

## Introduction, or What's Wrong with Literary Studies?

“All of it! – All! – Wilhelm!”  
 – *The Sorrows of Young Werther*

*Over the past few decades, humanists have insisted that it is important to resist generalizations.*

*The process of secularization associated with modernity has not spread as widely nor penetrated as deeply as Western humanists have tended to assume, not even in the academy.*

*Without its world, the human is merely another species on earth, testing itself against threats of its own creation and in the process becoming a force like nature (capable only of overt behavior) that jeopardizes its own existence.*

*Western European philology developed in the eighteenth century at much the same time that the notion of literature did.*

*The funny thing that happened to charm on its way to modernity was the disenchantment of the West.*

*Today, anything, it would seem, can be art.*

These statements, all drawn from recent academic research in literary studies, represent the tip of a giant methodological iceberg. They signal a much larger and much less well-known problem of what I will be calling the practice of generalization. “Generalization” is the rhetorical strategy whereby we move from partial evidence to knowledge claims about some larger group or category. Whether authors are describing “humanists,” “the human,” “Western philology,” “literature,” or simply “art,” what these statements all have in common is that they exceed, drastically and exorbitantly, the foundations of what has been observed to make claims about what is.

This Element is about wrestling with this problem, with the gap between what is observed and what is known or, better, believed to be known. While this is a problem for all fields of knowledge, not to mention everyday life, one of the beliefs that I want to explore in this Element is that textual evidence poses distinct challenges for the practice of generalization, of how we move from examples to knowledge. What difference does the weave of written language make when it comes to thinking about the relationship of parts to whole?

The problem of the reliability of textual evidence has emerged with particular urgency today. In what has come to be known as the “reproducibility crisis,” we are witnessing a major reassessment of truth claims being made across a number of scientific disciplines (Earp & Trafimow 2015; Spellman 2015). A preponderance of false positives is circulating within scholarly publications – things that are believed to be true but that are not reproducible over time. The increasing

failure to reproduce is a sign that the process of generalization, the validity of a belief based on some, but not all, of the evidence has not endured.

At the same time, we are also seeing a widespread concern with the veracity of popular information, which now goes by the colloquial name “fake news.” The increasing volume of written communication and the growing speed of circulation through social media have brought with them a similar crisis of verification. Whether in the domain of science or the news, the changing scale of information has foregrounded a basic and recurring problem surrounding the credibility of textual evidence. How can we believe what we read?

The growth of uncertainty surrounding textual evidence is emerging (probably not incidentally) at precisely the same moment as a vast array of new methods and technologies are appearing that claim privileged access to knowledge about texts. Whether under the heading of natural language processing, machine learning, artificial intelligence, or text and data mining, a variety of new techniques have developed over the course of the past decade that aim toward the empirical and quantitative understanding of texts. Scale and measurement are being proposed as an antidote: not simply to the problem of the new (and old) scale of texts but also as a means of repairing the evidentiary deficiencies surrounding the study of texts more generally. But are they the answer we have been looking for?

In plain terms, generalization involves the act of moving from tokens (instances of things) to types (a single representation of a more general category). For example, as a child I saw many trees around me. At a certain point, I developed from these instances the idea of a “tree.” As cognitive scientists have shown, generalization is crucial to any learning process (e.g., Erickson & Thiessen 2015; Thiessen 2017). It allows us to dispense with holding many individual representations in our mind in favor of a single, higher-order (more abstract) representation. Instead of maintaining all novels I have read in memory, I create a more general representation of what a novel is in my mind. There is a fundamental utility to generalization. It makes life easier to navigate. Generalizations are also cognitively enabling. They allow us to account for more of our experience with fewer rules.

But generalizations can also be dangerous. It should not be hard for readers to conjure up examples in their minds of disputes they have had with a friend, partner, or child that turned on the issue of an unfair generalization (“that is a total generalization!”). Similarly, one can imagine examples where the persistence of a belief about a particular group (in terms of racial, ethnic, or gender stereotypes, for example) can drive negative behavior toward that group. This is what researchers are concerned about in the case of false positives (i.e., unwarranted generalizations) that can lead to poor social policies (such as unnecessary medical tests) or poor judgments by individuals (such as whether to hire a person or even engage in violence against others). Poor generalizations can

make life harder, and potentially more precarious, either for ourselves or for those around us.

The practice of generalization is thus an essential aspect of both social life and knowledge production. How we move from particular observations to more general statements about the world – and the discrepancy between parts and wholes in this process – has an essential bearing on our understanding, our beliefs, and the actions and policies that reflect those beliefs.

This Element sets out to better understand the process of generalization as it pertains to the study of texts. First, what does it mean to generalize about texts and literary texts in particular? What do generalizations look like in literary scholarship? Second, to what extent is this practice happening in our field? How prevalent are the generalizing claims like the ones I cited earlier? Are there particular types of people, journals, or fields where it is happening more (or less) often? And when it is happening, what is the nature and scale of generalization? How can we better understand its scope and qualities? Finally, what, if anything, should we do about it?

My own awareness of the problem of generalization emerged along with the reproducibility crisis in the sciences that began almost a decade ago. Despite a long-term investment in quantitative methods, serious issues were surrounding the credibility of claims being made across a number of different fields. The antidote to the fallibility of generalization had produced its own new set of problems. Numerous reforms were proposed to make scholarship more sound, and yet those calls were not part of the conversation in the humanities. There was a sense of immunity in our field from the evidentiary and methodological problems that surrounded the reproducibility crisis.

Scholars have identified a number of potential causes surrounding the crisis of reproducibility, from small sample sizes, to excessive researcher degrees of freedom, to an absence of transparency with respect to the research process. The more I read around in literary scholarship, including my own past work, the more I began to see how the very problems that were undermining the credibility of claims in the sciences were endemic to literary studies as well. I began to marvel at the discrepancies between researchers' claims, the amount of observations used to make those claims, and the opacity of the methods used to arrive at critical judgments. I will never forget the day I came across this passage by Fredric Jameson:

Having made these initial distinctions, let me now, by way of a sweeping hypothesis, try to say what all third-world cultural productions seem to have in common and what distinguishes them radically from analogous cultural forms in the first world. All third-world texts are necessarily, I want to argue, allegorical, and in a very specific way: they are to be read as what I will call

national allegories, even when, or perhaps I should say, particularly when their forms develop out of predominantly western machineries of representation, such as the novel. (1986, p. 69)

Here Jameson was speaking for *every single* creative work ever produced outside of NATO countries and assigning to them a single distinguishing feature. How could this possibly be true? And how did Jameson think his methods of reading a few books endowed him with the authority to make such sweeping, unverifiable claims? The more I read, the more I saw how pervasive statements like Jameson’s were. And in a perverse kind of logic, it was as though the more sweeping and less well justified the claim, the more prestige one could accrue. Did people know what they were doing? And if so, why weren’t they concerned?

This Element argues that we should be concerned. By “we” I mean literary scholars across different national literatures, as well as scholars across the humanities for whom methods of case-based research remain their field’s primary research method. Generalization is not some esoteric practice that we need not worry about. On the contrary, it is an essential, one might even say existential, scholarly practice that until now has remained all but invisible in critical debates in the humanities. The failure to generalize *well* puts at risk nothing less than our credibility as scholars and cultural commentators.

When it comes to textual evidence, generalization is a problem that can cut two ways. On the one hand, the gap between what has been observed (the examples) and the claims being made (the generalizations) can be so large as to defy belief – as when Jameson speaks for all third-world cultural production or researchers make claims about “the West” or “art” or “the human” and use just a handful of examples to justify their claims. On the other hand, if we take claims to particularity seriously – that we ought to traffic in knowledge of particulars, not generalized beliefs – our knowledge runs the risk of becoming so specific as to be socially meaningless. When we write a whole book or even a single chapter about a single novel, whom or what is this for? What is the social value of knowledge of a single thing?

To better understand the prevalence and nature of generalization within literary studies, this Element applies new techniques of machine learning and natural language processing to study a sample of more than 16,000 statements drawn from recent scholarship in three fields of study (history, sociology, and literary studies).<sup>1</sup> It tries to offer readers one way of structuring an algorithmically informed investigation of texts, in particular as a team-based, interdisciplinary, and even intergenerational undertaking. How can we use machine

<sup>1</sup> All data and code can be found at the following repository: <https://doi.org/10.6084/m9.figshare.12669329.v1>

learning to study complex linguistic phenomena such as generalization? This Element provides a workflow for how to do so, focusing on issues of data collection, conceptual definition, data annotation, and model validation as core aspects of computationally informed textual study.

Based on these methods, the Element's empirical section (Part II) provides evidence for both the consistency and prevalence of generalizing statements within literary studies. Despite calls to the value of particularity that one may hear in the humanities, according to the evidence provided here, scholars in literary studies continue to make generalizing statements with surprising regularity. According to the data presented here, *an estimated one out of every two or three sentences* in the framing of research articles is predicted to be a generalization, on par with a far more quantitative field like sociology. Delving more deeply into the data, we can also see how the conceptual scales at which researchers are working are both vast and apparently growing.

The findings uncovered here raise a host of challenging questions, not just for literary studies but also for the humanities more generally. Given the glaring discrepancies between the scale of claims being made and the amount of evidence being used to support these claims, what should we do about it?

This is obviously a challenging question to answer, and I can only begin to sketch some possible ways forward here. In the discussion section (Part III), I entertain but ultimately eschew two simpler and more straightforward arguments: (1) we should either *not* be generalizing at all because this is not what literary scholarship does or (2) we should *only* be generalizing because this is the very definition of scholarship (in the German sense of *Wissenschaft*, a body of knowledge) and therefore we require a massive overhaul of our methods to move largely away from case-based research. As I will discuss at greater length, the latter is an idealization that overlooks the irreducibility of textual meaning, that there can be no final validity of what a text or group of texts *means*. There are basic and fundamental limits to the universality of truth claims that can be made when it comes to questions of textual meaning and our ability to model and measure them. This is the first “problem of textual evidence” that I invoke in my title.

At the same time, the former position (that we should not be generalizing at all) is an idealization that overlooks our own behavior. As this Element will demonstrate, the practice of generalization is widespread in the scholarly literature. The generalization that we do not generalize is, according to the data presented here, a very poor generalization. But the resistance to the value of generalization itself also overlooks a vibrant tradition of scholarship within the field's history where the arrival at generalizable knowledge about texts is an essential aspect of the project of textual understanding. From Lorenzo Valla and

early humanists to the nineteenth-century rise of hermeneutics to today's textual forensics (Iqbal et al. 2013), avoiding misinterpretation has been a central value of the field of textual study for centuries. At a time when misinformation has become a primary social ill, developing methods that generate confidence in textual meaning is arguably more important than ever. This is the second sense by which I mean the "problem of textual evidence." How can we develop methods that move beyond the often incredible gaps between observation and interpretation that are operative in the field today?

In place of these two more extreme positions, I argue instead for two more moderate ways forward, each of which is premised on differing notions of "openness." The first asks us to move from questions of evidentiary sufficiency – when is some evidence *enough* – to questions of evidentiary transparency – how much of the research process has been explicated? Given that we can never have all of the evidence before us, in either qualitative *or* quantitative scholarship, how can we make more explicit all of the tools, techniques, and procedures that were undertaken to arrive at our general statements about the world? My aim here is to move past a two-cultures approach to evidence, where quantitative and qualitative methods operate and are valued according to different evidentiary criteria. Drawing on recent work that has emphasized principles of transparency surrounding scholarship in the name of "open science" (Willinsky 2006; Stodden, Guo, & Ma 2013; Foster & Deardorff 2017; Vincente-Saez & Martinez-Fuentes 2018), I argue that these principles need to be embodied more fully within humanistic disciplines like literary studies. While we have made (some) strides in the area of open access with respect to publishing, we have not addressed the accessibility and openness of our methods, evidence, and modes of judgment in equal measure. Using this Element's approach as a model, I provide suggestions for ways that *all* scholarship in literary studies can implement more open evidentiary procedures. If we are going to continue to make generalizations about how documents in the world work at the rate at which we are currently making them, we need to address the evidentiary and methodological elisions of our methods. The result will ideally be a good deal more moderation with respect to the claims being made in case-based research, what we might call a new culture of limitation.

The second moderate pathway argues that instead of seeing all literary scholarship as invested only and ever in making generalized claims about how texts work, that we take seriously the idea of the contingency of textual interpretation whose roots go back at least as far as Immanuel Kant. No amount of sampling or predictive modeling will ever fully explicate once and for all the question of textual meaning. While the problem of the irreducibility of truth

claims is endemic to all fields of knowledge (Reichenbach 1930; Popper 1935), a long tradition of research in literary studies has drawn attention to the way the solicitation of multiple and subjective responses to texts is not a problem to be minimized but the fundamental purpose of the discipline. As I will show in this section, in addition to the prevalence of generalization within the field, we can also find evidence for the ways in which scholarship in literary studies is indeed more conceptually and rhetorically open than work in the sciences. This is the second notion of openness that I want to emphasize here, not in the sense of transparency and visibility, but as a form of novelty and unknownness. To return to my earlier question, what is the value of engaging deeply with single things, this section argues that it can serve as a mechanism for producing a more interrogative and expansive mind-set.

Rather than seeing this practice as the *only* value of textual study however – as oppositional to the project of generalization and knowledge production – my argument is that we would benefit by envisioning these two strains of research in more complementary ways. Seeing and celebrating the critical project of particularism as something partial rather than universal will not only help us better account for what we are actually doing in our scholarship. It will also help us see the relational value between these two methodological modes. The conceptual openness of literary scholarship – its exploration of the hypothetical, from possible worlds to potential ideas – can be an ideal companion to inform more nuanced, creative, and flexible generalizations about the textual worlds that we study. I want us to see the value of *mutuality* between the conceptual openness of literary scholarship and the methodological openness of generalizing about how the world works.

To be clear, my goal in writing this Element is not to put forward a normative account of the discipline of literary studies, a single vision of what literary studies is and does. It is too diverse a thing for a single account. Rather, my aim is to better understand the alignment (or in this case misalignment) between our discursive practices and epistemic ideals, what we say in our work and what we hope to know through it. In this, I see it as part of a long tradition of work in disciplinary self-assessment, what James Evans and Jacob Foster (2011) have called a form of “metaknowledge.” Work by John Guillory (1993), Gauri Viswanathan (1989), Rachel Burma and Laura Heffernan (2012), Merve Emre (2018), and Chad Wellmon and Paul Reitter (forthcoming, 2021) are just a few of the many valuable works that have used extensive archival research to highlight the ways in which institutional and political contexts shape disciplinary behavior. In using a computational approach, my hope is that this Element can serve as a model for future large-scale studies of disciplinary behavior, following in the footsteps of important precedents (Goldstone & Underwood

2014; Wellmon & Piper 2017; Degaetano-Ortlieb & Piper 2019; Bourrier & Thelwall 2020).

Many of the recent debates surrounding the future of literary studies have often been technologically inflected – whether one can or cannot arrive at insights about literature using one kind of technology (computation) or another (bibliographic). Or they have been invested in trying to define the particular nature of knowledge unique to literary studies as distinct (or not) from other disciplines. My goal here is to move us away from discussions about the kind of knowledge we produce to a study of the *relationships* between the kinds of evidence we use and the discursive procedures that make that evidence visible to others.

After several years of thinking about this problem and discussing it with others, I have also become aware just how complicated and nuanced the problem is. The practice of generalization is not something that can be solved within a single book or from a single point of view. Indeed, it cannot be “solved” in any strict sense. It represents an ongoing state that requires regular and collective action, from the changing habits of individual researchers to the policies of editors and editorial boards with respect to peer review. The work of generalization is by definition never finished. My hope with this Element is to start an important conversation.

## I Theory

### 1 Generally Speaking

#### *From Exemplarity to Estimation*

In the paper that would ultimately gain him admission to the French Academy of Sciences, “Mémoire sur la probabilité des causes par les évènements” (1774), Pierre Simon Laplace asked whether given a large number of astronomical observations that all differed from one another, there was a way to derive the true, or at least most confident, location of a planet traversing the heavens. Laplace’s aim was to develop a mathematical model to understand the relationship of observation to error, that is, to think probabilistically about the nature of events in the world.

Laplace’s work represented a watershed moment in the history of ideas. It attempted for the first time to formalize an understanding of the distribution of observations as a means of thinking about accuracy and truth. Rather than rely on a single best example, Laplace’s method attempted to consider a *set* (or in modern-day parlance a *sample*) of observations and produce a single best estimate. While all observations were not imagined to be equally valid (the



further away they were from some imagined midpoint, the more likely they were to be erroneous, or so Laplace theorized), neither was any one of them categorically assumed to be “the best.” It marked an important beginning in the process of formalizing uncertainty and error as *part* of knowledge and truth rather than as enemies of it. Instead of emphasizing a single observation, Laplace’s model provided a method to understand the relationship of observations to each other. It replaced exemplarity with estimation.

Laplace’s initial mathematical solution to estimating the distribution of error was, it turned out, in error (Stigler 1986). Understanding effective ways to model the distribution and potential error of observations was a complex mathematical problem that would gradually evolve over a number of years. Like all scholarly debates, it entailed an active social process of argument and publication. Laplace’s own work built on earlier investigations by Johann Lambert, Joseph Lagrange, and Thomas Bayes. The German mathematician Carl Friedrich Gauss would eventually provide the solution to what came to be known as the “normal” distribution in 1809.

Over the course of four decades, then, roughly equivalent to one academic career, a profoundly new way of understanding nature had emerged. Instead of relying on a single “best” observation – a model based on an authoritative observer – scholars were putting into place a new model that took into account numerous observations to infer an “ideal” observation. In one sense, both practices involved forms of idealization. The exemplary example was an ideal version of its kind, just as the inferred estimate based on a sample of observations was. But according to the latter model, the accuracy of one’s insight depended on an understanding of the *relationship* of observations to each other rather than on a single, arbitrarily representative body.

The story of how such probabilistic thinking would gradually penetrate numerous fields of knowledge over the course of the nineteenth century is by now well known (Porter 1986), though much of its twentieth- and twenty-first-century continuation still remains to be written. Also well known is the extent to which scholars in the humanities have largely, though not exclusively, resisted this transition.<sup>2</sup> Literary scholars have continued to overwhelmingly rely on methods that depend on an authoritative observer choosing the “best” observation(s) to prove some larger point.

---

<sup>2</sup> In a sample of the first 100 articles listed in the Modern Language Association database from different journals filtered according to the subject heading “English literature” on April 28, 2020, only 1 article used quantitative methods. In the journal *American Sociological Review* for 2016 drawn from this Element’s sample, by contrast, 43 of 47 articles were quantitative in nature. Our working understanding of the prevalence of quantitative and computational work in the field of literary studies is thus ~ 1 percent of research outputs. Future work can provide more accurate estimates.

The new availability of increasing amounts of textual data along with a variety of computational methods to analyze that data has allowed literary scholars to begin to follow in the footsteps of Laplace and utilize the same practices of generalizability as other disciplines. So why haven't we embraced these methods more broadly? And should we be doing so?

*The Curse of Lorenzo Valla*

In 1440, Lorenzo Valla completed his *Oration on the Falsely-Believed and Forged Donation of Constantine*. As the title of his self-described *opusculum* or “little work” announced, Valla claimed to have proven that the decree composed by the Emperor Constantine in the fourth century that transferred authority over the Roman Empire to the pope was false. As Valla wrote to one of his interlocutors, “Why did I write about the Donation of Constantine? . . . For the sake of truth” (cited in Camporeale 1996, p. 9.).

Valla's oration is consistently held up as one of the primary examples of the intellectual contributions of Renaissance humanism, what Donald Kelley later called “the historical revolution” in textual criticism (1970). As scholars have long pointed out, doubt about the Donation was hardly new (Bowersock 2008). What was new was the historical method that Valla applied to disprove the document's authenticity. Valla took aim at the language and logic of the decree by showing how incongruous it was with the surrounding documents of its age. “You say that within the first few days the Senate, the nobles, and the satraps [i.e. provincial governors], as though already Christians, passed decrees with the Caesar for honoring the Roman church! What! . . . Whoever heard of satraps being mentioned in the councils of Romans?” (Valla 2008, p. 85). Valla's extensive knowledge of ancient documents allowed him to identify these linguistic anachronisms, just as his sensitivity to language allowed him to ascertain when words belonged to a source that could not have been known by Constantine: “Constantine is made to arrogate himself the titles of God, and to try to imitate the language of the sacred scriptures, which he had never read” (Valla 2008, pp. 91–93). For Valla, detailed attention to the style and vocabulary of documents, what we would now call techniques of “close reading,” allowed him to, if not prove the actual date of composition, then at least falsify claims surrounding the document's date of composition and thus its institutional validity.

Valla's case has become celebrated as a heroic instance of when close reading was successfully used to challenge the claims of power. For Valla's disciples, his methods would become a foundation of humanist scholarship (Nauta 2009). From the subsequent editorial work of Erasmus, who sought to produce reliable