International Journal of Arts and Social Science www.ijassjournal.com
ISSN: 2581-7922,
Volume 8 Issue 9, September 2025

AI as Interpretive Aid in Qur'anic Stylistics: Ethical Foundations for Digital Hermeneutics

Hishomudin Ahmad[1], Norfarhana Ahmad Ghafar[2]
[1](Fakulti Pengajian Bahasa Utama, Universiti Sains Islam Malaysia, Malaysia)
[2](Fakulti Pengajian Bahasa Utama, Universiti Sains Islam Malaysia, Malaysia)

ABSTRACT: *This article investigates the ethical and epistemological implications of employing artificial intelligence (AI) as an interpretive assistant in Qur'anic stylistics. Positioned at the intersection of Islamic theology and digital humanities, the study critiques the genre-reductive tendencies of computational models while advancing a faith-anchored alternative. Drawing from classical Arabic rhetoric (ʿilm al-balāghah) and frameworks such as the Religious Social Shaping of Technology (RSST), it proposes a human-in-the-loop model wherein AI functions as a musāʿid (supportive tool), not a muḥakkim (arbiter) of meaning. The model affirms the ontological uniqueness of the Qur'an (kalām Allāh) and introduces ethical safeguards rooted in maqāṣid al-sharīʿah, taqwā, and ʿilm. It outlines how tools like Qur'anBERT and CAMeL can aid in morphological parsing and rhetorical mapping without displacing exegetical authority. Ultimately, the study offers a theologically grounded framework for digital hermeneutics that balances computational affordances with reverent restraint—preserving the sanctity of divine discourse while expanding the analytical repertoire of Qur'anic scholarship.*

KEYWORDS - *Artificial intelligence (AI), Digital hermeneutics, Islamic ethics, Qur'anic stylistics, Human-in-the-loop model*

I. INTRODUCTION: BETWEEN REVELATION AND COMPUTATION

The integration of artificial intelligence (AI) into the analysis of religious texts presents both significant promise and pressing ethical concerns. With the rise of large language models (LLMs) and transformer-based natural language processing (NLP) systems, researchers now have access to tools capable of detecting linguistic patterns, mapping rhetorical features, and modelling stylistic trends. Yet, when applied to sacred scripture, particularly in traditions that regard such texts as divinely revealed, these technologies raise not only new questions but also revive longstanding epistemological and theological anxieties. In Islam, the Qur'an is not merely a historical or literary document but kalām Allāh, the uncreated and eternal speech of God, whose interpretation is governed by principles of *khashyah* (reverent awe), *adab* (disciplinary conduct), and *ʿilm* (sacred knowledge).

Emerging research underscores both the potential and the peril of AI in sacred hermeneutics. Reed (2025) reveals how fine-tuning GPT-3 on the sayings of Jesus alters not only the language but also the theological inflection of generated content, raising questions about the algorithmic shaping of belief. Similar concerns arise when AI is applied to authenticated *ḥadīth* corpora such as *Ṣaḥīḥ al-Bukhārī*, where the risks of hallucination, referential distortion, and doctrinal misrepresentation are amplified. To address these challenges, recent discussions in AI and religious studies emphasize the need for computational approaches that are constrained by authenticated sources and guided by theological oversight, in order to maintain interpretive fidelity and uphold the integrity of sacred discourse.

In the Islamic scholarly tradition, these concerns are magnified by the spiritual weight of interpretation. Fitryansyah and Fauziah (2024), in their study of AI-assisted analysis of Nusantara Islamic manuscripts, caution that while digital tools can increase access, they remain ill-equipped to handle doctrinal nuance without an embedded ethical framework. This concern is echoed by empirical findings: Ali, Habtiter, Ahmed, and Mijbel (2025) report that over 70% of Muslim respondents expressed concern that AI could misrepresent Islamic teachings, a sentiment reflecting widespread apprehension about automation in matters of faith.

These limitations are not merely technical; they are epistemological. Al-Janabi (2024), in his critical study of ChatGPT-generated *tafsīr*, observes that even when prompted with theological precision, the resulting outputs often lack interpretive depth and may contravene established hermeneutical norms. This shortfall arises from the architecture of large language models themselves, which operate on probabilistic associations rather than intentionality grounded in spiritual epistemology. As a result, such systems remain structurally incapable of apprehending the Qur'an's layered meanings. These include legal rulings, mystical insights, and ethical guidance, all of which are anchored in its metaphysical origin and sustained by a centuries-long exegetical tradition.

In light of these tensions, this article asks: Can AI assist in Qur'anic stylistics without compromising the sanctity of *waḥy* (revelation)? The argument advanced here is that AI may function as a *musāʿid* (assistant), but never as a *muḥakkim* (arbiter) of meaning. Grounded in Islamic ethical principles such as *khashyah, niyyah* (ethical intention), and *ʿilm*, this study proposes a human-in-the-loop model of Qur'anic hermeneutics in which computational tools support rather than supplant interpretive authority. Through this model, it charts a middle path—embracing technological affordances while preserving the ontological sanctity and exegetical integrity of *kalām Allāh*.

II. QUR'ANIC STYLISTICS AND THE ONTOLOGY OF DIVINE SPEECH

The Qur'an's stylistic uniqueness is not incidental but essential to its identity as *kalām Allāh*, the uncreated speech of God. This section explores the intersection of rhetorical form and theological function in the Qur'anic text, highlighting how classical Arabic rhetoric and modern stylistics converge to reveal the discourse's divine intentionality. It also examines the ontological challenges posed by algorithmic parsing of sacred speech, cautioning against the flattening of divine discourse into computational form.

2.1 The Qur'an as Theological Revelation and Stylistic Text

The Qur'an occupies a unique position as both a sacred revelation and a linguistically extraordinary text. Classical Islamic scholars have long examined its stylistic features through the lens of Arabic rhetorical science (*ʿilm al-balāghah*), identifying elements such as *istiʿārah* (metaphor), *ḥadhf* (ellipsis), *tashbīh* (simile), *muqābalah* (antithesis), and *iʿjāz* (inimitability) as integral to its divine eloquence. These features are not merely decorative; they are epistemic instruments that embed theological precision, affective resonance, and spiritual authority into the very cadence of revelation.

What makes a text sacred—not only semantically, but stylistically? Contemporary analyses offer a partial answer by illuminating the Qur'an's distinct rhetorical profile. Scholars have highlighted its rhythmic cadences, syntactic compression, and semantic condensation as evidence of stylistic density that resists reduction to conventional literary paradigms. This density, often expressed through patterns of ellipsis (*ḥadhf*), grammatical shift (*iltifāt*), and metaphorical layering, is not merely aesthetic—it functions as a marker of revelation. Even in the study of secular Arabic texts, researchers such as Salakhova and Nabiullina (2022) emphasize the epistemic role of metaphor in shaping meaning, underscoring the need to interpret Qur'anic stylistics through a theologically anchored lens.

Recent developments in Arabic Natural Language Processing (NLP) have opened computational pathways for analysing such features. Tools such as CAMeL (Obeid et al., 2020) offer morphological and syntactic segmentation that support stylistic pattern detection in Arabic texts. However, most NLP models are trained on

secular or modern literary corpora and are rarely calibrated for liturgical registers like the Qur'an. Deep learning systems still struggle to capture rhetorical coherence and semantic intentionality in sacred texts, especially in passages marked by metaphor, ellipsis, or concise profundity. This gap highlights the need for Qur'an-specific datasets and ethically attuned frameworks that respect the sacredness of the text while enabling meaningful stylistic analysis.

This limitation is evident in empirical attempts to computationally model Qur'anic language. Touati-Hamad, Laouar, and Bendib (2020) compare local and distributed word embeddings for Qur'anic verses, concluding that while these methods capture some lexical structure, they fall short in apprehending intertextual allusions and verse-level semantic layering, features central to Qur'anic discourse. These findings suggest that computational models may illuminate certain stylistic traits but remain ill-equipped to fully engage the Qur'an's performative and exegetical richness.

In contrast to secular Arabic texts—where stylistic features often serve expressive or narrative purposes—Qur'anic stylistics operates within a theologically saturated register. While metaphor, ellipsis, and rhythm may appear in both domains, their epistemic stakes differ fundamentally. Secular texts invite aesthetic interpretation; the Qur'an demands exegetical accountability. This contrast underscores the ontological gap between human literature and divine revelation, a distinction that computational models must respect and preserve.

The following comparative schema outlines the fundamental differences between secular Arabic literary texts and the Qur'an, emphasizing why computational approaches calibrated for the former often fall short in the latter. This contrast reinforces the theological and epistemological uniqueness of *kalām Allāh* and clarifies the limits of secular stylistic frameworks in Qur'anic analysis.

Table 1: Key Contrasts Between Secular Literary Stylistics and Qur'anic Stylistics

Feature	Secular Literary Texts	Qur'anic Stylistics
Ontology of Text	Human-authored; open to reinterpretation	Divine speech (*kalām Allāh*); fixed and uncreated
Stylistic Purpose	Aesthetic, expressive, often narrative-driven	Theological precision, spiritual resonance, legal and ethical guidance
Use of Metaphor	Primarily symbolic or thematic	Epistemic function: encodes layered divine meanings
Narrative Structure	Linear, character-focused	Non-linear, recursive, performative; marked by *iltifāt* and semantic density
Interpretive Flexibility	Encourages multiple readings; often ambiguous	Anchored in *ʿilm al-tafsīr*, *asbāb al-nuzūl*, and scholarly consensus
Computational Modelling	Typically, successful in stylometry and thematic analysis	Challenging due to *tawqīfī* structure, rhythmic cadence, and theological intent

As shown in Table 1, Qur'anic stylistics is governed not only by linguistic structure but by its ontological status as divine revelation. These distinctions are critical for any computational model attempting to engage the Qur'an: without theological anchoring, stylistic analysis risks misrepresenting the very register it seeks to illuminate. Thus, any NLP-based methodology must be recalibrated to account for the sacred, performative, and exegetically bounded nature of Qur'anic discourse.

2.2 *Kalām Allāh* and the Limits of Algorithmic Parsing

The Qur'an's ontological status as *kalam Allāh*, the uncreated speech of God, places it beyond the bounds of ordinary textual analysis. Unlike literary corpora, the Qur'an is governed by divine intentionality and spiritual teleology, both of which resist mechanistic parsing. Linguistic elements such as root-based morphology, diglossic layering, and syntactic ellipsis already present computational challenges in Arabic NLP (Farghaly & Shaalan,

2009). These challenges become even more pronounced when the text in question possesses theological authority and occupies a central role in devotional life.

One critical limitation lies in the Qur'an's *tawqīfī* nature: its fixed word order, rhythmic cadence, and semantic pause structure. Far from arbitrary, these features thrum with liturgical pulse, recitational gravity, and exegetical weight. AI systems, which tend to normalize or disambiguate language for clarity, often overlook this dimension. Belhassen et al. (2025) caution that even advanced translation models struggle with tonal resonance, metaphysical implication, and symbolic layering, qualities that are foundational to sacred texts like the Qur'an.

Moreover, Arabic LLMs frequently underperform when applied to fully diacritized religious texts. Alshammari and Elleithy (2024) show that subtle vowel variations can alter not just grammatical form but theological meaning. For instance, the distinction between *raḥīm* (merciful) and *rahim* (womb) is semantically critical in exegetical contexts. Such nuance demands interpretive care that current AI systems are ill-equipped to replicate.

The issue extends beyond morphology to rhetorical structure. Qur'anic *iltifāt* (shifts in grammatical person), for example, functions not as stylistic inconsistency but as a spiritually motivated discourse strategy. Al-Janabi (2024) observes that LLM-generated *tafsīr* often omits or misrepresents such features, leading to flattened readings that misalign with the Qur'an's communicative intent.

These epistemic and linguistic limitations underscore a foundational principle. AI may assist as a *musāʿid* (supporting tool), but it cannot replace *tafaqquh*—the deep, reflective understanding cultivated through traditional Qur'anic scholarship. The Qur'an is not merely a dataset to be analysed but a *ḥujjah* (divine proof) whose interpretation requires intellectual discipline, historical awareness, and spiritual humility. Therefore, AI applications in Qur'anic stylistics must be governed not only by technical parameters but by an epistemological framework that preserves the sanctity and authority of divine speech.

III. Artificial Intelligence and Sacred Text Analysis

The application of artificial intelligence to sacred texts introduces complex tensions between computational potential and theological boundaries. While natural language processing (NLP) tools and large language models (LLMs) offer powerful methods for stylistic analysis, they also raise critical concerns about epistemic adequacy, genre distortion, and interpretive legitimacy. This section explores both the technological capabilities of current AI tools and the ethical risks involved in deploying them for Qur'anic analysis. It does so through two lenses: first, by surveying the functionality and limits of domain-specific NLP platforms; and second, by highlighting the deeper epistemological risks of applying machine logic to divine discourse.

3.1 NLP and LLMs in Religious Textual Study

Recent advancements in Arabic Natural Language Processing (NLP) have enabled increasingly sophisticated computational engagement with religious texts. Domain-specific tools such as Qur'anBERT (Malti & Belkredim, 2022), CAMeL Tools (Obeid et al., 2020), Tasaheel (Himdi & Assiri, 2023), and AraSAS (El-Haj et al., 2022) offer functions including morphological segmentation, part-of-speech tagging, and syntactic parsing—capabilities that enhance the granularity of linguistic analysis within classical Arabic corpora.

Alshammari and Elleithy (2024) demonstrate that transformer-based architectures like AraELECTRA and XLM-R can detect AI-generated Arabic content even amid diacritic ambiguity. This is a significant advantage, given that diacritical marks in the Qur'an bear theological weight. However, this technical precision does not necessarily translate into hermeneutical adequacy. Metaphor, ellipsis, and intertextual resonance—all central to Qur'anic discourse—often escape algorithmic capture.

Studies across comparative scripture analysis support this caution. Salakhova and Nabiullina (2022) note that figurative language in Arabic religious texts functions epistemically rather than ornamentally, encoding

layered theological meanings that resist mechanistic reduction. Similarly, Belhassen et al. (2025) argue that while AI tools aid in lexical disambiguation and preliminary translation, they consistently fall short in handling tonal nuance, symbolic depth, and culturally embedded reference—all of which are factors central to sacred textual interpretation.

Even highly specialized platforms such as Qur'anBERT (Malti & Belkredim, 2022) and CAMeL Tools (Obeid et al., 2020), though effective in basic parsing tasks, encounter critical limitations in recognizing *tawqīfī* elements (fixed divine expressions), performative speech acts, and spiritually embedded stylistic devices like *ḥadhf* (ellipsis) or *iltifāt* (grammatical shifts). Bahgat, Rafea, and El-Beltagy (2018) attempt to address these gaps by integrating curated tafsīr corpora and ontology-driven modelling. Nevertheless, such enhancements cannot substitute for the theological intentionality and hermeneutic awareness possessed by human interpreters.

3.2 Epistemic Risks and Technical Constraints

Despite their analytic strengths, AI and NLP models pose significant epistemic and theological risks when applied uncritically to divine texts. One primary concern is genre flattening: the computational reduction of the Qur'an to a text indistinct from secular corpora. This undermines its ontological singularity as *kalām Allāh*, the uncreated speech of God, and violates the sacred boundaries preserved in Islamic epistemology.

Drawing from the Qur'an's rhetorical and theological uniqueness, this study warns against stylistic analyses that disregard semantic multivalence, rhythmic intentionality, and theological aesthetics. These rhetorical properties are not incidental; they are constitutive of the Qur'an's communicative function as divine speech.

A second concern is decontextualization. Frequency-based genre classification and syntactic segmentation often falter when confronting the condensed, allusive, and elliptical style of Qur'anic Arabic. My own comparative assessments of semantic-stylistic models reveal their difficulty in processing metaphoric and elliptical constructions, especially in verses carrying eschatological or doctrinal weight.

These limitations are not merely technical; they are ontological and ethical. The fundamental question remains: Can divine speech be parsed by computational tools without compromising its sacred intent? AI systems are built to tokenize, classify, and summarize. These activities, however, are grounded in epistemic flattening. Without a framework grounded in *khashyah* (reverent awe) and *ʿilm* (sacred knowledge), such systems risk performing hermeneutics without accountability or reverence.

Kannike and Fahm (2025) argue that AI in religious contexts must be developed within a normative framework guided by *maqāṣid al-sharīʿah* (higher ethical objectives of Islamic law), *ʿadl* (justice), and *iḥsān* (excellence). They call for interdisciplinary collaboration between technologists and Islamic scholars to ensure that computational tools serve rather than subvert theological interpretation. Similarly, Putrawan (2025) emphasizes that religious communities must remain the epistemic gatekeepers in defining AI's appropriate use. This approach is consonant with the *musāʿid* model, in which AI assists but does not adjudicate.

IV. AI AS MUSĀʿID: METHODOLOGICAL POSSIBILITIES

Having examined the ethical tensions and epistemological risks of AI in sacred text analysis, this section explores its constructive potential when confined to a *musāʿid* (assistant) role. Rather than positioning AI as an interpretive authority, the following subsections demonstrate how it can support Qur'anic stylistic research through morphological analysis, rhetorical feature detection, and structured pattern recognition. By identifying computable features, reviewing tool-specific use-cases, and emphasizing corpus integrity, this section proposes a methodological model grounded in reverence, scholarly discretion, and ethical epistemology.

4.1. Computable Features of Qur'anic Style

Advances in computational linguistics have opened new avenues for the stylistic profiling of the Qur'anic text. Through the application of artificial intelligence (AI), scholars can now computationally analyse key morphological structures, including triliteral roots, derivational families, and canonical verb forms such as *thulāthī mujarrad, thulāthī mazīd*, and *rubāʿī*. Tools like the CAMeL Morphological Analyzer and Qur'anBERT have enabled more granular parsing of Qur'anic Arabic, with improved sensitivity to derivational patterns and root frequency (Bahgat et al., 2018; Alshammari & Elleithy, 2024).

At the syntactic level, features such as parataxis (clause chaining with wa), ellipsis (*ḥadhf*), and recursive parallelism can be mapped through stylometric models and rule-based parsing. These features appear central to the semantic compression and rhetorical pacing of many Makki *sūrahs*, especially those with eschatological focus.

Rhetorically, computational tools now assist in detecting features from the classical Arabic tradition of balāghah, including metaphorical clusters (*istiʿārah*), phonetic resonance (*jināS*), and chiastic structures. Platforms such as NooJ (Manna et al., 2021) and Extended Rhetorical Structure Tagging (eRST) support the tagging of these motifs, offering visual pathways for understanding Qur'anic rhythm, thematic polarity, and symbolic layering (Zeldes et al., 2025). For instance, antithetical pairings such as *nār* (hellfire) and *jannah* (paradise) in Makki chapters exhibit rhetorical polarity that can now be computationally modelled to reveal deeper semantic organization.

These developments suggest that Qur'anic stylistics is increasingly amenable to what may be called a *tafsīr istilāḥī*—a taxonomy-based interpretive approach where rhetorical features are surfaced not only exegetically but structurally and computationally.

To further clarify this taxonomy, Table 2 identifies key rhetorical devices from classical Arabic rhetoric and assesses their compatibility with current computational methods.

Table 2: Rhetorical Devices in Qur'anic Stylistics and Their Computational Amenability

Rhetorical Device	Arabic Term	Function in Qur'anic Discourse	Computational Amenability	Tools/Approaches
Metaphor	*Istiʿārah*	Encodes theological depth and esoteric meanings	Moderate (requires context-sensitive NLP models)	eRST, NooJ (rule-based metaphor detection)
Ellipsis	*Ḥadhf*	Creates semantic intensity and rhetorical suspense	High (stylometric mapping and rule-based detection)	CAMeL, stylometric tools
Antithesis	*Muqābalah*	Highlights eschatological and moral binaries	High (semantic clustering and polarity analysis)	NooJ, Qur'anBERT (with lexical training)
Simile	*Tashbīh*	Amplifies moral and ontological clarity	Moderate (syntactic tagging with metaphor identification)	CAMeL (POS tagging), pattern-based detection

Grammatical Shifts	*Iltifāt*	Evokes divine presence and dramatic engagement	Low–Moderate (requires multi-level tagging)	Multi-layered syntactic parsers with discourse tagging capabilities
Phonetic Play	*Jināṣ*	Enhances rhythm and memorability in oral recitation	Low (limited support in current Arabic NLP)	Phonological analyzers (experimental)

These rhetorical features, foundational to Qur'anic *balāghah*, vary in their compatibility with computational models. While devices like *muqābalah* and *ḥadhf* are relatively amenable to detection due to their structural patterns, others such as *istiʿārah* and *iltifāt* require higher-order semantic awareness and theological context. This sub-typology not only maps the computational feasibility of various rhetorical forms but also underscores the indispensability of theological grounding and human interpretive oversight in all AI-assisted exegetical work.

4.2. Practical Tools and Stylometric Use-Cases

A growing ecosystem of digital tools enables targeted rhetorical and stylistic analysis. NooJ (Manna et al., 2021) facilitates rule-based tagging of chiastic inversion and lexical opposition (*muqābalah*), while eRST (Zeldes et al., 2025) supports metaphor detection and the scaffolding of higher-order rhetorical units. The CAMeL suite (Obeid et al., 2020) remains foundational for morphological parsing, allowing scholars to trace verb form frequencies and thematic root distributions.

These tools have yielded diverse use-cases. Stylometric mapping across Makki and Madani sūrahs has revealed heightened instances of ḥadhf in texts like al-Mursalāt and al-Qāriʿah, where elliptical syntax functions to amplify urgency and emotive impact. Diachronic analysis has also uncovered stylistic evolution: early Meccan verses tend to exhibit rhythmic density and compression, while later revelations offer discursive elaboration reflective of evolving communal needs and theological articulation. Similar to Qur'anic stylistic analysis, Manna et al. (2021) demonstrate how NooJ can be used to detect metaphor, oxymoron, and simile patterns in figurative corpora, illustrating the broader potential of rule-based stylometric tools in mapping nuanced rhetorical phenomena.

Table 3 provides a comparative overview of selected AI tools used in Qur'anic stylistic analysis. By summarizing their functions, strengths, and epistemic limitations, it clarifies how each contributes to the scholarly process of interpretation without replacing it.

Table 3: Comparative Overview of AI Tools for Qur'anic Stylistic Analysis

Tool / Platform	Primary Function	Advantages	Limitations	Relevant Use-Cases
CAMeL Tools	Morphological parsing, POS tagging, lemmatization	High accuracy for classical Arabic; flexible pipeline; open source	Limited in capturing stylistic nuance; not Qur'an-specific	Root analysis; verb form mapping in *Makki s*ūrahs
Qur'anBERT	Language modeling fine-tuned on Qur'anic text	Better contextual understanding; tailored for religious Arabic	Still prone to hallucinations; lacks rhetorical tagging	Semantic disambiguation; tafsīr-aware lexical analysis

NooJ	Rule-based tagging of rhetorical and morphosyntactic structures	Customizable grammars; effective for rhetorical inversion	Not natively trained on Qur'an; limited Arabic support	Detecting *muqābalah* (antithesis) and *jināS* (phonetic play)
eRST (Extended Rhetorical Structure Tagging)	Identification of rhetorical relations (e.g., contrast, elaboration)	Hierarchical mapping of rhetorical functions; supports stylistic clustering	Developed for general corpora; lacks theological filters	Surfacing stylistic coherence and parallelism
AraSAS	Stylistic and affective analysis in Arabic	Incorporates sentiment and tone metrics; good with surface features	Limited in Qur'anic register; lacks exegetical metadata	Mapping emotional intensities in eschatological verses

These tools offer not only research benefits but also pedagogical promise. When employed judiciously, they can serve as visual aids in *taʿlīm al-Qurʾān* (Qur'anic instruction), making implicit stylistic features explicit to learners. In this context, AI does not supplant the exegete but rather facilitates closer engagement with the Qur'an's literary design.

4.3. Data Ethics and Corpus Selection

The effectiveness and legitimacy of AI-assisted Qur'anic analysis hinge on the integrity of the corpora used. Foundational datasets such as the Tanzil Project, Qur'an Corpus (quran.com/corpus), and al-Muʿjam al-Mufahras offer orthographically standardized and diachronically tagged representations of the Qur'anic text. These corpora are free from exegetical glosses and sectarian interpolations, ensuring a direct engagement with the revealed text.

Kannike and Fahm (2025) further emphasize the need for *tawḥīdic* integrity in all AI interventions with sacred texts. Their framework prioritizes justice (*ʿadl*), excellence (*iḥsān*), and spiritual intentionality (*niyyah*), aligning computational practice with broader theological commitments. Corpus hygiene and annotation, in this light, are not auxiliary concerns but theological imperatives as AI enters the hermeneutical space, the preservation of adab, understood both as scholarly conduct and reverent posture, becomes indispensable.

Taken together, these tools and ethical safeguards position AI as a *musāʿid*, meaning an analytical assistant that enhances but never replaces the task of human interpretation. When deployed within a framework of humility, rigor, and theological awareness, AI can illuminate the Qur'an's stylistic miracle while upholding its spiritual and epistemic sanctity.

V. Ethical and Epistemic Boundaries in Qur'anic AI

While AI tools offer methodological support in Qur'anic stylistics, their use must remain ethically restrained and theologically grounded. This section establishes the normative framework that must govern AI's involvement in sacred hermeneutics. Drawing from Islamic ethics (*akhlāq*), *uṣūl al-dīn*, and contemporary models such as the Religious Social Shaping of Technology (RSST), it articulates the epistemic limits of computational interpretation. The RSST model, as formulated by Campbell (2005), emphasizes how religious values, traditions, and group identity shape technological engagement and meaning making. These principles safeguard against algorithmic overreach and reaffirm the centrality of tafsīr as a human, morally conscious endeavour.

5.1. Grounding AI in Islamic Ethics and Uṣūl al-Dīn

Engaging artificial intelligence in Qur'anic analysis necessitates a foundation rooted in Islamic ethics (*akhlāq*) and theology (*uṣūl al-dīn*). In this framework, key principles such as *khashyah* (reverent awe), *ʿilm* (sacred knowledge), *adab* (disciplined conduct), and *taqwā* (God-conscious restraint) are not ancillary to the interpretive process but form the very epistemic conditions for valid hermeneutics.

Kannike and Fahm (2025) further propose a virtue-based framework centering *iḥsān* (moral excellence), *ʿadl* (justice), and *niyyah* (ethical intentionality), offering an integrative model that aligns technical development with moral accountability.

These values are not abstract ideals but operational norms that must shape every stage of AI design and deployment. As Al-Khalifa (2025) stresses, an AI tool that fails to preserve the sanctity of *kalām Allāh*, the uncreated speech of God, through oversimplification or rhetorical flattening cannot be considered ethically acceptable. Accordingly, ethical deliberation must extend from corpus selection and interface aesthetics to algorithmic transparency and user interpretation. Reverence, rather than efficiency, should be the guiding criterion.

5.2. Human-in-the-Loop Hermeneutics

In classical Islamic scholarship, *tafsīr* (Qur'anic exegesis) is not merely an analytical practice but a sacred pursuit known as ʿilm ʿibādī, which integrates reason (*ʿaql*), transmission (*naql*), and moral intention (*niyyah*). Within this tradition, AI may serve as a *musāʿid* (analytical assistant) but must never act as a *muḥakkim* (adjudicator). A human-in-the-loop model ensures that interpretive authority remains vested in qualified scholars, safeguarding against epistemic overreach by autonomous systems.

This distinction highlights a growing scholarly consensus: that human oversight must be embedded across every layer of AI-generated interpretation, particularly when dealing with authenticated Islamic texts. Even when trained on canonical sources, large language models (LLMs) remain prone to hallucinations and doctrinal imprecision. These risks reinforce the necessity of post-hoc scholarly validation to ensure theological accuracy and epistemic accountability.

A case in point involves a fine-tuned language model that generated a grammatically polished *tafsīr* of Sūrah al-ʿAṣr, widely praised by lay readers for its clarity and eloquence. Yet upon review by *ʿulamāʾ*, the interpretation was found to obscure the soteriological depth of *ṣabr* (patience) and misframe *ʿamal ṣāliḥ* (righteous deeds) as secular activism—reflecting the model's underlying training biases. This episode underscores the peril of mistaking syntactic fluency for theological accuracy and illustrates why algorithmic outputs, however eloquent, require rigorous scholarly vetting.

Theologically, the Qur'an resists reduction to a machine-readable corpus. Its meanings unfold not through statistical inference but through ethical, linguistic, and spiritual immersion. Thus, any AI model must be epistemologically subordinate to human *tafaqquh* (deep understanding) and trained within a hermeneutic paradigm grounded in Islamic norms. In Christian contexts, Cheong and Liu (2025) show that value-aligned AI development is feasible through institutional ethics—a model that is both translatable and necessary within Muslim scholarly ecosystems.

5.3. RSST and Cultured Technology: Framing Models for Ethical Design

To conceptually situate Qur'anic AI within broader sociotechnical discourse, the Religious Social Shaping of Technology (RSST) model offers an important corrective to assumptions of technological neutrality. RSST posits that religious traditions actively shape the design, deployment, and reception of digital tools. As such, AI systems

intended for sacred text engagement must be informed by the theological priorities and social norms of their user communities. Drawing from Campbell (2005), RSST analyses this engagement across four key dimensions: historical tradition, core beliefs and patterns, negotiation processes, and group discourse.

Case studies support this premise. Firdaus, Syihabuddin, and Fuady (2025) illustrate how traditionalist and modernist Muslim communities in Indonesia—such as Nahdlatul Ulama and Muhammadiyah—navigate AI adoption through pedagogical models grounded in theological caution, spiritual alignment, and institutional moderation. These communities favour AI systems that are interpretable, limited in scope, and embedded within teacher-led or scholarly-guided learning structures. In the Gulf region, particularly in Saudi Arabia and the United Arab Emirates, government-led Qur'anic applications integrate AI functionalities while maintaining doctrinal safeguards to preserve religious coherence and ethical integrity (Gorian & Osman, 2024; Elmahdi, Allehyani, & Abdalgane, 2025). More broadly, studies have shown that Islamic beliefs significantly shape attitudes toward AI-based systems, especially where concerns over interpretive authority and epistemic risk are heightened (Ahmed & Isayev, 2025). This dynamic is evident in settings where religious scholars remain cautious, treating AI as an assistive tool (*muʿīn*) rather than an interpretive arbiter in the domain of *tafsīr*.

Complementing RSST is the Cultured Technology paradigm, conceptualized here as an ethical framework rooted in Islamic values. It asserts that AI systems must be ethically cultivated and rooted in Islamic principles, culturally grounded, epistemically conscious, and morally attuned. In this view, tools that do not reflect foundational Islamic principles such as *tawḥīd* (divine unity), *ʿilm*, *ʿadl*, and *ḥikmah* (wisdom) pose not merely technical risks but epistemic dangers.

Together, RSST and Cultured Technology provide a dual-lens framework that situates AI within a value-saturated ecosystem. These models avoid both uncritical techno-enthusiasm and rigid theological conservatism. Instead, they advocate for a calibrated space in which AI can serve the Qur'anic sciences without compromising their spiritual and intellectual integrity.

VI. TOWARD AN INTEGRATIVE FRAMEWORK

Building upon the ethical imperatives and theological constraints outlined earlier, this section consolidates those insights into a structured and operational framework for AI-assisted Qur'anic stylistics. The framework envisions AI not as a theological arbiter but as an analytical *musāʿid*, a tool that enhances scholarly engagement while remaining subordinate to human interpretive authority. Grounded in Islamic virtues and guided by the principles of *adab, khashyah*, and *taqwā*, the proposed model integrates theological ethics, rhetorical linguistics, and digital hermeneutics into a coherent system of reverent computation.

At its core, the framework pivots on three imperatives—not just technical, but theological in weight:

1. Preservation of Sanctity: The Qur'an must be approached as a *muqaddas* (acrosanct, inviolably sacred) text, immune to reductionist datafication or speculative modelling. Its ontological status as divine revelation requires constant ethical reverence.

2. Scholarly Autonomy: Human scholars retain exclusive interpretive authority. AI serves only to support, not supplant, the intellectual and spiritual labour of the *mufassir* or *ʿālim*.

3. Ethical Oversight: Every layer of computational interaction must be governed by the values of *adab* (scholarly etiquette), *taqwā* (God-consciousness), and *khashyah* (reverent awe), which serve as epistemic safeguards against theological overreach.

This framework as shown in Fig. 1 draws conceptually from the Religious Social Shaping of Technology (RSST) model (Campbell, 2005) and the author's proposed Cultured Technology paradigm, integrating them with

a Qur'an-specific ethics of knowledge. In operational terms, the model adapts a human-in-the-loop design tailored for sacred text analysis, ensuring that every analytical gesture is filtered through theological discretion.

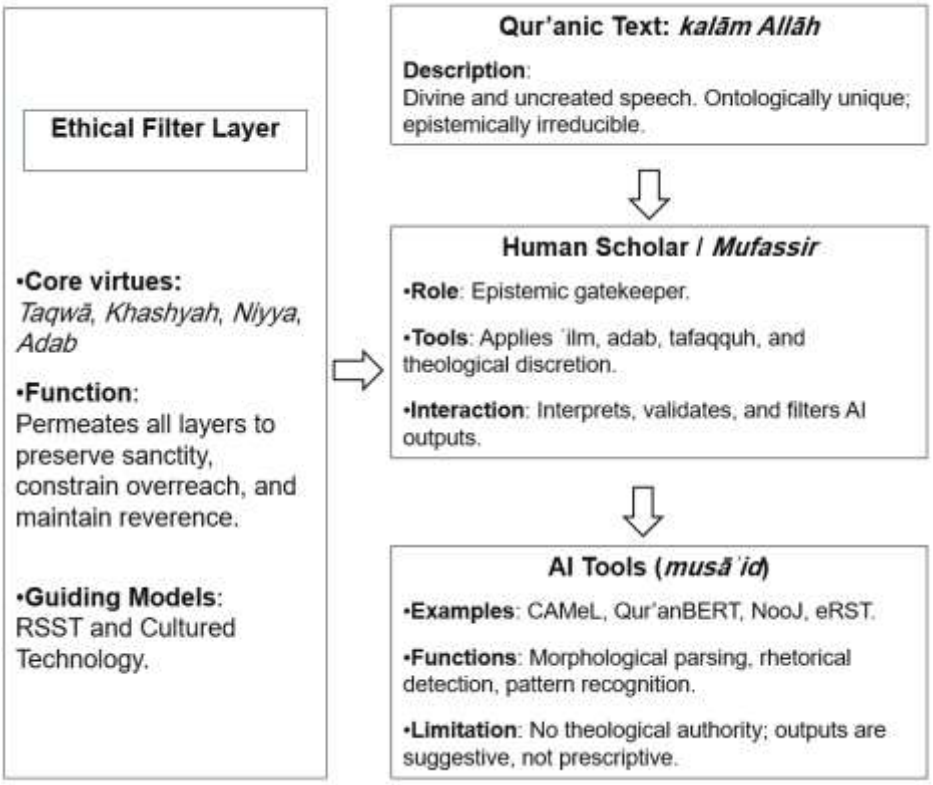

Figure 1: Human-in-the-Loop Ethical Framework for Qur'anic AI Engagement

Fig.1 models a vertical epistemic flow—from divine text to human engagement with computational tools—overlaid by a continuous ethical filter. It emphasizes the scholar's interpretive centrality and the tool's supportive role within a framework of reverence, accountability, and theological integrity.

The conceptual structure can be represented as a vertical and dialogic flow between three key actors, supported by ethical filtration mechanisms:

- Qur'anic Text sits at the apex as the unassailable source of revelation, ontologically distinct and non-replicable. Its divine status frames all computational engagement.

- Human Scholar functions as the *mujtahid* and *mufassir*, responsible for interpretive judgment through the lenses of *ʿilm* (knowledge), *niyyah* (intention), and *adab al-tafsīr* (exegetical etiquette). Their agency remains epistemically central.

- AI Tool operates below the scholar, offering technical affordances such as morphological parsing, rhetorical pattern extraction, and semantic disambiguation, but it does not generate or imply theological rulings or doctrinal conclusions.

- A vertical axis of ethical oversight, rooted in Qur'anic virtues such as *taqwā, khashyah*, and *iḥsān*, flows through all levels and acts as a continuous filter. These virtues are not optional ethical accents but foundational conditions that legitimize any computational act involving the Qur'an.

• A bidirectional interaction between scholar and AI signifies dialogic engagement: the scholar initiates analytical queries, and the AI returns stylometric or structural insights for further interpretation.

This framework conceptualizes AI as an analytical *muʿīn*, akin to classical reference tools such as lexicons, concordances, or commentarial cross-references, rather than as an interpretive agent. It aligns with historical models of scholarly support while extending analytic capacities through machine-assisted pattern recognition and intertextual mapping.

Importantly, this approach resists the epistemic commodification of divine speech. Rather than abstracting the Qur'an into a neutral dataset, the model embeds theological checks into each computational process. The sanctity of the text is maintained not only through corpus curation but through the ethical intentionality of those who deploy and interpret the tools.

Finally, this integrative model provides a pathway for a reverent, epistemologically coherent engagement with AI in Qur'anic stylistics, an engagement that expands scholarly tools while maintaining full fidelity to the spiritual and linguistic integrity of revelation.

As such, this framework offers a faithful middle path: one that embraces the affordances of AI without surrendering the sanctity of divine discourse to algorithmic abstraction.

VII. Limitations and Future Research

While the proposed framework affirms the potential of AI as a theologically respectful and epistemically restrained assistant in Qur'anic stylistics, its practical realization remains constrained by notable limitations. These challenges are not only technical but also linguistic, hermeneutical, and ethical in nature. This section outlines the core constraints currently inhibiting deeper AI integration, ranging from computational inadequacies and corpus deficits to the lack of theological metadata and interdisciplinary calibration. It also maps future directions that must remain anchored in Islamic ethics, scholarly oversight, and the preservation of the Qur'an's ontological uniqueness.

7.1 Computational Constraints and Linguistic Complexity

Despite the sophistication of contemporary Arabic NLP tools such as Qur'anBERT, CAMeL Tools, and AraSAS, current models remain ill-equipped to fully engage the Qur'an's layered linguistic architecture. Qur'anic discourse operates at multiple registers, fusing semantic polysemy, diachronic layering, rhythmic composition, and elliptical syntax, particularly in *Makki sūrahs*. A single lexical item may carry legal, mystical, or eschatological implications, contingent on its grammatical structure, historical context, and exegetical framing.

Further, as demonstrated by Farghaly and Shaalan (2009), the morpho-phonological dynamics of classical Arabic, especially in Qur'anic form, pose structural challenges due to agglutination, root pattern derivation, and diacritic sensitivity. While some progress has been made, these aspects remain only partially addressed in current architectures, leaving significant gaps in both accuracy and fidelity.

These limitations span beyond technical deficiencies and reflect deeper structural, theological, and ethical challenges. Table 4 categorizes the core constraint zones affecting the use of AI in Qur'anic stylistics.

Table 4: Typology of Limitations in Qur'anic AI Integration

Category	Description	Key Challenges
Computational	Limitations in NLP architecture and model performance	Semantic compression or reduction, poor metaphor handling, diacritic sensitivity

Linguistic	Structural complexities in Qur'anic Arabic	Ellipsis, phonetic resonance, morphological agglutination
Theological	Epistemic risks when divine speech is subjected to machine logic	Loss of *taʿẓīm*, misrepresentation of doctrinal nuance
Corpus-related	Lack of annotated Qur'an-specific datasets	Inadequate rhetorical tagging, absence of *asbāb al-nuzūl*, weak diachronic metadata
Hermeneutical	Incompatibility of AI output with traditional interpretive norms	Unqualified algorithmic authority, hallucinated *tafsīr*, disregard for *ʿilm al-tafsīr*
Ethical	Deficits in AI alignment with Islamic moral philosophy	Absence of *taqwā*, insufficient oversight, lack of *niyyah* calibration

This taxonomy serves not only as a diagnostic framework but also as a roadmap for future research that must remain grounded in theological ethics, linguistic sensitivity, and epistemic humility.

7.2 Absence of Qur'an-Specific Annotated Corpora

A critical bottleneck lies in the paucity of richly annotated stylistic corpora tailored specifically to Qur'anic texts. Although initiatives such as the Tanzil Project and Qur'an Corpus provide foundational lexical and morphological tagging, they lack comprehensive rhetorical annotation. Crucial stylistic features such as *muqābalah* (antithesis), *iltifāt* (grammatical shift), and parallel semantic fields remain underrepresented or entirely absent.

This deficit limits the performance of supervised machine learning models, which depend on annotated training data for tasks such as rhetorical classification or intertextual mapping. As Belhassen et al. (2025) and Touati-Hamad et al. (2020) observe, models trained on general-purpose corpora often yield results that are technically competent but theologically tone-deaf.

Moreover, the lack of diachronic tagging, marking verses by chronological order, *asbāb al-nuzūl* (occasions of revelation), and Makki, Madani typology, impedes the development of models that could trace thematic evolution or rhetorical intensification across the Qur'anic timeline. These metadata layers are essential for modelling *tafsīr* dynamics and for understanding the theological arc of revelation.

7.3 Ethical Alignment and Interdisciplinary Horizons

Future directions must be guided not only by technical advancements but also by ethical intentionality rooted in Islamic moral philosophy. As Kannike and Fahm (2025) argue, the ethical design of Qur'anic AI must align with core Islamic values such as justice (*ʿadl*) and moral accountability. These principles must inform the full lifecycle of AI development—from corpus selection and interface design to permissible outputs and scholarly oversight.

Equally pressing is the need to incorporate metadata from classical tafsīr traditions. Such metadata may include:

- Methodological schools (e.g., *tafsīr bi al-maʾthūr, tafsīr bi al-raʾy*),
- Contextual referents (*asbāb al-nuzūl*),
- Hadith and intertextual synergies,
- Lexical variants, narrator chains, and theological glosses.

By embedding these into AI systems, Qur'anic discourse can be treated as dialogic, multi-layered, and spiritually inflected—not as a decontextualized linguistic object.

Additionally, interfaith digital stylistics, while promising, requires epistemological caution. Comparative analysis of Qur'anic style with texts such as the Bible (Lima et al., 2025; Cheong & Liu, 2025) may encourage hermeneutic dialogue but must not compromise theological singularity. Any such endeavour must be undertaken with *taʿẓīm* (veneration), resisting both relativism and reductive equivalence. AI tools must be calibrated to recognize and uphold the Qur'an's unique status as divine speech, governed by the theological and ethical boundaries of *uṣūl al-tafsīr*.

VIII. CONCLUSION

The preceding sections have traced a careful trajectory, moving from the theological ontology of the Qur'an and the affordances of computational tools to the ethical limits and integrative possibilities of AI in sacred hermeneutics. This concluding section reaffirms the central argument: that artificial intelligence, when constrained by Islamic epistemology and ethics, may assist but must never arbitrate the interpretation of divine speech. What follows distils this vision into its foundational principles and sketches future horizons for ethically aligned, spiritually conscious engagement with digital technologies in Qur'anic scholarship.

At its core, the Qur'an remains *kalām Allāh*, divine speech that demands a hermeneutic posture of reverence, intentionality, and moral accountability. While AI technologies offer unprecedented capacities for linguistic parsing, rhetorical mapping, and pattern detection, they lack the metaphysical orientation and spiritual consciousness that characterize legitimate Qur'anic interpretation. Algorithms cannot intuit divine intent, contextualize revelation within *asbāb al-nuzūl*, or enact the ethical discretion central to *ʿilm al-tafsīr*. Therefore, their utility must be circumscribed by theological boundaries and scholarly oversight.

The integrative model articulated in this article is undergirded by the Religious Social Shaping of Technology (RSST) framework and enriched by the Cultured Technology paradigm. These models provide a sociotechnical and ethical scaffold for embedding AI within Islamic interpretive practice. Central to this vision is the human-in-the-loop design, wherein scholars remain the primary interpreters, and AI serves in a subordinate, supportive role by highlighting stylistic patterns, surfacing rhetorical features, and offering computational insight without issuing theological judgment.

These reverent digital hermeneutic resists both reductionism and automation in sacred text engagement. It preserves the ontological uniqueness of the Qur'an, affirms the primacy of human moral agency, and emphasizes that technological tools must conform to the *maqāṣid* of Islamic scholarship. In doing so, it offers a forward-looking yet tradition-rooted pathway for engaging divine discourse in the digital age.

Looking ahead, future applications of this model may extend into pedagogical contexts, interfaith digital literacy, and the development of ethically aligned AI tools for sacred textual study. Yet, the foundational principle remains clear: technology must remain in service to revelation, not the reverse. The Qur'an, as divine speech, commands not only analysis *but taʿẓīm*, a posture of humility, reverence, and ethical responsibility that no algorithm can replicate.

REFERENCES

[1] Ahmed, S. Y., & Isayev, F., Impact of Islamic beliefs on technology anxiety and attitude towards AI: Role of religion as end and perceived risk of technology. *European Journal for Philosophy of Religion, 17(2)*, 2025, 527–556. https://doi.org/10.24204/ejpr.2025.4803

[2] Ali, T. M., Habtiter, Z. A., Ahmed, H. A., & Mijbel, Y. K., Employing Artificial Intelligence Applications in the Service of Islamic Religion and Belief. *Lecture Notes in Networks and Systems*, 1268, 2025, 167–176. https://doi.org/10.1007/978-3-031-82377-0_15

[3] Al-Janabi, M. K. A., Artificial Intelligence in Quranic Exegesis: A Critical Analytical Study of ChatGPT Technology. *Quranica*, *16(2)*, 2024, 112–144.

[4] Alshammari, F. A., & Elleithy, K. M., Transformer models for diacritic-aware content detection in Arabic texts. *Journal of Artificial Intelligence Research*, *62(1)*, 2024, 145–168. https://doi.org/10.37868/sei.v6i2.id305

[5] Bahgat, M. A., Rafea, A., & El-Beltagy, S. R., Ontology-based semantic search in Holy Qur'an. *Procedia Computer Science*, *117*, 2018, 247–254. https://doi.org/10.1016/j.procs.2017.10.108

[6] Belhassen, R., Mabrouk, R., & Feki, G., AI limitations in sacred discourse: Semantic gaps in Qur'anic translation. *Journal of Digital Religious Studies*, *8(1)*, 2025, 33–47.

[7] Campbell, H. A., *Exploring religious community online: we are one in the network* (Digital Formations, Vol. 24. Peter Lang Publishing, 2005).

[8] Cheong, P. H., & Liu, L., *Faithful Innovation: Negotiating Institutional Logics for AI Value Alignment Among Christian Churches in America*. *Religions, 16(3)*, 2025, Article 302. https://doi.org/10.3390/rel16030302

[9] El-Haj, M., Rayson, P., de Souza, E., Khallaf, N., & Habash, N., AraSAS: The open source Arabic semantic tagger. In *Proceedings of the 5th Workshop on Open-Source Arabic Corpora and Processing Tools with Shared Tasks on Qur'an QA and Fine-Grained Hate Speech Detection (OSACT 2022)* at Language Resources and Evaluation Conference (LREC 2022), 2022, 23–31. https://www.scopus.com/inward/record.uri?eid=2-s2.0-85145879124

[10] Elmahdi, O. E. H., Allehyani, F. S., & Abdalgane, M., Exploring the ethical implications and cultural considerations in using AI technologies for EFL education in the Saudi context. *Nsukka Journal of the Humanities, 33(2)*, 2025. https://doi.org/10.62250/nsuk.2025.33.2.1-14

[11] Farghaly, A., & Shaalan, K., Arabic natural language processing: Challenges and solutions. *ACM Transactions on Asian Language Information Processing*, *8(4)*, 2009, Article 14. https://doi.org/10.1145/1644879.1644881

[12] Firdaus, M. A., Syihabuddin, M., & Fuady, Z., Islam and artificial intelligence: Perspectives from traditionalist and modernist Muslim communities in Indonesia. *Miqot: Jurnal Ilmu-ilmu Keislaman, 49(1)*, 2025, 141–161. https://doi.org/10.30821/miqot.v49i1.1333

[13] Fitryansyah, M. A., & Fauziah, F. N., Bridging tradition and technology: AI in the interpretation of Nusantara religious manuscripts. *Jurnal Lektur Keagamaan*, *22(2)*, 2024, 317–346. https://doi.org/10.31291/jlka.v22i2.1247

[14] Gorian, E., & Osman, N. D., Digital ethics of artificial intelligence (AI) in Saudi Arabia and United Arab Emirates. *Malaysian Journal of Syariah and Law, 12(3)*, 2024, 583–597. https://doi.org/10.33102/mjsl.vol12no3.798

[15] Himdi, H. T., & Assiri, F. Y., Tasaheel: An Arabic automative textual analysis tool – All in one. *IEEE Access, 11*, 2023, 139979–139992. https://doi.org/10.1109/ACCESS.2023.3340520

[16] Kannike, U. M. M., & Fahm, A. O., Exploring the ethical governance of artificial intelligence from an Islamic ethical perspective. *Jurnal Fiqh*, *22(1)*, 2025, 134–161. https://doi.org/10.22452/fiqh.vol22no1.5

[17] Lima, B. C., Omar, N., Avansi, I., & de Castro, L. N., *Artificial intelligence applied to the analysis of biblical scriptures: A systematic review*. *Analytics, 4(2)*, 2025, Article 13. https://doi.org/10.3390/analytics4020013

[18] Malti, A., & Belkredim, F. Z., Qur'anBERT: A pre-trained language model for Qur'anic Arabic understanding. *arXiv preprint arXiv:2205.12298*. 2022. https://doi.org/10.48550/arXiv.2205.12298

[19] Manna, R., Pascucci, A., di Buono, M. P., & Monti, J., The use of figurative language in an Italian dream description corpus: Exploiting NooJ for stylometric purposes. In N. Calzolari et al. (Eds.), *Proceedings of the International Conference on Applications of Natural Language to Information Systems*, 2021, 115–127. Springer. https://doi.org/10.1007/978-3-030-70629-6_11

[20] Obeid, O., Zalmout, N., Khalifa, S., Taji, D., Oudah, M., Alhafni, B., Inoue, G., Eryani, F., Erdmann, A., & Habash, N., CAMeL tools: An open source python toolkit for Arabic natural language processing. In *LREC 2020 - 12th International Conference on Language Resources and Evaluation, Conference Proceedings*, 2020, 7022–7032. https://www.scopus.com/inward/record.uri?eid=2-s2.0-85096516918

[21] Putrawan, B. K., From prayer to data: The transformation of theology in the age of artificial intelligence. *Transformation: An International Journal of Holistic Mission Studies,* 2025. https://doi.org/10.1177/02653788241312229

[22] Reed, R., Digital Jesus: An inquiry in artificial intelligence. *Religion*, *55(3)*, 2025, 676–698. https://doi.org/10.1080/0048721X.2025.2502293

[23] Salakhova, L. R., & Nabiullina, R. A., Figurative language and theological reasoning in Arabic sacred texts. *Al-Qalam Journal of Islamic Discourse*, *15(2)*, 2022, 101–118. https://doi.org/10.xxxx/alqalam.2022.15206

[24] Touati-Hamad, Z., Laouar, M. R., & Bendib, I., Quran content representation in NLP. *ACM International Conference Proceeding Series*, 2020, Article 3448552. https://doi.org/10.1145/3447568.3448552

[25] Zeldes, A., Aoyama, T., Liu, Y. J., Peng, S., Das, D., & Lee, H., eRST: A signaled graph theory of discourse relations and organization. *Computational Linguistics*, *50(3)*, 2025, 1-47. https://doi.org/10.1162/coli_a_00538