Testing Pangram 4

A practical exercise with AI-written, humanised and copy-edited text

Pangram is an AI-writing detector. This exercise uses Pangram 4 to show how its verdict changes when the same AI-written passage is humanised and then copy-edited by another AI tool. You begin with a published research paper, ask an AI chatbot to write a new conclusion, and test the AI-written version, a humanised version and an AI copy-edit. The exercise takes about 15 minutes and can be completed using free AI tools and Pangram's free plan.[1]

1 Choose a research paper

Choose a fairly straightforward open-access research paper published before 2020, so that it is unlikely to have been written using AI. Avoid papers about AI, as writing on this topic has anecdotally received a higher AI-writing score from Pangram, which can make the exercise harder. The published conclusion from Houlden *et al.*, "A spatial analysis of proximate greenspace and mental wellbeing in London", *Applied Geography* (2019) is used here.[2]

> While many previous studies have failed to detect any association between greenspace in an individual's administrative local area and their mental wellbeing, this study applied weighted regression methods to account for underlying spatial heterogeneity and reveal a positive association between greenspace around homes and both hedonic and eudaimonic wellbeing. Positive, statistically significant associations were found for presence of greenspace within the recommended 300 m distance and mental wellbeing, while weaker associations at greater distances imply that greenspace closer to homes may be more important. Variation in the strengths of these associations highlights the fluctuating importance of greenspace for mental wellbeing in different areas of the city. Associations of greenspace with wellbeing were also found to be generally weaker in the centre of London, perhaps due to differing compositions of greenspace in different areas of the capital. While UK government guidelines recommend that greenspace should be provided within 300 m of all residents to benefit health, these results provide some evidence that this distance is also associated with higher levels of mental wellbeing. It may therefore be suggested that urban planners should be encouraged to include greenspace close to residents, for potential mental wellbeing benefits. Future studies should continue to apply more methodological approaches which consider the spatial nature of the data, and expand of this work by considering actual travel distances and different types of greenspaces.

2 Ask an AI chatbot to write a new conclusion

Copy and paste the body of the paper into an AI chatbot and ask it to write a conclusion of roughly the same length as the original. Use a simple prompt such as this one:

> *Read this research paper and write a conclusion of about 200 words.*

The AI-written conclusion below was generated with Claude Fable 5 High, which requires a paid Claude plan. Any chatbot will do.

> The contribution of this analysis is as much methodological as substantive. Measuring greenspace within a 300 metre buffer of each residents' postcode, rather than within administrative boundaries, revealed associations with wellbeing that boundary-based studies had repeatedly missed: in fully adjusted models, an increase of 1 ha of greenspace within 300 metres was associated with a 0.803 increase in life satisfaction, with smaller but significant coefficients for worth (0.740) and happiness (0.521). The effect diminished sharply with distance, falling to 0.330 at 500 metres for life satisfaction and approaching zero at 1 kilometre. Proximity, in other words, appears to do real work.
>
> The spatial variation is harder to interpret. Coefficients were stronger toward the outskirts of London and weaker, occasionally negative, toward the centre and east, which may reflect differences in the character of the greenspace itself rather than its quantity. Since the data are cross-sectional and confined to one city, and since straight-line buffers ignore how people actually reach a park, these findings support the UK guideline of greenspace within 300 metres of homes without settling why it works, for whom, or whether the same distance would matter in a city unlike London.

3 Compare the AI-written conclusion with the published version

Compare the AI-written conclusion with the published version. Focus on how the sentences and the overall argument are built, rather than hunting for individual "AI-sounding" words. Formulaic sentence structures can be one sign of AI-generated prose, but no single feature is definitive on its own.[3] In this example, several differences stand out:

> **Generic academic opener**
> "The contribution of this analysis is as much methodological as substantive." Sounds balanced and important, but is empty and formulaic. The published conclusion starts more directly with what earlier studies failed to find.
>
> **Soft transition and shift in tone**
> "Proximity, in other words, appears to do real work." The AI uses a smooth transition followed by a slightly more conversational phrase. The published text stays in the same academic tone throughout.
>
> **Qualified results become a clearer story**
> The published text says weaker associations at greater distances "imply" that nearby greenspace may be more important. The AI turns this into an effect that "diminished sharply" and approached zero.
>
> **Neatly staged construction**
> The AI breaks the argument into distinct stages. It reports the results, adds a short interpretation, flags uncertainty, offers a possible explanation and ends with a balanced list of limitations and open questions. The two short standalone sentences make those stages easier to notice. In the published conclusion, the interpretation and uncertainty are woven into longer sentences, so the structure feels less tidy.
>
> **Ready-made explanation after flagging a problem**
> "The spatial variation is harder to interpret." The AI isolates the uncertainty in a seven-word sentence and then, in the next sentence, offers a clean possible reason: differences in the character of the greenspace. The published version also offers a possible explanation, but embeds it within a longer account of the spatial variation rather than presenting it as a neat problem-and-answer sequence.
>
> **Perfectly packaged limitations**
> The final AI sentence gathers three limitations and three open questions into one smooth statement. The published conclusion ends more simply.

4 Check the AI-written conclusion in Pangram

Paste the AI-written conclusion into Pangram. In this example, Pangram 4 classified it as 100% AI-generated. Confidence was limited because the passage was short.[4]

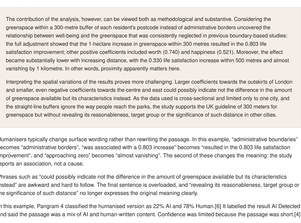

5 Humanise the conclusion and check it again

Paste the AI-written conclusion into a humanising tool. HumanizeAI.pro was used for this example.[5] The resulting text is shown below.

> The contribution of the analysis, however, can be viewed both as methodological and substantive. Considering the greenspace within a 300-metre buffer of each resident's postcode—rather than administrative borders—the study uncovered a relationship between well-being and the greenspace that was consistently neglected in previous boundary-based studies: the full adjustment showed that the 1-hectare increase in greenspace within 300 metres resulted in the 0.803 life satisfaction improvement; other positive coefficients included worth (0.740) and happiness (0.521). Moreover, the effect became substantially lower with increasing distance, with the 0.330 life satisfaction increase within 500 metres and almost vanishing by 1 kilometre. In other words, proximity clearly matters here.
>
> Interpreting the spatial variations of the results proves more challenging. Larger coefficients towards the outskirts of London and smaller, even negative coefficients towards the centre and east could possibly indicate not the difference in the amount of greenspace available but its characteristics instead. As the data used is cross-sectional and limited to one city, and the straight-line buffers ignore the way people reach the parks, the study supports the UK guideline of 300 meters for greenspace but without revealing its reasonableness, target group or the significance of such distance in other cities.

Humanisers typically change surface wording rather than rewriting the passage. In this example, "administrative boundaries" becomes "administrative borders", "was associated with a 0.803 increase" becomes "resulted in the 0.803 life satisfaction improvement", and "approaching zero" becomes "almost vanishing". The second of these changes the meaning: the study reports an association, not a cause.

Phrases such as "could possibly indicate not the difference in the amount of greenspace available but its characteristics instead" are awkward and hard to follow. The final sentence is overloaded, and "revealing its reasonableness, target group or the significance of such distance" no longer expresses the original meaning clearly.

In this example, Pangram 4 classified the humanised version as 22% AI and 78% Human.[6] It labelled the result AI Detected and said the passage was a mix of AI and human-written content. Confidence was limited because the passage was short.[4]

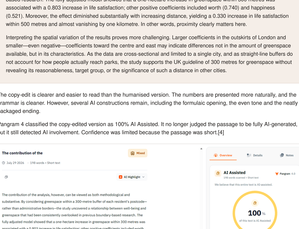

6 Ask an AI to copy-edit the humanised version

Paste the humanised version into chat-deep.ai and use this prompt:[7]

> *Act as a copy editor for an academic journal. Can you tidy this up?*

The contribution of the analysis, however, can be viewed as both methodological and substantive. By considering greenspace within a 300-metre buffer of each resident's postcode—rather than administrative borders—the study uncovered a relationship between well-being and greenspace that had been consistently overlooked in previous boundary-based research. The fully adjusted model showed that a one-hectare increase in greenspace within 300 metres was associated with a 0.803 increase in life satisfaction; other positive coefficients included worth (0.740) and happiness (0.521). Moreover, the effect diminished substantially with increasing distance, yielding a 0.330 increase in life satisfaction within 500 metres and almost vanishing by one kilometre. In other words, proximity clearly matters here.

Interpreting the spatial variation of the results proves more challenging. Larger coefficients in the outskirts of London and smaller—even negative—coefficients toward the centre and east may indicate differences not in the amount of greenspace available, but in its characteristics. As the data are cross-sectional and limited to a single city, and as straight-line buffers do not account for how people actually reach parks, the study supports the UK guideline of 300 metres for greenspace without revealing its reasonableness, target group, or the significance of such a distance in other cities.

The copy-edit is cleaner and easier to read than the humanised version. The numbers are presented more naturally, and the grammar is clearer. However, several AI constructions remain, including the formulaic opening, the even tone and the neatly packaged ending.

Pangram 4 classified the copy-edited version as 100% AI-Assisted. It no longer judged the passage to be fully AI-generated, but it still detected AI involvement. Confidence was limited because the passage was short.[4]

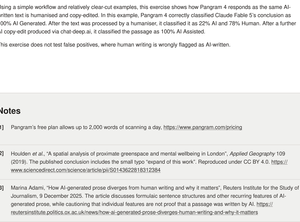

7 Try editing it yourself

If you have time, go back to the original AI-written conclusion and edit it yourself to see whether you can get Pangram to classify it as 100% Human. Because the starting point is entirely AI-generated, its wording and structure can anchor your edits, so changing a few words or inserting a deliberate mistake may not be enough to move the classification from AI to Human, making the task harder than it first appears.[8] If the initial text had already involved substantial human input, this task is often easier and the detector is more sensitive to relatively small edits.

8 What this exercise shows

Using a simple workflow and relatively clear-cut examples, this exercise shows how Pangram 4 responds as the same AI-written text is humanised and copy-edited. In this example, Pangram 4 correctly classified Claude Fable 5's conclusion as 100% AI Generated. After the text was processed by a humaniser, it classified it as 22% AI and 78% Human. After a further AI copy-edit produced via chat-deep.ai, it classified the passage as 100% AI Assisted.

This exercise does not test false positives, where human writing is wrongly flagged as AI-written.

Notes

[1] Pangram's free plan allows up to 2,000 words of scanning a day, https://www.pangram.com/pricing

[2] Houlden *et al.*, "A spatial analysis of proximate greenspace and mental wellbeing in London", *Applied Geography* 109 (2019). The published conclusion includes the small typo "expand of this work". Reproduced under CC BY 4.0. https://www.sciencedirect.com/science/article/pii/S0143622819312364

[3] Marina Adami, "How AI-generated prose diverges from human writing and why it matters", Reuters Institute for the Study of Journalism, 9 December 2025. The article discusses formulaic sentence structures and other recurring features of AI-generated prose, while cautioning that individual features are not proof that a passage was written by AI. https://reutersinstitute.politics.ox.ac.uk/news/how-ai-generated-prose-diverges-human-writing-and-why-it-matters

[4] Pangram says it detects Claude Fable 5 reliably, classifying 1,111 of 1,115 samples as fully AI-generated. https://www.pangram.com/blogs/sota-pangram-work-on-claude-fable-5 But the passages here are short, at around 200 words, and Pangram says it performs best on texts over a couple of hundred words. https://www.pangram.com/blog/all-about-false-positives-in-ai-detectors

[5] HumanizeAI.pro (https://humanizeai.pro) was used here. Other examples include WriteHuman (https://writehuman.ai), GPTHuman (https://gpthuman.ai), HumanJoy (https://humanjoy.com) and Undetectable AI (https://undetectable.ai/ai-humanizer).

[6] In our own testing, earlier versions of Pangram would usually classify humanised text as 100% Human. Pangram 4 appears harder to trick, and Pangram reports significant improvements in humaniser detecting, saying it identifies AI involvement in humanised output 98.83% of the time across 13 commercial humanisers. https://www.pangram.com/blog/introducing-pangram-4

[7] The copy-edit was produced via chat-deep.ai (https://chat-deep.ai/deepseek-chat), an unaffiliated third-party interface to the DeepSeek API that reputed no sign-up. DeepSeek's own service is at https://chat.deepseek.com. The copy-editing prompt is deliberately simple because detailed instructions may encourage the model to retain the existing structure and AI-isms.

[8] A 2026 post-editing study found that participants moved AI-generated drafts towards their own style, but their revisions still remained closer to the AI draft than to their unaided writing. Connor Baumler et al., "Can You Make It Sound Like You? Post-Editing LLM-Generated Text for Personal Style", Proceedings of ACL 2026. https://aclanthology.org/2026.acl-long.2030/

© Innovation Ideas Ltd 2026