



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/243311/>

Version: Accepted Version

Article:

Kousha, K. and Thelwall, M. (2026) Which abstract stylistic features fool ChatGPT research evaluations? *Scientometrics*, 131 (9). pp. 6497-6518. ISSN: 0138-9130

<https://doi.org/10.1007/s11192-026-05756-1>

© 2026 The Authors. Except as otherwise noted, this author-accepted version of a journal article published in *Scientometrics* is made available via the University of Sheffield Research Publications and Copyright Policy under the terms of the Creative Commons Attribution 4.0 International License (CC-BY 4.0), which permits unrestricted use, distribution and reproduction in any medium, provided the original work is properly cited. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

Which abstract stylistic features fool ChatGPT research evaluations?

Kayvan Kousha

Statistical Cybermetrics and Research Evaluation Group, Business School, University of Wolverhampton, UK. <https://orcid.org/0000-0003-4827-971X>; k.kousha@wlv.ac.uk

Mike Thelwall

School of Information, Journalism and Communication, University of Sheffield, UK. <https://orcid.org/0000-0001-6065-205X>; m.a.thelwall@sheffield.ac.uk

Large Language Models (LLMs) have the potential to be used to support research evaluation and have a moderate capability to estimate the research quality of a journal article from its title and abstract. This paper assesses whether there are abstract textual features unrelated to research quality that may influence ChatGPT's scores. Using a dataset of 99,277 journal articles submitted to the UK-wide Research Excellence Framework (REF) 2021 assessments, we calculated several readability indicators from abstracts and correlated them with ChatGPT scores and departmental REF scores (all fields) and individual article scores (health and life sciences). From the results, linguistic complexity and length correlated positively with both ChatGPT scores and a departmental average proxy for REF expert scores in most broad fields (Units of Assessment). They were also more strongly associated with ChatGPT research quality scores than with the proxy for REF expert scores in many subject areas, although this may have been due to the influence of the averaging process in the proxy for some. In the health and life sciences, where individual article scores were available, syllables per word and words per sentence had stronger correlations with individual ChatGPT scores than with individual article quality scores. Although cause-and-effect was not tested, these results suggest that ChatGPT may be more likely than human experts to reward linguistic complexity, with a potential bias towards longer and less readable abstracts in many fields. The apparent preference of LLMs for complex language is an undesirable feature for practical applications of LLMs for research quality evaluation, unless solutions can be found.

Keywords: Large Language Models; ChatGPT; AI bias; Research evaluation; Abstract readability; Research Excellence Framework (REF).

Introduction

The quality of published academic research is repeatedly assessed in the context of appointments, promotions, and national or local departmental research evaluations such as the UK's Research Excellence Framework (REF). These processes consume a lot of time, and the results can be inaccurate, especially when the evaluators lack the time or expertise to properly assess the outputs assigned to them. Historically, citation-based indicators have been used to support or replace expert evaluations, but several different Large Language Models (LLMs) seem to give results that rank more closely with expert scores (Thelwall, 2025c; Thelwall & Yaghi, 2025; Thelwall & Mohammadi, 2026). Nevertheless, the reason why LLM-based quality predictions align moderately well with human scores is not known, especially because they seem to be more effective when LLMs are fed with just article titles and abstracts than when they are fed with full texts (Thelwall, 2025ab). A similar pattern occurs for REF impact case studies, where ChatGPT scores based on titles and summaries correlate more

strongly with expert scores than when using full texts (Kousha & Thelwall, 2025). This raises the possibility that LLMs are using proxy signals in the sense of leveraging document attributes unrelated to research quality. One explanation could be that abstracts usually include the main claims about originality and significance, while full texts may include details that distract the models (Thelwall, 2025ab). It is also possible that human reviewers use titles and abstracts when forming initial impressions of research quality. Hence, stronger correlations based on titles and abstracts may reflect both possible model reliance on proxy signals and partial alignment with human evaluation practices. Identifying such attributes is therefore necessary to help understand this key LLM property as well as to assess whether authors might take simple steps to manipulate or fool LLM-based research evaluation systems.

Previous analyses of LLMs for research quality evaluation have tended to focus on obtaining the highest alignment with expert scores. Nevertheless, a few studies have tried to identify or test potential sources of bias. A regression analysis of ChatGPT scores for all REF2021 Units of Assessment (UoAs) found ChatGPT to give higher scores to research from some countries, articles in higher impact journals, and articles with longer abstracts. Of these, the last is most concerning since it seems to be unrelated to the quality of the paper. In contrast, some countries invest more in research and journals with higher citation rates may be expected to, on average, publish better research. Average scores also varied substantially between fields, suggesting a topic bias (Thelwall & Kurt, 2025). A comparison of terms occurring disproportionately in high and low scoring tourism titles/abstracts found that ChatGPT tended to score research with complex statistical methods higher and survey research lower (Nunkoo & Thelwall, 2025). For library and information science, LLMs seemed to give higher scores to papers with more technical terminology, theory, statistics, experiments and algorithms, and the opposite for abstraction and research involving libraries, students, or surveys (Zhu et al., 2026; Thelwall, 2025d). It is not clear whether some or all of these reflect underlying quality differences (as judged by human experts), but they are at least suggestive of potential AI bias. From wider research about LLMs for other tasks, it is known that LLMs can be influenced by the credibility or nationality of the source of the information if they are told it (Germani & Spitale, 2025; Vasu et al., 2025) but may not be influenced much by linguistic styles (Kundu & Barbosa, 2024).

Recent evidence shows that arXiv paper abstracts are becoming more complex and less readable after the release of ChatGPT (Alsudais, 2025; Bao et al., 2025), and that there has been a rapid increase in LLM-associated terms in academic articles across abstracts and full texts (Kousha & Thelwall, 2026). If LLM-based evaluation is sensitive to these features, this may encourage authors to adopt LLM writing styles to obtain higher scores, raising concerns that AI-supported research assessment may reward writing style rather than research quality. Despite the above findings and the research reviewed below, little is known about the relationship between measurable textual features of abstracts, such as abstract length, sentence length, vocabulary complexity, and readability indicators, and LLM scoring. Presentation quality may influence how experts understand and evaluate a paper, but linguistic complexity and abstract length are not explicit REF scoring criteria. Hence, this study tests whether these measurable abstract features are associated with ChatGPT scores, rather than assuming that they are either direct measures of research quality or unrelated to research quality.

In this article, abstract stylistic features are measurable properties of abstract wording and presentation, such as word count, sentence length, syllables per word, and readability scores. Readability here means the estimated ease or difficulty of reading an abstract, calculated using standard formulae described in the Methods section. Research evaluation means assigning quality scores to research outputs through expert scores. AI bias is used in a narrow sense to mean a systematic association between ChatGPT scores and abstract stylistic features that are not explicit REF research quality criteria.

Research questions

To investigate whether abstract features are associated with LLM research quality scores and whether these associations differ from REF expert scores across subjects (UoAs), this article addresses the following research questions:

- **RQ1:** Which textual features of abstracts (e.g., readability, linguistic complexity, and length) associate with LLM scores?
- **RQ2:** Are these associations stronger for LLM scores than for REF expert scores?
- **RQ3:** Are there differences in the answers to the above across subjects (UoAs)?

Background

Although this paper focuses on research quality, there is little relevant research about the relationship between research quality and abstract linguistic characteristics. Thus, this section reviews relevant research with citation data instead.

Abstract length and citation impact

Abstracts summarise the aims, methods, and findings of academic research. Along with the title, they are usually the first part of a paper read to decide whether a study is relevant. Because of this important characteristic, many scientometrics studies have used abstracts as the main source of textual data to investigate associations with research quality or citation impact. Some features of abstracts, such as their length, linguistic complexity, and readability, may influence how frequently papers are cited (Kousha & Thelwall, 2024).

Several studies have found that longer abstracts are associated with higher citation impact, perhaps because longer abstracts contain more information, helping readers to understand key aspects of the research and subsequently making the papers more citable. For example, a study of The Lancet articles (1997–1999) found that highly cited papers had much longer abstracts that contained 2.5 to seven times more words than poorly cited papers (Kostoff, 2007). A large study of one million abstracts across eight science fields found that shorter abstracts (with fewer words and sentences) were associated with fewer citations, although shorter sentences were associated with more citations in mathematics and physics, suggesting disciplinary differences (Weinberger et al., 2015). Another large study of 4.3 million papers from over 1500 journals also found that abstract length positively correlated with citation counts in almost all journals (Sienkiewicz & Altmann, 2016). Other studies across different subjects have similar findings (Didegah & Thelwall, 2013; Hafeez et al., 2019; Sohrabi & Iraj, 2017; Van Wesel et al., 2014). However, experimental evidence suggests that longer abstracts do not necessarily increase reader attention (Helbach et al., 2025). Hence, in the

context of research quality assessment, it is not clear whether experts or other evaluators would be influenced by abstract length when judging research quality.

Abstract readability and citation impact

Many studies have investigated whether more readable abstracts (e.g., Flesch reading ease score) associate with higher citation counts, finding that more complex (less readable) abstracts receive more citations, albeit with some disciplinary differences. For example, abstracts of highly cited information science articles are less readable (as measured by the Flesch Reading Ease score) (Gazni, 2011). A very large study of 4.3 million journal articles found that abstract complexity (measured by the Gunning Fog Index) positively associated with citation counts (Sienkiewicz & Altmann, 2016). Another study found that abstract readability, measured by the Flesch Reading Ease score, negatively correlated with citation counts in Biology and Biochemistry, while no significant association was found in Chemistry or Social Sciences (Didegah & Thelwall, 2013). Similarly, a study comparing highly cited and uncited Web of Science papers across multiple scientific fields found that abstracts of highly cited articles were longer, more linguistically complex, and less readable (Hu, Wang, & Deng, 2021). For articles on emerging technologies, less readable abstracts, measured using several readability indicators (the Flesch Reading Ease, Flesch–Kincaid Grade Level, and Gunning Fog Index), associated with higher citation impact (Ante, 2022). A study of 71,628 abstracts in Language and Linguistics journals (1991–2020) also found that lower abstract readability associated with higher citation counts, although the relationship was weak (Wang, Liu, & Zhou, 2022). A smaller-scale study of 550 articles published in Science also reported that abstract readability broadly related to online attention (Altmeteric scores) received by articles (Jin et al., 2021).

Methods

The research design was to obtain LLM research quality scores and stylistic characteristics of abstracts, such as readability, linguistic complexity, and length, for a large set of academic journal articles across different subjects and identify any relationship between the two. The second part of the research design was to repeat the process for expert scores instead of LLM scores, to assess whether there were any differences between experts and LLMs in the patterns found. Such differences would be suggestive of AI bias rather than indirect factors associating with research quality. The third part of the research design (not addressing a research question but added for context) involved replacing LLM scores with citation impact indicators to assess whether stylistic characteristics of abstracts were similarly or differently related to citation impact.

Dataset

The only recent large-scale dataset of published journal articles from all fields with an associated quality indicator is the UK REF2021 output set. This is a collection of research outputs that were submitted for evaluation. Each output had been assigned a research quality score on an integer scale of 1 to 4 by at least two field experts from the appropriate UoA broadly field-based grouping (<https://2021.ref.ac.uk/publications-and-reports/panel-criteriaand-working-methods-201902/>). REF output scores are expert judgements of research quality on the 1* to 4* REF scale. They assess originality, significance, and rigour, ranging from

nationally recognised quality at 1* to world-leading quality at 4*. REF scores are allocated by human expert review rather than calculated from a fixed formula based on textual features, citation indicators, or other bibliometric indicators (REF, 2019). Whilst the individual output scores have been destroyed, the number of outputs awarded each of the four score levels for each “submission” is public. These public score profiles allow an average REF score to be calculated for each submission which is used here as a proxy quality score. Articles submitted by multiple departments were instead given the average of the average scores of the submitting departments. This is not a replacement for individual expert scores but is the only available large-scale REF-based quality indicator for the full dataset. In the absence of department-level biases, switching to the proxy indicator does not affect the direction of the correlation, only its strength (Thelwall, 2025c, Thelwall & Yaghi, 2025). Thus, a correlation with the proxy indicator can give evidence of a relationship and a lower bound for the strength of that relationship.

A “submission” can be thought of as approximating a department or school in a single university, although it does not have to be. The term “department” will be used here for convenience. Each output was assigned the average score from the department that submitted it, as a proxy quality rating (or the mean of all submitting departments, if an article was submitted by multiple authors). Since every department in REF2021 had multiple different scores, this is inaccurate in all cases. Nevertheless, since individual output scores correlate with departmental scores, in the absence of bias, this serves to dampen correlations but not to invalidate any positive correlations. Hence, the correlations in this study should not be interpreted as showing that ChatGPT or REF scores are based on abstract length or linguistic complexity since they only suggest that these abstract stylistic features are associated with the scores.

Individual scores have been generated for 500 articles of each of UoAs 1 to 6 for a previous paper (Thelwall & Mohammadi, 2026) and these were also used here. They were obtained by the second author of the current article and represent his estimate of the REF quality score of these articles individually.

A spreadsheet of outputs assessed was obtained from the REF2021 website (<https://results2021.ref.ac.uk/all-submission-data>) and cross-referenced with a spreadsheet of departmental scores (<https://results2021.ref.ac.uk/profiles/export-all>), via the institution and submission (department) columns. An undesirable feature of the scores is that they are fully public and may have been included in the LLM training data. Nevertheless, it seems unlikely that they were because the data is in Excel format and therefore not a natural LLM input. In addition, cross-referencing the two spreadsheets is needed to associate an article with a score, which also seems unlikely. Finally, previous tests have shown that the API version of ChatGPT rarely knows the department or even journal of an article from its title and abstract (Thelwall, 2026), so is unlikely to be able to lookup REF-related quality information about articles.

For this dataset, in each UoA, non-article outputs (e.g., books, chapters) and the articles with the shortest 10% of abstracts were discarded because the study focused on standard journal articles rather than all types of REF outputs. Many short abstracts were from short form articles (e.g., letters, correspondence) or were incomplete and were less likely to contain enough information for robust ChatGPT quality scoring. There were too many articles to check individually, especially for the difficult decision of deciding whether an article was a short form

contribution, and so the 10% heuristic was used as a blunt instrument to remove unsuitable articles. This exclusion followed the previous REF2021 ChatGPT scoring study from which the LLM scores were taken (Thelwall, 2025c). Because the 10% rule was applied within each UoA, it should not remove whole fields, although it reduces the range of abstract lengths analysed. The final dataset size was 99,277 unique journal articles without short abstracts (106,843 including duplicates).

Citation rates

A field and year normalised citation rate was calculated for each article as an additional indicator. The Normalised Log-transformed Citation Score (NLCS) was used, which adjusts for the field and year of publication as well as for the highly skewed nature of citation data (Thelwall, 2017). For this calculation, all citation counts were first transformed with $\ln(1+x)$ to reduce skewing, and then each article's (log transformed) citation count was divided by the average log-transformed citation count of all articles from the same field and year (averaging the values from multiple fields when relevant). The data for this were taken from Scopus in 2024, giving at least 3 years of citations for each paper. The fields used for the normalisation were the All Science Journal Classification (ASJC) scheme used by Scopus. This process reduces age and field biases in citation data.

LLM scores

Scores were obtained from multiple LLMs, all recycled from a previous paper (Thelwall, 2025c). These scores were used because the previous study provided a large-scale REF-based set of LLM evaluations across many disciplines and included several ChatGPT models. The details are in that paper but, in brief, the models were ChatGPT-4o, its smaller sibling ChatGPT-4o mini, and ChatGPT-5 mini. For each model, article titles and abstracts were submitted to ChatGPT through its Application Programming Interface (API). The default OpenAI API settings (e.g., temperature) were used.

The scores were based on a zero-shot system prompt describing the task with the instructions originally given to the human reviewers in REF2021, requesting a score for the article based on its title and abstract. The scores follow the REF2021 scale of 1* (nationally recognised) to 4* (world leading) (for exact prompts, see: Thelwall, 2025c). In all cases except one (described next), the scores were the average of five separate requests because averaging has been shown to increase the correlation with expert scores (Thelwall, 2025abc).

The final case without five repetitions used a simple structured prompt requesting separate comments on the three quality components of originality, rigour, and significance, followed by a score for each one, and then an overall comment and an overall score at the end (for exact prompts, see: Thelwall, 2025c). This is similar to the process that experts might be requested to use, except that the user prompt states that fractional scores (e.g., 3.5*) are allowed, whereas the human experts had to assign integers: 1*, 2*, 3* or 4*.

For the primary analysis, all scores were averaged to give a single number representing the combined estimate from both ChatGPT-4o, ChatGPT-4o mini and ChatGPT-5 mini. The previous study had found this combination to give the best overall quality ranking, in terms of the strength of correlation with expert scores (Thelwall, 2025c). The increased power of averaging older and newer models might derive from the additional iterations or the combination of perspectives from the different models.

Readability features

Abstract readability and textual characteristics of abstracts in the dataset were calculated with a Python script based on the *textstat* library (<https://pypi.org/project/textstat/>), following previous studies that used this approach to assess readability and linguistic complexity of abstracts (Alsudais, 2025; Bao et al., 2025; Porwal & Devare, 2024). For each abstract, several standard readability indicators were extracted: Flesch Reading Ease, Flesch–Kincaid Grade Level, Gunning Fog, SMOG, and ARI. Abstract length was also calculated since this has previously been shown to associate with LLM scores (Thelwall & Kurt, 2025). The main analysis focuses on the Flesch Reading Ease and Flesch–Kincaid Grade Level measures because they are widely used readability indicators in scientometrics research. The Flesch Reading Ease score estimates how easy a text is to read based on average sentence length and syllables per word, with higher scores indicating easier texts (Flesch, 1948). The Flesch–Kincaid Grade Level uses the same two linguistic components but expresses the result as a U.S. school grade level required to understand the text (Kincaid et al., 1975). The main analysis focuses on the Flesch–Kincaid Grade Level because it is widely used. Additional readability indicators (Flesch Reading Ease and Gunning Fog) are reported in the appendix (Figures S1 and S2, respectively).

Textual features

An enormous range of textual indicators can be calculated. For this study, only indicators investigated in previous citation-analysis studies were used to keep the calculations manageable and to avoid adding many indicators speculatively. Testing too many indicators without a theoretical justification risks finding spurious positives through testing volume (the Bonferroni familywise error rate issue). On this basis, the following indicators were chosen for the main analyses: abstract length (word count), words per sentence, and syllables per word. These represent key aspects of abstract style, as used in standard readability indicators, and length. To avoid redundancy, the many closely related indicators (e.g., characters per word, characters per sentence, or total number of characters) are not reported.

Analysis

Previous studies with LLM scores of academic outputs have tended to find that the individual scores are usually wrong, with a high preference for 3*, but that the useful information is the ranking of the scores, which tends to align moderately well with expert scores (slightly weaker with a quality proxy) (Thelwall, 2025c). Thus, the only relevant measure of usefulness is rank correlation. Whilst there are several different rank correlations, Spearman was chosen in preference to Kendall's tau because the latter does not penalise error size. To report statistical significance, bootstrapped 95% confidence intervals were calculated for all Spearman correlations. These are not shown in the graphs because they would be difficult to read with error bars.

Results

Main Panel A: Health and life sciences (article-level comparison)

For the health and life sciences (UoAs 1 to 6) for which there are 500 individual scores per UoA (n=3,000, including duplicate articles that are in multiple UoAs), it is possible to directly compare the ChatGPT scores (rather than departmental averages) and individual article scores

with readability indicators. Overall, the most consistent Spearman correlations are for syllables per word and words per sentence and the Flesch-Kincaid Grade Level combination of the two. In all cases the ChatGPT correlations are more positive than the REF expert correlations, suggesting that ChatGPT may be more susceptible to these factors, although causation is unproven. For example, in Clinical Medicine the Spearman correlation between Flesch–Kincaid Grade Level and ChatGPT scores was 0.329 compared with 0.124 for REF expert scores. Similar differences occur for Public Health (0.218 vs 0.037), Allied Health Professions (0.291 vs 0.072), and Psychology (0.294 vs 0.137). ChatGPT scores were also more strongly correlated with the underlying linguistic features that form this measure, such as words per sentence (0.306 vs 0.150 in Clinical Medicine) and syllables per word (0.220 vs 0.039). Similar patterns occurred for other indicators of lexical complexity, including letters per word (0.234 vs 0.023).




Figure 1. ChatGPT or expert REF scores correlated against text complexity indicators for REF2021 health and life sciences UoAs (n=500 for each UoA).

Smaller vs. Larger dataset analysis for UoAs 1 to 6

Because individual REF scores were not available for the larger dataset within Main Panel A (UoAs 1 to 6; n=39,832 articles), departmental average REF scores were used as a research quality proxy for individual papers. To assess whether the patterns in the 500-article per unit analysis (with individual scores) replicate with the proxy indicator, the analysis was repeated using the larger dataset for the six health and life science UoAs with the departmental average REF scores. Tables 1 to 3 compare the correlations between abstract linguistic complexity indicators and evaluation scores for the 500-paper dataset with individual scores and the larger dataset with departmental averages. Across all six health and life science UoAs, the correlations with ChatGPT scores are consistently higher than those for REF scores in both datasets. This pattern remains broadly similar even when REF scores are represented by departmental average proxies in the larger dataset. This suggests that analyses based on departmental average REF scores can be used to make inferences equivalently to analyses based on individual scores.

Table 1. Correlations between the Flesch-Kincaid Grade Level and research quality scores in the 500-paper sample and the larger dataset.

UoA	ChatGPT (500 papers)	REF individual scores	N	ChatGPT (Larger dataset)	REF departmental score	N
Clinical Medicine	0.220	0.039	500	0.273	0.062	9,602
Public Health	0.056	-0.067	500	0.045	0.009	3,821
Allied Health Professions	0.216	0.041	500	0.201	0.055	9,337
Psychology	0.247	0.116	500	0.195	0.085	7,881
Biological Sciences	0.049	0.035	500	0.135	0.052	6,123
Agriculture	0.348	0.215	500	0.326	0.208	3,068

Table 2. Correlations between the words per sentence and research quality scores in the 500-paper sample and the larger dataset.

UoA	ChatGPT (500 papers)	REF individual scores	N	ChatGPT (Larger dataset)	REF departmental score	N
Clinical Medicine	0.306	0.150	500	0.266	0.039	9,602
Public Health	0.233	0.071	500	0.191	-0.013	3,821
Allied Health Professions	0.239	0.077	500	0.225	0.049	9,337
Psychology	0.208	0.075	500	0.160	0.000	7,881
Biological Sciences	0.030	-0.005	500	0.021	-0.032	6,123
Agriculture	0.188	0.055	500	0.174	0.088	3,068

Table 3. Correlations between syllables per word and evaluation scores in the 500-paper sample and the larger dataset.

UoA	ChatGPT (500 papers)	REF individual scores	N	ChatGPT (Larger dataset)	REF departmental score	N
Clinical Medicine	0.329	0.124	500	0.353	0.069	9,602
Public Health	0.218	0.037	500	0.185	0.003	3,821
Allied Health Professions	0.291	0.072	500	0.285	0.071	9,337
Psychology	0.294	0.137	500	0.244	0.057	7,881
Biological Sciences	0.058	0.019	500	0.113	0.013	6,123
Agriculture	0.342	0.168	500	0.322	0.177	3,068

Analysis across all 34 REF UoAs

Based on the comparison between ChatGPT and REF expert correlations in Main Panel A (Figure 1), three indicators were selected for further analysis across all 34 REF UoAs: Flesch–Kincaid Grade Level, words per sentence, and word count. These indicators were chosen because they capture different aspects of abstract writing style. Flesch–Kincaid Grade Level represents overall linguistic complexity, words per sentence reflects syntactic complexity, and word count measures overall abstract length, which is widely used in scientometric research.

Additional analyses using stylistic indicators with different linguistic characteristics, including Gunning Fog, SMOG, average letters per word, and sentence count, syllabus count are reported in the Appendix.

Flesch–Kincaid Grade Level

The Flesch–Kincaid Grade Level is a weighted combination of syllables per word and words per sentence that estimates the educational grade level required to understand a text. Higher scores indicate more linguistically complex abstracts and therefore a higher level of reading difficulty (Kincaid et al., 1975). Overall, 23 of the 34 UoAs had statistically significant positive correlations between ChatGPT scores and Flesch–Kincaid Grade Level, while one unit had a statistically significant negative association and the remaining UoAs had no statistically significant relationships, possibly due to small sample sizes and/or weak relationships. This indicates that linguistically more complex abstracts tended to receive higher ChatGPT scores. However, REF expert departmental average scores statistically significantly positively correlated with grade level in only 9 UoAs out of 34, following the same direction as ChatGPT, although these correlations were generally much smaller (Figure 2). Six UoAs also had statistically significant negative correlations between grade level and REF expert scores, including Politics and International Studies, Education, and several humanities subjects.

In several biological and physical sciences, ChatGPT scores had stronger positive associations than REF expert scores with abstract complexity features. For example, in Engineering (ChatGPT: 0.215; REF: 0.073), Chemistry (0.212; 0.083), and Biological Sciences (0.113; 0.013). These results suggest that ChatGPT scores were more strongly associated with linguistic complexity in abstracts, whereas REF expert scores were generally weakly related or unrelated to this readability measure.

Across many social science subjects, ChatGPT scores were more likely to increase with higher grade levels, whereas REF expert scores were often weakly or not significantly related to this measure. In several humanities subjects, REF expert scores also had weak and inconsistent associations with Flesch–Kincaid Grade Level, with some significant negative correlations in Modern Languages and Linguistics (–0.163), English Language and Literature (–0.109), and Philosophy (–0.112), suggesting that more linguistically complex abstracts were not favoured by expert evaluations in these subjects.




Figure 2. Spearman correlations between Flesch–Kincaid Grade Level and ChatGPT and REF departmental average expert scores across the 34 REF Units of Assessment.

Abstract length

In Figure 3, 29 of the 34 UoAs had significant correlations between ChatGPT scores and abstract length (number of words), including 22 positive and 7 negative associations. Positive correlations indicate that longer abstracts tended to receive higher ChatGPT scores, whereas negative correlations suggest the opposite. Negative correlations occurred in several biomedical and STEM fields, such as Clinical Medicine (-0.311) and Biological Sciences (-0.188), whereas positive associations appeared in many social science and humanities subjects, including History (0.218), English Language and Literature (0.194), and Philosophy (0.214). The negative association between abstract length and ChatGPT scores in some biomedical and STEM subjects, especially Clinical Medicine, may be related to the style of abstracts in these fields. Medical journals often require structured abstracts (e.g., Aims, Participants, Methods, Results) and have strict word limits to summarise key findings clearly. However, shorter abstracts do not necessarily mean they are easier to read. As shown in Figure

2, Clinical Medicine also had the strongest positive association between linguistic complexity and ChatGPT scores. In other words, ChatGPT may be more likely to assign higher scores to shorter but more linguistically complex abstracts.

In contrast, departmental average REF expert scores had statistically significant associations with abstract length in only 15 UoAs and most of these were negative (12 negative and 3 positive), indicating that shorter abstracts were often associated with higher average expert scores in these subjects.




Figure 3. Spearman correlations between abstract length (word count) and ChatGPT and REF departmental average expert scores across the 34 REF Units of Assessment.

Words per sentence

Words per sentence partially reflects the syntactic complexity of the text in terms of average sentence length. Many medical and Science, Technology, Engineering and Maths (STEM) UoAs had statistically significant positive correlations between ChatGPT scores and words per sentence, and these were generally stronger and more frequent than those for REF departmental average expert scores (Figure 4). However, in several arts and humanities subjects, REF expert scores had negative associations with sentence length while ChatGPT scores had little or no relationship, for example in Modern Languages and Linguistics (-0.198) and English Language and Literature (-0.135), suggesting that some REF expert evaluations may favour clearer writing with shorter sentences.




Figure 4. Spearman correlations between words per sentence and ChatGPT and REF expert scores across the 34 REF Units of Assessment.

Syllables per word

Syllables per word is a simple vocabulary complexity indicator. Higher values generally indicate longer or more technical words in the text. In most subjects, abstracts with more syllables per word tended to receive higher ChatGPT scores compared with REF expert departmental average scores. Across the 34 REF UoAs, 28 had statistically significant positive correlations between ChatGPT scores and syllables per word, whereas expert scores had statistically significant correlations in only 11 UoAs (Figure 5). For example, positive correlations between syllables per word and ChatGPT scores occurred in Clinical Medicine (0.273), Physics (0.262), Chemistry (0.236), and Engineering (0.210), whereas REF expert correlations with syllables per word in these subjects were much weaker. Similar patterns also occurred in several social sciences and humanities subjects, including Philosophy (0.211), History (0.190), Education (0.148), and English Language and Literature (0.145), where REF correlations were generally

weak, close to zero, or negative. This suggests that ChatGPT scores may be more sensitive than expert evaluations to vocabulary complexity in abstracts.




Figure 5. Spearman correlations between syllables per word and ChatGPT and REF expert scores across the 34 REF Units of Assessment.

Discussion

Comparison between ChatGPT scores and citation impact

Although citation impact is not a direct measure of research quality, it was included as an additional article-level indicator for comparison. Citation indicators are widely used in research evaluation and scientometrics, and previous REF2021 research found that field-normalised citation scores were positively associated with REF expert scores across all 34 UoAs, although the strength of the association varied substantially between fields, from very weak to strong (Thelwall et al., 2023). As a further analysis, ChatGPT scores and the Normalised Log-transformed Citation Score (NLCS 2024) were correlated with abstract readability (Flesch–Kincaid Grade Level). Both ChatGPT scores and NLCS were collected at the individual article level, allowing a direct comparison between the two indicators.

Figure 6 compares the correlations between ChatGPT and citation impact scores with Flesch–Kincaid Grade Level across the 34 REF Units of Assessment. In several biomedical and STEM subjects, ChatGPT scores were clearly associated with higher linguistic complexity, but the same was not true for citation impact: correlations were generally weaker and usually negative. This suggests that LLM-based evaluations are usually more sensitive to abstract complexity than citation-based indicators of research impact.




Figure 6. Spearman correlations between ChatGPT scores and citation impact (NLCS 2024) with Flesch–Kincaid Grade Level across the 34 REF Units of Assessment.

Limitations

This study has several data, methods, and interpretation limitations. The journal articles are from a single country and there may be different patterns for other countries, and particularly for those where English (the dominant language of science) is not commonly spoken. The articles are also preselected for quality because each department tries to submit its best work for the evaluation. This is likely to have reduced the strength of the correlations. Removing articles with the shortest abstracts (to remove short form contributions, which might be

expected to systematically attract lower scores) may also have reduced the strength of the abstract length correlations. Another limitation is that other LLMs, prompting strategies, or textual indicators may have yielded different results.

Unfortunately, individual REF expert scores were not available for the larger dataset used in Figures 2–4. Instead, departmental average REF scores were used as a proxy for individual papers. This approach reduces variation in the expert scores assigned to articles because multiple papers from the same department receive the same average score. As a result, correlations between abstract characteristics and REF scores may be weaker than if individual expert scores were available. This is unlikely to change the general direction of the correlations, but it may have reduced their strength. Some correlations that were non-significant may therefore have reached statistical significance if the original individual scores had been available. However, the direct comparison for UoAs 1 to 6 suggests that the reduction would not be large. Nevertheless, comparisons between ChatGPT and REF correlations should be interpreted cautiously, as ChatGPT scores were available at the article level whereas REF scores represent departmental averages. Another limitation is that the analysis focuses on key textual and readability features, such as abstract length and linguistic complexity, and does not consider other potentially relevant characteristics such as abstract structure (e.g., background, aims, methods, and conclusions), or methodological detail. These additional textual features may also influence LLM scoring. The findings may also differ for other cloud-based or locally hosted LLMs or for different prompting strategies which needs future investigations.

Finally, the study used correlation analyses which do not show causal relationships. For example, longer or more linguistically complex abstracts may not directly cause higher ChatGPT scores but may instead be associated with other characteristics of academic writing that influence the model's evaluation. Future research could test whether the same associations appear for other LLMs and when additional textual and linguistic features are included.

Conclusion

Although cause-and-effect has not been proven, the results suggest that abstracts with complex styles impress ChatGPT more than they impress expert evaluators, in the sense of ChatGPT tending to assign them higher scores. This is a plausible finding in the context of the text pattern matching of LLMs. For example, it may have previously ingested many unpublished badly written and poor manuscripts by hopeful authors, therefore “learning” to associate plain grammar with weak research. It is also possible that its historical data, with deliberately complex language being highly regarded, has given the same effect. More generally, it may have identified that outside academia humans can be impressed by complex language.

In answer to RQ1, the most consistent positive associations were for linguistic complexity indicators, especially Flesch-Kincaid Grade Level, words per sentence, and syllables per word. Abstract length was also associated with LLM scores, but the direction of the association varied by subject. In answer to RQ2, these associations were generally stronger and more frequent for LLM scores than for REF expert scores. This is evident with direct quality scores for the health and life sciences (UoAs 1 to 6) and is also suggested by the results for the remaining UoAs with proxy quality scores. In answer to RQ3, the patterns differed across UoAs,

with stronger positive associations for linguistic complexity indicators in many health, science, and engineering subjects, and more mixed patterns in social sciences and humanities.

The findings suggest the possibility that some textual characteristics of abstracts could be intentionally modified to obtain higher LLM evaluation scores. If LLM systems tend to give higher scores to features such as longer or more linguistically complex abstracts, these characteristics could potentially be adjusted by authors to increase their scores. If LLMs were used to support research evaluation processes such as peer review or national assessments like the REF, this could create the motivation to modify writing style rather than improve the underlying research quality. The behavioural consequences of such motivations are not clear yet. Hence, identifying which textual features influence LLM evaluation is important to understand potential risks of manipulation and to support the responsible use of LLMs in research assessment. Nevertheless, most of the correlations between LLM scores and abstract stylistic features are low to modest, suggesting that these features are unlikely to determine LLM evaluation scores by themselves. However, they may still influence the evaluation of individual papers when LLM scores are used to support research assessment. Given the results, it seems desirable to reduce the apparent impact of language complexity on LLM scores so that (a) authors are not rewarded for their language instead of their research quality and (b) authors do not try to write in a more complex style to impress LLMs. This might be achieved by incorporating prompt instructions to ignore language style or by using LLMs trained exclusively on relevant academic texts, if possible. Whilst in theory, the problem could be reduced by assigning a score penalty to complex texts to reduce their LLM score rank, this may be undesirable interference in the evaluation process and perhaps difficult to get agreement about.

Acknowledgement

This study is partly funded by the Economic and Social Research Council (ESRC), UK (UKRI2101).

Data availability: The data that support the findings of this study are available from the corresponding author upon reasonable request.

Declarations

Competing interests: The first and second authors are members of the Distinguished Reviewers Board of Scientometrics.

References

- Alsudais, A. (2025). Exploring the change in scientific readability following the release of ChatGPT. *Journal of Informetrics*, 19(3), 101679.
- Ante, L. (2022). The relationship between readability and scientific impact: Evidence from emerging technology discourses. *Journal of Informetrics*, 16(1), 101252.
- Bao, T., Zhao, Y., Mao, J., & Zhang, C. (2025). Examining linguistic shifts in academic writing before and after the launch of ChatGPT: a study on preprint papers. *Scientometrics*, 130(7), 3597-3627.

- Didegah, F., & Thelwall, M. (2013). Which factors help authors produce the highest impact research? Collaboration, journal and document properties. *Journal of Informetrics*, 7(4), 861- 873.
- Flesch, R. (1948). A new readability yardstick. *Journal of applied psychology*, 32(3), 221. https://comp311.wordpress.com/wp-content/uploads/2010/11/flesch_rudolph.pdf
- Gazni, A. (2011). Are the abstracts of high impact articles more readable? Investigating the evidence from top research institutions in the world. *Journal of Information Science*, 37(3), 273-281.
- Germani, F., & Spitale, G. (2025). Source framing triggers systematic bias in large language models. *Science Advances*, 11(45), eadz2924.
- Hafeez, D. M., Jalal, S., & Khosa, F. (2019). Bibliometric analysis of manuscript characteristics that influence citations: A comparison of six major psychiatry journals. *Journal of Psychiatric Research*, 108, 90-94.
- Helbach, J., Wandscher, K., Pieper, D., & Hoffmann, F. (2025). Readers' attention to shorter versus longer abstracts of systematic reviews: a randomised controlled trial. *BMJ Evidence-Based Medicine*.
- Hu, H., Wang, D., & Deng, S. (2021). Analysis of the scientific literature's abstract writing style and citations. *Online Information Review*, 45(7), 1290-1305.
- Jin, T., Duan, H., Lu, X., Ni, J., & Guo, K. (2021). Do research articles with more readable abstracts receive higher online attention? Evidence from Science. *Scientometrics*, 126(10), 8471-8490.
- Kincaid, J. P., Fishburne Jr, R. P., Rogers, R. L., & Chissom, B. S. (1975). Derivation of new readability formulas (automated readability index, fog count and Flesch reading ease formula) for navy enlisted personnel (No. RBR875). <https://stars.library.ucf.edu/cgi/viewcontent.cgi?article=1055&context=istlibrary>
- Kostoff, R. (2007). The difference between highly and poorly cited medical articles in the journal *Lancet*. *Scientometrics*, 72, 513-520.
- Kousha, K., & Thelwall, M. (2024). Factors associating with or predicting more cited or higher quality journal articles: An Annual Review of Information Science and Technology (ARIST) paper. *Journal of the Association for Information Science and Technology*, 75(3), 215-244.
- Kousha, K., & Thelwall, M. (2025). Assessing the societal influence of academic research with ChatGPT: Impact case study evaluations. *Journal of the Association for Information Science and Technology*, 76(10), 1357-1373.
- Kousha, K., & Thelwall, M. (2026). How much are LLMs changing the language of academic papers after ChatGPT? A multi-database and full text analysis. *Scientometrics*. preprint arXiv:2509.09596.
- Kundu, A., & Barbosa, D. (2024). Are large language models good essay graders? *arXiv preprint arXiv:2409.13120*.
- Nunkoo, R. & Thelwall, M. (2025). Evaluating the quality of tourism research using ChatGPT. *Current Issues in Tourism*. <https://doi.org/10.1080/13683500.2025.2596265>
- Porwal, P., & Devare, M. H. (2024). Scientific impact analysis: Unraveling the link between linguistic properties and citations. *Journal of Informetrics*, 18(3), 101526.

- REF (2019). Panel criteria and working methods (REF 2019/02). <https://2021.ref.ac.uk/publications-and-reports/panel-criteria-and-working-methods-201902/index.html>
- Sienkiewicz, J., & Altmann, E. G. (2016). Impact of lexical and sentiment factors on the popularity of scientific papers. *Royal Society Open Science*, 3(6), 160140.
- Sohrabi, B., & Iraj, H. (2017). The effect of keyword repetition in abstract and keyword frequency per journal in predicting citation counts. *Scientometrics*, 110(1), 243-251.
- Thelwall, M., Kousha, K., Stuart, E., Makita, M., Abdoli, M., Wilson, P., & Levitt, J. (2023). In which fields are citations indicators of research quality?. *Journal of the Association for Information Science and Technology*, 74(8), 941-953.
- Thelwall, M. & Yaghi, A. (2025). In which fields can ChatGPT detect journal article quality? An evaluation of REF2021 results. *Trends in Information Management*, 13(1), 1-29.
- Thelwall, M. (2017). Three practical field normalised alternative indicator formulae for research evaluation. *Journal of informetrics*, 11(1), 128-151.
- Thelwall, M. (2025a). Evaluating research quality with Large Language Models: An analysis of ChatGPT's effectiveness with different settings and inputs. *Journal of Data and Information Science*, 10(1), 7-25. <https://doi.org/10.2478/jdis-2025-0011>
- Thelwall, M. (2025b). Is Google Gemini better than ChatGPT at evaluating research quality? *Journal of Data and Information Science*, 10(2), 1–5. <https://doi.org/10.2478/jdis-2025-0014> with extended version here: <https://doi.org/10.6084/m9.figshare.28089206.v1>
- Thelwall, M. (2025c). In which fields do ChatGPT scores align better than citations with research quality? arXiv preprint arXiv:2504.04464.
- Thelwall, M. (2025d). Can ChatGPT replace citations for quality evaluation of academic articles and journals? Empirical evidence from library and information science. *Journal of Documentation*, 81(4), 1078-1094.
- Thelwall, M. (2026). Do Large Language Models know basic facts about journal articles? *Journal of Documentation*. <https://doi.org/10.1108/JD-11-2025-0330>
- Thelwall, M., & Kurt, Z. (2025). Research evaluation with ChatGPT: Is it age, country, length, or field biased? *Scientometrics*, 130(10), 5323-5343. <https://doi.org/10.1007/s11192025-05393-0>
- Thelwall, M., & Mohammadi, E. (2026). Can small and reasoning large language models score journal articles for research quality and do averaging and few-shot help? *Scientometrics*. arXiv preprint arXiv:2510.22389.
- Van Wesel, M., Wyatt, S., & ten Haaf, J. (2014). What a difference a colon makes: How superficial factors influence subsequent citation. *Scientometrics*, 98(3), 1601-1615.
- Vasu, S. S. M., Sheth, I., Wang, H. P., Binkyte, R., & Fritz, M. (2025). Justice in judgment: Unveiling (hidden) bias in LLM-assisted peer reviews. arXiv preprint arXiv:2509.13400.
- Wang, S., Liu, X., & Zhou, J. (2022). Readability is decreasing in language and linguistics. *Scientometrics*, 127(8), 4697-4729.
- Zhu, Y., Haunschild, R., Ye, J., & Xie, H. (2026). Can generative AI effectively perform quality evaluation within social sciences? A case study in library and information science. *Scientometrics*, 1-21.

Appendix: Additional readability analysis




Figure S1. Spearman correlations between **Flesch Reading Ease** with ChatGPT and REF expert scores across the 34 REF Units of Assessment. Higher Flesch Reading Ease scores indicate easier-to-read abstracts (Flesch, 1948). Positive correlations indicate that more readable abstracts tended to receive higher scores, whereas negative correlations indicate that less readable (more linguistically complex) abstracts tended to receive higher scores.




Figure S2. Spearman correlations between **Gunning Fog** readability scores and ChatGPT and REF expert scores across the 34 REF Units of Assessment. Higher Gunning Fog scores indicate more linguistically complex abstracts. Positive correlations indicate that more complex abstracts tended to receive higher scores, whereas negative correlations indicate that simpler abstracts tended to receive higher scores.