

VLCAP at ImageEval 2025 Shared Task: Multimodal Arabic Captioning with Interpretable Visual Concept Integration

Passant Elchafei

Ulm University, Germany
passant.elchafei@uni-ulm.de

Amany Fashwan

Alexandria University, Egypt
amany.fashwan@alexu.edu.eg

Abstract

We present VLCAP, an Arabic image captioning framework that integrates CLIP-based visual label retrieval with multimodal text generation. Rather than relying solely on end-to-end captioning, VLCAP grounds generation in interpretable Arabic visual concepts extracted with three multilingual encoders, mCLIP, AraCLIP, and Jina V4, each evaluated separately for label retrieval. A hybrid vocabulary is built from training captions and enriched with about 21K general domain labels translated from the Visual Genome dataset, covering objects, attributes, and scenes. The top- k retrieved labels are transformed into fluent Arabic prompts and passed along with the original image to vision–language models. In the second stage, we tested Qwen-VL and Gemini Pro Vision for caption generation, resulting in six encoder–decoder configurations. The results show that mCLIP + Gemini Pro Vision achieved the best BLEU-1 (5.34%) and cosine similarity (60.01%), while AraCLIP + Qwen-VL obtained the highest LLM-judge score (36.33%). This interpretable pipeline enables culturally coherent and contextually accurate Arabic captions.

1 Introduction

In today’s digital age, images are everywhere, shared across social media platforms, embedded in news articles, used in education, e-commerce, and communication. With the rapid growth of visual content, images have become a dominant form of information exchange and expression. This widespread presence highlights the need for intelligent systems that can understand, interpret, and describe visual content effectively (Gendy and Patel, 2024).

Image captioning is the task of automatically generating syntactically correct and semantically meaningful sentences that describe an image’s content. It plays a vital role in bridging the gap be-

tween visual content and natural language. Equipping machines with the ability to interpret and describe visual information offers numerous benefits, including enhanced information retrieval, support for early childhood education, assistance for visually impaired individuals, and applications in social media, among others. While understanding the content of an image may seem effortless, even for children, it remains a significant challenge for computers (Eljundi et al., 2020).

Vision-language (VL) models have significantly advanced image understanding and captioning tasks in English and other high-resource languages (Zhang et al., 2024). However, Arabic image captioning remains underexplored, particularly for culturally rich datasets requiring grounded and interpretable visual understanding. End-to-end generation models often fail to capture the fine-grained semantics, contextual nuances, and socio-cultural cues inherent in Arabic visual scenes.

This paper presents our submission to the ImageEval Shared Task, specifically the Image Captioning Models Evaluation Subtask. The objective of this subtask is to develop Arabic image captioning models capable of generating culturally relevant and contextually accurate image descriptions. To address this task, we propose VLCAP, a modular Arabic captioning pipeline that integrates visual label reasoning with multimodal generation. Unlike prior work, which mainly adapts English datasets or relies on end-to-end models, our approach explicitly grounds captioning in Arabic visual concepts. In the first stage, we conduct three separate experiments using CLIP-based encoders AraCLIP (Al-Barham et al., 2024), mCLIP (Chen et al., 2023), and Jina V4 (Günther et al., 2025) to extract the top- k Arabic visual labels for each image. These labels act as interpretable anchors representing “what is seen.”

In the second stage, we use these extracted labels to construct enriched prompts, which are then

combined with the original images and fed into two different vision–language models in separate experiments: Qwen-VL (Bai et al., 2023) and Gemini Pro Vision (Anil et al., 2023). This setup enables us to systematically evaluate which combination of label extraction model and caption generation model yields the most culturally aligned, semantically accurate, and fluent Arabic captions. By decoupling visual recognition from linguistic description, VLCAP improves both cultural relevance and model transparency.

The rest of this paper is organized as follows: Section 2 reviews the background and related work on Arabic image captioning and the dataset used in our study, while Section 3 presents VLCAP, our proposed Arabic vision–language captioning system, and describes the system overview. Section 4 discusses the results of our experiments. Finally, Section 5 concludes the paper.

2 Background

The ImageEval 2025 Shared Task focuses on evaluating image captioning models for the Arabic language, with two main subtasks: (1) Building an open-source dataset of images with culturally appropriate, naturally written Arabic captions, supporting the development of Arabic-native image captioning resources and (2) Automatic generation of Arabic captions for given images. We participated in Subtask 2: Image Captioning Models Evaluation, which requires participants to generate captions for a set of images in Arabic (Bashiti et al., 2025).

In Subtask 2, the input is a single image, and the output is a short, descriptive caption in *Arabic*. Captions are submitted in CSV format, with each row containing the image_id and the generated caption. The generated captions are evaluated against a hidden reference set using a combination of automatic metrics ROUGE, BLEU and LLM-as-a-judge to assess semantic similarity and overall quality. The images cover a broad range of everyday scenes, enabling the models to learn both literal and contextually enriched descriptions. The dataset is culturally relevant, reflecting both Modern Standard Arabic and occasional dialectal expressions.

The dataset used in this shared task is entirely in Arabic and consists of high-quality, manually curated captions. It includes a training set of 2,718 images with human-authored captions, a validation set of 76 images with undisclosed gold-standard

captions, and a test set of 753 images, also with undisclosed gold-standard captions.

Research on Arabic image captioning has been steadily growing, although it still lags behind progress in English captioning. Early work in the field primarily involved adapting English datasets by translating captions into Arabic or creating Arabic versions of existing datasets like Flickr8k and MS-COCO. (Al-muzaini et al., 2018) developed an Arabic image captioning dataset and implemented a model using a convolutional neural network (CNN) for image feature extraction and a recurrent neural network (RNN) with an LSTM decoder for sentence generation. These models laid the foundation for subsequent developments in Arabic captioning. (Eljundi et al., 2020) proposed an end-to-end model that directly transcribes images into Arabic text. They developed an annotated dataset for Arabic image captioning (AIC). They also developed a base model for AIC that relies on text translation from English image captions. The two models are evaluated with the new dataset, and the results show the superiority of their end-to-end model.

More recent studies have focused on building specialized datasets and improving model architectures. (Al-Malki and Al-Aama, 2023) built ‘ArabicFashionData’ dataset, which contains labeled images of clothing items. Using this data, researchers developed an attention-based encoder-decoder model that achieved a high BLEU-1 score of 88.52. (Za’ter and Talafha, 2022) highlighted the lack of standardized Arabic benchmarks and proposed unified datasets for evaluating multi-task learning approaches using pre-trained word embeddings, which showed moderate improvements in caption quality.

Transformer-based models have also gained traction. (Emami et al., 2022) explored Arabic image captioning using deep bidirectional transformers by integrating pre-trained language models into the generation process. (Alsayed et al., 2023) expanded on this by analyzing the impact of text preprocessing tools like CAMEL Tools and various image encoders such as ResNet152. Their experiments demonstrated substantial improvements in BLEU-4 scores up to 148% and their best-performing model outperformed existing approaches by 379%.

Vision-language models have introduced further advancements. The VIOLET model (Mohamed et al., 2023) combines a vision encoder with a Gemini text decoder. It leverages an automated

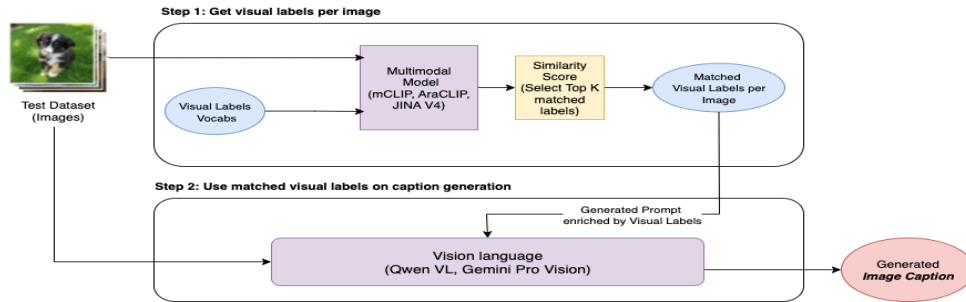


Figure 1: **VLCAP system overview.** The framework operates in two stages: (1) Arabic visual labels are retrieved by computing image–text similarity with a multilingual multimodal encoder (mCLIP, AraCLIP, or Jina V4) against a curated label vocabulary; (2) the retrieved labels are inserted into an Arabic prompt, which together with the original image, is passed to a vision–language model (Qwen-VL or Gemini Pro Vision) to generate the final caption.




Figure 2: **Visual Labels Vocab Builder.** Construction of the Arabic visual label vocabulary: Most frequent content words extracted from the training captions and augmented with general-domain visual concepts translated from the Visual Genome dataset, producing the final vocabulary used for label retrieval.

method to collect Arabic caption data from English sources, resulting in strong performance, including a CIDEr score of 61.2 on a manually annotated Arabic dataset.

In another recent approach, (Elbedwehy and Medhat, 2023) experimented with combining visual features extracted from powerful image encoders like SWIN, ConvNeXt, and XCIT, alongside Arabic pre-trained language models such as CAMeLBERT and MARBERTv2. This feature fusion strategy significantly improved the fluency and accuracy of the generated captions, outperforming earlier models.

3 System Overview

Unlike prior work that mainly adapts English datasets or relies on end-to-end models, our approach explicitly grounds captioning in Arabic visual concepts. We present **VLCAP**, a modular Arabic image captioning framework that decouples visual label extraction from caption generation to enhance cultural alignment, semantic accuracy, and interpretability. The system operates in two main stages as shown in Figure 1, supported by a comprehensive Arabic visual vocabulary constructed through the Visual Label Builder (Figure 2).

3.1 Arabic Visual Vocabulary Construction

As a preliminary step, we built a vocabulary of Arabic visual labels from two sources: (1) most frequent content words extracted from the training captions after removing Arabic stopwords, numbers, and very short tokens, and (2) an augmented set of 21,000 high-frequency visual concepts, covering objects, attributes, and scenes, translated from the Visual Genome dataset (Krishna et al., 2017) and adapted to Arabic cultural contexts. This vocabulary serves as the foundation for all subsequent label matching.

3.2 Visual Label Extraction

During inference (Figure 1), for each input image, multimodal similarity scores are computed between image embeddings and vocabulary label embeddings using three CLIP-based encoders: AraCLIP, mCLIP, and Jina V4. Cosine similarity ranking is applied to select the top- k matched labels per image, which serve as interpretable visual representations of “what is seen.” These matched labels are stored for later use in caption generation. Three separate experiments, one for each CLIP-based encoder, are conducted to determine the most effective label extractor.

Model	BLEU-1 Mean	Cosine Similarity Mean	LLM Judge Score
Base Model (Bashiti et al., 2025)	16.98	58.46	30.82
Gemini+mCLIP	5.34	60.01	33.05
Gemini+AraCLIP	4.25	58.89	36.33
Gemini+Jina V4	4.49	57.81	34.80
Qwen+mCLIP	5.20	58.39	23.49
Qwen+AraCLIP	4.57	57.19	31.40
Qwen+Jina V4	4.17	57.03	30.35

Table 1: Performance comparison of our CLIP-augmented captioning system (Gemini and Qwen combined with mCLIP, AraCLIP, and Jina V4) against the Base Model.

Participant	Cosine Similarity Mean
Base Model	58.46
VLCAP (Ours)	60.01
Averroes (Saeed et al., 2025)	58.55
phantom_troupe (Horaira et al., 2025)	57.48
ImpactAi (Al-Qasem and Hendi, 2025)	56.22
Codezone Research Group (Bichi et al., 2025)	38.30

Table 2: Cosine Similarity Mean scores for participating teams.

Participant	LLM Judge Score
Base Model	30.82
Averroes	33.97
VLCAP (Ours)	33.05
phantom_troupe	31.43
ImpactAi	26.55
Codezone Research Group	15.14

Table 3: LLM Judge Score results for participating teams.

3.3 Prompt-Guided Caption Generation

The top-ranked labels (typically 25–30) are used to construct an Arabic prompt of the form: “ باستخدام العناصر التالية: [top-ranked selected labels]. ”. This prompt, along with the input image, is fed into a vision–language model. We experiment with two models: Qwen-VL and Gemini Pro Vision. The resulting captions are grounded in both the visual content and the matched labels, ensuring cultural relevance and semantic accuracy.

3.4 Combination Analysis Evaluation

By pairing each of the three label extractors with both caption generation models, we evaluate a total of six configurations. The evaluation focuses on identifying the optimal combination for producing culturally aligned, semantically accurate, and fluent Arabic captions.

4 Results

The official evaluation of shared task submissions employed three complementary methods: Cosine Similarity, which quantifies lexical closeness between generated and reference captions after Arabic-specific normalization and TF–IDF n -gram comparison; LLM-as-a-Judge, using OpenAI’s GPT-4o to assess semantic accuracy, relevance, and fluency under reproducible conditions; and Manual Evaluation, where 5% of the test set was human-rated on cultural relevance, conciseness, completeness, and accuracy.

The results in Table 1 demonstrate that our system, which integrates CLIP-based visual label detection with Qwen and Gemini, consistently outperforms the base model (Bashiti et al., 2025) across all evaluation metrics. While the base model achieves the highest BLEU-1 score due to its direct captioning pipeline, it lags behind in semantic similarity and human-preference evaluations. For the Gemini model, mCLIP yields the strongest BLEU-1 and cosine similarity means, reflecting closer alignment with ground-truth captions and semantic coherence. Notably, AraCLIP, despite lower BLEU-1 and cosine similarity scores, achieves the highest LLM Judge Score, indicating that captions generated with AraCLIP labels are often judged as more contextually relevant or human-preferred. Jina V4 provides balanced performance across all metrics for Gemini. For Qwen, mCLIP again ranks highest in BLEU-1 and cosine similarity, but its

Participant	Cultural Relevance	Conciseness	Completeness	Accuracy
VLCAP	2.57	3.17	2.67	2.97
Averroes	3.63	3.43	2.60	2.80
Phantom Troupe	3.40	3.27	2.33	2.40
Codezone Research Group	1.10	2.03	1.47	2.03
ImpactAi	3.13	2.73	1.77	1.97

Table 4: Manual evaluation results based on Cultural Relevance, Conciseness, Completeness, and Accuracy.

LLM Judge Score is relatively low, suggesting that Qwen’s outputs are less favored by human evaluators compared to Gemini. Conversely, AraCLIP and Jina V4 improve Qwen’s LLM Judge Scores, highlighting the role of the CLIP-based label extractor in shaping user-perceived caption quality.

Overall, these results confirm that our CLIP-augmented system enhances performance beyond the baseline, particularly in semantic similarity and human preference, with the choice of CLIP model exerting a stronger influence on caption quality than the downstream vision–language model itself.

In this shared task, our system achieved the highest performance in the Cosine Similarity metric, ranking first among all participating teams (Table 2), and secured second place in the LLM-as-a-Judge evaluation (Table 3), reflecting strong results in both semantic adequacy and fluency. In the manual evaluation Table 4, our captions ranked first in Completeness (2.67%) and Accuracy (2.97%), while placing second in Cultural Relevance (2.57%) and Conciseness (3.17%). These outcomes underscore the system’s ability to generate Arabic captions that are accurate, comprehensive, and semantically faithful, while maintaining competitive performance in cultural appropriateness and conciseness.

5 Conclusion

In this work, we introduced VLCAP, a modular Arabic image captioning framework that separates visual label extraction from caption generation to enhance cultural alignment, semantic accuracy, and interpretability. Through six experiments combining three CLIP-based encoders with two vision–language models, we found that Gemini Pro Vision + mCLIP delivered the strongest lexical and semantic performance (*BLEU-1*: 5.34, *cosine similarity*: 60.01), whereas Qwen-VL + AraCLIP achieved the highest LLM-based human-alignment score (36.33). These outcomes demonstrate that VLCAP’s decoupled design allows tuning for dif-

ferent evaluation priorities and provides a transferable approach for culturally aware captioning in other low-resource languages.

References

- Muhammad Al-Barham, Imad Afyouni, Khalid Al-mubarak, Ashraf Elnagar, Ayad Turkey, and Ibrahim Hashem. 2024. [Araclip: Cross-lingual learning for effective arabic image retrieval](#). In *Proceedings of The Second Arabic Natural Language Processing Conference*, pages 102–110. Association for Computational Linguistics.
- Rasha Saleh Al-Malki and Arwa Yousuf Al-Aama. 2023. [Arabic captioning for images of clothing using deep learning](#). *Sensors*, 23(8).
- Huda Al-muzaini, Tasniem N., and Hafida Benhidour. 2018. [Automatic arabic image captioning using rnn-lstm-based language model and cnn](#). *International Journal of Advanced Computer Science and Applications*, 9.
- Rabee Al-Qasem and Mohannad Hendi. 2025. [Impactai at imageeval 2025 shared task: Region-aware transformers for arabic image captioning: A case study on the palestinian narrative](#). In *Proceedings of the Third Arabic NLP Conference (ArabicNLP 2025)*. Co-located with EMNLP 2025, Nov 5–9.
- Ashwaq Alsayed, Thamir M. Qadah, and Muhammad Arif. 2023. [A performance analysis of transformer-based deep learning models for arabic image captioning](#). *J. King Saud Univ. Comput. Inf. Sci.*, 35(9).
- Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew Dai, Anja Hauth, Katie Millican, David Silver, Melvin Johnson, Ioannis Antonoglou, Julian Schrittwieser, Amelia Glaese, Jilin Chen, Emily Pitler, Timothy Lillicrap, Angeliki Lazaridou, and Orhan Firat. 2023. [Gemini: A family of highly capable multi-modal models](#).
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. [Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond](#).
- Ahlam Bashiti, Alaa Aljabari, Hadi Hamoud, Md. Rafiul Biswas, Bilal Shalash, Mustafa Jarrar, Fadi Zaraket,

- George Mikros, Ehsaneddin Asgari, and Wajdi Zaghoulani. 2025. ImageEval 2025: The First Arabic Image Captioning Shared Task. In *Proceedings of the Third Arabic Natural Language Processing Conference (ArabicNLP 2025)*. Co-located with EMNLP 2025, November 5–9.
- Abdulkadir Shehu Bichi, Ismail Dauda Abubakar, Fatima Muhammad Adam, Aminu Musa, Auwal Umar Ahmed, Abubakar Ibrahim, Khadija Salihu Aua, Aisha Mustapha Ahmed, and Mahmud Said Ahmed. 2025. Codezone research group at imageeval 2025 shared task: Arabic image captioning using blip and m2m100: A two-stage translation approach for imageeval 2025. In *Proceedings of the Third Arabic NLP Conference (ArabicNLP 2025)*. Co-located with EMNLP 2025, Nov 5–9.
- Guanhua Chen, Lu Hou, Yun Chen, Wenliang Dai, Lifeng Shang, Xin Jiang, Qun Liu, Jia Pan, and Weping Wang. 2023. **mCLIP: Multilingual CLIP via cross-lingual transfer**. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13028–13043, Toronto, Canada. Association for Computational Linguistics.
- Samar Elbedwehy and T. Medhat. 2023. **Improved arabic image captioning model using feature concatenation with pre-trained word embedding**. *Neural Comput. Appl.*, 35(26):19051–19067.
- Obeida Eljundi, Mohamad Dhaybi, Kotaiba Mokadam, Hazem Hajj, and Daniel Asmar. 2020. **Resources and end-to-end neural network models for arabic image captioning**. pages 233–241.
- Jonathan Emami, Pierre Nugues, Ashraf Elnagar, and Imad Afyouni. 2022. **Arabic image captioning using pre-training of deep bidirectional transformers**. In *Proceedings of the 15th International Conference on Natural Language Generation*, pages 40–51, Waterville, Maine, USA and virtual meeting. Association for Computational Linguistics.
- Wen Gendy and Dularia Patel. 2024. **Advancements in computer vision: A comprehensive survey of image processing and interdisciplinary applications**. *Academic Journal of Science and Technology*, 13(2):28–34.
- Michael Günther, Saba Sturua, Mohammad Akram, Isabelle Mohr, Andrei Ungureanu, Sedigheh Eslami, Scott Martens, Bo Wang, Nan Wang, and Han Xiao. 2025. **jina-embeddings-v4: Universal embeddings for multimodal multilingual retrieval**.
- Muhammad Abu Horaira, Farhan Amin, Sakibul Hasan, Md. Tanvir Ahammed Shawon, and Muhammad Ibrahim Khan. 2025. Phantomtroupe at imageeval shared task: Multimodal arabic image captioning through translation-based fine-tuning of llm models. In *Proceedings of the Third Arabic NLP Conference (ArabicNLP 2025)*. Co-located with EMNLP 2025, Nov 5–9.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David Shamma, Michael Bernstein, and Fei-Fei Li. 2017. **Visual genome: Connecting language and vision using crowdsourced dense image annotations**. *International Journal of Computer Vision*, 123.
- Abdelrahman Mohamed, Fakhraddin Alwajih, El Moatez Billah Nagoudi, Alcides Alcoba, and Muhammad Abdul-Mageed. 2023. **Violet: A vision-language model for arabic image captioning with gemini decoder**. pages 1–11.
- Mariam Saeed, Sarah Elshabrawy, Abdelrahman Hagrass, Mazen Yasser, and Ayman Khalafallah. 2025. Averroes at imageeval 2025 shared task: Advancing arabic image captioning with augmentation and two-stage generation. In *Proceedings of the Third Arabic NLP Conference (ArabicNLP 2025)*. Co-located with EMNLP 2025, Nov 5–9.
- Muhy Eddin Za’ter and Bashar Talafha. 2022. **Benchmarking and improving arabic automatic image captioning through the use of multi-task learning paradigm**.
- Jingyi Zhang, Jiaying Huang, Sheng Jin, and Shijian Lu. 2024. **Vision-language models for vision tasks: A survey**. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(8):5625–5644.