COMMENT

Check for updates

Quantifying the quantitative (re-)turn in historical linguistics

Barbara McGillivray[1✉] & Gard B. Jenset [2]

Historical linguistics is the study of language change and stability, of the history of individual languages, and of the relatedness between languages. In spite of numerous acknowledgements, the adoption of quantitative methods in historical linguistics is still far from being mainstream and it falls below the level of other branches of linguistics. This comment considers the adoption of quantitative methods in recent historical linguistics research, and compares a study on 2012 publications with a similar study conducted seven years later. This comment argues for the advantages of a wider adoption of quantitative methods among historical linguists, and considers various reasons for the relatively slow progress in this direction. It also clarifies when quantitative methods are not the preferred route.

[1] King's College London, London, UK. [2] Springer Nature, London, UK. ✉email: Barbara.mcgillivray@kcl.ac.uk

Why should we talk about quantitative historical linguistics?

Historical linguistics is the academic field that studies language change and language stability, explores the history of individual languages, and identifies the relatedness between languages (Harrison, 2003, p. 214). This can involve investigating how different aspects of languages such as grammar, sound, or meaning changed over the course of the history of a language and across languages (diachronic analyses), reconstructing the pre-history of languages and language families, and tracing words' etymologies. Historical linguistics also covers the study of the languages of the past (historical languages) from a synchronic viewpoint, i.e. at a single point in time (see Campbell, 2021, among others).

Historical linguistics has been data-centric since its beginnings. Labov (1972, p. 100) acknowledged that historical linguistics makes the best use of "bad data", referring to the numerous gaps in the evidence available to historical linguists. The empirical basis of historical linguistics has also long been recognised by other scholars: Rydén (1980, p. 38) wrote that the "study of the past […] must be basically empirical", Fischer (2004, p. 57) that "[t]he historical linguist has only one firm knowledge base and that is the historical documents", a status also recognised by Penke and Rosenbach (2007, p. 1) more recently.

In spite of these numerous acknowledgements, the adoption of quantitative methods in historical linguistics is still far from being mainstream and it falls below the level reached by other branches of linguistics. For example, Joseph (2008, p. 687) notes that, while linguistics has always been an empirical field, "the bar [seems to have been raised] on the nature of the evidence we work with", noting, in particular, an increase in the reliance on corpus data. Similar arguments are put forward by Winter (2022), Kortmann (2021), and Brinton et al. (2021), among others. The fact that quantitative methods in historical linguistics are underused is a serious limitation because quantitative methods offer researchers the opportunity to test theoretical hypotheses that have been proposed on many historical linguistics phenomena. Moreover, quantitative methods can fruitfully complement qualitative example-based research: large- or medium-scale multivariate data analyses have the potential to provide descriptions of multidimensional phenomena where different factors are at play, which is a fairly typical situation in historical linguistics.

Going beyond historical linguistics into the broader field of digital humanities and cultural heritage studies, the recent availability of large cultural datasets (many of which are in textual form), coupled with breakthroughs in computational research (particularly machine learning, natural language processing, and scientific data analysis), have renewed excitement about the so-called "computational turn" in humanities research, concerned with applying and/or developing computational methods to answer research questions in the humanities (McGillivray et al., 2020). This trend is further supported and strengthened by the Open Science movement, which has brought issues of open data and reproducibility to the front of the scientific debate, also trickling into digital humanities discourse (cf. e.g. McGillivray et al., 2022).

Alongside specific examples of quantitative studies that can advance the field, it is also important to articulate a quantitative framework for *doing* quantitative historical linguistics research. In Jenset and McGillivray (2017) we introduced a corpus framework that aims to provide the methodological and epistemological "scaffolding" to bridge the gap between conceptual considerations and concrete quantitative techniques. In this comment, we present the results of a quantitative analysis of articles published in historical linguistics journals: based on this analysis, we argue for the importance of the wider adoption of quantitative methods in

historical linguistics. This study updates the study we presented in Jenset and McGillivray (2017, pp. 25–35), which we will refer to as the "2012 study". Our 2012 study focussed on 62 articles published in six historical linguistics journals in 2012 and found that 29% of the papers analysed were corpus-based, 40% were quantitative (as opposed to 80% of general linguistics articles in the study by Sampson, 2005), and corpus studies were more likely to adopt quantitative methods. Following the "chasm" model of technology adoption proposed by Moore (1991), and with appropriate caveats, this result pointed to historical linguistics being in an early phase of adoption of quantitative methods, with less than half of researchers adopting them and therefore falling into the category of "early majority". In contrast to historical linguistics, the evidence for general linguistics points to this field having progressed to the full adoption of quantitative methods. A few years later, we wanted to check if the situation had changed and if the trend towards more quantitative studies had stopped or continued.

Analysis

The aim of the analysis is to provide a snapshot of the field of historical linguistics today compared with the recent past. Following our 2012 study, the only previous quantitative study that has analysed the distribution of quantitative studies in historical linguistics journals, and to keep the task manageable, we selected six historical linguistics journals according to the following criteria (Jenset and McGillivray, 2017, p. 27):

1. Research journals (thus excluding monographs, edited books, and yearbooks);
2. Journals published in the English language;
3. Journals focussing specifically on historical linguistics and/or language change;
4. Journals that had a general scope (thus excluding specific subfields of historical linguistics such as historical pragmatics);
5. Linguistics journals (thus excluding interdisciplinary journals).

We based our methodology on this previous study to provide a longitudinal perspective on its findings. Therefore, we selected the same list of peer-reviewed academic journals we chose in Jenset and McGillivray (2017, pp. 25–35) according to criteria 1–5 above. The list of journals selected is:

– Diachronica
– Folia Linguistica Historica
– Journal of Historical Linguistics
– Language Dynamics and change
– Language variation and change
– Transactions of the Philological Society

We analysed all 63 research articles published in the journals listed above in 2019. The number of articles analysed is very close to the number we analysed in the 2012 study (62). We recognise that the size of this sample is rather limited, but we have decided to not expand the dataset further for a number of reasons. First, this analysis provides an empirical illustration of our argument, in line with the aims of a comment paper in contrast with a research paper. Second, as stated above, we kept the same selection criteria as our 2012 study to ensure a longitudinal perspective. Third, we carried out a statistical analysis that can measure the size of the effects detected and reveal whether indeed there is sufficient evidence for a statistically significant result. We selected all relevant research articles from the journal issues in question, excluding non-primary research such as editorials, comments, book reviews, and descriptions of software tools. We also

excluded a very small number of articles that were not historical or diachronic, as well as introductions to special issues.

We read each article to collect the following information: the type of evidence base used in the paper (digital corpora, word lists, examples, etc.) and the statistical techniques used for the analysis if any (*t*-tests, regression models, principal component analysis, etc.). We then classified the articles across two dimensions: corpus-based vs. non-corpus-based and quantitative vs. non-quantitative.

A paper was described as being corpus-based if the authors used a corpus (or a subset of a corpus) as the main evidence source of their research. In other words, the study had to use a machine-readable collection of historical natural language data which is published or at least accessible by others (even if not freely). Therefore, studies based only on word lists, private resources, purpose-built collections not available to the academic community, or other language resources such as dictionaries were not considered to be corpus-based. The type of data used in corpus-based studies in our sample included existing corpora such as Lip (Lessico di frequenza dell'italiano parlato) or portions of them, annotated corpora such as treebanks, and corpora of elicited utterances from fieldwork. The type of data used in non-corpus-based studies includes historical dictionaries, texts quoted in previous literature and examples from texts and manuscripts.

We considered a study to be quantitative if its conclusion relied on quantitative evidence, for example by including statements about the frequency of a given construction or set of items, testing a hypothesis quantitatively in some form or another, or measuring the statistical significance of a phenomenon such as a correlation between two variables. Phylogenetic studies, although they did not tend to use corpus frequency data in our sample, were considered quantitative because they compute distances between linguistic features. The techniques used in the quantitative studies range from simple percentages to chi-square tests and *t*-tests, to random forest, and regression models, including mixed effect models. Thus, the criterion for what we consider a quantitative study is not the presence of numbers in the article, nor is the definition as we operationalise it linked to any specific statistical technique. The only criterion we considered was whether the conclusion or main line of argumentation relied upon quantification in some form. We interpreted the absence of such quantification, determined by a close reading of each article, as indicating a qualitative study.

It is important to note that the two dimensions, corpus-based vs. non-corpus-based and quantitative vs. non-quantitative, although often correlated are nevertheless independent. A study may be corpus-based using qualitative methods, for example, if it relies on examples drawn from a corpus without presenting a quantitative analysis of them. On the other hand, a study may be quantitative without being corpus-based, for example, if it uses other evidence sources like phylogenetic research.

The articles covered a wide range of topics and linguistic subfields, from language typology and language classification to historical phonology, morphology, syntax, semantics and lexicon. The languages analysed include Latin, ancient Greek, Gothic, English, Medieval French, Eastern Tukanoan, Ecuadorian Siona, Vera'a, Spanish, Bantu languages, Japanese, Russian, old Saxon, Sanskrit, Celtic, Indian Punjabi, Italo-Romance languages, Dutch, German, and Grico.

One of the reviewers pointed out a potential risk of bias in our quantitative analysis, given that Transactions of the Philological Society (TPS) has a scope that might disproportionately attract studies of less attested and resourced languages, hence limiting the potential for quantitative analysis. However, this does not seem to be the case in our data. Of the 19 TPS articles in our

Table 1 Classification of the 2019 articles, classified according to whether they were corpus-based or not, and quantitative or not.

	Qualitative	Quantitative	Total
Not corpus-based	18 (29%)	14 (22%)	32 (51%)
Corpus-based	9 (14%)	22 (35%)	31 (49%)
Total	27 (43%)	36 (57%)	63 (100%)

sample we found that 13 were done on relatively well-attested and resourced languages: English (including Old and Middle English), Middle French, Middle Dutch, Old High German, Latin, Middle Norwegian and Old Irish.

Table 1 shows the number of articles in each category, alongside the percentages over the total number of articles (63). Of the articles analysed, 27 (43%) were qualitative and 36 (or 57%) were quantitative. Compared with the results from our 2012 (Jenset and McGillivray, 2017, pp. 25–35), we notice an *increase in the number of quantitative articles* (57% vs. 40%) *and corpus-based articles* (49% vs. 29%). In other words, the split qualitative/quantitative split seems to have changed in favour of quantitative studies and the same happened in favour of corpus-based studies.

The majority (22 out of 31) of corpus-based articles are also quantitative, while the majority of those that are not corpus-based (18 out of 32) are qualitative. Out of the quantitative studies, 22 (or 61%) were corpus-based and 14 (or 39%) were not. The association between these two dimensions was statistically significant, as per a chi-squared test ($\chi^2 = 5.73$, $p < 0.05$, $\varphi = 0.29$). This is similar to the 2012 study, which also found a statistically significant association between corpus-based and quantitative studies ($\chi^2 = 14.79$, $p << 0.05$, $\varphi = 0.49$) but a larger effect size as measured by the φ coefficient. Both chi-squared tests are exact tests, without Yates' continuity correction. Our original manuscript reported Yates' corrected results (the default in R), but as a reviewer pointed out, Yates' correction can be overly conservative. In our case, applying Yates' correction resulted in a non-significant result for the 2019 data. All expected frequencies in the table are above five, meaning that the conditions for using an uncorrected test, as reported above, are met by Yates' own criteria (Hitchcock, 2009). For completeness, we also ran Fisher's exact test on the 2019 data, which also showed a significant result ($p < 0.05$, OR = 3.34). Although the statistically significant association between corpus-based and quantitative studies persists from 2012 to 2019, the degree of association between them is less strong.

The 95% confidence intervals for quantitative articles in the 2012 study and in this study are presented in Table 2. These confidence intervals show that 95% of the observations from the underlying population of articles from the journals from which the sample was taken (if this is representative) would fall between 44% and 70% for the 2019 data and between 28% and 52% for the 2012 data. A binomial test shows a statistically significant difference between the two samples ($p << 0.05$).

To summarise our findings:

- Quantitative studies have gone from 40% of the sample in 2012 to 57% in 2019.
- Qualitative, corpus-based papers have increased since 2012. Qualitative, non-corpus papers have seen a decline.
- There is still a significant association between corpus use and quantitative methods, but the strength of the association between them, as measured by the φ coefficient, has decreased from 0.49 in 2012 to 0.29 in 2019.

Table 2 95% confidence intervals for the percentage of quantitative papers in the 2012 study (Jenset and McGillivray, 2017) based on a sample of historical linguistics articles published in 2012 and in this study using a sample of articles published in 2019.

	Percentage of quantitative papers	95% confidence interval
Jenset and McGillivray (2017)—2012 sample	40%	[28%,52%]
This study—2019 sample	57%	[44%,70%]

Arguments in support of quantitative methods in historical linguistics

Not all historical linguistics research can (or should) be quantitative. For certain linguistic phenomena, we simply do not have (enough) data to conduct statistical investigations. In other areas, non-quantitative computer-assisted methods, such as phylogenetic trees or networks, are more suitable (List, 2021). And in some areas, notably historical phonology and morphology, the traditional approach is in many cases not just the best but the only method available.

Nonetheless, it is clear from our data that 2019 has seen a statistically significant increase in the proportion of articles using quantitative methods compared to 2012. The increase from 40% to 57% represents a growth of 42.5%. In our opinion, this growth is a good thing, because this methodological alignment between synchronic and diachronic linguistics can facilitate other types of alignment and help break down the artificial distinction once introduced by Saussure (Pierce and Boas, 2019). However, the 42.5% growth in quantitative papers must be seen in its proper context. Firstly, the growth unfolds over a period of 7 years, meaning that the compound annual growth is only about 5%. For comparison, 5% of our 2019 sample is about 3 papers. This suggests that the growth might be a gradual one, rather than an abrupt shift, although a year-by-year analysis would be required to rule out the possibility of any sudden jumps. In other words, although historical linguistics articles have seen considerable growth in quantitative methods compared to 2012, the field remains behind, or at least not conclusively level with, that of linguistics as a whole.

This raises the interesting question of what is a reasonable, or appropriate, level of quantitative studies in historical linguistics. This is a question that cannot be answered prescriptively, if at all. In this article, we restrict ourselves to observing firstly that in general, an increase in the adoption of quantitative methods is desirable both to open new avenues of research and to facilitate alignment with synchronic linguistics. Secondly, we observe that the proportion of quantitative studies is growing, which suggests that historical linguists see value in conducting and publishing more quantitatively oriented research.

This ties in with broader trends in the humanities, where recent years have seen a number of textbooks in quantitative methods aimed specifically at researchers and students in the humanities. Examples include Tilton (2015), Lemercier and Zalc (2019), McGillivray and Toth (2020) and Karsdorp et al. (2021). A similar proliferation of quantitative methods textbooks can be found in linguistics in this period, with recent examples from cognitive linguistics (Winter, 2022), psycholinguistics (Rij van et al., 2020), and sociolinguistics (Macaulay, 2009).

There does not seem to be a similar publication burst of quantitative methods textbooks specifically for historical linguistics. Perhaps this is because some techniques, e.g. regression modelling, can be taught equally well with synchronic data and because some textbooks, such as Baayen (2008) and Johnson (2008) include chapters relevant to historical linguistics. However, it is also noteworthy that a popular historical linguistics textbook such as Campbell (2021) devotes its chapter on quantitative historical linguistics almost exclusively to criticism (Jenset and McGillivray, 2017, p. 86), suggesting at least some degree of resistance to their adoption in the field.

However, it is also worth considering the differences between historical linguistics and synchronic linguistics. The quantitative trend in linguistics generally seems driven partly by criticism of previous reliance on introspection in linguistics and partly by better access to quantitative or quantifiable data, such as web data or experimental data obtained via websites such as Amazon's Mechanical Turk (Winter, 2022). However, although web data might be an interesting source of data for some diachronic studies, historical linguistics in its proto-typical sense is cut off from these data sources. There are no native speakers of Old English or Latin to be recruited from Mechanical Turk. Instead, historical linguists must, by necessity, make use of the various types of evidence available, whether textual evidence or the present-day languages themselves, as related entities produced by a historical process. This, then, might constitute a form of absolute limit on the degree to which quantitative methods can be applied in historical linguistics. To be clear, a complete ban on qualitative studies in historical linguistics would be both futile and undesirable (Jenset and McGillivray, 2017; Kortmann, 2021). However, we believe the field should be striving towards a high degree of adoption of quantitative methods, to the extent possible, for these reasons of transparency, reproducibility, code and data sharing on a larger scale, as well as methodological alignment with linguistics in general and ultimately other adjoining fields.

Whence quantitative historical linguistics?

Based on our analysis, it seems clear that historical linguistics is undergoing, or has undergone, a quantitative turn, similar to linguistics in general (Winter, 2022; Brinton et al., 2021; Kortmann, 2021; Pierce and Boas, 2019; Janda, 2013; Joseph, 2008). It is difficult to judge if we have reached some natural or optimal level of application for quantitative methods in historical linguistics, or if there is still room to increase the proportion of quantitative studies further. Ultimately, that is a question for the future. However, after taking stock of where we are, it seems to us that a few clear challenges for the future can be formulated.

Firstly, there is the question of the quantitative turn itself, and its inclusion in historical linguistics. It should be noted that although the proportion of quantitative studies in our sample has increased, we should not forget the qualitative side of quantitative methods. Not all quantitative methods are equally informative or well-adapted to historical linguistics. As a consequence, we see room for moving away from the classical null hypothesis tests, towards more advanced methods that can better account for the context of the data. Null-hypothesis tests are sometimes useful (we have used them here, for instance) but they can be problematic with historical data (Jenset and McGillivray, 2017, p. 96), and although multilevel/mixed-effects regression models have gained a firm foothold in historical linguistics, there is probably room for further adoption of such models in particular, and generally for a broader repertoire of techniques suited for specific research questions. Even if the

quantitative method has been thoughtfully picked, it doesn't automatically follow that its use is well integrated with the linguistic problem at hand. The results, in our experience, are often studies where the conclusions and the quantitative analysis do not support each other. Kortmann (2021) discusses the same problem from a general linguistic point of view and argues (correctly in our view) that linguistic questions should lead the way in selecting the appropriate methods. This sounds (and is) reasonable but it is potentially challenging. Firstly, it requires a wider overview of the available statistical methods as well as a deeper conceptual understanding of what they do. It is also challenging since it might break with community norms, both for researchers and journal reviewers, and editors.

Next, there is clearly a set of open questions for historical linguistics in general that is not limited to quantitative historical linguistics, but which quantitative approaches to historical linguistics must also inevitably grapple with. For example, using quantitative methods in themselves does not automatically address the problem that multiple explanations and hypotheses might be compatible with the observed historical data (Jenset and McGillivray, 2017, p. 47). Roberts et al. (2020) present an interesting supporting tool to deal with this problem, which we find interesting and encouraging, but insufficient on its own to address this problem. Instead, we will probably need an even closer alignment of theory, hypotheses, data, and methods. Another such general problem is data quality, with gaps and various forms of historical preservation bias (geographical, social, gender-based, etc.) as prominent examples. Again we can find interesting partial technological solutions, such as imputation techniques for missing data, and simulation experiments including agent-based modelling (Stevens and Harrington, 2022; Harrington et al., 2019). Yet despite these promising technical advances we still see the greatest gains stemming from a closer engagement between theory, methods, and the available, imperfect, data.

We also think that historical linguistics could stand to gain from making better use of the data already available. In some cases, this would undoubtedly require the development of more natural language processing (NLP) tools for historical language varieties, to allow further enrichment of historical data, in addition to what already exists (Jenset and McGillivray, 2017). Another way in which the existing data could be better leveraged is by further enriching it with human annotation. Numerous such projects exist, and many historical linguists have undoubtedly done much annotation work that could, and should be shared with colleagues, e.g. as open datasets and described in a data paper. However, we believe there is also a benefit in unlocking annotated historical data that are, in our experience, too often difficult to integrate with current quantitative modelling platforms and techniques. A quantitative analysis of syntactically annotated data (Taylor, 2020), e.g. chapters 6 and 7 of Jenset and McGillivray (2017), will often require programming skills (or else very lengthy manual re-recording of annotations) to extract the rich, detailed information needed to perform multivariate regression analyses. New tools such as TreeNet (Jenset, 2022) would partially help, but a combination of training historical linguists in coding or having research teams with more diverse skills seems inevitable. As such, this challenge speaks not only to current researchers in historical linguistics but also to the coming generations that they will be training.

Data availability
The datasets analysed during the current study are available in the Dataverse repository: https://doi.org/10.7910/DVN/IIHRZ3.

References
Baayen RH (2008) Analyzing linguistic data: a practical introduction to statistics using R. Cambridge University Press, Cambridge

Brinton LJ, Honeybone P, Kortmann B, Seoane E (2021) 25 years of English Language and Linguistics: a celebration and analysis. English Lang Linguist 25(4):677–685

Campbell L (2021) Historical linguistics: an introduction. Edinburgh University Press, Edinburgh

Fischer O (2004) What counts as evidence in historical linguistics? Stud Lang 28(3):710–740

Harrington J, Gubian M, Stevens M, Schiel F (2019) Phonetic change in an Antarctic winter. J Acoust Soc Am 146(5):3327–3332

Harrison S (2003) On the limits of the comparative method. In: Joseph BD, Janda RD (eds.) The handbook of historical linguistics. Blackwell, Malden, MA, pp. 213–243

Hitchcock DB (2009) Yates and contingency tables: 75 years later. Electron J Hist Probab Stat 5:1–14

Janda LA (2013) Quantitative methods in cognitive linguistics: an introduction. In: Janda LA (ed) Cognitive linguistics—the quantitative turn. De Gruyter, Mouton, Berlin/Boston, pp. 1–32

Jenset GB (2022) TreeNet—a computational system for discovering constructions in historical parsed corpora. In: International Conference In Historical Linguistics 25 (Book of abstracts). Oxford https://ichl.ling-phil.ox.ac.uk/abstracts/209

Jenset G, McGillivray B (2017) Quantitative Historical Linguistics. A corpus framework. Oxford University Press, Oxford

Johnson K (2008) Quantitative methods in linguistics. Blackwell

Joseph BD (2008) The Editor's Department: last scene of all. Language 84(4):686–690

Karsdorp F, Kestemont M, Riddell A (2021) Humanities data analysis: case studies with Python. Princeton University Press

Kortmann B (2021) Reflecting on the quantitative turn in linguistics. Linguistics 59(5):1207–1226

Labov W (1972) Some principles of linguistic methodology. Language Soc 1(1):97–120

Lemercier C, Zalc C (2019) Quantitative methods in the humanities: an introduction. University of Virginia Press

List JM (2021) Computer-assisted approaches to historical language comparison. Habilitation Thesis, Friedrich Schiller University, Jena

Macaulay RKS (2009) Quantitative methods in sociolinguistics. Bloomsbury

McGillivray B, Alex B, Ames S, Armstrong G, Beavan D, Ciula A et al (2020) The challenges and prospects of the intersection of humanities and data science: a White Paper from The Alan Turing Institute. figshare https://doi.org/10.6084/m9.figshare.12732164.v5

McGillivray B, Toth G (2020) Applying language technology in humanities research. Design, application, and the underlying logic. Palgrave Macmillan

McGillivray B, Marongiu P, Pedrazzini N, Ribary M, Wigdorowitz M, Zordan E (2022) Deep Impact: a study on the impact of data papers and datasets in the humanities and social sciences. Publications 10(4):39. https://doi.org/10.3390/publications10040039

Moore GA (1991) Crossing the chasm: marketing and selling high-tech products to mainstream customers. Harper Business, New York

Penke M, Rosenbach A (2007) What counts as evidence in linguistics? An introduction. In: Penke M, Rosenbach A eds What counts as evidence in linguistics. John Benjamins, Amsterdam, pp. 1–49

Pierce M, Boas HC (2019) Where was historical linguistics in 1968 and where is it now?. In: Pierce M, Boas HC (eds) New directions for historical linguistics. Brill, pp. 1–41

Rij van J, Vaci N, Wurm LH, Feldman LB (2020) Alternative quantitative methods in psycholinguistics: implications for theory and design. In: Pirrelli V, Plag I, Dressler WU (eds) Word knowledge and word usage. A cross-disciplinary guide to the mental lexicon. De Gruyter Mouton

Roberts AK et al. (2020) CHIELD: the causal hypotheses in evolutionary linguistics database. J Language Evol 5(2):101–120. https://doi.org/10.1093/jole/lzaa001

Rydén M (1980) Syntactic variation in a historical perspective. In: Jacobson S (ed) Papers from the Scandinavian symposium on syntactic variation, Stockholm, 18–19 May 1979. Almqvist & Wiksell, Stockholm, pp. 37–45

Sampson GR (2005) Quantifying the shift towards empirical methods. Int J Corpus Linguist 10:10–36

Stevens M, Harrington J (2022) Individual variation and the coarticulatory path to sound change: agent-based modeling of/str/in English and Italian. Glossa 7(1):1–34.

Taylor A (2020) Treebanks in historical syntax. Annu Rev Linguist 6:195–212

Tilton L (2015) Humanities data in R. Exploring networks, geospatial data, images, and text. Springer

Winter B (2022) Mapping the landscape of exploratory and confirmatory data analysis in linguistics. In: Tay D, Pan M (ed) Data analytics in cognitive linguistics: methods and insights. De Gruyter Mouton, Berlin, Boston, pp. 13–48

Competing interests

GBJ is employed by Springer Nature. BMG declares no competing interests.

Ethical approval

This article does not contain any studies with human participants performed by any of the authors.

Informed consent

This article does not contain any studies with human participants performed by any of the authors.

Additional information

Correspondence and requests for materials should be addressed to Barbara McGillivray.

Reprints and permission information is available at http://www.nature.com/reprints

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.