



HAL
open science

Intertextuality and narrative context: digital narratology?

Damien Nelis, Christopher W Forstall, Lavinia Galli Milić

► **To cite this version:**

Damien Nelis, Christopher W Forstall, Lavinia Galli Milić. Intertextuality and narrative context: digital narratology?. 2017. hal-01480773

HAL Id: hal-01480773

<https://inria.hal.science/hal-01480773>

Preprint submitted on 1 Mar 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Intertextuality and narrative context: digital narratology?

Damien Nelis^{1*}, Christopher W. Forstall², Lavinia Galli Milic^{1*}

¹University of Geneva, Switzerland

² University of Notre Dame, USA

*Corresponding author: damien.nelis@unige.ch

Abstract

Within the context of an interdisciplinary research project studying allusion in first century Latin epic, we examine some of the shortcomings of current digital methods for detecting poetic allusions through text reuse. We discuss the role that textual features above and beyond word identity often play in signalling intertextual connections to human readers, in particular, narratological structures and typical scenes. We suggest that such features can be captured by performing clustering and/or classification on word frequencies in paragraph-sized chunks of text, and propose a possible pipeline for fusing similarity measurements based on these “thematic” features with existing metrics produced by the intertext search tool *Tesserae*. While work testing the stability of the proposed features is ongoing, we report some promising initial results based on k-means clustering, and describe plans for future work.

keywords

Latin epic; intertextuality; narratology; clustering

INTRODUCTION

This paper describes ongoing work to improve the results of digital intertextual search algorithms in Latin epic poetry by supplementing existing phrase-based matching with a feature based on narrative structures. Scholars of intertextuality, both digital and non-digital, are already aware of the limits of purely verbal repetition as a proxy for poetic allusion and other literary gestures that make reference to earlier texts. In our view, going beyond simple lexical reuse should necessarily include consideration of the text at a thematic and a contextual level. In the endeavour to bring these additional levels of analysis to bear upon existing digital search procedures, one important first step is the development of prototype methods for automatically extracting information about narrative structure from the text. A second, equally important aspect is the comparison of current computational intertextual search results with the products and process of traditional philological reading methods.

Our team is made up of classicists working with a combination of traditional philological and Digital Humanities methods to study the literary history of classical epic. The present project is devoted to trying to measure the complexities of text reuse in Latin epic poetry of the Flavian period (69-96 CE), contextualised by the larger frame of classical epic tradition. This work is funded by a Swiss National Science Foundation research grant which permits us to collaborate with Neil Coffee's *Tesserae* group at the University at Buffalo. This partnership gives our classicist team in Geneva access to a wider range of computational competences than would otherwise be available, and provides a unique opportunity to compare the critical practices of non-digital philologists with ongoing development and design work in an established computational intertextual search project. One potential benefit of such a partnership is that close comparison of the priorities and reading habits of the digital and non-

digital teams will not only lead to improvements in each, but ideally also provide occasion for reflection and insight into the very nature of intertextual reading.

I BACKGROUND

1.1 Verbal reuse and Digital Classics

For several decades now, classical scholars have been able to search for verbal connections between Latin texts by searching for specific terms in databases such as *Poetria nova* [Mastandrea and Tessarolo, 2001], the *Packard Humanities Institute* corpus [PHI], the *Bibliotheca Teubneriana Latina Online*, or the Brepolis corpus and search tools [CDS] [LLT-A] [LLT-B].¹

But the last decade has witnessed a decisive increase in the sophistication of computational approaches to intertextual study, and this rapid expansion is transforming our vision of the role played by digital tools. More than a simple index or concordance, software is now seen as able to intervene in research at all scales and even to contribute to the theoretical definition of what constitutes a meaningful intertext [Trillini and Quassdorf, 2010] [Coffee et al., 2012]. For example, the *Musisque Deoque* [MQDQ]² co-occurrence search allows the researcher to compare one particular poem to the whole poetic corpus and can identify metrical patterns as well as words.³ *Filum* performs similar one-to-many searches using edit distance between character strings rather than a dictionary of Latin words, allowing words with different lemmata but similar spellings to match [Filum] [Chauduri and Dexter, 2016]. Marco B uchler’s [TRACER] is an open-source, general-purpose text re-use detection tool which, although it is not associated with a particular corpus, has been used with success in several classics projects. It provides several flexible word n-gram features out of the box, with the option to employ stemmers and more complex natural language processing tools if the research has them available. *Tesserae*, also open-source, provides a core classical corpus and several pre-configured search tools via a web interface [Tesserae]. By default *Tesserae* matches on lemma bi-grams⁴, a relatively small feature in the field of text reuse. This fine-grained sensitivity allows *Tesserae* to catch the often very brief allusions characteristic of Latin poetry, but at a cost: it also returns a large number of false positives which must be filtered out based on additional textual features. The parameters of this filtering, as well as the considerations involved in designing a model of allusion around text re-use, are discussed in [Coffee et al., 2012].

An important difference between today’s digital searches and their precursors is that they assume a more exploratory mode of research: whereas earlier tools required an a priori term or phrase of interest to search for (“how many works use the word *arma*?”), these permit more open-ended queries (“what similarities exist between the *Aeneid* and the *Amores*?”) and aim explicitly at the discovery of the unexpected. As the kind of questions posed to computational tools begins to look more like the subjective, interpretive questions that are of greatest interest to scholars, the successes (and importantly also the failures) of these algorithmic approaches can shed new light on the behaviour and experience of readers [Scheirer et al., 2016].

¹The [PHI], [BTL] and Brepolis search tools are available online by subscription only.

²This open-access tool is an improved version of [Mastandrea and Tessarolo, 2001].

³Metrical features are more fully developed in a sister project [Pedecerto].

⁴Additional, experimental options allow search on character-based and semantic features, as well as a many-to-many Text search, an LSA search tool and the Tri-gram visualizer, see [Tesserae] under “Other tools.”

We principally use the *Tesserae* search tool for the digital side of our research. *Tesserae*'s own evaluation of their search tool's performance was, at least initially, largely based on a comparison between book 1 of Lucan's *Bellum Civile* and the whole Vergil's *Aeneid*, using lemma bi-grams as the only features. "These works were chosen," its authors argue, "because they were long enough to give representative results, and well studied enough to allow comparison with traditional approaches" [Coffee et al., 2012]. The choice to focus on epic poetry is further justified, in our opinion, because this corpus is recognized as being intensely intertextual on the verbal level, and literary allusion in this field has always been seen as fundamental to literary analysis.

At the same time, although lemma bi-grams may represent "the most basic and common form of intertextuality" in this corpus, as [Coffee et al., 2012] claim, it is also clear that lemma identity alone is not sufficient to characterize intertextuality. [Wills, 1996] describes a large set of features which may operate in combination with lexical repetition or even without it to build intertextual links between Latin poems. [Bamman and Crane, 2008] have highlighted the need for digital intertextual search tools to take into account multiple features including word order, syntax, metre and sound similarities as well as the semantic similarity, in order to capture "indirect references" as well as direct quotations.⁵The *Tesserae* team themselves estimate, based on their Lucan-Vergil benchmark, that at least one quarter of the allusions considered significant by commentaries cannot be caught by word-based matching alone[Coffee et al., 2012].

We assume that every example of text-reuse recovered by digital means occurs within a narrative structure that is sometimes far more likely than individual verbal reminiscences to trigger associations among readers, as this has been suggested by major contributors to the study of intertextuality in the world of classical scholarship, including [Knauer, 1979], [Juhnke, 1973], and [Nelis, 2001]. In our view, elements of this level of prominence should be taken into account as a criterion in the automatic detection of intertextuality. It is the goal of our project to explore, on the digital side, how we might add one or more narrative features to *Tesserae*'s existing lemma bi-gram search, and, on the philological side, what sort of results we expect to see from a narrative-sensitive tool, above and beyond what is available now.

1.2 Allusion and Intertextuality

The meaning and scope of the term "intertextuality" are contested. When it was used for the first time by Julia Kristeva in 1966, the term had essentially a linguistic-semiotic dimension, extending well beyond the narrow world of classical literature: "Every text," she wrote, "constructs itself as a mosaic of quotations, every text is the absorption and transformation of another text" [Kristeva, 1967]⁶. Yet the term has most frequently seen service in the context of literary criticism and philology, and almost from the moment it was coined, scholars have continued to modify, challenge, and very frequently to refine and reduce the scope of this

⁵ The authors categorized the allusions in two types: the "direct allusion" which can easily be identified because is a sort of quotation and the "indirect reference," a very common form of allusion in Latin texts, where the referent text has so lexically changed that it is hardly uncovered by a simple string search. Other efforts in this direction are expended by the German [eTRACES] project, whose goal is to create new tools for detecting "winged words and quotations."

⁶"Tout texte se construit comme une mosaïque de citations, tout texte est absorption et transformation d'un autre texte."

notion.⁷ Classicists in particular have engaged with its theoretical definition and its application as an interpretive practice, focusing not only on the ways in which texts interact with one another but also on the respective roles played by authorial intent and by readerly interpretation. Even within classics views diverge significantly on the relative power and freedom of the author and the reader in performing the work of making meaning out of the links between literary texts.⁸

In this paper we assume that there are different degrees of intertextuality, ranging along a gradient of subtlety. At one end of this scale is obvious quotation of the *source* (earlier) text, verbally pointed out in the *target* (later) text; at the other end lies allusion and other literary artifice, where source and target texts merge into each other. Our primary interest, here, is in the latter. We consider such allusions to be part of the literary design of the author and we are interested in understanding how they enrich the meaning of the target text for its audience. We work with the idea of a so-called “ideal reader,” that is, a reader who is capable of decoding the allusive strategies and we assume that awareness of references to earlier texts is a fundamental element in the whole process of literary communication.⁹

Our approach also takes into account multiple tiers of text re-use, that is, analysis of texts which contain layers of similarity to several earlier texts, and potentially bear meaningful references to all of these predecessors simultaneously. Scholars call this kind of intertextuality “window reference,” “double allusion,” or “two-tier allusion,” that is “alluding to a model both directly and through an intermediary” [Hinds, 1997, 9 and 151][Nelis, 2001, 5]. Each of the model-texts has its own respective narrative forms and sequences, and yet in the juxtaposition of the target and its sources, a pattern emerges which sheds new light on their interrelationship. This will be illustrated in the example below (Section II).

1.3 “Heroic” epic poetry

Our Flavian epic corpus consists of Statius’ *Thebais* and *Achilleid*, Valerius Flaccus’ *Argonautica*, and Silius Italicus’ *Punica*. These are long poems on mythical (Statius and Valerius) and historical (Silius) subjects, composed in Latin hexameters, each work comprising multiple “books” of many hundreds of lines.¹⁰

It is widely recognized that this poetry marks an important stage in the history of Latin intertextuality. Yet, for most of the nineteenth and twentieth centuries, the work of these authors was generally relegated to an imaginary second division, considered inferior to their predecessors in the Augustan age, especially Vergil. The Flavian poets were at best seen as representative of a “silver” age, at worst as proof of an irreversible trend towards hopeless decadence in Latin literature. In the latter part of the twentieth century, however, this dismissive approach began to evolve significantly, as a series of important studies

⁷For examples, see [Jenny, 1976], [Riffaterre, 1979], [Riffaterre, 1981], and [Genette, 1982]. For a summary history of these terms and concepts see [Allen, 2011]; for their use with respect to Classical literature, see [Pucci, 1998], [Edmunds, 2001], and [Baraz and van den Berg, 2013].

⁸Scholars have tried variously to categorize the multiple kinds of text-reuse, mostly referring to the shape or the function of these practices. Such attempts include [Conte, 1974], [Barchiesi, 1984], [Conte and Barchiesi, 1989], [Hinds, 1998], [Edmunds, 2001], [Thomas, 1989], and [Fowler, 1997].

⁹Amongst the scholars working in the field of the reception theory, see [Riffaterre, 1979], [Riffaterre, 1981], [Iser, 1985], and [Eco, 1979].

¹⁰Statius’ *Achilleid* was interrupted by the author’s death partway through book 2, but clearly was intended by the author to be much longer—he claims in book 1 that the poem will recount the entire life of Achilles.

inaugurated a radically new evaluation of these texts.¹¹ A relevant role in this process of renewal has been played by the appreciation of the sophistication of these authors' allusive techniques, and there can be little doubt today that in the field of complex intertextuality much work remains to be done on the corpus of surviving texts at our disposition.

Although within the Graeco-Roman literary system any text's network of intertextual connections will inevitably extend beyond the bounds of a single literary genre, for the purpose of our article we will limit our focus to relationships between texts within the same generic framework, i.e. epic poetry. In our computational analysis, we consider the three full-length epics of our Flavian authors, plus three canonical Latin predecessors of comparable size and thematic structure, from which our authors of interest all draw source material for their intertextual connections: the *Aeneid* of Vergil (written *circa* 19 BCE), the *Metamorphoses* of Ovid (ca. 14 CE), and the *Bellum Civile* of Lucan (ca. 65 CE).

There is no simple, comprehensive definition of epic poetry because of the organic nature of the genre and its long history.¹² Nevertheless, we feel confident that scholars are broadly in agreement as to the following heuristics: that an epic poem is a narrative text, considerable in scale, structured in sections or books, written in hexameters; further, that it consists of a story, told by an omniscient narrator, whose action takes place in a mythical or historical past made up of gods, heroes and/or historical characters. Within this framework, which encompasses some very different works of classical literature, the three Flavians can claim to belong to the corpus of what we can call in a rather old-fashioned definition "heroic" epic, whose ancestors are obviously the Homeric *Iliad* and the *Odyssey*. "Homer has shown in which rhythm the successes of kings and leaders, and ominous wars can be written," prescribed the Roman poet Horace in his *Art of Poetry*.¹³ Recognizing its debt to the genre's originator, post-Homeric epic has continuously imitated and challenged the Homeric poems, which have lastingly dominated the literary landscape of Greece and Rome both directly and through the mediation of the scholarship of the Hellenistic age.

1.4 Narrative structure and epic type-scenes

In describing the scope of phenomena we hope to characterize, we use the collocation "narrative structure" in a very broad sense, including features as different as the articulation of the poem in books and scenes, the alternation of multiple types of "discourse" (narrative, speeches, *excursus*, similes, apostrophes, *ekphraseis* and so on), and the disposition of repeated story elements often called "type-scenes," or "typical scenes." It has been recognized for a long time now that Homeric poems are structured by large-scale story patterns [Edwards, 1987]¹⁴ as well as by "recurrent block[s] of narrative with an identifiable structure, such as a

¹¹ Over the last thirty years or so, one of the most valuable and exciting trends in the study of Flavian epic has been the publication of numerous commentaries on individual books. Some of these have achieved almost classic status in our field—among them, notably, [Dewar, 1991], [Smolenaars, 1994], [Zissos, 2008], and [Littlewood, 2011]. It is important to note also that these commentaries are building on a long tradition of earlier commentaries going back to the Renaissance (as investigated in detail for the case of Statius' *Thebaid* by [Berlincourt, 2013] and on such valuable modern contributions as [Juhnke, 1972].

¹² For instance on the discussion about the 'epic identity' of Ovid's *Metamorphoses* and, in general, the features which constitute the essence of the genre, see [Hinds, 2000].

¹³ *Res gestae regumque ducumque et tristia bella / quo scribi possent numero, monstravit Homerus* (Hor. *ars poetica* 73-4).

¹⁴ This idea has seen application beyond the bounds of epic of course—for example by [Levi-Strauss, 1955]. More specifically to the Homeric tradition, [Edwards, 1987] gives one list of such story patterns

sacrifice, the reception of a guest, the launching and beaching of a ship, the donning of armour” [Edwards, 1992]. This latter definition encompasses a set of concrete, stereotypical actions, taking place in passages approximately the size of a paragraph. It corresponds to the units called “themes” by [Lord 1960], and the English equivalent of [Arend, 1975]’s *Typische Szenen*.

Albert Lord identified the use of type-scenes as a defining characteristic of oral-formulaic composition, noting analogous structures in the poetry of contemporary oral cultures outside the Graeco-Roman tradition [Lord, 1960]. Edwards who supports the theory of an oral composition of the Homeric poems has pointed out [Edwards, 1992] that these type-scenes may be assimilated to the Homeric formulae first identified with the oral style by [Parry, 1928]. Recently, scholars including [Rubin, 1996] and [Minchin, 2001] have presented evidence from cognitive science that certain narrative structures native to or informed by oral tradition may be shaped by the way the brain optimises memory and perception, and [Lyne, 2016] extends at least some of the underlying intertextual mechanisms more explicitly to non-oral literature.

Whatever the origins of these structures in Homer and in the oral pre-history from which he emerged, it is important for our purposes to emphasize that within the framework of a mainly written culture—that is, Hellenistic and Roman literature—these scenes come to play a different, highly literary role, becoming the privileged places for confrontation and consequently for intertextual connections. Following [Edwards, 1992], who claims that “in narratological terms, an amplified type-scene is not necessary to the ‘story’, . . . but is part of the ‘discourse,’” we assume that these scenes can be considered as equivalent to story-telling devices in an essentially narratological framework.¹⁵

II CASE STUDY: APOLLONIUS RHODIUS, VERGIL, AND VALERIUS FLACCUS

We will illustrate the role that similarities in narrative structures can play in marking intertextual relationships with a passage from the Flavian epic poet Valerius Flaccus, who wrote the story of the Argonauts in 8 books.¹⁶ It is well established [Ripoll, 1998] [Lieberman, 1997, xxxii-xlvi], and indeed has been accepted by scholars at least since the beginning of the 20th century [Stroh, 1905], that Valerius’s poem is profoundly influenced by the *Aeneid*. For instance Philip Hardie points out that Valerius “shows himself . . . Virgilian . . . in his use of symmetry and repetition to give form to his poem” [Hardie, 1990, 5-9]. At the same time, the storyline and structure of Valerius’ poem are very close to his Hellenistic Greek predecessor, Apollonius of Rhodes, who wrote his own *Argonautica* in four books sometime during the

(withdrawal, devastation, eternal return [order into disorder, then return to order], loss of hero’s best friend, quarrel between heroes, rescue of a man/hero by another one...).

¹⁵ Even within the oral context, there exists a diversity of opinion as to the level of control and artistic freedom exerted by the poet in deploying narrative structural elements. Thus [Minchin, 2001, 39-42] argues that typical scenes in Homer proceed from an artistic development of existing cognitive functions, and stresses the creative originality of the artist. For a large-scale narratological perspective, see [Bal, 1978]. As epic is a narrative genre and also an ancestor, so to speak, of the novel, it is a frequent subject of narratological theory—see, for instance, [de Jong, 2014]’s Chapter 6, “Narratology and Epic,” and her bibliography. As de Jong has pointed out, “in fact, narratology can be said to have started in antiquity, when a number of central concepts were developed” [de Jong, 2014, 3].

¹⁶ In fact the poem was unfinished, or at any rate it has come to us in fragmentary shape. Nevertheless we agree with [Schetter, 1959], who thinks that the poem was designed to be in 8 books and that we have missed in the worst case a few hundred lines.

third century BCE. As Andrew Zissos puts it, if “Valerius makes Virgil his principal model on the linguistic, conceptual and thematic levels,” then Apollonius is his primary source “on the level of the plot” [Zissos, 2009, 508].¹⁷

In particular, it is now clear that Valerius is organizing his text in two different halves (4 + 4 books) repeating the shape of the Vergilian *Aeneid* (6 + 6 books) and, at the same time, following the thematic bipartition of Apollonius Rhodius' *Argonautica* (2 books on voyage + 2 books on love, i.e. the meeting between Jason and Medea and its consequences). Consider the parallel thematic sequences which cluster around the poems' respective midpoints, schematically represented in Table 1 below. These three poems share not only a macro-structure (i.e. sea voyage, love story) but also a number of much more specific type-scenes, including departure, arrival, burial and invocation to the Muse, in each case laid out in more or less the same order at the turning-point of each epic.

Apollonius <i>Argonautica</i> 2-3		Vergil <i>Aeneid</i> 7	Valerius Flaccus <i>Argonautica</i> 5	
2.815-898	Mariandyni; death and burial of Idmon and Tiphys; Ancaeus chosen as helmsman.	7.1-7	<i>BOOK DIVISION</i>	<i>BOOK DIVISION</i>
			Death and burial of Caieta; departure .	5.1-70a
2.899-1259	Departure and voyage along southern coast of the Black Sea; Argonauts pass Carambis and Prometheus.	7.8-24	Voyage along the coast; Trojans pass Circe's land.	5.70b-176
2.1260-1286	Night and arrival in the Phasis; prayer of Ancaeus.	7.25-36	Dawn and arrival in the Tiber.	
	<i>BOOK DIVISION</i>	7.37-106	Invocation of the Muse Erato and the situation in Latium.	5.177-216
3.1-5	Invocation of the Muse Erato.			Evening and arrival in the Phasis; prayer of Jason.
3.6-166	Divine intervention: Hera, Athena, Aphrodite, and Eros; love.	7.107-147	Meal; prayer of Aeneas; sacrifice.	5.217-277
				Invocation of a Muse (<i>dea</i>) and the situation in Colchis.
3.167-212	Argonauts make their way to the city and palace of Aietes.	7.286-640	Trojans make their way to the city and palace of Latinus.	5.278-295
				Divine intervention: Juno and Minerva; war.
				5.296-328
				Argonauts make their way to the city and palace of Aietes.

Table 1. Valerius Flaccus' *Argonautica* and two of its principal intertextual sources share a common sequence of thematic segments, grouped around the mid-point of the narrative.

¹⁷On Valerius Flaccus' intertextuality in general see [Hershkowitz, 1998, 35-104]. For further discussion of these issues, see [Nelis, forthcoming].

By this close alignment, Valerius Flaccus is telling his reader that he has built his story against both the *Aeneid* and the Greek *Argonautica*, and moreover that he has himself read the relationship between Vergil and Apollonius. In acknowledging and extending the existing allusive relationship between his two sources, he invites the reader to make sense of this intertextual background when trying to understand these sections of the poem. Unfortunately, the digital search tools currently available, such as *Tesserae*, being based on text re-use as the model for allusion, will not succeed in highlighting this kind of extremely meaningful relations because the texts are verbally connected at a very thin level, sharing only one word or lemma.

Let us look for a moment at one specific allusive locus within this sequence of interrelated passages. All three texts interrupt the narrative at its midpoint with an invocation of the Muse, a traditional epic gesture signalling a significant new direction in the story:

Apollonius Rhodius *Argonautica* 3.1
Εἰ δ' ἄγε νῦν, Ἑρατώ,
“Come on, now, Erato”

Vergil *Aeneid* 7.37
Nunc age . . . , **Erato**,
“Come on, now, Erato”

Valerius Flaccus *Argonautics* 5.217
Incipe **nunc** cantus alios, **dea**,
“Begin now, goddess, other songs”

Vergil addresses the Muse Erato in what is not only a translation of Apollonius' Greek, but very nearly a word for word transliteration. Valerius Flaccus signals his awareness of the Vergil-Apollonius link with a more oblique reference, avoiding the Muse's name and addressing her simply as “goddess.” At the same time, in a kind of one-upsmanship, he exaggerates Vergil's delay of the invocation after the central book division—what was for Apollonius the first line of his poem's second half becomes line 37 in Vergil and 217 in Valerius. To a human reader engaged in close reading, the three invocations to the Muse are unmistakably related. Yet a lemma-based automatic reading would lose the match, if not between Apollonius and Vergil,¹⁸ certainly between Valerius Flaccus and the other two. What is needed in order to catch the significance of Valerius' faint verbal repetition is sensitivity to the surrounding context—the alignment of the surrounding themes, proximity to the book division, and the identity of the discursive mode.¹⁹ The question posed by the computational

¹⁸It might not be strictly appropriate to call any automatic match between a Greek source and a Latin target text simple “lemma” matching, since the lemmata are obviously different. While such cross-language matching is not reliably provided by the tools online today, an experimental version Greek-Latin search is available through [Tesserae]. This search in fact is based around the notion that each lemma in Greek ought to have some close equivalent in Latin, and, having calculated the best “translation” of each Greek word, then proceeds as a lemma-based search [Gawley et al., 2014]. In theory, much more sophisticated machine translation techniques exist, and we hope to see their application to this problem soon.

¹⁹As Jeffrey Wills shows, “[a]lthough allusions through content or context are obvious enough,” the combination of multiple features can strengthen connections where any one feature might not be sufficient [Wills, 1996, 18]. This includes in particular “positional marking (in relationship to line- or

component of this project is to what degree these narrative elements can be successfully extracted from the text.

DIGITAL METHODS

3.1 General goals

The digital aspect of this research has two related goals. The principal question is whether what we call “thematic features” can be successfully integrated into the automatic scoring of intertextual search results on a purely verbal level. That is, does similarity of theme, narratological structure, or other contextual features contribute in a systematic way to the reader’s perception of a verbal allusion’s literary significance? If so, this kind of similarity can be used in the way we currently employ the frequency and proximity of specific matching words to predict the relevance of results [Forstall et al., 2015]. A closely related, secondary goal is to extract thematic features from the text automatically, that is, unsupervised or semi-supervised tagging of the corpus for at least the most important type scenes, narrative devices, and so on. The work reported here pertains principally to this latter, secondary goal; even so, its development and its evaluation always take place within the larger theoretical context of evaluating allusivity.

It is important to note that in theory our principal goal, predicting allusive significance, does not depend on achieving the secondary goal of extracting thematic features automatically. In particular, it is plausible from the point of view of the *Tesserae* scoring system to employ a set of thematic tags assigned manually. Moreover, although complete classification of the themes of the classical corpus by hand represents a monumental task, it is not entirely out of the question, as it would be, for example, in a living and growing contemporary canon. For at least some subsets of Greek and Latin literature, in fact, such hand-tagging is being attempted. Thus, the *Epische Bauformen* project, directed by Christiane Reitz at the Universität Rostock, for example, is embarking on a mission to create a typology and catalogue of the type scenes of Greek and Latin epic poetry over the next three years, and envision making the entire catalogue available afterwards in a machine-accessible format [EB]. Likewise, the collaborative interdisciplinary project *Memorata Poetis*, bringing together researchers at several Italian universities, is currently engaged in manually tagging the major themes in a corpus of Latin, Greek, Italian and Arabic epigrams [MP]. This database, which counts Paolo Mastandrea among its principal investigators, is currently online and searchable, although still in progress, and is apparently to be made compatible with the [MQDQ] allusion search engine.

At the same time, the ability to automatically extract such information from the texts with much less human effort would provide an important complement to such databases. First, although as classical philologists we have the good fortune to work on relatively small and closed corpora, to devote this level of attention to the entire extant classical canon would require more resources than our field can afford at present. Indeed, one of the lessons of the digital humanities is that the greatest potential to learn something new about our field often comes from the examination of those more marginal and less-frequently studied texts that are often the last to receive careful treatment by hand.

book-boundaries)” [Wills, 1996, 22]. Putting these together, “sequential order seems another possible element of this grammar of allusion” [Wills, 1996, 27].

Second, it is important that we look beyond the borders of our discipline, and make the case that our research has relevance to the study of literature more generally. We would hope that the search for intertextuality would be equally interesting when applied to much larger, open corpora, where manual tagging was not a feasible option. At least some branches of intertextual theory specifically argue that literary activity takes place in the vast, open system of our total semiotic environment, including texts outside the genre in question, everyday language, and even non-linguistic sign systems. [Fowler, 1997], for example, notes the growth of the view that, “intertextuality is a property of language-and of semiotic systems in general-not simply of literature.” A method that could be generalized to something larger than our own proprietary, manually-prepared corpus would be potentially valuable to others, and at the same time invite collaboration and comparison that could benefit and add context to the present study.

Finally, although the automation of feature-extraction is proposed as secondary to and in support of the principal goal of improving our scoring functionality, we believe that, if successful, it also represents an important end in itself, namely the modeling of readerly behavior. To the extent that we can use surface textual features to identify where readers are most likely to identify coherent and transcendent themes, we will better understand how such themes evolve out of the continuous, but ephemeral, stream of poetic composition, performance and transmission. This is a particularly pertinent line of research in the study of classical epic, which draws so much, even in its later and highly literary forms, from the oral tradition. While it is clear that the Roman poets, who in emulating one another also emulated Homer, felt that type scenes were a necessary and defining feature of the epic genre, nevertheless the question of how closely two scenes must resemble each other in order to belong to the same “type” was for them, as it is for us, a subjective one. In attempting to model such judgments, to identify the textual features on which that subjectivity is nourished, we shine new light on what it is to read and to make meaning from these works.

Work on the automatic classification of type scenes is still exploratory, but we plan a two-part program, comprising supervised and unsupervised approaches. At the moment we are undertaking the beginning of the unsupervised approach, which consists in four discrete steps: sampling, feature extraction, dimensionality reduction, and clustering. At the same time we are preparing hand-tagged data for training the supervised approach, also to be used for testing of both approaches.

3.2 Sampling

Sampling has traditionally been a thorny problem in the analysis of literary texts because the questions posed are often particularly finely-grained, as here, while at the same time the amount of text available is often small in comparison with other stylometric applications such as forensics or broader cultural studies. In particular, [Eder, 2015] cautions that for Latin poetry, samples of several thousand lines are necessary in order to extract a reliable stylometric signal from a work, and also underlines the need for random rather than sequential sampling to ensure that the samples are representative. While we take these considerations very seriously, they are somewhat complicated by the nature of our object of study. The bulk of such stylometric feature extraction aims to characterize a unique, and relatively constant signal associated with a particular, author, work, or genre. Here, however, we are interested in bringing to light the internal heterogeneity of individual works, and, by contrast, the stability across the corpus of the different component structures employed in their composition. We therefore continue to use passages of consecutive lines in building our samples, since a

defining feature of the epic typical scene is that it is in fact one, coherent passage of text.²⁰ Our samples must necessarily be small, as the scenes we hope to identify are often fewer than 100 verse lines in length. At the same time, the degree to which automated feature extraction methods will tolerate such fine-grained sampling is an open question. [Scheirer et al., 2016] considered the use of topic modelling for allusion in Latin and offered some positive results tempered by cautionary examples. The authors found that two popular implementations of Latent Dirichlet Allocation, Gensim [Řehůřek and Sojka, 2010] and MALLETT [McCallum, 2002], while potentially useful at capturing something very like the type of similarity we are looking for, became unstable at sample sizes below several hundred lines, but that Gensim's implementation of Latent Semantic Indexing produced stable results for passages of the size considered here [Scheirer et al., 2016].

The other important consideration in sampling is the irregularity of the passages we hope to find. Our type scenes vary widely in length, and may be found anywhere within the texts we are searching. With too-large, and fixed-size samples, there is a danger that scenes falling between two samples will be so diluted by the surrounding material that their signals are unrecognizable. One solution would be an ability to accommodate irregularly sized samples, and ideally to identify the borders automatically through a bootstrapping process where samples were successively re-aligned so as to maximize the differences in their signals. A second option is to automatically detect editorial indications of scene change, such as whitespace between paragraphs, punctuation marks indicating speech boundaries, and the numbering of books and verses. Much of this information is already in machine-readable form in the TEI XML documents, downloaded from the Perseus Digital Library [Perseus], which form the basis of our corpus.²¹ It must be noted, however, that use of this information resituates our unsupervised approach rather under a "semi-supervised" rubric.

For the work discussed here we have used an entirely unsupervised approach to sampling, taking samples of 50 consecutive verse lines. As a step towards the bootstrap method proposed above, we created 10 independent runs by changing the sampling window offset in five-line increments (i.e. with sample 1 beginning, respectively, at line 1, 6, 11, and so on). Each of these runs was treated independently through the remaining steps of the experiment—feature extraction, dimensionality reduction, and clustering. After classification, the results across all runs were compared on a line-by-line basis. Our hypothesis was that those lines which tended to have more stable classifications across multiple runs would represent regions of the text more firmly embedded within typical thematic material, while lines whose classification varied with the sampling offset would be more likely to represent transitional regions or otherwise less-stereotypical material. While in initial experiments the most stable areas have tended to correspond to typical scenes, we don't yet have enough evidence to claim that the method works generally, or to take the next steps towards fully automating the identification of irregular sample bounds.

²⁰ Homer famously portrays all action as sequential—[Scott, 1921, 158-159] goes so far as to claim that oral style made it "impossible for the poet to picture events as taking place simultaneously." While this seems unfair to the narrative sophistication of the *Odyssey* in particular, and while Homer's Hellenistic and Roman successors continue to refine and develop the handling of epic time, it is safe to assume that we will not encounter the type of interleaved narration of independent scenes found, for example, in modern novels.

²¹ Preliminary work described here used XML versions of the texts downloaded directly from the [Perseus] website between 2010 and 2014; we are currently updating our texts to correspond with the corpus available on GitHub at <https://github.com/PerseusDL/canonical-latinLit>. The specific text files we used are also available on GitHub at <https://github.com/cwf2/dh2016>.

In comparison to this process, we have also tested, in a different but related study, a sample set defined by editorial scene boundaries, as taken from a modern published commentary. Compared with fixed-length samples, in that experiment the samples defined by editorially defined boundaries tended to show stronger and more heterogeneous stylistic signatures [Forstall and Galli Milić, 2015].

3.3 Feature extraction

The features we are using for this work are relatively simple, bag-of-words lexical features based on term frequency / inverse document frequency (tf-idf) weights of the lemmatized Latin words. We exclude a stoplist of words that occur in more than half of all samples; although the precise list can vary based on the sampling, this tends to amount to the top 50-60 lemmata. While for the most part these are function words (*et, qui, in, sum, hic, tu,* and the like), a stoplist based solely on the number of samples in which words appear also includes generically important, but ubiquitous, content words such as *arma, bellum, magnus,* and *deus*. It is not inconceivable that such words, though present throughout the work, play special roles in defining particular type scenes, and should thus be included. This remains to be tested in future experiments. Finally, *hapax legomena* are also excluded. Intuitively, this seems an obvious step, but it is worth noting that in fact the pointed use of Homeric *hapax legomena* was an established intertextual gesture for later authors. It is possible, in fact, that a word found nowhere else in a specific canon nevertheless has a set of correspondences with other rare words in analogous passages, words which share not a common lemma or meaning but merely their rarity—what [Wills, 1996, 22] calls “frequency marking.”

3.4 Dimensionality reduction

Each sample was thus characterized by the tf-idf weights of all remaining lemmata. Again, while the specific number of terms varied somewhat with the sampling, this was on the order of seven or eight thousand for our corpus of six Latin epics. We used principal component analysis (PCA) to reduce this to the first 500 principal components. The use of PCA greatly reduced compute time for classification; we are still in the process of testing whether it reduces or improves classification accuracy over the original feature vectors, as well as the optimal number of dimensions for the reduced features. It would seem natural that methods from the field of topic modeling should be appropriate to this study. While initially wary of the small sample size, we are encouraged by the results of [Scheirer et al., 2016] and more recently [Köntges, 2016], and we are currently working to adapt these methods to our own pipeline.

3.5 Clustering

Samples in a given run were clustered by the resulting feature vectors using k-means clustering. We had no *a priori* expectation as to the number of clusters, and so we attempted to determine the most appropriate value for *k* experimentally. We tested values between 2 and 12. For each value of *k*, we repeated the clustering process 15 different times (Figure 1). Because initial cluster centers were selected randomly, the resulting assignment of samples to clusters could differ among the 15 repetitions. We tested agreement between repetitions by doing pairwise comparisons of the 150 resulting cluster assignments (10 different runs, with 15 repetitions for each, Figure 2). The group assigned to a given sample was applied to each of the lines that made up that sample—considering lines rather than the samples allowed us to compare the tendency of specific lines cluster together even among runs with different sampling offsets. To measure agreement between any two clusterings we used the adjusted Rand index, which gives a value between 0 and 1 for two sets of sample labels, 1 representing

perfect agreement, and 0 representing the level of agreement expected by chance. The stability of results for each value of k was approximated by an average adjusted Rand index for all pairwise comparisons. Preliminary results suggest that five or six clusters represent the limits of this method; higher values of k produce increasingly unstable classifications (Figure 3). However, this process was not combined with the bootstrapping method for determining sample boundaries. It is possible that with more precisely defined samples, a greater diversity of clusters could be reliably distinguished.

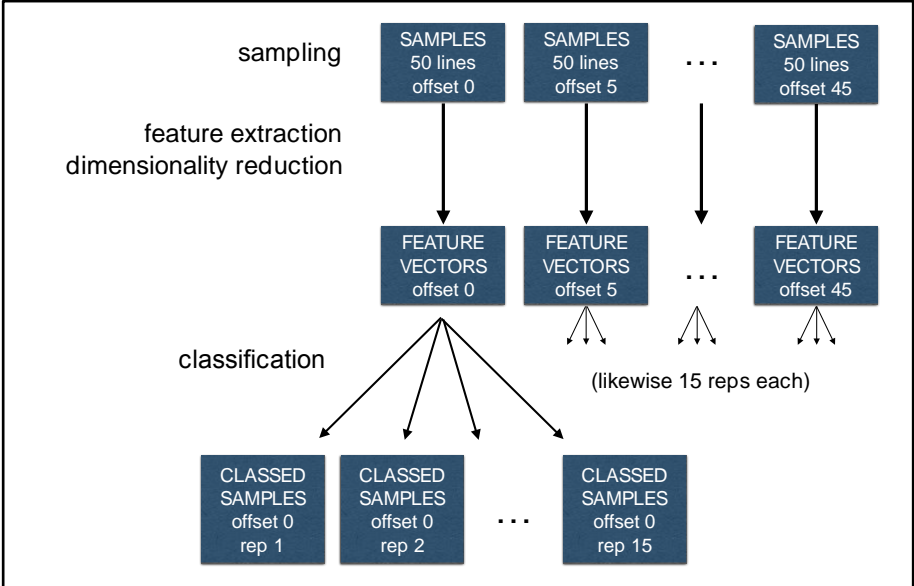
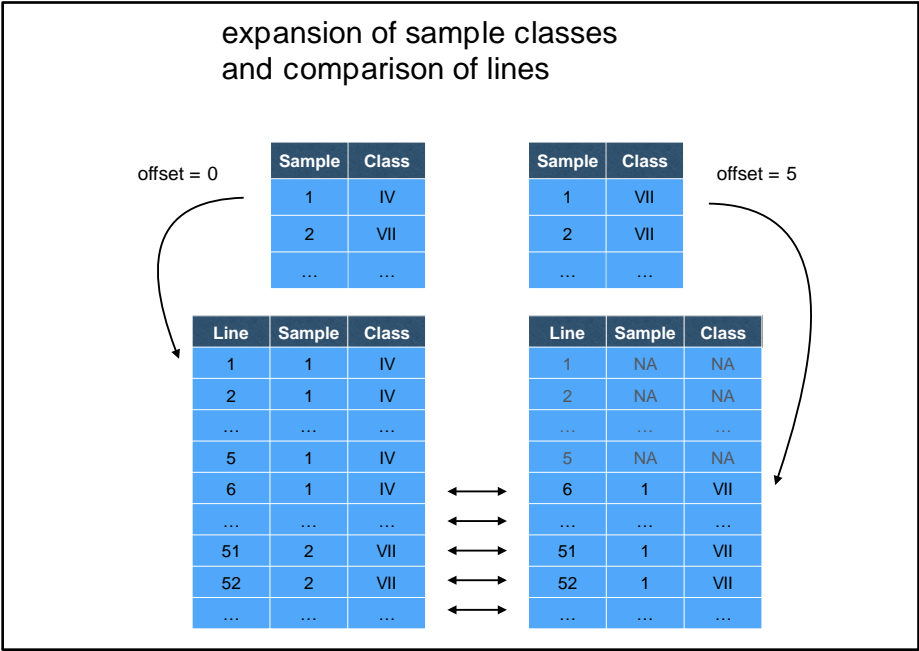


Figure 1. Overview of sampling and clustering. Sampling was repeated in 10 independent runs, increasing the sampling window offset in increments of five lines. For each run, clustering was performed in 15 repetitions. In total, 150 different cluster assignments were produced.

Figure 2. Comparison of clustering results. The cluster assigned to a given sample was extended to each of the 50 lines that comprised the sample. Pairwise agreement between clusterings, on a line-by-line basis, was measured using the adjusted Rand index.



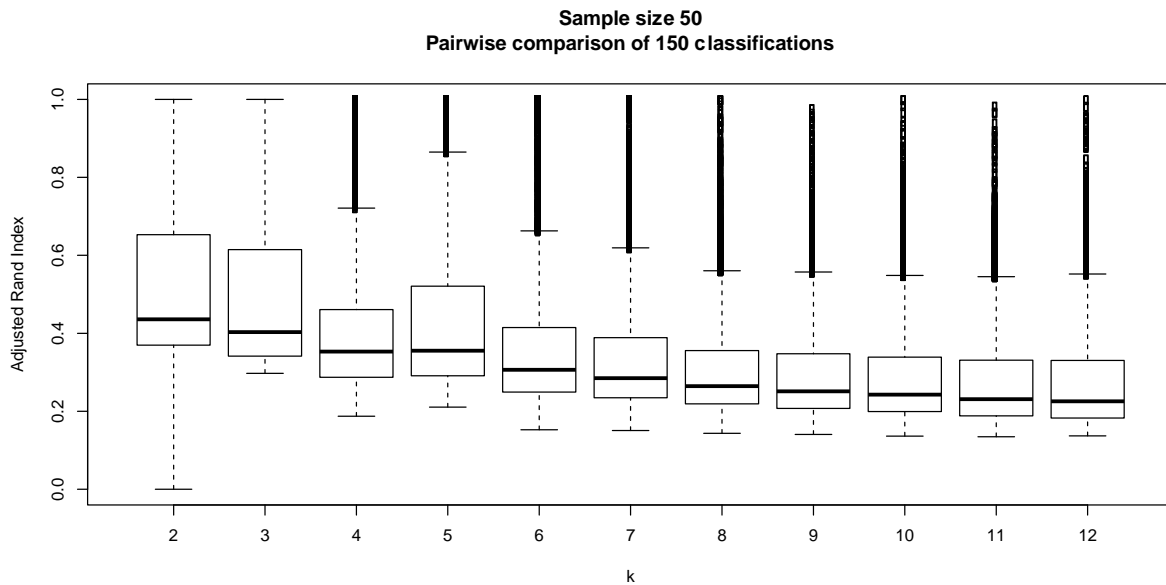


Figure 3. Sample-by-sample stability of repeated k-means clustering, for varying values of k . For each value of k we measured the adjusted Rand index for every possible pair of 150 different clusterings. Each column thus represents 11175 values of this metric. Mean agreement decreases with increasing k ; in initial experiments we preferred $k=6$ as a compromise between stability and precision.

3.6 Authorship signal adjustment

The initial clustering experiments demonstrated that even with small samples of consecutive verse lines, and even with a significant stoplist of most frequent words removed, the predominant factor in clustering was authorship. Figure 4 shows the distribution of samples within the reduced feature space after PCA; even with only two principal components it can be seen that samples fall into distinct regions according to author. In particular, Ovid's *Metamorphoses*, Lucan's *Bellum Civile*, and Silius Italicus' *Punica* are localised in separate portions, around the periphery of the space. Vergil's *Aeneid*, Statius' *Thebaid*, and Valerius Flaccus' *Argonautica* overlap somewhat more in this figure, but still can be largely distinguished using further principal components.

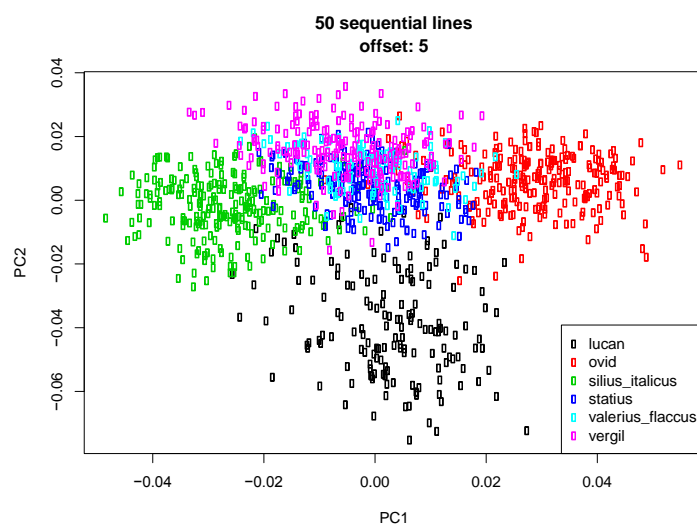


Figure 4. Distribution of samples within the PCA feature space. The symbols here show true authorship for each sample. It is immediately apparent that samples by the same author tend to be near one another in the feature

space. For Lucan, Ovid, and Silius Italicus in particular, the first two principal components alone allow for almost complete separation by author.

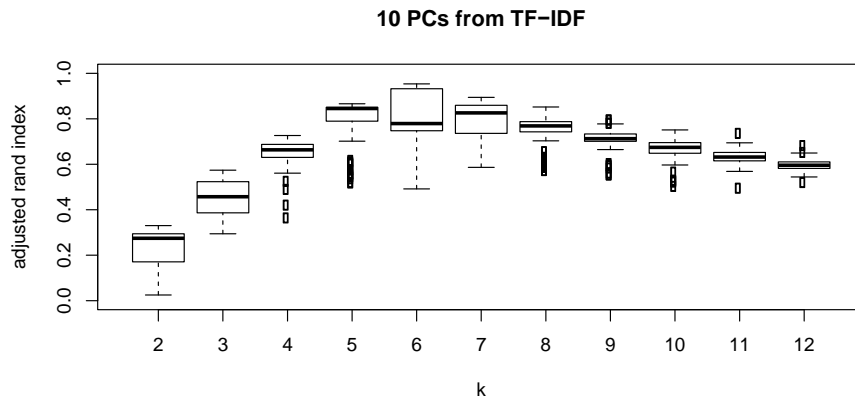


Figure 5. Agreement between clustering by k-means and true authorship, by k . For a wide range of k , unsupervised assignment to clusters corresponds highly with true authorship; this effect peaks around $k=6$, the true number of authors. Each sample was represented by the first 10 principal components derived from tf-idf weights, as described above. It is clear that automated clustering for this feature set primarily groups samples by author.

The results of k-means clustering of these samples correlates strongly with true authorship (Figure 5). In particular, agreement between clustering and true authorship peaks at values of k around 6, the number of authors in the corpus.

In order to compensate for the effects of authorship, we created a “mean author signal” for each work, subtracting from every sample in that work the mean vector for the work as a whole. This had the effect of centering each work within the feature space, while retaining the heterogeneity of its component samples (Figure 6). This was done after feature reduction, and the new, author-adjusted feature set was then put through the clustering stage as described above. This procedure eliminated the correlation between clustering and true authorship shown in Fig. 5 above.

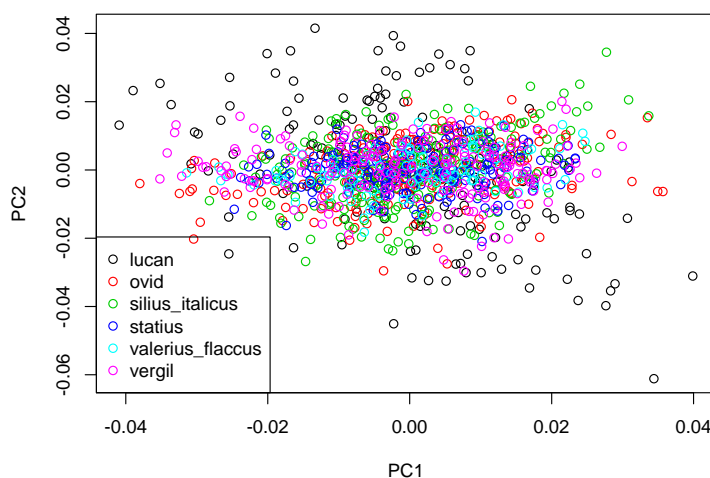


Figure 6. The same feature space as Fig. 4, but with samples adjusted for mean author signal. Each cloud of samples representing a single author’s work has been independently recentered, so that all authors overlap, while preserving the relative distances between samples within each author.

3.7 Verse-by-verse stability measurements

Our hypothesis is that certain passages in these poems share stereotypical language that identifies them as type scenes belonging to specific, canonical types—*tempest at sea*, *catalogue of warriors*, *lion among the herds*, etc.—while the bulk of each poem is dominated by original, creative content. Our samples are taken from across the corpus, and recentered by author. The feature space now represents the internal heterogeneity of these poems, the diversity of language which occurs even within each author. It is our hope that in this space we will find small clusters of samples that partake of common thematic language, but it is also our expectation that most of the samples will not cluster reliably, since the work of the poet is generally to “make it new,” as Ezra Pound put it, rather than simply to rearrange static thematic building blocks.

We test this hypothesis by comparing the stability of the labels assigned by k-means clustering line-by-line through each work. As in the selection of k , we use stability across multiple re-clusterings as a clue to the coherence of the underlying clusters. Passages from different works that partake of a common, special vocabulary connected with a particular typical scene ought to be placed reliably in the same cluster across many repetitions of the sampling and clustering process. Although the lines of such a passage may be grouped differently as the sampling window moves from run to run, if the passage is long enough, we should see consistency overall when we consider a large number of trials at different offsets.

As above, we apply the cluster label for each sample individually to the lines that make it up, allowing us to compare labels between runs even when they have different offsets. We then measure the degree to which each line of the corpus tends to be assigned to a reliable cluster in randomly selected pairs of trials. Of the 150 different sets of cluster labels we have for the verse-lines of the corpus, we choose two at a time and use the R package *mclust*'s `classError()` function to create the best possible alignment of cluster labels between the two trials. Those lines whose labels cannot be harmonized are labeled as misclassified by `classError()`. Over a large number of such comparisons, we calculate the rate at which each line of the corpus is assigned to the same cluster in both trials, according to the harmonized labels. We hope that this success rate will serve as a clue to regions of stereotypical language.

3.8 Hand tagging

In further work, we will systematically compare the consistency of cluster labels line by line to manual scene labels assigned by human readers. A definitive, “gold standard” database of scene labels is for the moment out of our reach, and thus validation of these results remains a work in progress. Our current local database of hand-tagged type assignments was developed entirely by two post-doctoral researchers, and limited to a somewhat arbitrary repertory of scene types, of which those with the most numerous tagged examples are *tempest* and *battle* scenes. While such a limited approach, relying on the personal judgments of a very small set of readers, has proven sufficient for prototyping to this point, we will ultimately need a more diverse data set for robust training and testing. One potential source is the *Epische Bauformen* project mentioned above [EB]. Two other sources to which we have turned in the past are the work of graduate students, performed for credit in seminars, and published commentaries which, at least in some cases, contain internal classification schemes or habits of description consistent enough to be amenable to coding and digital transcription.²²

²² See <http://tesseract.caset.buffalo.edu/blog/collected-benchmark-sets/> for datasets used by Tesseract, including student-entered data contributed by our own group.

To a certain degree, the influence of commentaries is already present in the passages tagged so far, in that the classification of scene types is loosely based on that of [Edwards, 1987] and the boundaries of scenes are in most cases determined by the editorial choices of the editions used. That is, the majority of the researchers' attention is focused on evaluating which of the scene types best suits the passage under consideration, the boundaries of which are considered closely only in cases where ambiguity is apparent. In a future, larger-scale phase of data collection, this methodology will have to be more strictly controlled, especially as the broadening of the readership whose judgements are consulted will necessarily increase the complexity of this subjective process.

IVPRELIMINARY RESULTS

Results to this point have shown limited success, suggesting that the unsupervised method for extracting thematic features has potential, but also that it will likely be improved by closer study of the ways in which our human readers recognize typical scenes. Considering the tempest scenes which represent our most complete human-tagged benchmark at this point, we identified the longest example of this type for each of the authors in our corpus. Of these six scenes, five were flagged as containing lines consistently assigned to the same cluster across multiple trials (Figures 7-11) while one showed no difference in cluster stability from the surrounding material (Figure 12).

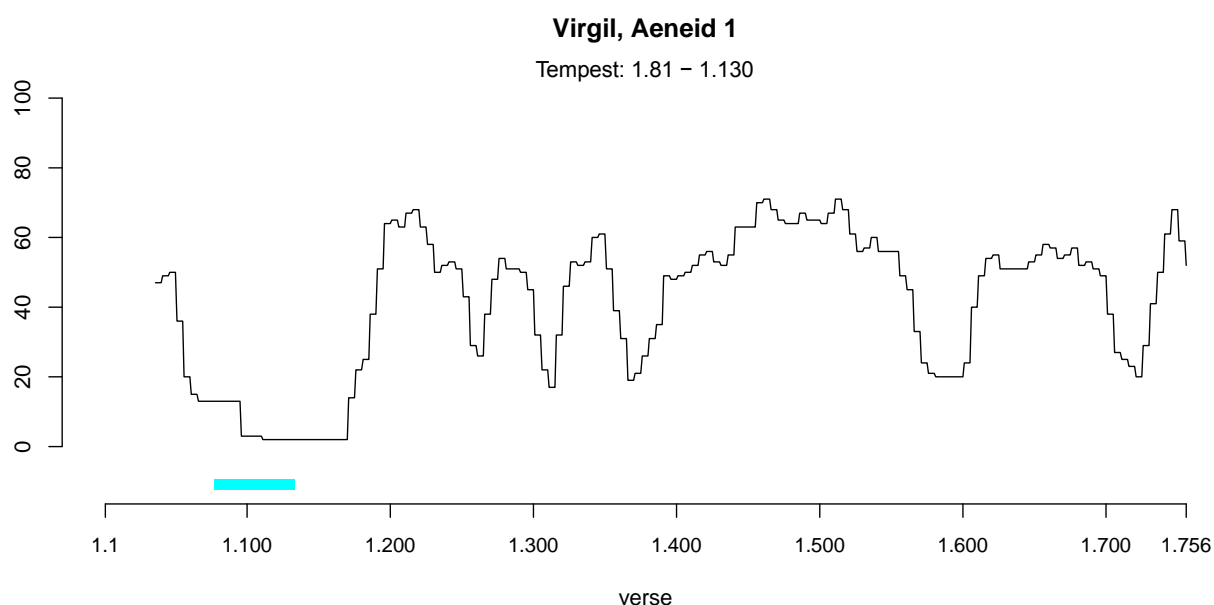


Figure 7. Line-by-line cluster stability across a random selection of 100 trials, for book 1 of Vergil's *Aeneid*. While most trials disagreed about the clusters to which most lines were assigned, nevertheless a few passages were assigned to the same clusters in significant numbers of trials. In particular, a group of lines coinciding with the tempest scene of this book were clustered together consistently across trials, regardless of the sampling window offset. In this plot, the line shows the number of times two randomly selected trials agreed as to the assignment of a certain line, out of 100 such comparisons. The colored bar at bottom shows our ground truth, the lines independently assigned to the tempest scene by human annotators. This is the classic scene in which Aeolus, king of the winds, lets loose a storm at sea at Juno's command in order to hinder and punish the Trojans.

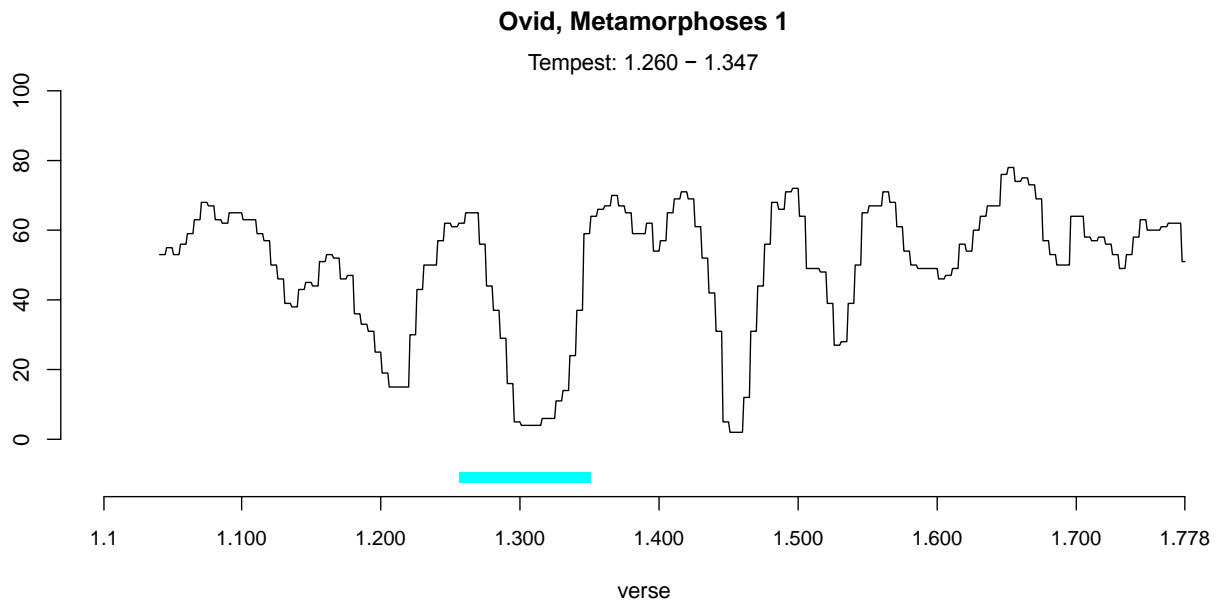


Figure 8. Line-by-line stability measurements, as in Figure 7, detail showing book 1 of Ovid’s *Metamorphoses*. Again, the great majority of the poem did not show any consistency across trials, but a few small passages showed high stability. The scene marked by our human readers here was the story of the deluge, in which Jupiter sets out to destroy the world by flood. Of the six definitive tempest scenes selected for our benchmark, this is in some ways the furthest from the classic epic ideal; nevertheless, its allusive connections to e.g. *Aeneid* 1.81-1.30 (Fig. 7) are well established [Anderson, 1997, ad loc.].

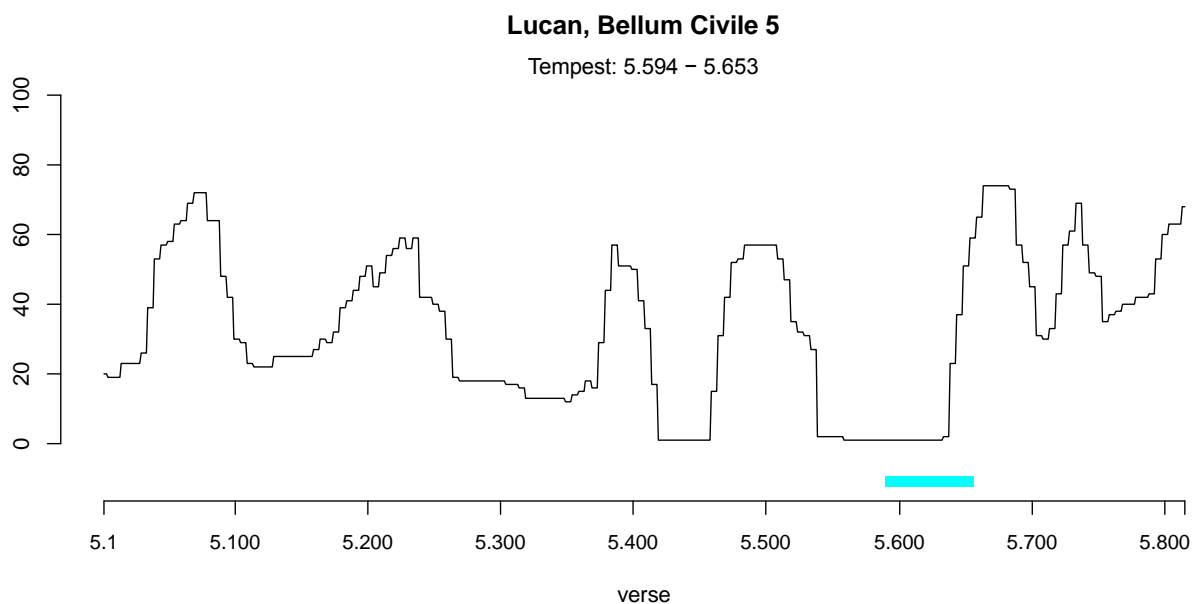


Figure 9. Line-by-line cluster stability, as above. Detail showing the most substantial tempest scene in Lucan’s *Bellum Civile*, lines 5.594-5.653. In this passage, Julius Caesar attempts a crossing of the Adriatic, alone in a tiny boat, but is turned back by a storm and high seas. As above, the preselected passage roughly corresponds with an unusually consistent cluster stability, while most of the poem’s verse lines fluctuate much more.

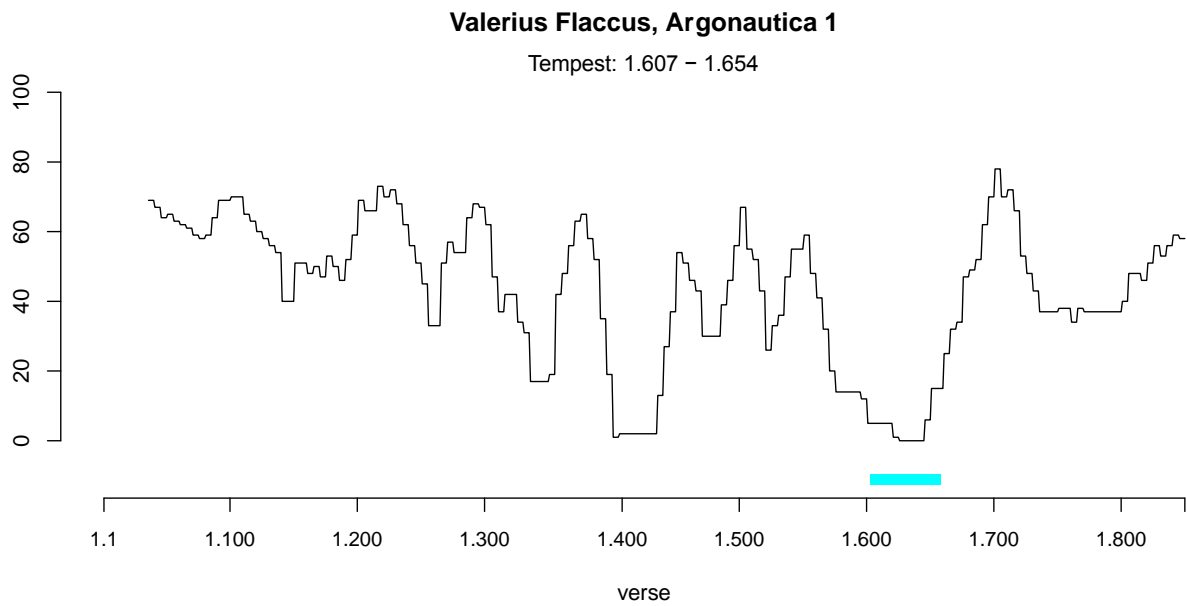


Figure 10. Line-by-line cluster stability, as above. Detail showing the pre-selected tempest scene for Valerius Flaccus' *Argonautica*, lines 1.607-654. As in the *Aeneid* passage above, Aeolus frees the winds to batter the story's heroes—in this case Jason and the Argonauts rather than the Trojans, and at the winds' own urging rather than Juno's. Multiple allusions to the passages of Vergil, Ovid, and Lucan examined in Figs. 7-9 above, as well as with that of Statius below (Fig. 11) are noted by [Zissos, 2008] in his commentary on these lines.

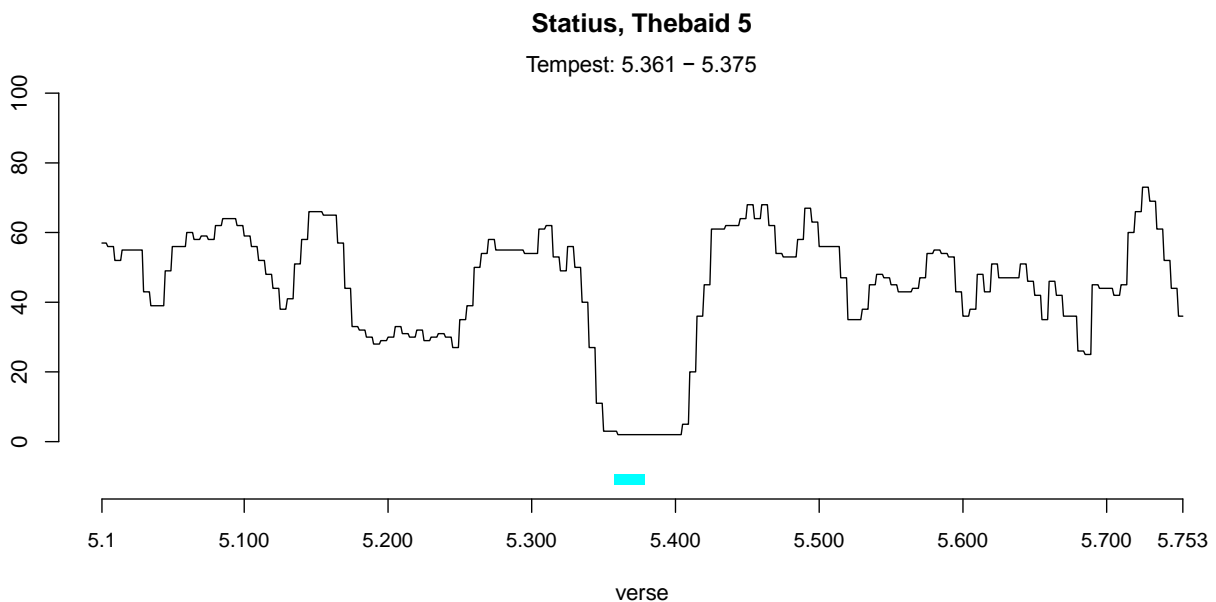


Figure 11. Line-by-line cluster stability, as in the preceding figures. Detail showing Statius *Thebaid*, lines 5.361-375, the longest reader-selected tempest scene for this poem. Again a storm besets the Argonauts, but here the scene is recounted from the point of view of the princess Hypsipyle, who watches from the walls of Lemnos.

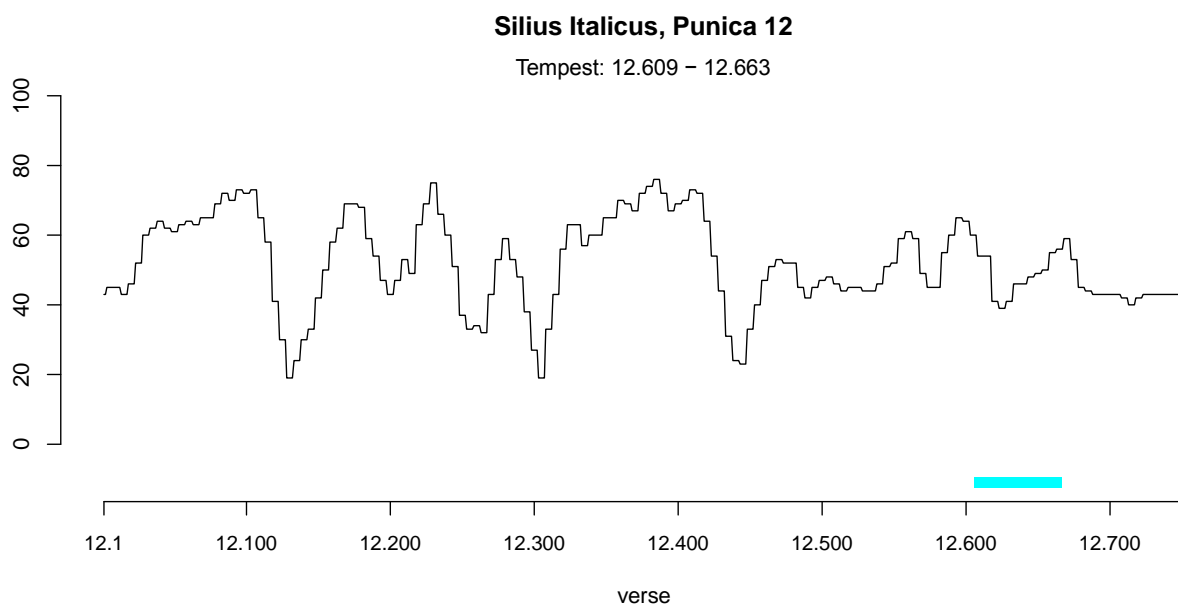


Figure 12. Line-by-line cluster stability, as in the preceding figures. Detail showing the pre-selected tempest scene for Silius Italicus' *Punica*, lines 12.609-12.663. Here, Hannibal and his troops threaten the gates of Rome, but Jupiter brings a terrible storm to the city's defense. Unlike the five previous cases, the passage determined to be the poem's most significant example of the tempest type scene was unremarkable in terms of cluster stability. Indeed, while a few brief sections exceeded this one in stability, none of the lines in this book showed a level of clustering consistency across trials comparable to that seen in the other five passages studied.

VFUTURE WORK

5.1 Continued work on theme extraction

In general, the poems showed great variation in stability from line to line. At the same time, over multiple comparisons, the boundaries of passages found to be stable were consistent. It would appear that the choice of tempests for the human tagging task was serendipitous: not only was the stability signal strong in five out of six cases, but in fact in most cases the pre-selected region of interest demonstrated the strongest stability in its respective book. Based on the graphs examined here, it seems unlikely that similarly striking success will be achieved with every one of the types that has been identified in catalogues such as [Edwards, 1987]. As work on the benchmark tags proceeds, we will be better able to put the results shown here in the context of a larger repertoire of thematic narrative elements.

The success of this method at flagging the passages independently selected by readers as worthy of interest belongs to the *recall* domain discussed above; the complementary ability in terms of *precision* can be measured by the degree to which other passages flagged as stable also turn out to contain typical thematic material. A related task will be to examine larger patterns of stability over the entire corpus, and to determine whether the simple percentage of misclassifications used here can be better calibrated to readers' impressions of the text.

5.2 Supervised approach

In the supervised approach, we begin from hand-tagged texts subdivided into unequal samples according to editorial indications and according to the subjective judgements of our human

taggers.²³ Although we have not yet implemented the remaining steps, we plan on implementing an analogous process to that found in the unsupervised method: feature extraction using tf-idf scores, followed by dimensionality reduction using PCA and adjustment for mean authorship signals.

We are considering a limited number of type scenes for the present, about a dozen, in order to test the viability of the general concept with thematic elements widely acknowledged as core to the genre, before wading into detailed discussion of exactly what constitutes a typical scene and how finely one ought to divide their classification. This also means the results will have a granularity roughly comparable to those from the unsupervised approach.

To model the human classification of the resulting features, we intend to use a linear support vector machine (SVM), a method with which we have some experience and which has produced successful results in literary applications in the past.²⁴

5.3 Application to text-reuse scoring

The success of our project's primary objective will be evaluated using cases of text reuse in Valerius Flaccus' *Argonautica*, ranked by readers familiar with the poem and its precedents for their allusive significance. This set of parallels will be generated using the Tesseract search engine, which returns instances where a verse line in Valerius reuses at least two words from a verse line in one of his predecessors. These are ranked by Tesseract according to the frequency and proximity of the matching words. Tesseract regularly employs a cutoff score to weed out the vast number of cases where reuse of small phrases has no literary significance. We will evaluate the performance of this automated ranking using independent measures of *recall* and *precision*. For us, *recall* indicates the ability of the scoring system to place specific instances of text reuse previously identified by readers as significant above an arbitrary score, which can then be used as the cutoff; while *precision* indicates the proportion of all results above the given threshold that are deemed interesting after the search is done and the cutoff chosen.

We do not attempt to combine precision and recall in a single metric such as the F1 measure. In fact, our measures of recall and precision are not entirely commensurate, since they deal with different populations. For us, recall answers the question, how many of the allusions *that we knew we wanted* did we actually get? It does not have any way to measure the serendipitous discovery of new, interesting intertexts, even though this is in fact a core function of search tools like Tesseract. Precision, on the other hand, answers the question, how many uninteresting cases did we have to sift through, even above the cutoff score? This is really more a measure of our patience with the machine than it is a measure of what proportion of all true negative results were excluded.

At bottom, the incommensurability of these two metrics comes from the subjective nature of intertextuality and literary interpretation: we consider it impossible in principle to define once and for all the set of all "true" allusions, since what constitutes an allusion depends largely on

²³We consulted the following editions for paragraph and scene breaks: for Vergil, the online edition of Itinera Electronica[IE]; for Ovid, [Miller and Goold, 1977] and [Miller and Goold, 1984]; for Lucan, [Bourguery, 1927] and [Bourguery and Ponchont, 1930]; for Valerius Flaccus, [Lieberman, 1997] and [Lieberman, 2002]; for Silius Italicus, [Nisard 1855]; and for Statius, [Nisard, 1878].

²⁴For example, [Zhao and Zobel, 2003] show the utility of SVMs in traditional authorship attribution tasks using stylistic features on literary texts. [Forstall et al., 2011] is one attempt by members of our team to broaden the literary application of linear SVMs to more open-ended, literary critical tasks.

the experience of the reader. What we can attempt to model, however, is the degree to which the algorithm fulfills our expectations (recall) and at the same time avoids what we don't want (precision).

In judging the success of the thematic features, then, we will take Tesseract's best recall and precision measures for searches on Valerius Flaccus as a baseline. We will then adjust Tesseract's scores by adding an extra term to account for whether the source and target phrases share thematic features. The best achievable recall and precision will again be evaluated, systematically testing various treatments of the thematic term and adjusting the cutoff. To the extent that better recall and / or precision are achieved, we will judge the new featureset to be useful.

Acknowledgments

The work described here was supported by the Fonds National Suisse de la Recherche Scientifique, Project No. 146976, "Intertextuality in Flavian Epic Poetry." We would like to thank our collaborators, the Tesseract project at the University at Buffalo, USA, directed by Neil Coffee.

References

- Allen G. *Intertextuality*. London: Routledge (London), 2011.
- Anderson W.S. *Ovid's Metamorphoses: Books 1-5*. University of Oklahoma Press (Norman), 1997.
- Arend W. *Die typischen Scenen bei Homer*. Weidmann (Berlin), 1975.
- Bal M. *De theorie van vertellen en verhalen: inleiding in de narratologie*. Coutinho (Muiderberg), 1978. English translation: *Narratology: Introduction to the Theory of Narrative*. University of Toronto Press (Toronto), 2009.
- Bamman D. and Crane G. "The Logic and Discovery of Textual Allusion." *Proceedings of the Second Workshop on Language Technology for Cultural Heritage Data (LaTeCH 2008)*. Marrakesh, Morocco, 2008.
- Baraz Y. and van den Berg C.S. "Intertextuality: Introduction." *American Journal of Philology* 2013:34:1-8.
- Barchiesi A. *La traccia del modello: Effetti omerici nella narrazione virgiliana*. Giardini (Pisa), 1984.
- Berlincourt V. *Commenter la Thébaïde (16e-19e) s.: Caspar von Barth et la tradition exégétique de Stace*. Brill (Leiden), 2013.
- Bourguery A. *Lucain. La Guerre civile. La Pharsale. Tome I: Livres I-V*. Les Belles Lettres (Paris), 1927.
- Bourguery A. and Ponchont M. *Lucain. La Guerre civile. La Pharsale. Tome II: Livres VI-X*. Les Belles Lettres (Paris), 1930.
- CDS CDS: Cross Database Searchtool. Brepolis Publishers NV. 2015. <http://clt.brepolis.net/cds/Default.aspx>. Accessed 28-11-2016.
- BTL *Bibliotheca Teubneriana Latina Online*. De Gruyter. <http://www.degruyter.com/view/db/btl>. 2014. Accessed 28-11-2016.
- Chaudhuri P. and Dexter J., "Bioinformatics and Classical Literary Study." arXiv:1602.08844, 2016.
- Coffee N., Koenig J.-P., Poornima S., Ossewaarde R., Forstall C.W., and Jacobson S. "Intertextuality in the Digital Age." *Transactions of the American Philological Association* 2012:142:383-422.
- Conte G.B. *Memoria dei poeti e Sistema letterario*. Einaudi (Torino), 1974. English translation: *The Rhetoric of Imitation: Genre and Poetic Memory in Virgil and Other Latin Poets*. Cornell University Press (Ithaca), 1986.
- Conte G.B. and Barchiesi A. "Imitazione e arte allusiva. Modi e funzioni dell'interstualità", in: Cavallo G., Fedeli P., and Giardina A. (eds): *Lo spazio letterario di Roma antica, vol. I: La produzione del testo*. Salerno (Rome), 1989. Pages 81-114.
- Dewar M. *Statius: Thebaid IX*. Clarendon Press (Oxford), 1991.
- EB Reitz C. and Finkmann S. (eds) *Epische Bauformen: Strukturen epischen Erzählens*. Universität Rostock. <http://www.epische-bauformen.uni-rostock.de/home/>. Accessed 28-11-2016.
- Eco U. *Lector in fabula: La cooperazione interpretativa nei testi narrativi*. Bompiani (Milan), 1979. English translation: *The role of the Reader: Explorations in the Semiotics of Texts*. University of Indiana Press (Bloomington), 1979.
- Eder M. "Does size matter? Authorship attribution, small samples, big problem." *Digital Scholarship in the Humanities* 2015:30(2):167-182.
- Edmunds L. *Intertextuality and the Reading of Roman Poetry*. Johns Hopkins University Press (Baltimore), 2001.
- Edwards M. *Homer: Poet of the Iliad*. Johns Hopkins Press (Baltimore), 1987.
- Edwards M. "Homer and Oral Tradition: The Type-Scene." *Oral Tradition* 1992:7(2):284-330.
- eTRACE <http://etraces.e-humanities.net/home-etraces.html>. Accessed 28-11-2016.
- Filum Filum – A tool for identifying verbal resemblances in literature. <https://qcrit.org/filum>. Accessed 21.02.2017.
- Forstall C.W., Jacobson S., and Scheirer W.J. "Evidence of Intertextuality: Investigating Paul the Deacon's *Angustae Vitae*." *Literary and Linguistic Computing* 2011:26(3):285-296.

- Forstall C.W., Coffee N., Buck T., Roache K., and Jacobson S. "Modeling the Scholars: Detecting Intertextuality through Enhanced Word-Level N-Gram Matching." *Digital Scholarship in the Humanities* 2015:30(4):503-515.
- Forstall C.W. and Galli Milić L. "Thematic features for intertextual analysis." *Digital Classicist Berlin Seminar Series*, 13 October 2015. <http://de.digitalclassicist.org/berlin/2015/10/13/Forstall>
- Fowler D. "On the Shoulders of Giants: Intertextuality and Classical Studies." *Materiali e discussioni per l'analisi dei testi classici* 1997:39:13-34. Reprinted in Fowler D. *Roman Constructions: Readings in Postmodern Latin*. Oxford University Press (Oxford), 2000. Pages 115-137.
- Gawley J., Forstall C.W. and Clark K. "Automating the Search for Cross-language Text Reuse." Paper presented at *Digital Humanities 2014*, Lausanne, Switzerland, July 9–11, 2014.
- Genette G. *Palimpsestes: La littérature au second degré*. Éditions du Seuil (Paris), 1982. English translation: *Palimpsests: Literature in the Second Degree*. University of Nebraska Press (Lincoln), 1997.
- Hardie P. "Flavian Epicists on Virgil's Epic Technique." In: Boyle A.J. (ed.) *The Imperial Muse: Ramus Essays on Roman Literature of the Empire: Flavian Epicists to Claudian*. Aureal Publications (Bendigo), 1990. Pages 3-20.
- Hershkowitz D. *Valerius Flaccus' Argonautica: Abbreviated Voyages in Silver Latin Epic*. Clarendon Press (Oxford), 1998.
- Hinds S. *The Metamorphosis of Persephone: Ovid and the Self-conscious Muse*. Cambridge University Press (Cambridge), 1997.
- Hinds S. *Allusion and Intertext: Dynamics of Appropriation in Roman Poetry*. Cambridge University Press (Cambridge), 1998.
- Hinds S. "Essential epic: Genre and gender from Macer to Statius." In: Depew M. and Obbink D. (eds) *Matrices of genre*. Harvard University Press (Cambridge, MA), 2000. Pages 221-244.
- IE *Itinera Electronica: du texte à l'hypertexte*. <http://agoraclass.fltr.ucl.ac.be/concordances/intro.htm#vir>. 2002. Accessed 28-11-2016.
- Iser W. *Der Akt des Lesens*. Wilhelm Fink (Munich), 1975. French translation: *L'Acte de lecture: Théorie de l'effet esthétique*. Mardaga (Brussels), 1985.
- Jenny L. "La stratégie de la forme." *Poétique* 1976:27:257-281.
- de Jong I.J.F. *Narratology and Classics*. Oxford University Press (Oxford), 2014.
- Juhnke H. *Homerisches in römischer Epik flavischer Zeit: Untersuchungen zu Szenennachbildungen und Strukturentsprechungen in Statius' Thebais und Achilleis und in Silius' Punica*. C. H. Beck (Munich), 1972.
- Knauer G.N. *Die Aeneis und Homer: Studien zur poetischen Technik Vergils mit Listen der Homerzitate in der Aeneis*. Vandenhoeck und Ruprecht (Göttingen), 1979.
- Köntges T. "Researchers to your Driving Seats: Building a Graphical User Interface for Multilingual Topic-Modelling in R with Shiny." *Digital Humanities 2016*, Kraków, Poland, 14 July 2016.
- Kristeva J. "Bakhtine, le mot, le dialogue, le roman." *Critique* 1967:239:438-465.
- Levi-Strauss C. "The structural study of myth." *The Journal of American Folklore* 1955:68(270):428-444.
- Liberman G. *Valerius Flaccus: Argonautiques: Tome I: Chants I-IV*. Les Belles Lettres (Paris), 1997.
- Liberman G. *Valerius Flaccus: Argonautiques: Tome II: Chants V-VIII*. Les Belles Lettres (Paris), 2002.
- Littlewood J. *A commentary on Silius Italicus Punica 7*. Oxford University Press (Oxford), 2011.
- LLT-A LLT-A: Library of Latin Texts – Series A. Brepolis Publishers NV. <http://clt.brepolis.net/llta/Default.aspx>. 2016. Accessed 28-11-2016.
- LLT-B LLT-B: Library of Latin Texts – Series B. Brepolis Publishers NV. <http://clt.brepolis.net/lltb/Default.aspx>. 2016. Accessed 28-11-2016.
- Lord A. *The Singer of Tales*. Harvard University Press (Cambridge, MA), 1960.
- Lyne R. *Memory and Intertextuality in Renaissance Literature*. Cambridge University Press (Cambridge), 2016.
- McCallum A. *MALLET: A machine learning for language toolkit*. <http://mallet.cs.umass.edu>. 2002.
- Mastandrea P. and Tessarolo L. *Poetria nova: a CD-ROM of Latin medieval poetry, 650-1250 A.D. (with a gateway to Classical and Late Antiquity Texts)*. Florence, 2010.
- Miller F.J. and Goold G.P. *Ovid: Metamorphoses: Books 1-8*. Harvard University Press (Cambridge MA), 1977.
- Miller F.J. and Goold G.P. *Ovid: Metamorphoses: Books IX-XV*. Harvard University Press (Cambridge MA), 1984.
- Minchin E. *Homer and the Resources of Memory: Some Applications of Cognitive Theory to the Iliad and the Odyssey*. Oxford University Press (Oxford), 2001.
- MP Mastandrea P., Sblendorio M.T., Perrelli R., Bossina L., Paolucci P., Marchetta A., Mordenti R., and Boschetti F. (eds) *Memorata Poetis: memoria poetica e poesia della memoria*. <http://www.memoratapoetis.it/public/index/index>. Accessed 28-11-2016.
- MQDQ Mastandrea P., Perrelli R., Biondi G., Zurli L., and Viparelli V. (eds) *Musisque Deoque: un archivio digitale di poesia Latina*. <http://www.mqdq.it/>. 2005. Accessed 28-11-2016.
- Nelis D. *Vergil's Aeneid and the Argonautica of Apollonius Rhodius*. Francis Cairns (Leeds), 2001.
- Nelis D. "Valerius Flaccus, Vergil and Apollonius Rhodius: the intertextual strategies of a Flavian poet." In: Fratantuano L. (ed) *Companion to Latin Epic ca. 14-96 CE*. Wiley-Blackwell (Malden MA), forthcoming.
- Nisard M. (ed) *Lucain, Silius Italicus, Claudien. Oeuvres complètes avec la traduction en français*. Didot (Paris), 1855.
- Nisard M. (ed) *Stace, Martial, Manilius, Lucilius Junior, Rutilius, Gratius Faliscus, Némésianus et Calpurnius. Oeuvres complètes avec la traduction en français*. Firmin-Didot (Paris), 1878.
- Parry M. *L'épithète traditionnelle dans Homère: Essai sur un problème du style homérique*. Les Belles Lettres (Paris), 1928. English translation: "The Traditional Epithet in Homer." In: Parry A. (ed) *The Making of Homeric Verse: The Collected Papers of Milman Parry*. Oxford University Press (New York), 1987. Pages. 1–190.
- Pedecerto Colombi E. (ed) *Pede Certo: metrica latina digitale*. Università di Udine. <http://www.pedecerto.eu/>. Accessed 28-11-2016.
- Perseus Crane G.R. (ed) *Perseus Digital Library*. Tufts University. <http://www.perseus.tufts.edu/>. Accessed 28-11-2016.

- PHI *Classical Latin Texts: A Resource Prepared by The Packard Humanities Institute*. <http://latin.packhum.org/>. Accessed 28-11-2016.
- Pucci J. *The Full-Knowing Reader: Allusion and the Power of the Reader in the Western Literary Tradition*. Yale University Press (New Haven), 1998.
- Řehůřek R. and Sojka P. "Software Framework for Topic Modelling with Large Corpora." In: *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*, 2010.
- Riffaterre M. "Sémiotique intertextuelle: l'interprétant." *Revue d'esthétique* 1979:1-2:128-150.
- Riffaterre M. "L'intertexte inconnu." *Littérature* 1981:41:4-7.
- Ripoll F. *La morale héroïque dans les épopées latines d'époque flavienne: Tradition et innovation*. Peeters (Louvain), 1998.
- Rubin D.C. *Memory in Oral Traditions: The Cognitive Psychology of Epic, Ballads, and Counting-out Rhymes*. Oxford University Press (New York), 1995.
- Scheirer W.J., Forstall C.W., and CoffeeN. "The sense of a connection: Automatic tracing of intertextuality by meaning." *Digital Scholarship in the Humanities* 2016:31(1):204-217.
- Schetter W. "Die Buchzahl der Argonautica des Valerius Flaccus." *Philologus* 1959:103:297-308.
- Scott J. *The Unity of Homer*. The University of California Press (Berkeley), 1921.
- Smolenaars H. *Status: Thebaid VII: A commentary*. Brill (Leiden), 1994.
- Stroh, H. *Studien zu Valerius Flaccus: besonders über dessen Verhältnis zu Vergil*. P.J. Pfeiffer (Augsburg), 1905.
- Tesserae Coffee N. (ed) *Tesserae*. The State University of New York at Buffalo. <http://tesserae.caset.buffalo.edu/>. 2010. Accessed 28-11-2016.
- Thomas R.F. "Virgil's Georgics and the Art of Reference." *Harvard Studies in Classical Philology* 1986:90:171-198.
- TRACER Buchler M. (ed) *TRACER: Text Reuse Detection Machine*. The Electronic Text Reuse Acquisition Project, Georg-August-Universität Göttingen. <http://www.etrapp.eu/research/tracer/>. 2016. Accessed 28-11-2016.
- Trillini, R.H. and S. Quassdorf "A 'key to all quotations'? A corpus-based parameter model of intertextuality." *Literary and Linguistic Computing* 2010:25(3):269-286.
- Wills J. *Repetition in Latin Poetry: Figures of Allusion*. Clarendon Press (Oxford), 1996.
- Zhao Y. and Zobel J. "Searching with style: Authorship attribution in classic literature." *Proceedings of the 30th Australasian Conference on Computer Science* 2003:62:59-68.
- Zissos A. *Valerius Flaccus' Argonautica Book 1*. Oxford University Press (Oxford), 2008.
- Zissos A. "Valerius Flaccus." In: Foley J.M. (ed.) *A Companion to Ancient Epic*. Blackwell (Malden), 2009. Pages 503-513.