, p. 2). These are never certain. Poe and Elliot (2019) connect this to the importance of the interpretation and use argument (IUA) advanced by Kane by foregrounding the need for ongoing recognition of the contingencies involved in validation processes. The concept of IUA supports iterative, evidence-driven writing analytic inferences with acknowledgment of their uncertain—or probabilistic—nature.

8.0 Assessment

A crucial consideration for both writing analytics uses addressed in the previous section and writing assessment design addressed in this section is the epistemic role of related choices. Just like choices about writing analytics study design and use, the design and use of writing assessments constitute what we (can) know from them. Each choice shapes what students demonstrate in their writing and what is concluded about their needs and abilities. On a basic level, having students write an argumentative essay versus an explanatory report, for example, will change the cognitive and social demands of organization, invention, and discourse that they confront. And any given analysis of said essay or report will influence what is observed about it. Primarily focusing on cohesion using natural language processing tools such as TAACO will

facilitate different inferences about what students know than primarily focusing on thesis statements using human readers. In turn, applying those inferences prescriptively will create different emphases and opportunities in students' learning and evaluation. Even the same focus, such as a definition of syntactic complexity offered by Biber and Gray (2010) and that offered by Arthurs (2018) will change the nature of how writing is assessed and understood. Any subsequent instructional efforts aimed to support that assessment will therefore emphasize particular choices related to syntax.

More broadly speaking, any assessment used over time ideally includes *construct evidence*, which forms a precondition to and ongoing part of "the traditional and emerging conceptual definition of writing" in a given context, *scoring evidence* that rules are applied accurately and consistently, and *decision evidence* that patterns over time support an interpretative argument to justify the use of a given assessment (Elliot, Deess, Rudniy, & Joshi, 2012, p. 293-294). In many first-year college writing programs, for example, a common construct is academic argumentative essay writing. To exemplify these concepts, let's imagine the goal of an example first-year college writing program was to support students' writing success in subsequent courses. For construct evidence, we would look for consistent evidence of academic argumentative essay tasks and/or expectations in the first-year writing courses and later courses. Scoring evidence would need to show that grades the students received on those tasks were consistently distributed across student groups (and thus not affording greater opportunity for success for certain groups). The decision evidence would need to show a correlation between students' writing performance on these argumentative essay tasks and their retention and performance in first-year and later writing courses. Depending on the program and assessment, our example might also need to include *extrapolation evidence*, e.g., showing that argumentative essay scores allowed predictions beyond test performance if those scores were used to place students in later writing courses that match their needs and abilities. Any assessed writing and analytic approach, then, bears the constraints and possibilities of all of these influences, in terms of the writing construct(s) assessed in a given assessment task, in terms of aspects of said constructs examined and emphasized in resulting analytic inferences, and in terms of how those inferences are used. Writing analytics can help us look for evidence of constructs, scoring, decision, and extrapolation, allowing us to modify assessments when such evidence does not point to consistency and fairness.

Consider some writing analytics studies that illustrate these ideas. Burstein et al. (2018) present a study of Writing Mentor™ (WM), a Google Docs add-on designed to help students improve their writing, that uses natural language processing (NLP) methods and related linguistic resources in four areas commonly emphasized in secondary and early-postsecondary academic writing instruction: use of sources, claims, and evidence; topic development; coherence; and knowledge of conventions. An assessment using this tool therefore prioritizes and sheds light on these aspects of writing (versus, for instance, lexical cohesion or syntactic complexity); furthermore, Burstein et al.'s study, which focuses in part on self-efficacy, then foregrounds the interpersonal domain of writing as part of understanding and guiding student

writing. (Tate and Warschauer's [2018] research note in the same issue of *The Journal of Writing Analytics* reviews several constructs used to measure self-efficacy.)

In a second example, Kumar et al.'s (2017) study uses an expert-graded corpus of student writing to create an automated evaluation score (AES) model based on a five-variable model (spelling accuracy, grammatical accuracy, semantic similarity, connectivity, lexical diversity) in order to predict the holistic scores of Grade 10 narrative essays. Their approach and use of findings are intended to be predictive and prescriptive. In other words, they focus on "how the tool can be useful in automatically identifying students scoring below a certain threshold when performing a writing activity" (p. 181). They recommend using this score to notify teachers "of students' poor performance," so that "teachers would be able to provide remedial interventions in time" as well as "formative feedback throughout a school year" in addition to using it as an advance indication of "students who are struggling with text production skills" related to the five variables as they tend to be valued in a particular assessment (in this case, the Grade 10 narrative exam). Indeed, Kumar et al. note that a limitation of their approach concerns "its omission to assess higher-order thinking skills that all writing constructs are ultimately designed to assess" (p. 177).

Here, Kumar et al.'s wording is useful in that it draws attention to the nature of writing assessments as constructed response tasks: any intervention based on writing analytics findings based on a constructed response task, for instance, evaluates and reinforces writing with respect to that task, and any analytic approach will be limited to the constructs on which it focuses. Written communication itself is always more intersubjective, context-influenced, and complex than writing analytic tools. On this note, writing analytic tools can likewise be used to assess assessments: to hold a mirror up to assessment processes in order to descriptively analyze what they entail. Ross and LeGrand (2017) provide an example study in which they assess writing constructs and propose an expanded view of inter-reader reliability based on writing analytics inferences. They conclude, for instance, that a formative, rubric-based approach to ePortfolio assessment that uses disciplinarily diverse raters can achieve medium-level rates of interrater agreement and reliability; but they also conclude that there are many potential concerns about construct validity and rater training that remain unresolved.

The three example studies above highlight both assessment *of* learning (AoL)—evaluation of summative tasks poised at the end of an arc of learning, such as a writing assignment at the end of a course or school year—and assessment *for* learning (AfL)—ongoing assessment that identifies learner strengths and weaknesses in an iterative fashion, such as in a writing portfolio. AoL is the most common approach to writing assessment, not least because institutions regularly use summative test performances to make inferences about student writing level and development. AfL emerged in response to concerns about the decontextualized, product-oriented nature of AoL. It focuses on formative assessment on the "edge" of student competence, using assessment to propel learning forward (Heritage & Wiley, 2018, p. 729). It aims for a "fundamentally different way of thinking about assessment" by seeing assessment as interactive and ongoing, rather than as something external to classrooms or at the end of student learning (p.

737). The newest approach, sometimes discussed as part of AfL, is labeled assessment *as* learning (AaL) and emphasizes self-assessment, self-reflection, self-evaluation, and metacognition in student learning and development. Practices that emphasize AaL include sharing of evaluation criteria with learners and learner evaluation of their work, learner identification of what they need to work on, and learner reflection on what steps they are taking in their learning. Sadeghi and Rahmati (2017) have suggested that with the emergence of writing analytics tools like corpus linguistics, assessment models that integrate AaL, AfL, and AoL offer promising, empirical ways to explore theoretical views of assessment in light of practical learner realities.

9.0 Ethical Considerations

WA is rife with ethical considerations that are the consequences of how we (and others) see our emerging research specialization. First, as in other research areas, we rely on data from others, whether those others be students, instructors, or workplace writers. Typically, however, this data is not considered on an equal basis with, say, data acquired from subjects in medical or other research. Second, we conduct our studies and projects for a variety of reasons that are governed by local context; some of these are internally motivated and some are both internally and externally motivated. Third, the research of WA locates in a hybrid area between research, as classically defined, and scholarship of discovery, application, engagement, teaching, and learning (Boyer, 1990).

A particular set of concerns deals with data acquisition and data safeguarding procedures. Often, the actual data used by WA researchers comes along with other personally identifiable data points. Data breaches of personally identifiable information/private information have the potential for tremendous harm, albeit harm not directly a result of the research—instead, a result of accidental disclosure of data. Therefore, any WA research project or assessment program will need clear procedures to prevent disclosure of either personally identifiable information about students/faculty or inadvertent disclosure of student writing. Some procedures will be independent of scale, for example, security requirements around password protection and physical security of data to prevent disclosure, and encryption of data to limit harm given a data breach. By contrast, de-identification of data may be easily managed by hand for a small effort, but requires automated means for a true enterprise-level effort to analyze writing. As Rudniy (2018) finds, however, automating de-identification of data is not necessarily possible given current machine learning or neural network techniques.

The specifics of data safeguarding will then likely vary from context to context, and so we suggest guiding principles for thinking about protecting data and privacy while pursuing legitimate pedagogical and intellectual goals. We think Hoel and Chen's (2018) principles for data protection for LA are applicable to WA:

- Privacy and data protection are achieved by negotiating data sharing with each student.

- Openness and transparency are essential and should be an integral part of institutional policies. How the educational institution will use data and act upon the insights of analysis should be clarified in close dialogue with the students.
- Big data will impact all society. Therefore, in negotiating privacy and data protection measures with students, workplace writers, and institutions, researchers should use this opportunity to strengthen their personal data literacies.

Key documents regarding treatment of human subjects, including *The Belmont Report* and the *Common Rule* (2018), were developed for the protection of those subjects of biomedical and behavioral research. Information potentially applicable to research construed as WA-related is mentioned in the following sections of the revised *Rule* under the Electronic Code of Federal Regulations (2018):

- §46.104(d) 1- 3: Exempt research
- §46.110: Expedited review procedures for certain kinds of research involving no more than minimal risk
- §46.111 (a): Criteria for IRB approval of research
- §46.116: General requirements for informed consent (particularly broad consent)

Although these documents describe principles and procedures designed to maximize

- beneficence (“first do no harm”),
- justice (“describe the risks and benefits equally and to disseminate any research findings, both bad and good”), and
- respect for subjects (“people are autonomous agents”),

they do so in the context of human subjects in clinical trials (Sims, 2010). It is incumbent upon current and future researchers in WA to determine how to apply these principles to their own studies.

Most Institutional Review Boards (IRBs) have developed their policies based on these foundational documents; consequentially, WA researchers looking for guidance from their IRBs will likely find that 1) many IRBs do not see a need to be involved in WA activities for purely local purposes, such as local program review and other internal mandates, and 2) that other than noting that data must be de-identified, IRBs provide minimal ethical guidance, even for research completed with the intent of future external publication. For now, many such considerations fall to writing analytics researchers as the research specialization develops. In discussing ethical concerns applicable to both learning and writing analytics, Palmquist (2020) finds that even as researchers in both areas point out the potential for appropriation of and erroneous inferences drawn from data, they also recognize the need to keep exploring responsible and effective uses of that data. As research in writing analytics continues, watchful consideration of whether data is used justly and fairly is essential.

Thus, in addition to considerations for data safeguarding and privacy, we add two principles related to the interpretive power of data selection and analysis:

- Data selection facilitates certain inferences and not others. Researchers and institutions must be critical and transparent about what is present and absent in any data set.
- Algorithms and other data sorting and analysis practices facilitate certain inferences and not others, including the potential for bias toward underrepresented populations, calling for what media studies researchers have labeled “algorithmic accountability.” Researchers and institutions must be aware of and transparent about data organization and analysis choices.

10.0 Conclusions

To this point, we find that writing analytics is an emerging, multidisciplinary applied area of study that has largely focused its efforts on a limited range of instructional situations. As a specialization that has as primary characteristics agility and malleability, it is imperative that those of us professing to conduct writing analytics research make the effort to define how our work fits the larger conversation of our area. This taxonomy provides us with a framework for doing so. It also provides us with a way to describe the research specialization to new/interested researchers in order to broaden the stakeholder pool and create a sustainable climate for research. To that end, we offer the following observations:

- Implementing a WA effort should be a deliberate process, considering the different WA implementation paradigms and their potential affordances and constraints, in context. This planning process should also be forward looking, and account for potential growth/change over time.
- Data is critical to writing analytics. By definition, analytics is about finding patterns in data, and so sufficient data in the form of student writing, instructor commentary, or other associated data points is a foundational requirement for any such effort. There are also critical entailments to data, around storage, access, and safekeeping. Thinking through data requirements is another essential pre-implementation process.
- Descriptive analytics is a potentially powerful way to engage and empower students in their acquisition of writing competencies, by allowing them to see genre and disciplinary patterns, and compare them to their own writing.
- Ethics are another critical aspect of writing analytics, and implementation of a program will require a deliberate effort to think about data safeguarding procedures, harm mitigation procedures, and privacy considerations. Such efforts will likely need to be reconciled with existing university ethics policies but made specific to local writing analytics contexts.

11.0 A Final Note: Ongoing Challenges to WA

As we have found in constructing this taxonomy, WA has three readily identifiable ongoing challenges.

First of these is that we continue to iteratively redefine the research specialization even as we expand the foci of the research specialization. Challenges to both iterative definition and expansion include explaining the impact of our work to and on our stakeholders, who include students, teaching faculty, and administrators, most immediately, as well as the larger university and other communities. Our dominant focus on application and emerging methodologies has to this point made it difficult to build even a small collection of foundational documents to assist those who would like to explore conducting research in this specialization.

Second, the above conditions have made it difficult to develop a pipeline of scholars/researchers to both grow and sustain the research specialization. One reason for this difficulty is that WA researchers use empirical and quantitative methods that have long been met with resistance by many in writing studies; consequently, even current graduate students in writing studies receive limited training at best in those methods. Another reason links back to the challenge for new researchers to find senior faculty mentors. Many of those principally involved in WA projects are researchers who also balance administrative appointments and have limited teaching and mentoring responsibilities.

Finally, although WA research has been conducted by individuals, as is customary in humanities scholarship, the multidisciplinary nature of WA would point to collaborative work as providing a larger portal into the specialization. The need for such teams is clear—collaborations could provide access to textual corpora, as well as knowledge in linguistics, computer science, statistics, and writing studies—and access to funding that is often reserved for only multidisciplinary or multi-institution projects. These types of access and knowledge are essential, not only for those immediately interested in conducting research in WA areas, but for graduate students and early-career faculty as well. Providing these new stakeholders with access to workshops/courses/concentrations in which to learn WA methods, as well as venues in which to present and publish, are key to ensuring that WA is not a temporary research exigency in early 21st century writing studies and related fields. Thus, we end with a call to existing and potential stakeholders to join us in building the next iteration of this taxonomy: to conduct research in the myriad of application areas that would benefit from WA, to contribute to emerging theoretical and ethical frameworks that can inform our research, and to seek ways to involve others in this growing research specialization.

Author Biographies

Susan Lang is Director of the Center for the Study and Teaching of Writing and associate professor of English at The Ohio State University. Her research examines aspects of writing program administration, writing analytics, and technical communication. Her work has been published in *College English*, *College Composition and Communication*, *Writing Program*

Administration, and *Technical Communication*, among others. She is the recipient of the 2016 Kenneth Bruffee Award for Best Article in *Writing Program Administration*.

Laura Aull is Director of the English Department Writing Program and associate professor of English at the University of Michigan. Her research focuses on intersections between applied linguistics, composition, writing assessment, and writing analytics, and most recently examines student writing and civil discourse. She is the author of *First-Year University Writing* and *How Students Write: A Linguistic Analysis*.

Bill Marcellino is a Senior Behavioral Scientist at the RAND Corporation, a non-profit public policy research institution. Bill is a Professor of Text Analytics within the Frederick S. Pardee RAND Graduate School and teaches qualitative research methods within John Hopkin's Global Security Studies program. He also serves on the editorial review board of *The Journal of Writing Analytics*.

References

- Arthurs, N. (2018). Structural features of undergraduate writing: A computational approach. *The Journal of Writing Analytics*, 2, 138–175.
- Aull, L. L. (2017). Corpus analysis of argumentative versus explanatory discourse in writing task genres. *The Journal of Writing Analytics*, 1, 1–47.
- Biber, D. (1988). *Variation across speech and writing*. Cambridge: Cambridge University Press.
- Biber, D., & Gray, B. (2010). Challenging stereotypes about academic writing: Complexity, elaboration, explicitness. *Journal of English for Academic Purposes*, 9(1), 2–20.
- Boyer, E. (1990). *Scholarship reconsidered: Priorities of the professoriate*. Princeton, NJ: The Carnegie Foundation.
- Burstein, J., Elliot, N., Klebanov, B. B., Madnani, N., Napolitano, D., Schwartz, M., ... & Molloy, H. (2018). Writing Mentor™: Writing progress using self-regulated writing support. *The Journal of Writing Analytics*, 2, 285–313.
- Collier, D., LaPorte, J., & Seawright, J. (2012). Putting typologies to work: Concept formation, measurement, and analytic rigor. *Political Research Quarterly*, 65(1), 217–232.
- Colvin, C., Dawson, S., Wade, A., & Gašević, D. (2017). Addressing the challenges of institutional adoption. In C. Lang, G. Siemens, A. Wise, & D. Gašević (Eds.), *Handbook of learning analytics* (1st ed., pp. 281–289). USA: Society for Learning Analytics Research.
- Cronbach, L. J., & Meehl, P. E. (1955). Construct validity in psychological tests. *Psychological Bulletin*, 52, 281–302.
- Crossley, S. A., & Kyle, K. (2018). Assessing writing with the tool for the automatic analysis of lexical sophistication (TAALES). *Assessing Writing*, 38, 46–50.
- Crossley, S. A., Kyle, K., & McNamara, D. S. (2016). The tool for the automatic analysis of text cohesion (TAACO): Automatic assessment of local, global, and text cohesion. *Behavior Research Methods*, 48(4), 1227–1237.
- Crossley, S., Liu, R., & McNamara, D. (2017). *Predicting math performance using natural language processing tools*. Paper presented at the Seventh International Learning Analytics & Knowledge Conference.

- Crossley, S. A., Russell, D. R., Kyle, K., & Römer, U. (2017). Applying natural language processing tools to a student academic writing corpus: How large are disciplinary differences across science and engineering fields? *The Journal of Writing Analytics*, 1, 48–81.
- Driscoll, D. L. (2009). Composition studies, professional writing and empirical research: A skeptical view. *Journal of Technical Writing and Communication*, 39(2), 195–205.
- Electronic Code of Federal Regulations. (2018). Retrieved from <https://catalog.data.gov/dataset/electronic-code-of-federal-regulations>
- Elliot, N., Deess, P., Rudniy, A., & Joshi, K. (2012). Placement of students into first-year writing courses. *Research in the Teaching of English*, 46(3), 285–313.
- Federal Policy for the Protection of Human Subjects ('Common Rule'), 45 C.F. R. §46. (2017). Retrieved from <https://www.hhs.gov/ohrp/regulations-and-policy/regulations/common-rule/index.html>
- Hardy, J. A., & Römer, U. (2013). Revealing disciplinary variation in student writing: A multi-dimensional analysis of the Michigan Corpus of Upper-level Student Papers (MICUSP). *Corpora*, 8(2), 183–207.
- Haswell, R. H. (2005). NCTE/CCCC's recent war on scholarship. *Written Communication*, 22(2), 198–223.
- Helberg, A., Poznahovska, M., Ishizaki, S., Kaufer, D., Werner, N., & Wetzel, D. (2018). Teaching textual awareness with DocuScope: Using corpus-driven tools and reflection to support students' written decision-making. *Assessing Writing*.
- Heritage, M., & Wylie, C. (2018). Reaping the benefits of assessment for learning: Achievement, identity, and equity. *ZDM*, 50(4), 729–741.
- Hoel, T., & Chen, W. (2018). Privacy and data protection in learning analytics should be motivated by an educational maxim—towards a proposal. *Research and Practice in Technology Enhanced Learning*, 13(1), 20.
- Kane, M. T. (2013). Validation as a pragmatic, scientific activity. *Journal of Educational Measurement*, 50(1), 115–122.
- Kumar, V., Fraser, S. N., & Boulanger, D. (2017). Discovering the predictive power of five baseline writing competences. *The Journal of Writing Analytics*, 1, 176–226.
- Kyle, K., & Crossley, S. A. (2015). Automatically assessing lexical sophistication: Indices, tools, findings, and application. *TESOL Quarterly*, 49(4), 757–786.
- Lancaster, Z. (2016). Expressing stance in undergraduate writing: Discipline-specific and general qualities. *Journal of English for Academic Purposes*, 23, 16–30.
- Marradi, A. (1990). Classification, typology, taxonomy. *Quality & Quantity*, 24(2), 129–157.
- Mislevy, R. J. (2018). *Sociocognitive foundations of educational measurement*. New York, NY: Routledge.
- Moxley, J., Elliot, N., Eubanks, D., Vezzu, M., Elliot, S., & Allen, W. (2017). Writing analytics: Conceptualization of a multidisciplinary field. *The Journal of Writing Analytics*, 1, v–xvii.
- Moxley, J., & Walkup, K. (2016). *Mapping writing analytics*. Presented at the 9th International Conference on Educational Data Mining (EDM 2016), Raleigh, NC, USA. Retrieved from <https://pdfs.semanticscholar.org/de36/8a95f91b5bef34443d2f2edc0c7e99b0bdb9.pdf>
- Nelson, W. (1965). A standard corpus of edited present-day American English. *College English*, 26(4), 267–273.
- Nelson, F. W., & Kučera, H. (1967). *Computational analysis of present day American English*. Providence, RI: Brown University Press.

- Palmquist, M. (2020). Learning analytics in writing instruction: Implications for writing across the curriculum. In L. E. Bartlett, S. L. Tarabochia, A. R. Olinger, & M. J. Marshall (Eds.), *Diverse approaches to teaching, learning, and writing across the curriculum: IWAC at 25*. Fort Collins, CO: The WAC Clearinghouse and University Press of Colorado. Retrieved from <https://wac.colostate.edu/books/perspectives/iwac2018>
- Peele, T. (2018). Is this too polite? The limited use of rhetorical moves in a first-year corpus. *The Journal of Writing Analytics*, 2, 78–95.
- Peña-Ayala, A. (2018). Learning analytics: A glance of evolution, status, and trends according to a proposed taxonomy. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 8(3), e1243.
- Poe, M., & Elliot, N. (2019). Evidence of fairness: Twenty-five years of research in *Assessing Writing*. *Assessing Writing*, 100418.
- Purcell, K., Buchanan, J., & Friedrich, L. (2013). *The impact of digital tools on student writing and how writing is taught in schools*. Washington, DC: Pew Research Center. Retrieved from <https://www.pewinternet.org/2013/07/16/the-impact-of-digital-tools-on-student-writing-and-how-writing-is-taught-in-schools/>
- Ringler, H., Klebanov, B. B., & Kaufer, D. (2018). Placing writing tasks in local and global contexts: The case of argumentative writing. *The Journal of Writing Analytics*, 2, 34–77.
- Ross, V., & LeGrand, R. (2017). Assessing writing constructs: Toward an expanded view of inter-reader reliability. *The Journal of Writing Analytics*, 1, 227–275.
- Rude, C. D. (2009). Mapping the research questions in technical communication. *Journal of Business and Technical Communication*, 23(2), 174–215.
- Rudniy, A. (2018). De-identification of laboratory reports in STEM. *The Journal of Writing Analytics*, 2, 176–202.
- Sadeghi, K., & Rahmati, T. (2017). Integrating assessment as, for, and of learning in a large-scale exam preparation course. *Assessing Writing*, 34, 50–61.
- Schum, D. (1994). *The evidential foundations of probabilistic reasoning*. Evanston, IL: Northwestern University Press.
- Shum, S. B., Knight, S., McNamara, D., Allen, L., Bektik, D., & Crossley, S. (2016). Critical perspectives on writing analytics. *LAK '16 Proceedings of the Sixth International Conference on Learning Analytics & Knowledge*, ACM, New York, 481–483.
- Sims, J. M. (2010). A brief review of the Belmont Report. *Dimensions of Critical Care Nursing*, 29(4), 173–174.
- Smith, K. B. (2002). Typologies, taxonomies, and the benefits of policy classification. *Policy Studies Journal*, 30(3), 379–395.
- Staples, S., & Dilger, B. (2018–). Corpus and repository of writing [Learner corpus articulated with repository]. Available at <https://crow.corporaproject.org>
- Tate, T. P., & Warschauer, M. (2018). Going beyond “That was fun”: Measuring writing motivation. *The Journal of Writing Analytics*, 2, 257–279.
- Turnitin. (2019). Turnitin for higher education. Retrieved from <https://www.turnitin.com/divisions/higher-education>
- Wright, G. T. (1974). The lyric present: Simple present verbs in English poems. *PMLA*, 89(3), 563–579.