

AI Detection Analysis Report

Generated: June 29, 2026 | Topic: Reconceptualising Authorship in the Age of Generative AI: A Theoretical Review

Client Action Summary

Overall AI-detection risk: Low (0%) — the text reads as human-written.

Confidence: High — based on the AI-detection model.

Sections to review: No section needs structural attention — only the optional phrasing refinements listed below.

Recommended actions:

- Vary how sentences begin.
- Strings of exactly three items can read as mechanical when they recur — a common AI-writing tic. Keep the two strongest items, or rephrase one triad as a full sentence, so the rhythm varies.

Section-by-Section Action Plan

Section	Status	What to improve
Abstract	OK (normal for this section type) — Low sentence variation is normal here.	No change needed.
Introduction	OK	No change needed.
Scope and Source Selection	OK	No change needed.
Authorship as Spectrum: From Binary to Continuum	OK	No change needed.
Legal and Institutional Frameworks Under Pressure	OK	No change needed.
Tensions and Debates	OK	No change needed.
Theoretical Synthesis	OK	No change needed.

Section	Status	What to improve
Position within the field	OK	No change needed.
Theoretical contribution	OK	No change needed.
Limitations	OK	No change needed.
Future Research	OK	No change needed.
Recap of the argument	OK	No change needed.
Concluding remarks	OK	No change needed.
References	OK (normal for this section type) — Low sentence variation is normal here.	No change needed.

Phrasing to Revise

Issue	Suggested fix
Many sentences begin with “the” (60 times)	Vary how sentences begin.
Many sentences begin with “this” (17 times)	Vary how sentences begin.
Many sentences begin with “when” (5 times)	Vary how sentences begin.
Many sentences begin with “section” (4 times)	Vary how sentences begin.
Many sentences begin with “if” (4 times)	Vary how sentences begin.
Repeated three-item lists, e.g. “analytic epistemology, posthumanist scholarship, and the sociology of” (41 times)	Strings of exactly three items can read as mechanical when they recur — a common AI-writing tic. Keep the two strongest items, or rephrase one triad as a full sentence, so the rhythm varies.

Executive Summary

This report analyses the document for AI-generated content indicators, comparing scores **before** and **after** the academic humanisation (polish) stage.

Metric	Before Polish	After Polish	Change
AI Detection Score	0%	0%	↓ 0pp
Burstiness	78%	78%	↑ 1pp

Metric	Before Polish	After Polish	Change
Perplexity Proxy	80%	80%	↑ 0pp

Understanding the Metrics

AI detection tools are not infallible. Different detectors use different methods and may produce contradictory results. This report provides indicators based on statistical analysis and a dedicated detection model, but results should be interpreted as guidance, not absolute truth.

AI Detection Score (chatgpt-detector-roberta model)

What it measures: This score is produced by a dedicated AI detection model (Hello-SimpleAI/chatgpt-detector-roberta), a RoBERTa transformer fine-tuned specifically to classify text as human-written or AI-generated. The model analyses linguistic patterns, word choice distributions, and contextual coherence to produce a probability score.

How to read it: 0% = certainly human-written, 100% = certainly AI-generated — the higher the score, the more likely the text is to be flagged as AI-generated. The risk levels below explain what each band means.

Risk Levels

Low Risk (0–30%): The text exhibits characteristics typical of human writing. Most AI detectors are unlikely to flag it.

Medium Risk (30–60%): The text has some AI-like patterns. Some detectors may flag portions. Consider reviewing flagged sections for natural phrasing.

High Risk (60–100%): The text shows strong AI-generation signals. Most detectors will likely flag it. Significant rewriting may be needed.

Burstiness

What it measures: Burstiness measures the variation in sentence lengths throughout the document. Human writers naturally produce “bursty” text — a short punchy sentence followed by a long analytical one, then a medium-length transition. AI models tend to generate sentences of similar length, resulting in monotonous, uniform rhythm.

How to read it: Technically, burstiness is the coefficient of variation (standard deviation / mean) of sentence word counts, normalised to a 0–1 scale. A burstiness of 0% means all sentences are the same length (AI-like). A burstiness of 60%+ indicates natural human variation.

IMPORTANT: low burstiness in some sections is **NORMAL**, not a defect. Sections built on a fixed structural template — Abstract, Research Objectives, Hypotheses, References, bulleted lists — are inherently low-burstiness even when written entirely by a human author. The metric measures sentence-length variance, and a 4-bullet list of research objectives shaped “To investigate... / To identify... / To evaluate...” will score around 25–40% by definition. Interpret the score in context of the section type; analytical sections (Discussion, Theoretical Background) are the meaningful places to look for low burstiness as an AI signal.

Perplexity Proxy

What it measures: True perplexity requires running the text through a language model to measure how “surprised” it is by each word. Since we avoid loading heavy ML models, this metric approximates perplexity using measurable proxies: vocabulary richness (how many unique words relative to total words), long-word usage, and sentence structure variety.

How to read it: Higher values indicate more diverse, less predictable text — characteristic of human writing. Lower values suggest formulaic, predictable patterns typical of AI output. This is a statistical approximation, not a true perplexity measurement.

Vocabulary Richness (Root TTR)

What it measures: The root type-token ratio counts unique words divided by the square root of total words. This corrects for document length — longer texts naturally repeat words more. AI-generated text tends to reuse the same vocabulary repeatedly, while human writers draw from a broader lexicon with more varied word choices.

How to read it: Values are reported as a raw score. Higher values indicate richer, more varied vocabulary. This metric feeds into the Perplexity Proxy composite score.

Limitations

- The AI detection score depends on the chatgpt-detector-roberta model, which may not cover all languages equally well. Results are most reliable for English text.
- Burstiness and perplexity proxy are statistical approximations. Short documents (< 500 words) produce less reliable scores.

- Detection technology evolves rapidly. Scores reflect the state of detection methods at the time of generation.
- Academic writing naturally shares some patterns with AI output (formal register, structured arguments). This can inflate scores.

Section-by-Section Analysis

The following table shows detection scores for each major section of the document.

Section	AI Score (Before)	AI Score (After)	Burstiness (Before)	Burstiness (After)
Abstract	0%	0%	61%	66%
Introduction	0%	0%	100%	100%
Scope and Source Selection	0%	0%	46%	48%
Authorship as Spectrum: From Binary to C	0%	0%	58%	56%
Legal and Institutional Frameworks Under	0%	0%	91%	89%
Tensions and Debates	0%	0%	59%	58%
Theoretical Synthesis	0%	0%	74%	77%
Position within the field	0%	0%	59%	59%
Theoretical contribution	0%	0%	84%	86%
Limitations	0%	0%	70%	71%
Future Research	0%	0%	82%	81%
Recap of the argument	0%	0%	100%	100%

Section	AI Score (Before)	AI Score (After)	Burstiness (Before)	Burstiness (After)
Concluding remarks	0%	0%	94%	94%
References	0%	0%	100%	100%

Note: Burstiness is a measure of sentence-length variance. Some section types are inherently low-burstiness — Abstract, Research Objectives, Hypotheses, References, and any bulleted list have fixed structural templates, so a 25–40% score for these sections reflects the format, not AI authorship. The metric is most informative on analytical sections (Discussion, Theoretical Background) where a human writer would naturally vary sentence rhythm.