AI-Assisted Translation of Qur'anic Text: Linguistic and Rhetorical Challenges in Surah al-Kawthar

Mahboobah Kaprakkadan

PhD Research Scholar, Department of Arabic, Anvarul Islam Women's Arabic College, Mongam, India

DOI: https://doi.org/10.47772/IJRISS.2026.102000049

Received: 15 October 2026; Accepted: 20 October 2026; Published: 19 November 2026

ABSTRACT

Artificial intelligence (AI), particularly neural machine translation and large language models, has created new opportunities for translating religious texts. However, translating the meanings of the Qur'an involves challenges that extend beyond grammatical accuracy and linguistic fluency. Qur'anic expressions are characterized by semantic richness, rhetorical features, contextual relationships, and interpretive dimensions that may not be fully represented by a single equivalent in another language.

This article examines the linguistic and rhetorical challenges of AI-assisted Qur'anic translation through a descriptive-analytical discussion of three expressions from Surah al-Kawthar: al-Kawthar (Qur'an 108:1), wa-nḥar (108:2), and al-abtar (108:3). Drawing on Qur'anic linguistic resources and interpretive traditions, the study explores how the translation of these expressions may preserve their basic meanings while leaving certain contextual, rhetorical, or interpretive dimensions insufficiently expressed. It also discusses the potential contribution of retrieval-augmented generation (RAG) to source-based translation and the importance of qualified human review.

The article proposes an evaluation framework incorporating semantic adequacy, linguistic acceptability, contextual and rhetorical fidelity, terminological consistency, source traceability, and interpretive transparency. Rather than reporting experimental results or ranking specific AI systems, it develops a conceptual framework for examining the challenges of AI-assisted Qur'anic translation and identifying directions for future empirical research. The study argues that reliable translation requires a clear distinction between the translated text and its explanatory interpretation, supported by verifiable sources and scholarly judgement.

Keywords: Qur'anic translation; artificial intelligence; large language models; Qur'anic rhetoric; retrieval-augmented generation; translation evaluation.

INTRODUCTION

Artificial intelligence (AI) has become increasingly important in modern translation. Neural machine translation systems can produce fluent translations, while large language models (LLMs) can follow instructions, rephrase sentences, explain terms, and suggest different ways of expressing ideas. These technologies have improved access to information in different languages and created new opportunities for sharing religious knowledge digitally. However, they have also raised questions about translation accuracy, reliability, and the authority of AI-generated interpretations.

These issues are especially important in translating the meanings of the Qur'an. In Islamic scholarship, a translation is generally considered a human effort to convey the meanings of the Arabic revelation in another language. It does not replace the original Arabic text or fully reproduce all its linguistic and rhetorical features. Translators must consider the different meanings of words, sentence structures, omitted expressions, emphasis, sound patterns, imagery, and interpretive choices. A translation may be grammatically correct and

easy to read but may still limit the meaning of an expression or fail to explain its context in Qur'anic commentaries.

Large language models make this issue more important because they can produce confident and well-organized answers even when their claims are not supported by reliable evidence. In religious translation, even a small change in wording may influence how readers understand a theological concept, a religious practice, or a rhetorical contrast. Therefore, the main question is not simply whether AI can translate the Qur'an, but how its output can be evaluated, corrected, and used responsibly.

This article examines the opportunities and limitations of AI-assisted translation of Qur'anic meanings. It focuses on the linguistic and rhetorical features of Qur'anic Arabic, uses Surah al-Kawthar as an illustrative case study, and discusses retrieval-augmented generation (RAG) and human review as parts of a responsible translation process. The article does not claim that all machine translations are inaccurate, nor does it present an experimental comparison that has not been conducted. Instead, it identifies possible risks in AI-assisted Qur'anic translation and proposes a framework for future empirical studies.

Research Problem, Questions and Objectives

The main problem discussed in this article is the possible gap between the fluency of AI-generated translations and their ability to convey the full meaning of the Qur'anic text. Translation systems usually produce target-language expressions based on patterns learned from large amounts of data. However, translating the Qur'an may require a deeper understanding of classical Arabic, established interpretive traditions, and the difference between translating a word and explaining its meaning through Qur'anic commentary.

The article addresses four research questions:

- Which linguistic and rhetorical features of Qur'anic Arabic present particular challenges for translation into other languages?
- How do the expressions al-Kawthar, wa-nḥar, and al-abtar illustrate the relationship between lexical meaning, rhetorical context, and Qur'anic interpretation?
- How can digital linguistic resources and retrieval-augmented generation support the translation of Qur'anic meanings, and what limitations remain?
- What criteria can be used to assess AI-assisted Qur'anic translation while distinguishing linguistic evaluation from theological judgement?

The objectives of this study are:

- To examine linguistic and rhetorical features of Qur'anic Arabic that may be difficult to represent fully in translation.
- To analyze three selected expressions from Surah al-Kawthar in relation to their lexical meanings, contextual functions, and interpretive dimensions.
- To discuss the potential contributions and limitations of digital linguistic resources and retrieval-augmented generation in AI-assisted Qur'anic translation.
- To propose an evaluation framework that combines linguistic and contextual criteria with source verification and qualified human review.

The study is conceptual and illustrative in nature. It does not aim to measure the performance of specific AI systems or establish experimentally that one translation method is superior to another.

METHODOLOGY AND SCOPE

The analysis of the selected expressions follows a qualitative, source-based procedure. First, each expression is examined in its Qur'anic context, with attention to its lexical and grammatical features. Second, relevant explanations from identified Qur'anic commentaries and linguistic resources are considered to clarify the interpretive dimensions of the expression. Third, the possible implications for translation are discussed in

terms of semantic adequacy, contextual and rhetorical fidelity, and the distinction between translated meaning and explanatory interpretation. These criteria provide an analytical structure for the discussion; they are not presented as scores derived from an empirical evaluation. Since the study does not collect outputs from specified AI systems or use independent expert ratings, its findings are interpretive observations rather than measured results of AI performance.

This study is limited to Surah al-Kawthar and general issues related to AI-assisted Qur'anic translation. Its observations cannot be applied to all Qur'anic passages, languages, or AI models without further research and testing.

Qur'anic Inimitability and the Challenges of Translation

The concept of iʿjāz al-Qurʾān refers to the inimitability of the Qur'an, understood within Islamic scholarship as a distinctive feature of divine revelation that human expression cannot reproduce.[1] Its discussion encompasses linguistic and rhetorical characteristics, the organization of Qur'anic discourse, and theological understandings of the Qur'an as divine speech. These dimensions are interconnected, although they should not be treated as identical or reduced to a single linguistic criterion.

From a linguistic perspective, Qur'anic expressions may combine lexical precision, grammatical structure, semantic depth, and contextual relationships within concise forms. Rhetorically, word order, emphasis, sound patterns, parallel structures, and figurative expression contribute to the organization and effect of the discourse. Translating such expressions requires decisions about which aspects can be conveyed directly and which may require explanation. A target-language equivalent may communicate the central meaning without reproducing every feature of the original expression.

The theological dimension requires a further distinction. Translation evaluation can examine semantic accuracy, rhetorical representation, and the preservation of contextual meaning. However, these assessments do not, by themselves, establish or disprove theological claims concerning the divine origin or inimitability of the Qur'an. Computational systems can assist in identifying linguistic patterns and comparing translations, but their outputs should not be treated as definitive judgements on theological questions.

In this context, iʿjāz al-Qurʾān provides an important background for examining the limits and possibilities of translation. The purpose of this article is not to claim that AI can measure or reproduce the full inimitability of the Qur'an. Rather, it investigates selected linguistic and rhetorical challenges and considers how AI-assisted translation can be evaluated more transparently through linguistic analysis, reliable interpretive sources, and qualified human review.

Linguistic and Rhetorical Challenges in Qur'anic Translation

Qur'anic Arabic presents several related challenges for translation. At the morphological level, the root, word pattern, prefixes, suffixes, and grammatical form of a word can influence its meaning. At the syntactic level, word order, grammatical relationships, omitted words, pronoun references, and emphasis may affect how a verse is understood. A translation that follows the original word order without understanding these relationships may give an incorrect meaning. On the other hand, explaining every implied meaning may make the translation longer and weaken its original rhetorical effect.

At the semantic level, a single Arabic word may have several meanings that cannot always be expressed through one word in another language. The context may clarify or add to these meanings, while Qur'anic commentaries may present more than one interpretation. The translator must decide whether to use a short

[1] Al-Jurjānī, ʿA. al-Q. (1995). Dalāʾil al-iʿjāz (M. al-Tanji, Ed.). Dar Al-Kitab Al-Arabi.

equivalent, add an explanatory note, or clearly separate the translation from further explanation. Each choice affects the readability and clarity of the translation.

Rhetorical features create additional challenges. The Qur'an uses parallel structures, sound patterns, contrasts, emphasis, imagery, and concise expressions to communicate meaning and influence the reader. Some of these features can be explained through notes, but their full effect may be difficult to reproduce in another language. AI systems may produce natural and fluent sentences without clearly distinguishing between the original wording, scholarly interpretation, and additional explanations included to improve readability.

However, these challenges do not prove that computer systems are unable to process classical Arabic. Their performance depends on several factors, including training data, model design, tokenization, access to reliable sources, task instructions, and evaluation methods. For example, the Quranic Arabic Corpus provides information about word forms, grammatical structures, and semantic relationships. This information can support both human and computer-based analysis of the Qur'an[2]. Although such resources make linguistic information easier to access, they cannot replace the need for careful scholarly judgement.

Surah Al-Kawthar as an Illustrative Case Study

Surah al-Kawthar is a short chapter, but its three verses provide useful examples of how Qur'anic expressions can have different linguistic, contextual, and interpretive meanings.

بِسْمِ ٱللَّهِ ٱلرَّحْمَٰنِ ٱلرَّحِيمِ ـ إِنَّآ أَعْطَيْنَٰكَ ٱلْكَوْثَرَ (١) فَصَلِّ لِرَبِّكَ وَٱنْحَرْ (٢) إِنَّ شَانِئَكَ هُوَ ٱلْأَبْتَرُ (٣)

In the name of Allah, the Most Gracious, the Most Merciful – (1) Indeed, We have granted you (O Muḥammad) abundant goodness. (2) So pray to your Lord and offer sacrifice (to Him alone). (3) Indeed, your enemy is the one cut off. [3]

{ In the name of Allah, the Most Compassionate, the Most Merciful.
1. Indeed, We have bestowed upon you Al-Kawthar (abundant goodness).
2. So pray to your Lord and offer sacrifice.
3. Indeed, it is the one who hates you who is truly cut off.}[4]

The following examples do not compare or rank the translations produced by specific AI systems. Instead, they explain the issues that should be considered in a responsible translation process.

Al-Kawthar (108:1)

The first verse states: "Indeed, We have granted you al-Kawthar" (Qur'an 108:1). In Qur'anic commentaries, the term "al-Kawthar" is associated with abundant goodness. Some reports and commentaries also connect it with a river in Paradise granted to Prophet Muhammad ﷺ. Translating the word as "abundance" conveys its general meaning, but this single word does not fully communicate the interpretations found in Qur'anic commentaries.

However, including one specific interpretation directly in the translated verse also requires care. The expression "al-Kawthar" illustrates the difficulty of representing semantic richness through a single target-language equivalent. A translation such as "abundant goodness" communicates a broad meaning, while retaining the Arabic term preserves a direct connection with the original expression but may leave readers unfamiliar with its interpretation. An explanatory note can introduce interpretations discussed in Qur'anic commentaries, provided that these are clearly distinguished from the wording of the translated verse. The analytical issue is therefore not simply whether one English equivalent is correct, but how a translation

[2] Ontology of Quranic Concepts. https://corpus.quran.com/ontology.jsp

[3] https://quran.com/al-kawthar?translations=20, Translation: Saheeh International

[4] AI-generated translation produced by ChatGPT on 20-October-2026

communicates the lexical meaning while allowing readers to access relevant interpretive dimensions without presenting one explanation as exhaustive.

Wa-nḥar (108:2)

The second verse states: "So pray to your Lord and sacrifice" (Qur'an 108:2). Different translations use different expressions for "wa-nḥar". The verb is related to "nahr", and Qur'anic commentaries discuss its meaning in the context of worship and devotion. A general translation may communicate the basic meaning but leave some of its religious and interpretive context unexplained; thereby this does not mean that a general translation is necessarily incorrect.

The expression "wa-nḥar" requires attention to both its linguistic form and its interpretive context. The rendering "sacrifice" communicates an established understanding of the command in this verse, but an analysis of translation should also consider how the relationship between prayer, devotion, and the act indicated by the verb is explained in relevant Qur'anic commentaries. Such explanations should not automatically be incorporated into the translated sentence as though every interpretive detail were explicitly expressed by the Arabic wording. A responsible translation can preserve a concise rendering while using a separate note to identify relevant scholarly interpretations. This distinction allows the translator to maintain readability without concealing the interpretive decisions involved.

Al-Abtar (108:3)

The final verse states: "Indeed, it is your hater who is cut off" (Qur'an 108:3), although published translations use different expressions. The word "al-Abtar" is associated with the idea of being cut off or severed. However, its meaning in the context of this surah involves more than its basic literal sense. A direct translation may preserve an important part of the word's meaning without fully expressing its rhetorical force in the verse.

The expression "al-abtar" illustrates the relationship between lexical meaning and rhetorical context. Its basic sense of being cut off or severed provides an important starting point, but the expression in Surah al-Kawthar also functions within a response to hostility directed at the Prophet. The translation must therefore be examined not only for its lexical adequacy but also for the way it communicates the contextual contrast and rhetorical force of the verse. Different target-language expressions may convey different aspects of this meaning, and their suitability should be assessed in relation to the Arabic text and relevant exegetical explanations. This example shows why matching Arabic words with target-language equivalents is not enough to evaluate translation quality. The evaluator must consider whether the translation is clear, whether it conveys the meaning of the verse in context, and whether any additional interpretation is supported by reliable sources.

Together, these examples show that responsible Qur'anic translation requires a clear distinction between the original Arabic text, the translated version, and explanations drawn from Qur'anic commentaries. AI can help generate possible translations, but selecting and explaining the most suitable expressions requires careful examination of reliable sources and scholarly judgement.

Large Language Models and Retrieval-Augmented Generation

Large language models (LLMs) can perform many translation-related tasks through natural-language instructions. They can suggest translations, explain grammatical forms, compare different expressions, and help users search large collections of texts. However, their explanations may include unsupported claims, mix information from different sources, or present uncertain interpretations with too much confidence. These problems are especially serious when readers cannot check the original Arabic text or the cited Qur'anic commentary.

Retrieval-augmented generation (RAG) is one possible way to improve the reliability and traceability of AI-generated content. In the framework introduced by Lewis and his colleagues[5], a language model uses an external retrieval system to obtain relevant documents or passages before generating an answer. For Qur'anic translation, these sources could include verified Qur'anic texts, identified editions of Qur'anic commentary, classical Arabic dictionaries, linguistic studies, and published translations with complete reference details.

A well-designed RAG system could be developed to show the source and location of each interpretive claim, distinguish direct quotations from paraphrased information, and avoid claims that are not supported by the retrieved sources. It could also record the versions of the sources used, allowing researchers to check and reproduce the results. These features would make the system easier to verify than one that produces explanations without showing its sources.

However, RAG does not guarantee correct results. The retrieval system may select an irrelevant passage, overlook another important interpretation, or retrieve information from an unreliable edition. The language model may also misunderstand or incorrectly present the information it receives. Therefore, the quality of retrieval and the quality of generated responses must be evaluated separately. The system should also indicate uncertainty and identify cases that require review by specialists. In Qur'anic translation, RAG should be considered a promising approach that needs further testing, rather than a complete solution to interpretive errors.

Digital resources such as the Quranic Arabic Corpus can also support this process by providing information about word forms, grammatical structures, and conceptual relationships[6]. However, these resources offer linguistic data and analytical tools; they do not automatically guarantee that an AI-generated translation accurately represents the theological and rhetorical meanings of the Qur'an.

Evaluating AI-Assisted Qur'anic Translation

Automatic evaluation methods are useful in machine-translation research, but their results must be understood according to the purpose of the translation. BLEU measures the similarity between a translated text and one or more reference translations by comparing groups of words and applying a penalty to overly short translations[7]. It allows researchers to compare translations systematically. However, it may give a low score to a correct translation that uses different wording or give a high score to a translation that shares similar words with the reference but does not fully convey the meaning.

COMET is a neural evaluation method that uses learned language representations. It was developed to measure translation quality in a way that better matches human judgements in general machine-translation tasks[8]. However, this does not mean that COMET or similar metrics can independently evaluate the theological and interpretive accuracy of Qur'anic translations.

The evaluation of Qur'anic translation requires a distinction between linguistic and theological questions. Linguistic evaluation can examine whether a translation preserves lexical meaning, grammatical relationships, contextual significance, and identifiable rhetorical features. It can also investigate whether explanatory additions are supported by relevant scholarly sources. However, these procedures cannot independently

[5] Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems, 33, 9459–9474. https://arxiv.org/abs/2005.11401

[6] Quranic Arabic Corpus. (n.d.). Documentation. https://corpus.quran.com/ documentation/

[7] BLEU: A method for automatic evaluation of machine translation. Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, 311–318. https://doi.org/10.3115/1073083.1073135

[8] COMET: A neural framework for MT evaluation. Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, 2685–2702. https://doi.org/10.18653/v1/2020.emnlp-main.213

determine theological questions concerning the divine origin or inimitability of the Qur'an. Automatic evaluation metrics may assist in identifying differences between translations, but similarity scores alone cannot establish interpretive correctness or theological adequacy. Accordingly, the framework proposed here treats linguistic and contextual analysis as components of translation evaluation while recognizing that theological questions require appropriate scholarly consideration.

A more suitable evaluation method should combine automatic measurements with expert assessment. This article proposes six criteria:

- Semantic adequacy: Whether the translation preserves the main meaning without adding unsupported meanings.
- Linguistic acceptability: Whether the translated text is grammatically correct and easy to understand.
- Contextual and rhetorical fidelity: Whether the translation reflects the meaning and function of the expression within its context.
- Terminological consistency: Whether important and repeated terms are translated consistently.
- Source traceability: Whether explanations and interpretive claims can be linked to identifiable and reliable sources.
- Interpretive transparency: Whether the translation clearly distinguishes the original wording from commentary and indicates uncertainty where necessary.

For future empirical research, each criterion should be clearly defined in a scoring guide before the translations are evaluated. The study should identify the selected verses, target languages, and published reference translations. It should also record the AI model and version, the date of use, the prompts, and the available settings. The reviewers should have relevant knowledge of Arabic linguistics and Qur'anic studies. The research should also explain the instructions given to reviewers and report the level of agreement between them. Automatic scores can provide additional evidence, but they should not replace expert evaluation.

This framework also recognizes that differences between translations are not always errors. Translators may use different expressions because they follow different translation methods or address different audiences. The main task is to distinguish acceptable differences from the loss of important meanings, changes in interpretation, unsupported additions, and inconsistent use of key terms.

Institutional Responsibility and Ethical Governance

The reliability of digital Qur'anic translation depends on more than the ability of an AI model. It also depends on the quality and origin of the source texts, the rules governing their use, editorial supervision, version control, and clear responsibility for the published translation. Institutions involved in printing, translating, and sharing the Qur'an digitally can help by providing verified texts and properly documented translations. Open linguistic resources can also support research, provided that researchers follow their terms of use and cite them correctly.

A responsible publication process should clearly state whether AI was used to prepare or revise a translation. It should identify the sources consulted, keep records of corrections made by human reviewers, and avoid presenting unreviewed AI-generated text as an authoritative translation. If an AI system produces explanations similar to Qur'anic commentary, editors should check every reference against the original source. Copyright, licensing, privacy, and the risk of spreading incorrect religious information must also be considered.

Responsibility should be shared among developers, researchers, and publishers, with clear duties for each group. Developers should explain the limitations of their systems and how they handle sources. Researchers should describe their methods clearly and avoid making claims that are not supported by evidence. Publishers should establish suitable review procedures, especially for sensitive religious content. Human review should not be treated as a simple final formality. It should be a careful and accountable process in which qualified reviewers compare the AI-generated translation with the original Arabic text and relevant scholarly sources.

A Proposed Human-In-The-Loop Workflow

Based on the above discussion, this article proposes a five-stage process for responsible AI-assisted Qur'anic translation.

- The translation team should identify the target audience, purpose, and language. It should also decide how to distinguish the translated text from explanatory notes.
- The team should prepare a verified set of sources, including the Arabic Qur'anic text, relevant linguistic resources, and clearly identified editions of Qur'anic commentaries.
- An AI system may be used to generate possible translations or alternative expressions. Where appropriate, it should also retrieve information from reliable sources.
- A qualified reviewer should examine the output for accuracy of meaning, grammatical correctness, rhetorical context, consistency in the translation of key terms, and reliability of references.
- The team should record all revisions, conduct a final review, and publish the translation with a clear explanation of the method used and its limitations.

This process is a proposal and has not yet been tested as an established standard. It should therefore be evaluated through a controlled study comparing different translation methods. For example, researchers could compare AI-generated translations without external support, AI-generated translations supported by retrieved sources, human translations, and AI translations revised by human experts.

The proposed workflow is a methodological recommendation rather than a tested technical system. Its potential value lies in organizing the translation process around verified source materials, explicit analytical criteria, traceable interpretive claims, and qualified human review. Whether this approach improves translation quality remains an empirical question. Future studies should therefore compare translations generated under clearly documented conditions, including systems operating without retrieved sources, systems using retrieval-augmented generation, and AI-generated translations revised by qualified reviewers. The resulting translations should be evaluated using predefined criteria and independently assessed by appropriately qualified experts. Such research would provide evidence about the contribution of source retrieval and human review without assuming in advance that either intervention necessarily improves every aspect of translation quality.

DISCUSSION

The analysis leads to three main observations:

- The main challenge in AI-assisted Qur'anic translation is not simply producing grammatically correct sentences. It also involves preserving word meanings, context, rhetorical features, and the interpretations developed in Islamic scholarly traditions.
- Digital linguistic resources and retrieval-based systems can make reliable sources easier to access and help readers check the information used. However, these tools cannot independently resolve all questions of interpretation.
- A reliable evaluation method must give an important role to expert judgement rather than depending only on automatic evaluation scores.

These observations suggest that AI should mainly be used as a supporting tool in Qur'anic translation, rather than as an independent authority. It can help prepare initial drafts, explore alternative translations of important terms, and access relevant reference materials more quickly. However, its output must still be reviewed by people with suitable knowledge of Qur'anic Arabic, relevant Qur'anic commentaries, and the target language. This position does not reject technological development. Instead, it explains the conditions needed to use AI in a transparent and responsible way.

The main limitation of this study is its illustrative nature. It does not provide a reproducible comparison of current AI systems. Therefore, it cannot establish error rates, rank different models, or prove that retrieval-

augmented generation (RAG) performs better than other methods in Qur'anic translation. The proposed evaluation framework should be understood as a starting point for future empirical research, not as a presentation of experimentally measured results.

Thus, the contribution of this study is to establish a conceptual basis for examining the relationship between Qur'anic linguistic and rhetorical features and the evaluation of AI-assisted translation. The three expressions from Surah al-Kawthar demonstrate how translation analysis can consider lexical meaning, contextual interpretation, and rhetorical function without treating them as interchangeable. The proposed evaluation criteria and human-in-the-loop workflow provide a structure that future empirical research can apply to actual AI-generated and published translations. However, the present study does not establish the comparative effectiveness of particular models or translation methods. Its findings should therefore be understood as analytical observations and methodological proposals that require further empirical investigation.

CONCLUSION AND RECOMMENDATIONS

Artificial intelligence (AI) offers useful opportunities for translating and sharing the meanings of the Qur'an. It can help prepare translation drafts, make the text more accessible in different languages, and support the use of digital linguistic resources. However, fluent language does not necessarily mean that a translation fully conveys the intended meaning or can be considered authoritative. Qur'anic translation involves decisions related to vocabulary, grammar, rhetoric, and Qur'anic interpretation. These decisions require clear methods and careful review by qualified experts.

The examples from Surah al-Kawthar show that a single translated word may express an important part of the original meaning while leaving out some contextual or interpretive details. The solution is not always to make the translated verse longer. In some cases, it is better to keep the translation simple and provide additional explanations in a separate note supported by reliable sources. Retrieval-augmented generation (RAG) may help readers identify and check sources, but it cannot guarantee that the retrieved information is relevant or correctly explained by the AI system.

This article recommends developing well-documented multilingual Qur'anic resources, establishing transparent methods for AI-assisted translation, and creating evaluation systems that combine automatic measurements with expert review. Future studies should clearly report their data selection, prompts, model versions, reference translations, and evaluation criteria. Cooperation among Qur'anic scholars, translators, computational linguists, and publishers will also be important.

Ultimately, AI-assisted Qur'anic translation should be approached as a process requiring linguistic analysis, reliable interpretive sources, and qualified human review. The examples examined in this study illustrate why translation quality cannot be judged solely by grammatical fluency or lexical similarity. The proposed framework offers a basis for examining semantic adequacy, contextual and rhetorical fidelity, source traceability, and interpretive transparency.

Nevertheless, the framework has not yet been empirically validated, and the present study does not establish the performance of any specific AI system. Future research should apply these criteria to documented AI-generated and published translations to determine how different translation methods represent the linguistic, rhetorical, and interpretive features of Qur'anic expressions.

REFERENCES

1. Quranic Arabic Corpus. (n.d.). Documentation. https://corpus.quran.com/ documentation/
2. Quranic Arabic Corpus. (n.d.). Ontology of Quranic Concepts**.** https://corpus.quran.com/ontology.jsp
3. Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems, 33, 9459–9474. https://arxiv.org/abs/2005.11401

4. Papineni, K., Roukos, S., Ward, T., & Zhu, W.-J. (2002). BLEU: A method for automatic evaluation of machine translation. Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, 311–318. https://doi.org/10.3115/1073083.1073135
5. Rei, R., Stewart, C., Farinha, A. C., & Lavie, A. (2020). COMET: A neural framework for MT evaluation. Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, 2685–2702. https://doi.org/10.18653/v1/2020.emnlp-main.213
6. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. Advances in Neural Information Processing Systems, 30. https://papers.nips.cc/paper/7181-attention-is-all-you-need
7. Saleh Ibrahim Al-R; Jihad Muhammad Al-Nasser, Surah al-Kawthar: A Stylistic Study, IUG Journal of Islamic Studies, 2023, v. 31, n. 3, p. 501, doi. 10.33976/IUGJIS.31.3/2023/21
8. Haqa Abdul Karim, The Beauty of Qur'anic Language in Surah al-Kawthar, Reader for Literary, Critical and Linguistic Studies, Vol. 5, No.3, Pg 146-157, 30/09/2022, https://asjp.cerist.dz/en/article/201449?utm