

Structured Prompting for Arabic Essay Proficiency: A Trait-Centric Evaluation Approach

Salim Al Mandhari¹, Hieu Pham Dinh², Mo El-Haj^{1,2}, Paul Rayson^{1,2}

¹Lancaster University, UK

²VinUniversity, Vietnam

s.m.almandhari, m.el-haj, p.rayson@lancaster.ac.uk

24hieu.pd, elhaj.m, paul.r2@vinuni.edu.vn

Abstract

This paper presents a novel prompt engineering framework for trait specific Automatic Essay Scoring (AES) in Arabic, leveraging large language models (LLMs) under zero-shot and few-shot configurations. Addressing the scarcity of scalable, linguistically informed AES tools for Arabic, we introduce a three-tier prompting strategy (standard, hybrid, and rubric-guided) that guides LLMs in evaluating distinct language proficiency traits such as organization, vocabulary, development, and style. The hybrid approach simulates multi-agent evaluation with trait specialist raters, while the rubric-guided method incorporates scored exemplars to enhance model alignment. In zero and few-shot settings, we evaluate eight LLMs on the QAES dataset, the first publicly available Arabic AES resource with trait level annotations. Experimental results using Quadratic Weighted Kappa (QWK) and Confidence Intervals show that Fanar-1-9B-Instruct achieves the highest trait level agreement in both zero and few-shot prompting (QWK = 0.28 and CI = 0.41), with rubric-guided prompting yielding consistent gains across all traits and models. Discourse-level traits such as Development and Style showed the greatest improvements. These findings confirm that structured prompting, not model scale alone, enables effective AES in Arabic. Our study presents the first comprehensive framework for proficiency oriented Arabic AES and sets the foundation for scalable assessment in low resource educational contexts.

Keywords: Arabic AES, Prompt engineering, Trait evaluation

1. Introduction

AES has emerged as a critical tool in educational assessment, enabling scalable and consistent evaluation of language performance. Although substantial progress has been made in AES for high-resource languages such as English, development for Arabic remains limited by the scarcity of annotated datasets, the linguistic complexity of Arabic, and the lack of tools assessing proficiency based on distinct linguistic skills. Existing Arabic AES systems predominantly focus on holistic scoring or content correctness, often neglecting nuanced aspects of proficiency such as organization, vocabulary richness, idea development, and stylistic appropriateness (El-Haj et al., 2024; Almujaivel et al., 2025). Ghazawi and Simpson (2024), for example, demonstrated the use of AraBART for Arabic AES, achieving high accuracy in content-based scoring. However, their focus was limited to correctness of responses, with no evaluation of traits or linguistic proficiency. Conversely, English AES research has advanced toward trait scoring and prompt-based evaluation using LLMs (Lee et al., 2024; Bashendy et al., 2024), yet these innovations have not been fully extended to Arabic due to data and methodological constraints.

To address this gap, our research investigates how prompt engineering techniques including zero-shot and few-shot configurations can be applied to LLMs for trait-specific scoring of Arabic student

essays. Rather than relying on fine-tuning or training custom models, we propose a flexible, rubric-informed prompting strategy designed to work directly with existing LLMs¹

Our contributions are as follows:

- Our study is the first to evaluate trait-specific prompt engineering strategies on the QAES dataset, using a linguistically aligned hybrid evaluation framework.
- We introduce a novel three-tier prompting framework for AES in Arabic, composed of: (1) Standard prompts for holistic evaluation, (2) Hybrid trait-based prompts, wherein specialist raters assess specific linguistic features (e.g., organization, vocabulary, tone, development), and (3) Rubric-informed prompts that guide the model using explicit scoring criteria for each trait.
- We implement few-shot learning for trait-level prompting, providing example-scored responses to calibrate the model's judgment. This is, to our knowledge, the first application of few-shot prompting for trait-specific scoring in Arabic.

¹Code and prompts are available at https://github.com/dinhieufam/Arabic_AES/tree/master.

- We evaluate our framework on the QAES dataset, a trait-annotated Arabic AES corpus, and demonstrate competitive alignment with human scores across multiple traits in zero and few-shot settings.
- We provide qualitative insights into LLM performance variability across traits and prompt types, revealing which linguistic features are most robustly evaluated by current models.

This work demonstrates the feasibility of prompt-based scoring as a scalable and adaptable alternative to traditional supervised AES models, particularly in low-resource settings. It also provides a foundation for the future development of LLM-driven educational tools tailored to the Arabic language.

2. Related Work

AES has evolved from rule-based and statistical models to transformer-based language models and advanced prompt engineering. This section categorizes prior work by modeling strategies, dataset development, and prompt-based innovations, emphasizing efforts in Arabic AES.

Early AES systems for Arabic utilized string and vector similarity methods. For instance, [Abbas and Al-qaza \(2014\)](#) developed a Vector Space Model (VSM) to compare 30 essays with reference texts, which was later extended with Latent Semantic Indexing ([Abbas and Al-qazaz, 2015](#)). While foundational, these approaches lacked granularity and linguistic depth. [Shehab et al. \(2018\)](#) expanded AES with algorithms such as N-gram and DISCO, achieving moderate accuracy (up to 82%) on short Arabic responses. [Lofy et al. \(2023\)](#) further enhanced scoring reliability using ensemble models including Random Forest and Adaboost.

Transformer models introduced contextual depth to AES. [Ghazawi and Simpson \(2024\)](#) demonstrated this by leveraging AraBART for Arabic AES, achieving 0.88 overall QWK score after training the model across domains using a dataset of over 2,000 essays. Their focus, however, was on evaluating answer correctness in subject-specific contexts rather than linguistic proficiency. Our work, in contrast, centers on assessing language skills through multi-trait analysis. Sentence-level semantic modeling has been effectively explored by [Ramesh and Sanampudi \(2024\)](#), who employed sentence-BERT embeddings with Bi-LSTM and CNNs to capture coherence and cohesion in English AES. The highest QWK score achieved was 0.75. [Settles](#)

[et al. \(2020\)](#) combined psychometric modeling (IRT and CAT) with linear ML to adaptively assess language ability in the Duolingo exam framework.

Prompt-based AES has gained traction as a fine-tuning-free method. [Lee et al. \(2024\)](#) introduced a zero-shot prompting framework called Multi-Trait Specialization (MTS), enabling LLMs like ChatGPT and LLaMA to achieve competitive trait-level scoring of the TOFEL and ASAP English datasets, gaining QWK scores of 0.43 and 0.35 respectively. The TRATES framework proposed by [Eltanbouly et al. \(2025\)](#) translated rubrics into guiding questions, improving the scoring alignment of traits using Starling and Gemma models to the average QWK score of 0.59 on the ASAP and ELLIPSE English datasets. [Mansour et al. \(2024\)](#) tested a four-tier prompt design on English AES using ChatGPT and LLaMA but found these underperformed compared to fine-tuned models, achieving QWK score around 0.57. They did not experiment with few-shot prompting or linguistic trait specialization. [Bashendy et al. \(2024\)](#) implemented feature-based (LR), holistic SOTA (NPCR), single-task learning (STL) and multi-task learning (MTL) models for both holistic and trait scoring using AraBERT on the QAES dataset. The highest QWK score was 0.26 achieved by LR, on the holistic scoring,. This finding can be considered as a baseline for our study.

Although no work has been found in literature on prompting design for Arabic proficiency scoring, our approach builds on these insights by proposing a three-tier prompting framework: standard, hybrid, and rubric-based for Arabic AES. We incorporate few-shot examples and trait-decomposed scoring, thereby addressing both scoring fidelity and pedagogical interpretability without model fine-tuning.

The Arabic AES domain has been limited by resource scarcity. QCAW, introduced by [Zaghouni et al. \(2024\)](#), provides a bilingual corpus of argumentative essays. QAES, built on QCAW, is the first publicly available Arabic dataset annotated by traits such as organization, vocabulary, and grammar ([Bashendy et al., 2024](#)). Trait rubrics were derived from official scoring guidelines by the Qatar University Testing Center. While the only study using QAES relied on logistic regression and AraBERT embeddings ([Bashendy et al., 2024](#)), it did not explore prompt-based or few-shot approaches.

3. Methodology

This study proposes a three-level prompting framework for Arabic AES using large language models (LLMs) in zero-shot and few-shot configurations. Each level is designed to improve linguistic fidelity

and alignment with trait-based scoring rubrics, without requiring any model fine-tuning. Table 1 illustrates the framework structure.

Level	Approach	Key Mechanism
1	Standard Prompting	Direct scoring of all traits via a structured zero-shot prompt
2	Hybrid Trait Prompting	Multi-agent simulation using domain-specialist raters
3	Rubric-Guided Few-Shot Prompting	Trait-specific prompts with rubric guides and scored examples

Table 1: Overview of Prompting Framework

The standard prompting level 1 initiates zero-shot prompting by instructing the LLM to assign a score for each of the seven linguistic traits directly based on the student’s essay. The prompt template calls for numeric trait scores (e.g., for organization, vocabulary, etc.) and a summed total score. No prior examples or rubrics are provided, and the prompt is general across essays.

In the hybrid trait prompting with specialized raters - Level 2, each essay is evaluated by five simulated raters, with each rater specializing in a distinct linguistic trait. Each rater is responsible for scoring only the rubric dimensions that fall within their area of expertise. Table 2 presents the five rater specializations. The rubric traits are then mapped to the corresponding rater scores as shown in Table 3, and the final score is computed by averaging the scores assigned by all relevant raters. A single rater may influence one or more rubric traits, depending on their specialization. This approach ensures that the language model concentrates on a specific linguistic skill during evaluation, thereby simulating the behavior of an expert rater. For instance, the vocabulary specialist focuses exclusively on vocabulary-related aspects and disregards other dimensions such as content development or essay organization. Equation 1 highlights that the score for each rubric is calculated:

$$S_j = \frac{1}{|R_j|} \sum_{i \in R_j} s_{ij} \quad (1)$$

Where:

- S_j : Final score for rubric j
- R_j : Set of raters assigned to rubric j
- s_{ij} : Score given by rater i for rubric j

This structure promotes focused evaluation while maintaining multi-perspective judgment

Rater	Specialization	Evaluation Focus
A	Organization & Coherence	Logical flow, paragraph transitions, structural clarity
B	Vocabulary & Lexical Variety	Word choice, lexical diversity, sophistication, repetition
C	Grammar, Spelling & Mechanics	Punctuation, syntax, spelling, readability
D	Content Development & Reasoning	Argument quality, elaboration, evidence use
E	Style, Tone & Contextual Appropriateness	Voice, stylistic consistency, audience alignment

Table 2: The Five Rater Specializations

Rubric	Assigned Raters
Organization	A, D, C
Vocabulary	B, E, C
Style	B, E, C
Development	D, A, B
Mechanics	C
Structure	A, B, C
Relevance	D, B, E

Table 3: Rubric-to-Rater Mapping

through overlapping expertise.

Moving to the last prompting level, The rubric-guided few-shot prompting level (level 3) introduces rubric-informed, trait-specific prompts paired with scored examples. For each trait, the model is prompted with a detailed evaluation guide in Arabic, along with a file containing three examples illustrating low, mid, and high scores (e.g., 1, 3, 5). The model compares the target essay to the rubric and examples and produces a trait score with justification in structured JSON format. This method enables greater alignment between model scoring and human evaluation standards, reinforcing clarity, consistency, and explainability across traits. Appendix A highlights the 3 prompting levels designed for the experiments.

Although level 1 and level 3 prompts have already been applied in the literature, they are new to Arabic proficiency scoring. The main novelty lies in Prompt 2, which is inspired by human assessment. This reduces the model load and simulates specialist raters. Unlike previous work such as Abbas and Al-qaza (2014), Abbas and Al-qazaz (2015), and Ghazawi and Simpson (2025), where the prompts are unpartitioned (all traits scored in one pass), our method distributes the responsibilities between the simulated raters. The appendix in the camera-ready version will highlight the 3 prompting levels designed for the experiments.

4. Experimental Setup

4.1. Dataset and Models

We evaluated our models using the QAES dataset (Bashendy et al., 2024). QAES extends the Qatar Corpus of Argumentative Writing (QCAW) by providing fine-grained scores for seven linguistic traits: Organization, Vocabulary, Style, Development, Mechanics, Structure, and Relevance. Trait scores range from 0–5 (except Relevance, which is 0–2), with a total possible score of 32, as shown in Table 4.

Essay	Org.	Vocab.	Style	Dev.	Mech.	Struct.	Rel.	Final
1	4	4	4	4	4	4	2	26
2	4	4	4	4	4	4	2	26
3	5	4	4	4	4	4	2	27
4	5	5	4	4	4	4	2	28

Table 4: Sample trait annotations for 10 essays in the QAES dataset (Bashendy et al., 2024).

Trait annotations were produced by trained human raters following Qatar University’s standardized writing rubrics. Inter-rater agreement was verified through Cohen’s kappa analysis, achieving about 0.72 as an average of the two tasks scored by the two main annotators. (Bashendy et al., 2024). We consider this substantial agreement as a threshold for acceptable model performance in our study. The dataset consists of 195 essays written by undergraduate students in response to argumentative prompts in educational settings, primarily focused on technology and communication. Although this dataset is quite small, it was chosen for this project because it is the only dataset available with annotated traits for the assessment of language proficiency in Arabic AES. This resource provides a rich foundation for both zero-shot and few-shot AES modeling in Arabic. Other synthetic data by Qwaider et al. (2025) was explored also but found less representative of authentic learner writing. To mitigate the limited dataset, we tested eight LLMs of varying scales across three prompting strategies, focusing on trends consistent across models rather than absolute scores.

Turning to the models, we tested our prompting framework on 8 LLMs based on: (1) model size variety, (2) different foundation bases, and (3) considering efficiency reporting in the literature. We used the corresponding checkpoints for each model available in Hugging Face. Consequently, we selected eight models as follows: ChatGPT-4 (Achiam et al., 2023); Fanar-1-9B-Instruct (Team et al., 2025); Jais-family-13B-Chat (Sengupta et al., 2023); ALLAM-7B-Instruct (Bari et al., 2024); Qwen1.5-1.8B-Chat (Bai et al., 2023); Qwen2.5-7B-Instruct (Ahmed et al., 2025); Qwen3-VL-8B-Instruct (Team, 2024); and LLaMA-2-7B-Chat-hf (Touvron et al., 2023). These models were selected to test a spectrum of

capabilities in open-source and proprietary LLMs with varying levels of Arabic support and instruction tuning.

4.2. Experimental Protocol

Our research used the experimental protocol for the three prompting levels described in Section 3. Each essay is evaluated under the same conditions and the predicted scores are compared against the final human annotations. All prompts are in Arabic with rubric-aligned instructions. Each level is implemented as follows:

- Level 1: Single-pass, full-trait zero-shot evaluation.
- Level 2: Hybrid multi-agent simulation with trait-specialist raters and rubric-based averaging.
- Level 3: Few-shot prompting with trait-specific rubrics and exemplar comparisons.

Predictions are stored per trait in structured JSON/CSV outputs. All models are evaluated using identical input texts.

For the evaluation metric, we used the Quadratic Weighted Kappa (QWK) to measure the agreement between the predicted models and the human gold scores. QWK is defined in Equation 2:

$$\kappa = 1 - \frac{\sum_{i,j} w_{i,j} O_{i,j}}{\sum_{i,j} w_{i,j} E_{i,j}}, \quad w_{i,j} = \frac{(i-j)^2}{(N-1)^2} \quad (2)$$

Where:

- $O_{i,j}$: observed agreement matrix
- $E_{i,j}$: expected matrix
- $w_{i,j}$: penalty for rating difference $|i-j|$
- N : number of possible score levels for each trait

This metric is widely regarded the AES standard and is used to assess the severity of misclassification on an ordinal grading scale. Its use was thoroughly examined by Doewes et al. (2023) and Ghazawi and Simpson (2024), who also discuss its sensitivity to rating scales and best practices for its interpretation.

5. Results and Analysis

In this section, we discuss our baseline performance by addressing various scoring dimensions and calculating the QWK score. We compare the

3 prompting levels across the 8 models to demonstrate how our prompt engineering logic effectively enhances the ability of LLMs to assess Arabic language proficiency using zero- and few-shot settings without fine-tuning. In our experiments, both rubric-based and total scores are considered for each essay to evaluate the performance of each model in terms of linguistic skill rubrics and overall assessment quality. Table 5 presents the QWK scores across all language traits and the total score for the 8 evaluated models under different prompting levels. It highlights model differences: stronger agreement on Organization, Vocabulary, Development, Mechanics, and Relevance by Fanar and Qwen3-VL. It also shows slight agreement on Style and Structure by the same models. However, smaller models such as Qwen1.5, Qwen2.5, and LLaMA-2, along with models like Jais and ChatGPT-4, exhibit consistently low agreement across all traits. ALLaM performs best among all models in evaluating Structure. Overall, the agreement levels can generally be described as low to fair.

5.1. Trait and Prompt-Level QWK Scores and Confidence Intervals

To provide insight beyond a single estimate of QWK scores, we calculate the confidence interval (CI) of all traits across the models and prompting levels applied. Figure 1 presents confidence intervals of each trait, along with its QWK score.

The CI offers a statistical range around the observed QWK score that likely contains the true agreement level between the model and human annotations. Specifically, we compute the 95% CI using non-parametric bootstrapping with 1,00 iterations per trait-model-prompt group. This method captures the variability in QWK scores due to the sample of essays evaluated. The CI bounds are calculated as:

$$CI_{95\%} = [QWK_{2.5}, QWK_{97.5}] \quad (3)$$

Where: $QWK_{2.5}$ and $QWK_{97.5}$ denote the 2.5th and 97.5th percentiles of the bootstrapped QWK score distribution, respectively (Efron and Tibshirani, 1993).

Key Observations:

- **Performance Trends:** Across traits, Fanar and Qwen3-VL consistently achieve the highest QWK scores, primarily under `prompt_level_3`. Vocabulary and Style evaluation in particular is more responsive to `prompt_level_1`. Their CIs are relatively narrow, indicating stable performance.
- **Prompting Effect:** All models show notable improvement from `prompt_level_1` to

`prompt_level_3`, confirming that more detailed prompts enhance scoring reliability.


- **Trait Variability:** Traits such as *Vocabulary* and *Relevance* yield higher agreement across models, whereas *Organization*, *Mechanics*, and *Development* exhibit lower QWK scores and wider CIs. Focusing on specific model performance, *Mechanics* appears to be the most responsive trait, with Fanar achieving the highest confidence interval upper bound of **0.41**.
- **Low Performing Models:** Smaller models such as Qwen1.5, Qwen2.5, and Jais perform inconsistently, often scoring near or below zero on several traits with large CIs, suggesting poor agreement and higher uncertainty.

Comparing our findings to the inter-rater agreement threshold and existing baselines, Table 6 summarizes the comparative results. The data show that prompt engineering yields a modest improvement in QWK scores, although the results remain considerably below the threshold. Additionally, our study introduces the use of confidence intervals for QWK in Arabic AES, achieving a peak CI upper bound of 0.41, which provides a more robust interpretation of model agreement.


To evaluate the effectiveness of different prompting strategies, we visualize the distribution of QWK scores across all models and prompt levels, along with their corresponding confidence intervals. Figure 2, 3, and 4 present bar plots summarizing QWK scores and 95% CIs aggregated over the seven linguistic traits and the total score, grouped by prompt level.

The results reveal a consistent and interpretable pattern: model performance improves as prompts become more structured and trait-specific. Prompt Level 1, which uses a single holistic instruction, exhibits the widest spread and lowest median scores. In contrast, Prompt Level 3, which leverages rubric-aligned, few-shot trait-specific instructions, achieves the narrowest confidence intervals and the highest QWK scores across most models. These findings underscore that prompting specificity, rather than model size, plays a more critical role in aligning with human-rated scoring.


The CI results reveal wide CI across several traits, particularly under Prompt Level 3. This indicates substantial uncertainty in the estimated QWK agreement and suggests high variability in the essay samples used to calculate CI. The wide intervals correspond to larger standard errors, reflecting instability in agreement estimates. Notably, the *Mechanics* trait exhibits comparatively narrower CI, suggesting more stable and consistent agreement.




(a) Development




(b) Mechanics



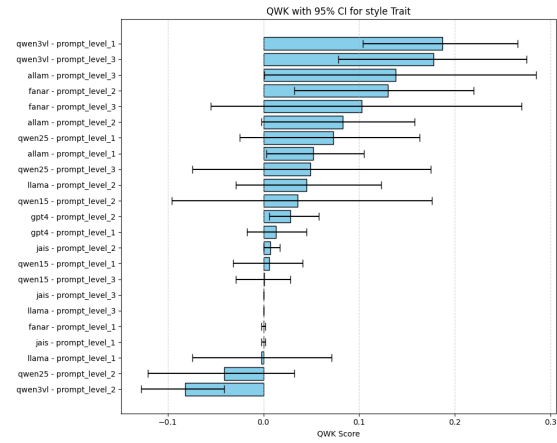
(c) Organization




(d) Relevance




(e) Structure



(f) Style



(g) Total



(h) Vocabulary

Figure 1: QWK Scores with 95% Confidence Intervals for All Traits

Model	Prompt Level	Org.	Vocab.	Style	Dev.	Mech.	Struct.	Rel.	Total
Qwen1.5-1.8B-Chat	L1	0.039	0.016	0.005	0.015	0.014	0.018	0.006	0.014
	L2	0.001	0.058	0.035	0.045	-0.017	0.024	0.115	0.045
	L3	-0.003	-0.033	0.001	0.002	0.040	-0.020	0.051	0.006
Qwen2.5-7B-Instruct	L1	0.015	0.016	0.074	0.015	0.067	0.007	0.00	0.021
	L2	0.157	-0.028	-0.040	0.004	-0.023	-0.000	-0.026	0.067
	L3	0.099	0.101	0.049	-0.010	0.064	0.072	-0.071	0.040
Qwen3-VL-8B-Instruct	L1	0.086	0.220	0.191	0.045	0.188	0.096	0.000	0.125
	L2	-0.059	-0.087	-0.082	-0.041	-0.035	-0.071	-0.014	-0.060
	L3	0.264	0.001	0.179	0.061	0.194	0.176	0.206	0.206
ALLaM-7B-Instruct-preview	L1	0.143	-0.037	0.051	-0.001	-0.031	0.015	0.021	0.039
	L2	0.052	0.092	0.084	0.102	0.035	0.072	0.180	0.118
	L3	0.071	0.205	0.141	0.193	0.117	0.198	0.141	0.264
Jais-family-13B-Chat	L1	-0.006	0.002	0.000	0.026	0.010	-0.006	0.077	0.007
	L2	0.006	0.004	0.006	0.009	0.006	0.003	0.003	0.008
	L3	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Fanar-1-9B-Instruct	L1	0.015	0.000	0.000	0.000	0.012	0.013	0.000	0.008
	L2	0.103	0.172	0.132	0.044	0.103	0.163	0.177	0.161
	L3	0.142	0.205	0.105	0.204	0.280	0.123	0.227	0.284
LLaMA-2-7B-Chat-HF	L1	0.000	0.030	-0.000	0.093	-0.001	0.047	0.073	0.033
	L2	0.015	0.039	0.044	0.091	-0.000	0.034	0.030	0.046
	L3	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
ChatGPT-4	L1	0.090	0.043	0.012	0.055	0.060	0.035	0.000	0.050
	L2	0.025	0.033	0.027	0.015	0.043	0.034	0.021	0.034
	L3	-	-	-	-	-	-	-	-

Table 5: Quadratic Weighted Kappa (QWK) scores for each language trait and total score across all models and prompt levels. GPT-4 results for Prompt Level 3 are omitted due to the high computational cost of running extended evaluations.

Metric	Baseline	Our Score	Threshold
QWK	0.26	0.28	0.72
CI Upper	-	0.41	-

Table 6: Comparison of QWK scores and confidence intervals against the threshold and baseline achieved by QAES.

This may be attributed to the surface-level nature of mechanical errors, which are less subjective than higher-level discourse traits. While reducing the confidence level to 90%, for example, or increasing the sample size would mechanically narrow the interval width, this would not resolve the underlying variability in agreement.

5.2. Interpretation and Discussion

These results demonstrate that prompt structure has a greater impact on AES performance than model size alone, particularly for complex linguistic traits. The most significant QWK improvements were observed with the hybrid prompting method operating in a zero-shot setting, as reflected in the following:




Figure 2: QWK Scores with 95% Confidence Intervals for Prompt Levels 1

- Most of the traits, Organization, Vocabulary, Style, and Relevance, as well as the total score, showed improvement in most LLM runs. In contrast, other traits such as Development, Mechanics, and Structure exhibited only moderate improvement.




Figure 3: QWK Scores with 95% Confidence Intervals for Prompt Level 2




Figure 4: QWK Scores with 95% Confidence Intervals for Prompt Level 3

- The Fanar-1-9B and ALLaM-7B models demonstrated the strongest instruction tuning and the most robust Arabic language coverage, with consistently better understanding and scoring in all traits.

When examining the effect of few-shot settings within Trait Prompting, the Structure trait was the only one to show notable improvement. In contrast, the impact on the other traits and the overall score was minimal.

Another notable observation is that models showing a performance drop across prompt levels also tend to have very low QWK scores. This supports the hypothesis that as prompt complexity increases, smaller models struggle more with comprehension. Interestingly, for these smaller models, scores for Prompt 2 are often higher than for Prompt 1. This could be because Prompt 1 requires 8 trait scores, while Prompt 2 requires only 5 (A, B, C, D, and E).

Additionally, for Prompt 3, both LLaMA and JAIS have a QWK score of 0, as they failed to generate meaningful outputs from which scores could be extracted. This further illustrates the difficulty these models face with more complex prompts.

Assessing language proficiency is considerably more complex than evaluating content correctness in AES. Ghazawi and Simpson (2025) reported a QWK of 0.45 for GPT-4 in zero-shot content scoring, while in our study, the same model achieved only 0.05 in Prompt Level 1 for trait-based proficiency scoring. This highlights the added difficulty of evaluating traits such as organization and development, which require deeper contextual and rubric-aligned understanding.

These findings support our hypothesis that prompt engineering, especially rubric-aligned and example-based prompting, can transform general-purpose LLMs into effective AES tools for Arabic, without any fine-tuning or domain adaptation.

6. Conclusion and Future Work

This paper proposed a prompt-based framework for trait-specific Arabic AES using large language models. Through a three-level prompting strategy, standard, hybrid, and rubric-guided, we showed that carefully designed prompts can elicit reliable trait-level scoring in zero and few-shot settings. Experiments on eight LLMs using the QAES dataset confirmed that prompt engineering improves performance, especially for small and medium models suitable for educational use where resources and privacy matter. Our results also reveal that traits such as development, organisation, and style remain challenging, requiring structured linguistic guidance rather than simple holistic scoring.

For future work, we plan to expand the dataset to improve robustness and generalisability and explore the transferability of our framework to non-Arabic AES tasks. We also aim to address methodological challenges such as rater disagreement and score aggregation in hybrid prompting. This work provides a foundation for scalable, low-resource AES solutions in Arabic and other underrepresented languages.

Limitations

Although our framework demonstrates strong performance in trait-specific Arabic AES, several limitations remain. The scarcity of annotated Arabic AES datasets restricts large-scale evaluation and limits linguistic and demographic coverage. In addition, reliance on proprietary models like GPT-4 incurs high cost, making large-scale deployment difficult in resource-constrained environments.

These limitations highlight the need for more accessible Arabic AES datasets and optimisation strategies for cost-effective implementation. Future work should also explore open-source alternatives and parameter-efficient adaptation techniques to improve practicality and reproducibility.

7. Bibliographical References

- Ayad R Abbas and Ahmed S Al-qaza. 2014. Automated arabic essay scoring (aaes) using vector space model (vsm). *Journal of Al-Turath University College*, 15:25–39.
- Ayad R Abbas and Ahmed S Al-qazaz. 2015. Automated arabic essay scoring (aaes) using vectors space model (vsm) and latent semantics indexing (lsi). *Eng. Technol. J.*, 33(3):410–426.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Imtiaz Ahmed, Sadman Islam, Partha Protim Datta, Imran Kabir, Naseef Ur Rahman Chowdhury, and Ahshanul Haque. 2025. Qwen 2.5: A comprehensive review of the leading resource-efficient llm with potential to surpass all competitors. *Authoria Preprints*.
- Sultan Almujaivel, Damith Premasiri, Tharindu Ranasinghe, Mo El-Haj, and Ruslan Mitkov. 2025. [Complex concept-based readability estimation from arabic curriculum](#). *ACM Trans. Asian Low-Resour. Lang. Inf. Process.* Just Accepted.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- M Saiful Bari, Yazeed Alnumay, Norah A Alzahrani, Nouf M Alotaibi, Hisham A Alyahya, Sultan Al-Rashed, Faisal A Mirza, Shaykhah Z Alsubaie, Hassan A Alahmed, Ghadah Alabduljabbar, et al. 2024. Allam: Large language models for arabic and english. *arXiv preprint arXiv:2407.15390*.
- May Bashendy, Salam Albatarni, Sohaila Eltanbouly, Eman Zahran, Hamdo Elhuseyin, Tamer Elsayed, Walid Massoud, and Houda Bouamor. 2024. Qaes: First publicly-available trait-specific annotations for automated scoring of arabic essays. In *Proceedings of The Second Arabic Natural Language Processing Conference*, pages 337–351.
- Afrizal Doewes, Nugthoh Kurdhi, and Akirati Saxena. 2023. Evaluating quadratic weighted kappa as the standard performance metric for automated essay scoring. In *16th International Conference on Educational Data Mining, EDM 2023*, pages 103–113. International Educational Data Mining Society (IEDMS).
- Bradley Efron and Robert J. Tibshirani. 1993. *An Introduction to the Bootstrap*. Chapman & Hall/CRC.
- Mo El-Haj, Sultan Almujaivel, Damith Premasiri, Tharindu Ranasinghe, and Ruslan Mitkov. 2024. [DARES: Dataset for Arabic readability estimation of school materials](#). In *Proceedings of the Workshop on DeTermIt! Evaluating Text Difficulty in a Multilingual Context @ LREC-COLING 2024*, pages 103–113, Torino, Italia. ELRA and ICCL.
- Sohaila Eltanbouly, Salam Albatarni, and Tamer Elsayed. 2025. Trates: Trait-specific rubric-assisted cross-prompt essay scoring. *arXiv preprint arXiv:2505.14577*.
- Rayed Ghazawi and Edwin Simpson. 2024. Automated essay scoring in arabic: a dataset and analysis of a bert-based system. *arXiv preprint arXiv:2407.11212*.
- Rayed Ghazawi and Edwin Simpson. 2025. How well can llms grade essays in arabic? *arXiv preprint arXiv:2501.16516*.
- Sanwoo Lee, Yida Cai, Desong Meng, Ziyang Wang, and Yunfang Wu. 2024. Unleashing large language models' proficiency in zero-shot essay scoring. *arXiv preprint arXiv:2404.04941*.
- Nourmeen Lotfy, Abdulaziz Shehab, Mohammed Elhoseny, and Ahmed Abu-Elfetouh. 2023. An enhanced automatic arabic essay scoring system based on machine learning algorithms. *CMC-COMPUTERS MATERIALS & CONTINUA*, 77(1):1227–1249.
- Watheq Mansour, Salam Albatarni, Sohaila Eltanbouly, and Tamer Elsayed. 2024. Can large language models automatically score proficiency of written essays? In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 2777–2786, Torino, Italia. ELRA and ICCL.
- Chatrine Qwaider, Bashar Alhafni, Kirill Chirkunov, Nizar Habash, and Ted Briscoe. 2025. Enhancing arabic automated essay scoring with synthetic data and error injection. *arXiv preprint arXiv:2503.17739*.

Dadi Ramesh and Suresh Kumar Sanampudi. 2024. Coherence-based automatic short answer scoring using sentence embedding. *European Journal of Education*, 59(3):e12684.

Neha Sengupta, Sunil Kumar Sahu, Bokang Jia, Satheesh Katipomu, Haonan Li, Fajri Koto, William Marshall, Gurpreet Gosal, Cynthia Liu, Zhiming Chen, et al. 2023. Jais and jais-chat: Arabic-centric foundation and instruction-tuned open generative large language models. *arXiv preprint arXiv:2308.16149*.

Burr Settles, Geoffrey T. LaFlair, and Masato Hagiwara. 2020. Machine learning-driven language assessment. *Transactions of the Association for Computational Linguistics*, 8:247–263.

Abdulaziz Shehab, Mahmoud Faroun, and Magdi Rashad. 2018. An automatic arabic essay grading system based on text similarity algorithms. *International Journal of Advanced Computer Science and Applications*, 9(3).

Fanar Team, Ummar Abbas, Mohammad Shahmeer Ahmad, Firoj Alam, Enes Altinisik, Ehsannedin Asgari, Yazan Boshmaf, Sabri Boughorbel, Sanjay Chawla, Shammur Chowdhury, Fahim Dalvi, Kareem Darwish, Nadir Durani, Mohamed Elfeky, Ahmed Elmagarmid, Mohamed Eltabakh, Masoomali Fatehkia, Anastasios Fragkopoulos, Maram Hasanain, Majd Hawasly, Mus'ab Husaini, Soon-Gyo Jung, Ji Kim Lucas, Walid Magdy, Safa Messaoud, Abubakr Mohamed, Tasnim Mohiuddin, Basel Mousi, Hamdy Mubarak, Ahmad Musleh, Zan Naeem, Mourad Ouzzani, Dorde Popovic, Amin Sadeghi, Husrev Taha Sencar, Mohammed Shinoy, Omar Sinan, Yifan Zhang, Ahmed Ali, Yassine El Kheir, Xiaosong Ma, and Chaoyi Ruan. 2025. [Fanar: An arabic-centric multimodal generative ai platform](#).

QwenLM Team. 2024. Qwen3-VL: Multimodal Large Language Model. <https://github.com/QwenLM/Qwen3-VL>. Accessed: October 20, 2025.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Wajdi Zaghouani, Abdelhamid Ahmed, Xiao Zhang, and Lameya Rezk. 2024. Qcaw 1.0: building a qatari corpus of student argumentative writing. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 13382–13394.

Appendix

A. Prompting Templates

The following templates represent the three common prompting designs used to run all the LLMs.

```
SCORING_PROMPT = """
```

```
You are an expert Arabic language evaluator. Your task is to assess the proficiency of an Arabic essay based on seven traits:
```

1. Organization (0-5): How well-structured and coherent is the essay?
2. Vocabulary (0-5): Does the writer use a rich and appropriate vocabulary?
3. Style (0-5): Is the writing engaging, fluent, and stylistically appropriate?
4. Development (0-5): Are ideas elaborated with sufficient details and examples?
5. Mechanics (0-5): Are grammar, spelling, and punctuation correct?
6. Structure (0-5): Does the essay follow proper syntactic structures?
7. Relevance (0-2): Does the essay address the given topic appropriately?
8. Final Score (0-32): The sum of all the scores.

```
Each trait should be scored on a scale from 0 (poor) to 5 (excellent), except for relevance which is scored on a scale from 0 (poor) to 2 (excellent).
```

```
The final score should be the sum of all the scores on a scale from 0 to 32.
```

```
Return ONLY this JSON object with your scores (replace X with actual numbers):
```

```
{  
  "organization": X,  
  "vocabulary": X,  
  "style": X,  
  "development": X,  
  "mechanics": X,  
  "structure": X,  
  "relevance": X,  
  "final_score": X
```

```
}
```

```
"""
```

Figure 5: Standard Prompt - Level 1

```

RATER_SPECIALIZATIONS = {
    "A": "Evaluate the essay's organization and how well ideas are connected. "
        "Consider logical flow, paragraph structure, and transitions. "
        "Provide a score between 0-5",
    "B": "Assess the vocabulary quality and lexical variety in the essay. "
        "Consider word choice, sophistication, and avoidance of repetition. "
        "Provide a score between 0-5",
    "C": "Evaluate grammar, spelling, punctuation, and mechanical accuracy. "
        "Identify errors and assess their impact on readability. "
        "Provide a score between 0-5",
    "D": "Analyze content development and reasoning quality. "
        "Consider depth of analysis, argument strength, and evidence use. "
        "Provide a score between 0-5",
    "E": "Assess style, tone, and contextual appropriateness. "
        "Consider voice, audience awareness, and stylistic effectiveness. "
        "Provide a score between 0-5"
}

RUBRIC_MAPPING = {
    "organization": ["A", "D", "C"],
    "vocabulary": ["B", "E", "C"],
    "style": ["B", "E", "C"],
    "development": ["D", "A", "B"],
    "mechanics": ["C"],
    "structure": ["A", "B", "C"],
    "relevance": ["D", "B", "E"]
}

```

Figure 6: Hybrid Trait Prompt - Level 2 (first part)

```

prompt = f"""
You are an Arabic essay scoring assistant. You will read a student's Arabic
essay and assign scores from 0 to 5 for the following five linguistic dimensions:

A: {RATER_SPECIALIZATIONS['A']}
B: {RATER_SPECIALIZATIONS['B']}
C: {RATER_SPECIALIZATIONS['C']}
D: {RATER_SPECIALIZATIONS['D']}
E: {RATER_SPECIALIZATIONS['E']}

Return ONLY this JSON object with your scores (replace X with actual numbers):
{{
    "A": X,
    "B": X,
    "C": X,
    "D": X,
    "E": X
}}
"""

```

Figure 7: Hybrid Trait Prompt - Level 2 (second part)

```

RUBRIC_GUIDES = {
  "organization": {
    "arabic": "organization",
    "guide": "1. Is the essay well organized?\n2. Is there a clear introduction, body, and conclusion?\n3. Is the paragraph sequence logical?",
    "scoring": "0-5"
  },
  "vocabulary": {
    "arabic": "vocabulary",
    "guide": "1. How varied is the vocabulary?\n2. Is there repetition or inappropriate usage?\n3. Is the vocabulary precise and meaningful?",
    "scoring": "0-5"
  },
  "style": {
    "arabic": "style",
    "guide": "1. Is the style appropriate and engaging?\n2. Is the expression smooth?\n3. Are the structures and techniques suitable?",
    "scoring": "0-5"
  },
  "development": {
    "arabic": "development",
    "guide": "1. Are the ideas supported with examples?\n2. Is there sufficient elaboration?\n3. Is the argument persuasive?",
    "scoring": "0-5"
  },
  "mechanics": {
    "arabic": "mechanics",
    "guide": "1. Are there grammatical or spelling errors?\n2. Is punctuation used correctly?",
    "scoring": "0-5"
  },
  "structure": {
    "arabic": "structure",
    "guide": "1. Are the sentences well-formed?\n2. Is the syntactic structure clear and controlled?",
    "scoring": "0-5"
  },
  "relevance": {
    "arabic": "relevance",
    "guide": "1. Does the essay address the topic of communication and technology?\n2. Does the overall idea align with the topic?",
    "scoring": "0-2"
  }
}

```

Figure 8: Rubric-Guided Few-Shot Prompting - Level 3 (first part)

```

return f"""
You are a linguistic evaluator specialized in assessing the [{rubric_info['arabic']}] skill in Arabic essays.

You will evaluate the essay based on this skill only.

Please follow these steps:
1. Read the essay carefully.
2. Refer to the following evaluation guide:
{rubric_info['guide']}
3. Assign a score from {rubric_info['scoring']}.
4. Provide clear justifications for your decision.

The rubric below includes:
- A description of what it measures
- Three example levels (scores: 1, 3, 5)
\\\\"
{example}
\\\\"
Use these examples to compare with the essay being evaluated and justify the score accordingly.

Essay:
\\\\"
{essay_text}
\\\\"

Respond in JSON format only, like this:
{{
  "score": X,
  "justification": "..."
}}
"""

```

Figure 9: Rubric-Guided Few-Shot Prompting - Level 3 (second part)