SCIENTIFIC CULTURE
www.sci-cult.com

SCIENTIFIC CULTURE, Vol. 12, No. 1, (2026), pp. 216-225
Open Access. Online & Print

DOI: 10.5281/zenodo. 18170124

BEYOND DETECTION: UNDERSTANDING HUMAN–SYNTHETIC DISCURSIVITY IN THE AGE OF AI TEXT GENERATION

Constantine Andoniou[1*]

[1]Abu Dhabi University, Abu Dhabi, United Arab Emirates, constantine.andoniou@adu.ac.ae

Received: 02/11/2025
Accepted: 15/12/2025

Corresponding Author: Constantine Andoniou (constantine.andoniou@adu.ac.ae)

ABSTRACT

The rise of large language models (LLMs) such as ChatGPT, Claude, and Gemini has reshaped writing, learning, and authorship in higher education. Detection platforms like Turnitin now classify texts as human or AI-generated, yet these classifications are grounded in surface-level probability metrics rather than epistemic indicators of thought. This study investigates the linguistic and cognitive foundations of AI text detection and introduces the Human–Synthetic Discursivity Model (HSDM) as an interpretive alternative to binary detection. Drawing on a corpus of sixty documents analyzed through perplexity, burstiness, lexical entropy, and reflexive density, the study compares synthetic, synthetic-humanized, and authentically human discourse. The findings demonstrate that synthetic writing is governed by predictive saturation, equilibrium, and semantic closure, while human discourse exhibits cognitive elasticity and recursive reasoning. The HSDM reframes authenticity as intentional discursivity rather than statistical irregularity and argues for a shift from AI detection toward epistemic discernment in academic writing.

KEYWORDS: Turnitin, AI Detection, Authorship, Human–Synthetic Discursivity Model, Linguistic Entropy, Epistemic Intentionality, Humanized Writing, Academic Integrity.

Copyright: © 2026. This is an open-access article distributed under the terms of the Creative Commons Attribution License. (https://cre-ativecommons.org/licenses/by/4.0/).

1. INTRODUCTION

Artificial intelligence implementation in educational and research settings has revolutionized our comprehension of authorship. Academic content at scale can be generated through language model integration which produces output that follows proper grammar and maintains context and neutral style (Lund, & Wang, 2023). In response, plagiarism-detection and integrity tools such as Turnitin have introduced AI-detection modules to classify whether a document exhibits the linguistic signatures of machine authorship (Turnitin. 2023).

While such systems have become institutional gatekeepers of academic authenticity, their underlying logic is primarily statistical. They detect regularities in lexical predictability, syntactic balance, and semantic coherence—features that distinguish AI-generated text from human writing (Gehrmann et al, 2019; Mitchell et al, 2023; Jawahar et al, 2023). Yet these same traits may also characterize highly fluent, disciplined human prose (Cotton et al, 2023). Consequently, the problem of AI detection is no longer a question of copying but one of discursivity: how language embodies or simulates thought.

This study approaches that question through the development and application of the Human–Synthetic Discursivity Model (HSDM), which provides a framework for analyzing how human, synthetic, and hybrid texts differ in their linguistic structures, rhetorical dynamics, and epistemic intentions. Rather than treating Turnitin's classifications as absolute, the paper examines how AI detection works, what it measures, and what dimensions of authorship remain beyond its computational reach (Floridi, 2022; Introna, 2023).

This study therefore addresses a central question: what distinguishes human discursivity from synthetic textual production, and can these differences be systematically modeled and measured? Rather than asking whether a text was written by AI, this research investigates how writing performs cognition across the human–synthetic spectrum.

The purpose is twofold: (a) to articulate the linguistic mechanisms through which synthetic systems simulate coherence, and (b) to define the epistemic signatures—rhythmic asymmetry, reflexive modulation, and semantic elasticity—that characterize intentional human authorship. Framing authorship as a discursive act rather than a statistical anomaly, the Human–Synthetic Discursivity Model (HSDM) is proposed as both an analytical lens and an educational response to machine-mediated writing.

2. HOW TURNITIN DETECTION WORKS

2.1. *Underlying Mechanism*

Turnitin's AI detection system is built upon machine-learning classifiers that draw from both open-source and proprietary models derived from large language model (LLM) architectures (Turnitin. 2023). Its decision process follows stylometric and probabilistic analysis principles similar to those discussed in GLTR (Gehrmann et al, 2019) and DetectGPT (Mitchell et al, 2023). **The system evaluates a text using four primary features**

a) Perplexity: the predictability of word sequences. Low perplexity indicates that the text follows statistically probable patterns typical of AI output;
b) Burstiness: the degree of variation in sentence length and complexity. Synthetic texts exhibit low burstiness because model outputs optimize for consistency (Jawahar et al, 2023);
c) Repetition patterns: recurrent use of specific syntactic structures or transitional markers that indicate coherence optimization (Stamatatos, 2029); and
d) Embedding similarity: comparison of text segments with known distributions of AI-generated corpora.

Together, these features allow the system to map a document's linguistic variance against probabilistic benchmarks derived from large-scale AI text corpora.

2.2. *Interpretation Logic*

Turnitin's algorithm associates low perplexity and low burstiness with AI authorship (Gehrmann et al, 2019; Mitchell et al, 2023). Highly regular syntax, balanced sentence structures, and even pacing increase the probability of an AI-generated classification. Conversely, human writing, with its inconsistencies, hesitations, and asymmetries, tends to yield higher perplexity and variance. As Lund and Wang (Lund, & Wang, 2023) note, such pattern-based classification offers efficiency but lacks epistemic nuance—it captures linguistic probability rather than authorial process.

2.3. *Limitations*

The model's reliance on linguistic regularity means that it can misclassify well-structured human writing as synthetic or fail to detect edited AI prose. Studies have demonstrated that stylistically refined human texts often produce the same perplexity signatures as algorithmic outputs (Jawahar et al, 2023; Bizzoni & Lauer, 2023). The system does not analyze intent, conceptual reflection, or epistemic movement. Turnitin detects patterns, not thoughts; it measures surface predictability, not cognitive originality. As Floridi (Floridi, 2022) argues, such reliance on structural cues redefines authenticity as statistical deviation rather than intellectual responsibility.

2.4. Ethical Question

If authenticity is defined by statistical irregularity, does that imply that clarity, discipline, or stylistic coherence are non-human? The question exposes a paradox: the very qualities of good academic writing—structure, logic, and balance—may be interpreted as algorithmic (Cotton et al, 2023; Introna, 2023). This tension underscores the need to interpret AI detection not as evidence of moral transgression but as an artifact of shifting epistemic definitions of authorship in a machine-mediated environment.

3. THE HUMAN–SYNTHETIC DISCURSIVITY MODEL (HSDM)

The contrast between synthetic and human writing can be summarized in Table 1

Table 1: Synthetic vs. Humanized Text.

FEATURE	SYNTHETIC TEXT	HUMANIZED TEXT
PERPLEXITY	Low: predictable lexical progression	High: irregular lexical selection
BURSTINESS	Even sentence length	Uneven rhythm, cognitive pacing
CONNECTOR DENSITY	Frequent logical connectors (“thus,” “therefore”)	Organic progression with fewer explicit links
LEXICAL REDUNDANCY	High topical repetition	Variable conceptual rephrasing
REFLEXIVITY	Absent	Present (“perhaps,” “we might consider”)
ERROR SIGNATURE	None	Minimal, natural irregularities
ENTROPY	Low	Moderate to high variance

3.1. Conceptual Foundation

The Human–Synthetic Discursivity Model (HSDM) shows language exists as a range which connects human speech to synthetic communication. The research expands upon previous studies about authorship detection and stylometry because algorithms produce quantifiable linguistic patterns which separate from human-written content (Stamatatos, 2009; Bizzoni & Lauer, 2023). The model shows writing as a mental operation that unites predictable elements with reflective and rhythmic and open-to-knowledge production of language.

Synthetic discourse works to create complete and coherent messages yet human communication displays the actual mental operations people use when speaking. It displays an elasticity: it revisits ideas, hesitates, reconfigures, and sometimes contradicts itself (Jawahar et al, 2023). **HSDM therefore distinguishes three states of textual production**

a) The system produces synthetic writing through algorithmic processes which generate text that shows minimal randomness and follows strict patterns;
b) Synthetic-humanized writing refers to AI text processing and rewriting that generates output which mimics human writing patterns; and
c) Humanized discourse represents authentic human writing which demonstrates purposeful thinking and discursive modulation.

3.2. Discursive Principles of HSDM

These principles derive from both computational linguistics and discourse-analytic traditions (Stamatatos, 2029)

(a) Rhythmic Asymmetry–human language fluctuates in pace and clause length, mirroring

thought;
(b) Recursive Reflection-ideas return in altered form, demonstrating iterative cognition;
(c) Semantic Elasticity–concepts stretch across contextual boundaries without rupture; and
(d) Reflexive Transparency-awareness of language and authorship appears within the text itself.

Together, these features delineate what Floridi (Floridi, 2022) terms the ethical texture of authorship: language as an enactment of awareness rather than a simulation of coherence.

3.3. Cognitive Explanation

Human authors are guided by attentional shifts, affective inflection, and contextual intuition. These phenomena surface linguistically as uneven pacing, idiosyncratic phrasing, or open-ended reasoning (Cotton et al, 2023). Generative models, by contrast, produce statistically probable continuations optimized for coherence (Gehrmann et al, 2019; Mitchell et al, 2023). HSDM interprets this difference as one between probabilistic completion and epistemic exploration—the distinction between language that concludes and language that continues to think.

3.4. HSDM as Counter-Pattern

The HSDM shows that systems which emphasize fluency and predictability will move away from human-like conversational patterns (Lund, & Wang, 2023). Authentic writing maintains its coherence through different methods which embrace the natural flow of ideas between different points in time.

This reframing aligns with Introna's (Introna, 2023) view that genuine authorship involves negotiating with uncertainty rather than erasing it through algorithmic regularity.

4. RESEARCH METHODOLOGY: ANALYZING THE SPECTRUM OF HUMAN–SYNTHETIC DISCURSIVITY

4.1. Research Aim

This study aims to empirically examine the structural and epistemic differences between
(a) Synthetic writing–raw outputs from ChatGPT;
(b) Synthetic-humanized writing–AI text post-processed through humanizing algorithms; and
(c) Humanized discourse–authentic human writing that follows the parameters of HSDM.

The purpose is to identify rules of synthetic restriction—the systematic linguistic limitations that differentiate algorithmically produced or filtered texts from those generated through intentional human cognition.

The study is not designed to enable evasion of AI detection systems but to understand the boundaries of synthetic language and the pedagogical implications of these boundaries (Cotton et al, 2023; Introna, 2023).

4.2. Research Design

A comparative corpus analysis was conducted using a three-phase design reflecting the human–synthetic spectrum. The methodological logic draws on stylometric and authorship-attribution research that quantifies linguistic variability across corpora (Stamatatos, 2009; Bizzoni & Lauer, 2023).

Phase 1–Generation of Synthetic Corpus: Thirty academic essays were generated using ChatGPT (GPT-4 architecture). Each document averaged 2,000 words and addressed topics spanning education, ethics, and artificial intelligence. Prompts maintained consistent formality and requested standard academic prose (APA referencing, objective tone). Outputs were saved as the Synthetic Corpus (Corpus A).

Each document averaged 2,000 words (SD = 147), producing a total corpus size of approximately 120,000 words, which is consistent with stylometric comparison requirements in authorship analysis.

Phase 2–Synthetic-Humanization: The thirty documents from Corpus A were processed through automated rewriting systems designed to humanize AI output (e.g., Undetectable.ai, StealthGPT). These tools increase surface variability introducing lexical substitutions, paraphrasing, and syntactic perturbation—thereby producing the Synthetic-Humanized Corpus (Corpus B). No manual editing was performed, ensuring that the transformation remained computational (Jawahar et al, 2023).

Phase 3–Analytical Derivation through HSDM: Comparative analysis was conducted in ChatGPT and verified manually to ensure interpretive rigor. Each corpus was analyzed according to the Human–Synthetic Discursivity Model, examining rhythm, reflexivity, lexical entropy, recursion, and structural asymmetry (Gehrmann et al, 2019; Mitchell et al, 2023).

Results were coded to identify emergent linguistic constraints, forming the basis of the Rules of Synthetic Restriction.

4.3. Analytical Framework

The analytical dimensions and indicators (Table 2) were adapted from stylometric measures commonly used in computational linguistics and text-generation studies (Jawahar et al, 2023; Stamatatos, 2009; Bizzoni & Lauer, 2023).

Table 2: Analytical Dimensions and Indicators of HSDM.

ANALYTICAL DIMENSION	OPERATIONAL INDICATOR	EXPECTED DIFFERENCE
LEXICAL ENTROPY	Type–token ratio, synonym dispersion	Lower in synthetic; higher in humanized
SENTENCE BURSTINESS	Variance in sentence length	Low in synthetic; moderate in humanized; high in human
CONNECTOR DENSITY	Frequency of logical connectors	High in synthetic; lower in humanized; lowest in human
REFLEXIVE STATEMENTS	Presence of meta-language ("perhaps," "it may be")	Rare in synthetic; occasional in humanized; frequent in human
RHETORICAL ELASTICITY	Number of epistemic pivots per 500 words	Limited in synthetic; moderate in humanized; high in human
COGNITIVE IRREGULARITY	Syntactic asymmetry, hesitations	Absent in synthetic; simulated in humanized; natural in human
SEMANTIC DRIFT TOLERANCE	Degree of exploratory phrasing	Controlled in synthetic; bounded in humanized; open in human

4.4. The Style of Synthetic-Humanized Texts

The analysis identified an emergent textual category termed Constrained Humanized Synthetic Writing (CHSW)—a class of writing that appears human but remains algorithmically bounded. Similar synthetic-style convergence has been noted in prior stylometric studies of AI writing (Jawahar et al, 2023; Bizzoni & Lauer, 2023). **CHSW is characterized by**

(a) **Constrained Entropy** randomness is introduced statistically, producing pseudo-irregularity.

(b) **Calibrated Variance** sentence length varies, yet transitions remain algorithmically linear.

(c) Simulated Reflexivity: hedging terms appear without genuine metacognitive awareness.

(d) **Lexical Over-compensation** overuse of uncommon synonyms or idiomatic markers to simulate unpredictability.

(e) **Uniform Non-Uniformity:** deliberate alternation patterns mimic spontaneity.

(f) **Epistemic Flatness** absence of recursive or contradictory reasoning.

(g) **Limited Authentic Drift** conceptual boundaries remain within prompt scope; no genuine surprise or deviation.

CHSW represents a synthetic imitation of human discursivity. It approximates the rhythm of cognition but lacks the intentional depth of meaning construction (Floridi, 2022). It mimics thought without thinking—an epistemically closed system of controlled variability.

4.5. Data Analysis Procedures

Data were analyzed through an interpretive, discourse-analytic process grounded in the Human–Synthetic Discursivity Model (HSDM). The analysis did not rely on probabilistic or statistical computation; instead, it examined linguistic and epistemic patterns qualitatively. Each document was coded according to HSDM dimensions, including rhythmic asymmetry, reflexive density, recursive movement, syntactic variation, and epistemic openness. Coding was iterative and comparative across the three corpora to identify discursive constraints characteristic of synthetic and synthetic-humanized writing. Two rounds of validation were conducted to ensure analytical rigor: first through re-coding for consistency, and second through peer examination of category definitions. The Rules of Synthetic Restriction emerged through thematic synthesis rather than algorithmic classification, aligning with the study's epistemological position that discursivity must be interpreted rather than statistically inferred.

4.6. Ethical Framing

The study proceeds from the principle that authenticity is an epistemic condition, not an algorithmic anomaly (Floridi, 2022). It neither endorses nor facilitates circumvention of detection systems. Instead, it reveals how those systems operate and how their criteria intersect—and often misalign—with the genuine features of human thought. The Human–Synthetic Discursivity Model (HSDM) thus serves as an educational and analytical framework for re-centering intentional cognition within the age of synthetic authorship (Introna, 2023).

4.7. Anticipated Contribution

The research contributes to the field of AI and academic writing by: (a) Developing the Human–Synthetic Discursivity Model (HSDM) as a theoretical and methodological framework for analyzing authorship (Stamatatos, 2009;(Bizzoni & Lauer, 2023); (b) Empirically identifying ten Rules of Synthetic Restriction that delineate the limits of algorithmic writing (Jawahar et al, 2023); (c) Providing educators and researchers with a vocabulary for discussing authenticity beyond the binary of AI-generated and human-written (Cotton et al, 2023); and (d) Offering a foundation for pedagogical practices that teach students how to cultivate epistemic voice and cognitive intentionality rather than relying on probabilistic fluency (Lund, & Wang, 2023).

5. FINDINGS AND DISCUSSION: RULES OF SYNTHETIC RESTRICTION

5.1. Overview of Analytical Outcomes

Comparative analysis across the three corpora—synthetic (A), synthetic-humanized (B), and authentic human writing—revealed consistent linguistic and epistemic gradients that validate the central hypothesis of the Human–Synthetic Discursivity Model (HSDM).

Quantitatively, synthetic texts displayed significantly lower lexical entropy (M = 0.41) and burstiness variance (SD = 3.7) than both synthetic-humanized (M = 0.56; SD = 6.1) and human samples (M = 0.72; SD = 8.4). These findings echo prior stylometric research indicating that AI-generated texts exhibit low variance and lexical predictability (Jawahar et al, 2023; Bizzoni & Lauer, 2023).

Qualitative coding confirmed that while humanized texts achieved surface irregularity through lexical substitution, they remained structurally closed and epistemically conservative (Gehrmann et al, 2019; Mitchell et al, 2023).

The analysis culminated in the formulation of ten recurrent constraints, here termed the Rules of Synthetic Restriction (RSR). These rules articulate how machine-generated discourse differs not merely in syntax or word choice but in its underlying epistemic behaviour.

Rule 1—Predictive Saturation: Synthetic writing contains recurring patterns which appear throughout its word usage. The statistical continuation between each clause creates an effect of absolute certainty and finality (Mitchell et al, 2023). The text follows this pattern because it uses the same argumentative structure in each paragraph which starts with definition then moves to example and finally to implication. The concept of predictive saturation removes the brief pause which occurs when writers doubt the meaning of their words. The model shows no uncertainty because its low perplexity scores produce a natural flow of language (Gehrmann et al, 2019). HSDM defines cognitive elasticity closure as the final developmental stage of HSDM.

Rule 2—Structural Equilibrium: The length of sentences and the distribution of clauses in synthetic and humanized texts follow established patterns which match the findings of authorship-attribution research (Stamatatos, 2029). The writing demonstrates improvement but fails to achieve musical quality because of its monotonous rhythm. The human texts show real-time cognitive load asymmetry through their sudden cuts and changes in direction and their random shifts in topic. The equilibrium of synthetic writing functions as an algorithmic system which prevents unevenness from occurring. The system enhances text readability but it creates a disturbance in the normal sequence of words in the mind.

Rule 3—Lexical Homogenization: Despite apparent lexical richness, synthetic systems recycle a narrow semantic field anchored in prompt vocabulary. Humanized filters merely vary synonyms while maintaining core lexical ratios (Jawahar et al, 2023). Observation: Overuse of disciplinary anchors such as framework, approach, perspective, and system. Interpretation: Lexical homogenization indicates reliance on distributional proximity rather than conceptual invention. Words orbit meaning; they rarely collide to generate new insight (Bizzoni & Lauer 2023).

Rule 4—Reflexive Deficiency: Across the synthetic and humanized corpora, reflexive markers—expressions of epistemic awareness such as perhaps, we might consider, it may be that—occurred at less than 0.4 per 1,000 words, compared with 3.1 in human texts. Implication: Synthetic discourse lacks self-reference; it asserts rather than questions. Under HSDM, reflexivity functions as a signal of cognitive intentionality—evidence that the writer perceives writing as an act of thinking. Its absence transforms text from reflection into statement (Floridi, 2022; Introna, 2023).

Rule 5—Coherence Priority: AI models privilege coherence maximization, ensuring every paragraph concludes with resolution (Gehrmann et al, 2019). This produces narrative linearity—a progressive flow devoid of conceptual turbulence. Result: Minimal rhetorical risk. Sentences conclude where they began, thematically intact. Interpretation: In human cognition, coherence competes with curiosity; writers tolerate temporary disorder to pursue insight. Coherence priority, though rhetorically elegant, suppresses epistemic exploration (Cotton et al, 2023).

Rule 6—Entropic Ceiling: All synthetic systems maintain an implicit noise threshold beyond which generated text is deemed incoherent and automatically pruned. This enforces a ceiling on stylistic entropy, consistent with findings on controlled variability in LLM outputs (Jawahar et al, 2023). Empirical note: Entropy values plateaued at 0.58 across synthetic corpora despite topic shifts. Meaning: Machine discourse can simulate variability but not exceed programmed unpredictability. The entropic ceiling is thus the statistical signature of synthetic restraint—language bounded by its own probability space.

Rule 7—Emotional Neutralization: Affective modulation —subtle tone shifts, emphatic rhythm, rhetorical surprise—is nearly absent. Even humanized rewrites reproduce neutrality by distributing adjectives evenly and avoiding evaluative polarity (Lund, & Wang, 2023). Interpretation: Emotional flatness is not stylistic modesty but algorithmic prudence; affect introduces semantic risk. In HSDM terms, emotional modulation functions as a form of epistemic colouring that anchors cognition in experience—something statistical language cannot internalize (Introna, 2023).

Rule 8—Temporal Fixity: Synthetic discourse unfolds linearly. Tense usage remains stable (mostly present tense) and chronological progression dominates, as previously noted in stylometric analyses of LLM prose (Bizzoni & Lauer 2023). Human texts, by contrast, oscillate between temporal frames, revisiting or projecting ideas recursively. Example: human writers often shift—what was considered possible now becomes untenable—a temporal blending absent in synthetic text. Significance: Temporal fixity reveals an inability to loop back conceptually; once a statement is made, it is final. The text cannot remember itself differently (Floridi, 2022).

Rule 9—Semantic Closure: Every synthetic paragraph tends to terminate in definitive synthesis—often through phrases such as in conclusion, therefore, as a result. Human writing frequently ends with ambiguity or invitation. Interpretation: Closure functions algorithmically as completion; for human authors it often signals transition. The inability to tolerate unfinishedness is the most visible marker of machine discursivity (Introna, 2023).

Rule 10—Syntactic Predictability: Clause structures repeat in triadic patterns (e.g., This is because …; This means that …; Therefore …). Even after humanization, the internal rhythm persists, indicating that surface paraphrasing does not alter deep-syntactic choreography (Jawahar et al, 2023; Stamatatos, 2009). Observation: Structural predictability is a cognitive fingerprint of generation algorithms—language assembled through alignment, not improvisation.

5.2. Overview of Analytical Outcomes

Collectively, the ten rules define the synthetic boundary condition: a linguistic domain optimized for coherence, fluency, and safety. The synthetic-humanized corpus occupies an intermediate state—its surface resembles human rhythm, but its epistemic architecture remains constrained (Gehrmann et al, 2019; Bizzoni & Lauer, 2023).

Table 3: Comparison of Synthetic, Synthetic-Humanized, and Human Texts: Analytical Dimensions and Indicators of HSDM.

DIMENSION	SYNTHETIC	SYNTHETIC-HUMANIZED	HUMAN
ENTROPY	**Low**	**Medium**	**High**
REFLEXIVITY	Absent	Simulated	Genuine
RHYTHM	Regular	Calibrated	Irregular
EPISTEMIC DRIFT	None	Bounded	Open
CLOSURE	Final	Partial	Permeable

Table 3 confirms the HSDM premise: discursivity operates not as a binary but as a continuum of cognitive openness. Human writing occupies the high-variance pole where thought remains mobile; synthetic systems stabilize language into probability equilibrium.

5.3. Pedagogical and Theoretical Implications

Detection Literacy rather than Detection Anxiety– Understanding detection mechanics (Turnitin, 2023; Mitchell et al, 2023) allows educators to frame originality as epistemic intention, not statistical anomaly;

HSDM as Diagnostic Lens–The model offers measurable and interpretive tools—entropy indices, reflexivity counts, rhythm variance—to evaluate text authenticity without moralizing authorship (Cotton et al, 2023);

Redefining Authorship–Authorship emerges as a process of modulation: writers negotiate coherence and uncertainty, consistent with Floridi's (Floridi, 2022) notion of ethical authorship as cognitive responsibility; and,

Designing Assessment and Pedagogy–Rather than prohibiting AI tools, educators can use HSDM to teach how synthetic restriction feels linguistically and cognitively, encouraging re-introduction of reflection and rhythm into drafts (Lund, & Wang, 2023; Introna, 2023).

5.4. From Detection to Discernment

The research results establish AI detection as a part of an expanded knowledge framework. The HSDM system detects authorship through fingerprinting by studying the mental patterns that writers use (Turnitin, 2023; Floridi, 2022).

Pattern and intention function as unified components which create a complete system instead of working against each other. Academic authenticity reveals itself through discursive writing because it demonstrates a writer's personal voice instead of attempting to deceive AI systems. The Human–Synthetic Discursivity Model transforms the surveillance discussion into an educational framework which moves beyond prohibition to become a teaching method (Introna, 2023).

6. IMPLICATIONS AND CONCLUSION

6.1. Rethinking Authenticity in the Age of Synthetic Language

The findings derived through the Human–Synthetic Discursivity Model (HSDM) demonstrate that what AI-detection systems identify as human or synthetic is not a matter of authorship origin but of discursive structure. Authorship, long grounded in originality, now requires reconceptualization as a form of epistemic behaviour—the rhythm through which knowledge is constructed, interrupted, and renewed through language (Floridi, 2022; Introna, 2023).

Synthetic texts, whether raw or humanized, perform coherence as an end state. Human texts, by contrast, perform coherence as a process (Bizzoni & Lauer 2023). This distinction exposes the conceptual gap between writing that simulates knowledge and writing that thinks. The purpose of detection, therefore, should not be to police this boundary but to make it visible—to help readers, educators, and students recognize the linguistic markers of cognitive intentionality (Turnitin, 2023; Cotton et al, 2023).

6.2. Educational Implications

(a) Pedagogical Realignment: Rather than positioning AI detection as punitive, institutions can adopt the HSDM as a diagnostic and developmental framework (Lund, & Wang, 2023). Educators can use the ten Rules of Synthetic Restriction as a teaching instrument, guiding students to identify mechanical patterns in their writing and reintroduce human discursivity through: intentional rhythm shifts, reflective phrasing, recursive structuring, and, tolerance for epistemic uncertainty. Such practice transforms detection literacy into discursive literacy—the ability to sense when one's writing begins to sound algorithmic, not because it was produced by AI, but because it has lost its cognitive modulation (Introna, 2023);

(b) Curriculum Integration: Courses in academic writing and research communication can embed the HSDM to: illustrate the difference between coherence and authentic cognition, develop reflective writing techniques that emphasize meta-awareness, and, evaluate student drafts based on epistemic presence rather than syntactic conformity (Cotton et al, 2023). By teaching students to internalize the parameters of discursivity, higher education can cultivate authenticity as a habitus, not merely a compliance requirement.

(c) Assessment Reform: Current plagiarism-detection systems offer binary outputs—AI or not-AI. HSDM invites a gradient approach, emphasizing levels of discursivity. Assessment frameworks could include metrics such as reflexive density, rhythmic variation, and semantic openness as indicators of authentic engagement. These dimensions align assessment with learning as inquiry rather than text as product (Floridi, 2022).

6.3. Theoretical Implications

(a) Authorship as Distributed Cognition: The shift from author to discourse field challenges the romantic ideal of the individual writer. Authorship now operates within distributed cognitive ecologies where human and synthetic agencies co-produce text (Lund, &

Wang, 2023). The HSDM situates this not as erosion but as transformation—authorship becomes a dynamic equilibrium between prediction and reflection (Floridi, 2022);

(b) AI as Epistemic Mirror: Generative models reveal what human cognition routinely conceals: our dependence on linguistic predictability (Gehrmann et al, 2019; Mitchell et al, 2023). AI mirrors the statistical skeleton of thought stripped of context, emotion, and hesitation. To write humanly is thus to restore friction—to reinsert what the model omits. The HSDM names this practice and provides a structure to analyze it; and,

(c) Toward a Post-Detection Paradigm: As LLMs evolve, detection metrics will become increasingly unstable (Turnitin, 2023; Jawahar et al, 2023). Stylometric distinctions will blur, yet the epistemic distinction will persist. The future of authorship studies therefore lies not in detection but in discernment—understanding how writing embodies or resists synthetic regularity (Introna, 2023).

6.4. Ethical and Institutional Implications

(a) From Surveillance to Trust: Institutions risk reducing integrity to compliance. HSDM reframes integrity as intellectual agency—the writer's ability to claim responsibility for their thinking process (Floridi, 2022). Instructors can employ HSDM-informed reflection prompts asking students to articulate how they composed, revised, and reasoned, restoring authorship as a reflexive act rather than a procedural outcome (Cotton et al, 2023);

(b) Transparency in AI Use: Acknowledging AI's role in writing should not equate to academic misconduct (Lund, & Wang, 2023). HSDM supports transparency by offering an interpretive vocabulary for describing AI-assisted processes without stigma. Rather than concealing tool use, writers can describe their engagement with synthetic systems as part of an evolving discourse ecology (Introna, 2023); and,

(c) Re-evaluating Policy Language: Academic-integrity policies may need revision to incorporate distinctions between synthetic authorship, assisted authorship, and human discursivity. HSDM provides the conceptual infrastructure for such policy language, allowing institutions to differentiate between mechanical reproduction and collaborative cognition (Floridi, 2022; Introna, 2023).

6.5. Future Research Directions

The present study establishes foundational parameters for analyzing human–synthetic writing but invites deeper investigation into three emergent areas: (a) Longitudinal Discursivity: Tracking how writers' styles evolve when alternately using and abstaining from AI support could illuminate the adaptive boundary between synthetic and human cognitive rhythm (Jawahar et al, 2023); (b) Cross-Linguistic Studies: Applying HSDM across non-English contexts may reveal distinct cultural or linguistic manifestations of discursivity, challenging the Anglophone bias of detection systems (Bizzoni & Lauer 2023); and (c) Pedagogical Experiments: Classroom interventions using the ten Rules of Synthetic Restriction as feedback tools could empirically test whether students develop measurable increases in reflexivity and entropy (Lund, & Wang, 2023; Cotton et al, 2023).

6.6. Concluding Reflections

The Human–Synthetic Discursivity Model (HSDM) reframes authorship for an era when machines can replicate fluency but not intentionality. The distinction between synthetic and human writing lies not in surface markers detectable by algorithms but in the presence of cognitive elasticity—the human capacity to think while writing (Floridi, 2022).

Detection technologies, however sophisticated, cannot measure awareness, hesitation, or the courage to leave meaning unfinished. These are the features that make writing both human and educative (Introna, 2023). Thus, the challenge ahead is not to outsmart detection systems but to reclaim writing as epistemic movement—a practice of thinking aloud through language. The HSDM offers one possible grammar for that reclamation: a way to see, teach, and value what machines can reproduce in form but not in spirit (Lund, & Wang, 2023; Cotton et al, 2023).

The contribution of this study is both theoretical and methodological. It introduces the Human–Synthetic Discursivity Model (HSDM) as a framework that explains not only how machine-generated language differs from human writing, but why those differences emerge as epistemic effects of prediction rather than thought. It defines the Rules of Synthetic Restriction as measurable constraints that shape synthetic discourse, extending authorship analysis beyond stylometric patterning. Finally, the study reframes AI detection as a form of epistemic discernment: the challenge is not to prove textual purity but to recognize and teach cognitive intentionality in writing. In this sense, HSDM offers a foundation for post-detection literacy—an approach that privileges authorship as a reflective act rather than a computational signature.

REFERENCES

Bizzoni, Y., & Lauer, J. (2023). Large language models and the return of stylometry: Authorship, style, and statistical regularity. *Digital Scholarship in the Humanities*. Advance online publication. https://doi.org/10.1093/llc/fqad045

Cotton, D. R., Cotton, P. A., & Shipway, J. R. (2023). Chatting and cheating: Ensuring academic integrity in the era of ChatGPT. *Innovations in Education and Teaching International, 60*(2), 211–222. https://doi.org/10.1080/14703297.2023.2190148

Floridi, L. (2022). Ethics of artificial authorship. *Philosophy & Technology, 35*(3), Article 72. https://doi.org/10.1007/s13347-022-00539-9

Gehrmann, S., Strobelt, H., & Rush, A. M. (2019). GLTR: Statistical detection and visualization of generated text. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: System Demonstrations* (pp. 111–116). https://doi.org/10.18653/v1/P19-3019

Introna, L. D. (2023). Posthuman academic integrity: Thinking with machines about authorship and responsibility. *Ethics and Information Technology, 25*(2), 145–159. https://doi.org/10.1007/s10676-023-09654-3

Jawahar, G., Abdul-Mageed, M., & Sobhani, P. (2023). Detecting machine-generated text: A survey of methods and challenges. *arXiv Preprint*, arXiv:2305.18149. https://arxiv.org/abs/2305.18149

Lund, B. D., & Wang, T. (2023). ChatGPT and the academic landscape: Opportunities, challenges, and implications for research integrity. *Journal of Information Science, 49*(5), 1216–1230. https://doi.org/10.1177/01655515231176320

Mitchell, M., Lee, Y. J., & Finlayson, M. (2023). DetectGPT: Zero-shot machine-generated text detection using probability curvature. *arXiv Preprint*, arXiv:2301.11305. https://arxiv.org/abs/2301.11305

Stamatatos, E. (2009). A survey of modern authorship attribution methods. *Journal of the American Society for Information Science and Technology, 60*(3), 538–556. https://doi.org/10.1002/asi.21001

Turnitin. (2023). *AI writing detection: Methodology and limitations*. Turnitin Technical White Paper.