

BcQLM: Efficient Vision-Language Understanding with Distilled Q-Gated Cross-Modal Fusion

Sike Xiang and Shuang Chen and Amir Atapour-Abarghouei*

Department of Computer Science, Durham University

{sike.xiang, shuang.chen, amir.atapour-abarghouei}@durham.ac.uk

Abstract

As multimodal large language models (MLLMs) advance, their large-scale architectures pose challenges for deployment in resource-constrained environments. In the age of large models, where energy efficiency, computational scalability and environmental sustainability are paramount, the development of lightweight and high-performance models is critical for real-world applications. As such, we propose a lightweight MLLM framework for end-to-end visual question answering. Our proposed approach centres on BreezeCLIP, a compact yet powerful vision-language encoder optimised for efficient multimodal understanding. With only 1.2 billion parameters overall, our model significantly reduces computational cost while achieving performance comparable to standard-size MLLMs. Experiments conducted on multiple datasets further validate its effectiveness in balancing accuracy and efficiency. The modular and extensible design enables generalisation to broader multimodal tasks. The proposed lightweight vision-language framework is denoted as BcQLM (BreezeCLIP-enhanced Q-Gated Multimodal Language Model). It offers a promising path toward deployable MLLMs under practical hardware constraints. The source code is available at <https://github.com/thico0224/BcQLM>.

1 Introduction

Multimodal learning integrates visual and linguistic modalities, which has become a key direction in building more general and human-like AI systems. By combining information from different modalities, such systems can better understand complex real-world scenarios and perform more advanced reasoning and decision-making tasks.

Among various multimodal tasks, Visual Question Answering (VQA) has emerged as a widely

adopted and representative benchmark. VQA requires models to comprehend an image, understand a corresponding natural language question and generate an appropriate answer (Antol et al., 2015). It poses significant challenges in cross-modal semantic alignment, contextual reasoning and fine-grained visual understanding (Anderson et al., 2018; Li et al., 2022; Zhang et al., 2021), and thus serves as an important indicator of multimodal intelligence.

To effectively address the inherent challenges of VQA and other complex multimodal tasks, mainstream approaches have predominantly adopted a two-stage architecture. Pretrained encoders such as CLIP (Radford et al., 2021), ViT (Dosovitskiy et al., 2021), and BERT (Devlin et al., 2019) are first used to extract semantic features from images and texts, providing strong unimodal representations and initial cross-modal alignment through large-scale pre-training. These features are then passed to large language models (LLM), which serve as decoders to perform reasoning and generate answers. This paradigm has been adopted by several representative systems, including BLIP-2, Qwen and LLaVA (Li et al., 2023; Bai et al., 2023; Li et al., 2024). These models integrate visual understanding with language modelling, which enables them to generate contextually relevant and accurate answers across a wide range of question types.

However, existing multimodal large models often suffer from large parameter sizes and high computational costs, limiting their deployment on edge devices. This hinders practical use in scenarios requiring efficient on-device inference, such as healthcare (Wang et al., 2023), remote education (Chu et al., 2025), disaster response (Kumbam and Vejre, 2024) and assistive technologies (Driess et al., 2023). In these settings, cloud-based solutions may be infeasible due to connectivity or latency issues (Wang et al., 2024), underscoring the need for lightweight yet capable multimodal

*Corresponding author.

systems.

To address this issue, we propose BcQLM (BreezeCLIP-enhanced Q-Gated Multimodal Language Model). BreezeCLIP is a lightweight vision–language encoder designed to significantly reduce computational costs while enabling efficient multimodal representation. It replaces the original BERT and ViT backbones in CLIP with compact transformer modules inspired by the inverted bottleneck design (Sandler et al., 2018). To train BreezeCLIP, we construct image–text pairs from the GQA dataset (Hudson and Manning, 2019), leveraging its rich scene graph annotations that provide detailed object relationships and spatial context. We adopt a dual training strategy: contrastive learning ensures visual and textual features are well aligned in a shared embedding space, while knowledge distillation from a CLIP teacher (Radford et al., 2021) transfers high-level semantic alignment to BreezeCLIP. This enables BreezeCLIP to achieve strong performance despite its lightweight architecture.

A lightweight decoder based on LLaMA-3.2-1B (Grattafiori et al., 2024) is first employed to generate answers conditioned on fused multimodal features, achieving strong performance with significantly reduced computational overhead. Building on this, we introduce a Q-Gated Cross-Modal Fusion Module that dynamically adjusts the contribution of visual and textual features according to the input question, which enables fine-grained, question-aware interaction between modalities. These components together constitute an efficient end-to-end multimodal question answering system designed for real-world, resource-constrained deployment scenarios.

Our primary contributions are thus as follows:

- (i) We propose BcQLM, involving a compact BreezeCLIP by distillation learning and a Q-Gated Cross-Modal Fusion Module. Comprehensive experiments demonstrate that our BreezeCLIP is able to effectively preserve vision-language alignment capabilities under a tiny model setting.
- (ii) We propose a Q-Gated Cross-Modal Fusion Module to enable fine-grained and adaptive multimodal fusion.
- (iii) Our BreezeCLIP only contains 1.2B parameters (10% of SoTA method), which achieves performance comparable to several standard-sized MLLMs with much higher parameter counts on VQA tasks.

2 Related Work

We consider prior work leading to recent advances in MLLMs, with connections to multimodal fusion techniques and vision-language tasks, such as VQA, highly relevant to our work.

2.1 Visual Question Answering

Visual Question Answering (VQA) is a prominent multimodal task that requires a model to answer natural language questions based on the content from an image. It requires both visual perception and language understanding, as well as reasoning across the two modalities. First introduced by Antol et al. (Antol et al., 2015), the VQA task has become a widely used benchmark for evaluating multimodal intelligence. Subsequent work by Goyal et al. (Goyal et al., 2017) extended the original dataset with balanced answer distributions, mitigating language bias and enabling more robust evaluation.

Early VQA approaches extract CNN-based image features and LSTM/GRU question embeddings, fusing them via concatenation, bilinear pooling or co-attention (e.g., MCB (Fukui et al., 2016)); later attention-based models like BAN (Kim et al., 2018) and Bottom-Up Top-Down (Anderson et al., 2018) enhance fine-grained reasoning and cross-modal alignment. These early methods achieved some success but relied on task-specific fusion and lacked generalisation. Recent work leverages pretrained large language models and generative architectures to enable broader reasoning.

2.2 Vision-Language Pretraining & LLM Generation

Advances in multimodal representation learning focused on aligning visual and textual modalities through joint embedding spaces. One of the most representative models, CLIP (Radford et al., 2021), learns such a space by contrastively training an image encoder and a text encoder on large-scale image-text pairs. Subsequent models such as ViLT (Kim et al., 2021), ALBEF (Li et al., 2021), and BLIP (Li et al., 2022) further enhance cross-modal alignment by incorporating attention-based fusion mechanisms and multi-task pretraining objectives. These models demonstrate strong transferability to a wide range of downstream tasks, including VQA, image-text retrieval (Xiao et al., 2025) and caption generation (Liu et al., 2025). Such dual-encoder and fusion-based models adopt architectures where

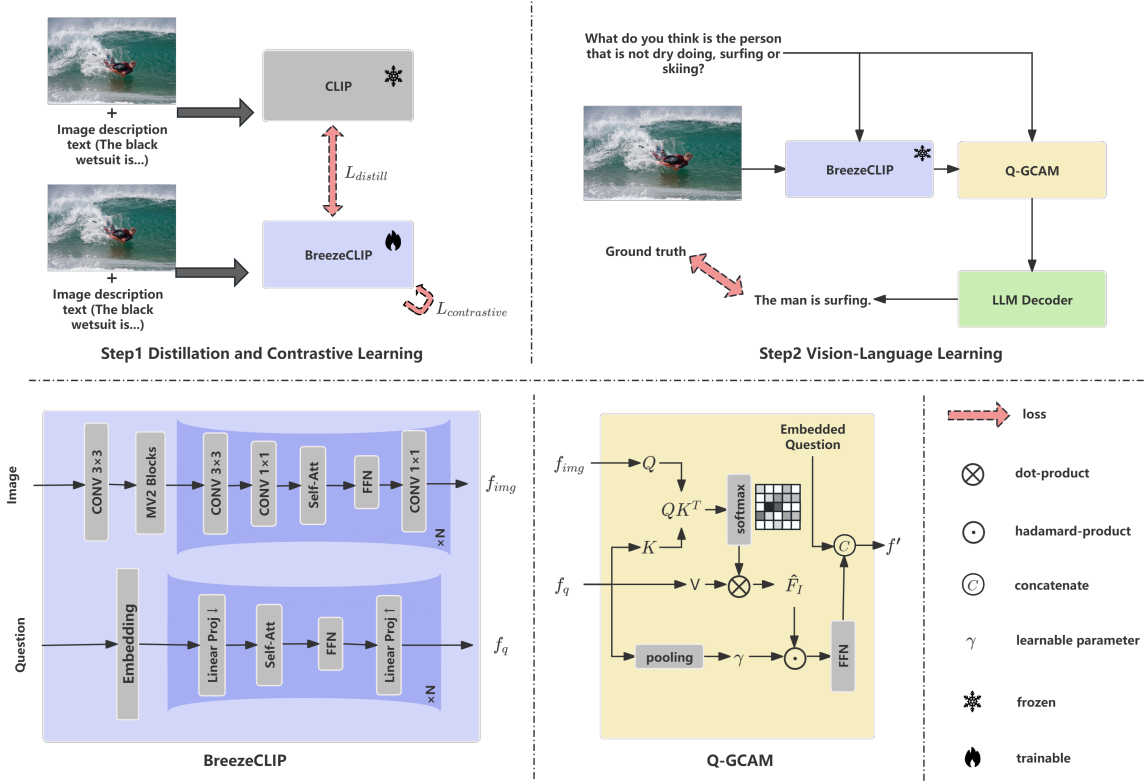


Figure 1: An overview of the proposed framework. Training consists of two stages: (1) BreezeCLIP encoder is trained using contrastive ($\mathcal{L}_{\text{contrastive}}$) and distillation loss ($\mathcal{L}_{\text{distill}}$) from the CLIP teacher; (2) the Q-GCAM module fuses image and question features to generate answers via a lightweight LLM decoder, optimised with cross-entropy loss (\mathcal{L}_{CE}). The bottom diagrams illustrate BreezeCLIP and Q-GCAM architectures.

visual and textual inputs are first encoded independently and then fused via attention modules or interaction layers. While they excel at semantic alignment, they remain limited in their ability to support complex reasoning or generate language that relies on common sense and broader world knowledge not captured by the training corpus.

In response to these limitations, recent studies have introduced pretrained LLMs into multimodal systems, marking a shift toward generative architectures. Unlike traditional approaches that rely solely on vision-language encoders, models such as BLIP-2 (Li et al., 2023) and LLaVA (Li et al., 2024) integrate visual features into frozen or lightly tuned LLMs, treating the language model as the central reasoning and generation component. Visual inputs are typically embedded as prompts or special tokens to guide multimodal interaction. These models demonstrate superior performance in open-ended and knowledge-intensive tasks like VQA and produce more coherent and contextually relevant answers. This shift from traditional dual-encoder models to LLM-centred generative architectures

brings improved reasoning and generation capabilities, but high computational costs and large model sizes remain major obstacles for practical deployment. Recent works such as AdapterCLIP (Gao et al., 2023) and LoRA (Hu et al., 2022) mitigate this issue via lightweight adapters or low-rank tuning, but still rely on large fixed backbones and show limited flexibility in complex generation tasks. Our work advances this direction by emphasising lightweight and efficient modelling through the introduction of BreezeCLIP, coupled with gated cross-modal fusion and a LLaMA-based decoder.

3 Method

The proposed BcQLM (BreezeCLIP-enhanced Q-Gated Multimodal Language Model) is a lightweight multimodal large language model designed for visual question answering under resource constraints, to support effective multimodal reasoning in resource-limited environments. The architecture comprises three main components. First, a compact vision-language encoder, BreezeCLIP, is proposed to extract semantic features from

both images and text. BreezeCLIP is trained using contrastive learning and knowledge distillation from a CLIP teacher to ensure efficient representation learning. Next, we propose a question-guided cross-modal attention module (Q-GCAM) to dynamically fuse visual and textual features based on the semantics of the input question. Finally, a lightweight language decoder based on LLaMA-3.2-1B generates the final answer. An overview of the framework is shown in Figure 1. It follows a two-stage design: (1) BreezeCLIP is first pretrained via distillation and contrastive learning; and (2) the pretrained BreezeCLIP is then used as a frozen feature extractor, working with Q-GCAM and an LLaMA-3.2-1B decoder to perform vision-language generation. This balances performance and efficiency, and enables BcQLM to support effective multimodal reasoning in resource-limited environments.

3.1 Breezeclip

To enable efficient vision-language representation learning within our framework, we design BreezeCLIP, a lightweight dual-encoder model that significantly reduces computational costs while effectively aligning visual and textual modalities.

Feature Encoder: The text encoder in BreezeCLIP adopts a lightweight Transformer to enable efficient language modelling in resource-constrained settings, as shown in Figure 1. The architecture adopts the same 24-layer Transformer structure as BERT-Large. To reduce the number of computation parameters while preserving model capacity, we incorporate a bottleneck structure (Sun et al., 2020) into each layer: token embeddings are first projected to lower dimensions, where attention and feed-forward computations are performed and then restored via a linear projection. Residual connections and layer normalisation are applied to ensure training stability. Furthermore, we adopt a factorised embedding strategy to reduce the size of the vocabulary embedding matrix and apply layer-wise knowledge distillation to preserve semantic representation quality. The hidden states from the last Transformer layer are retained as token-level features without pooling, preserving fine-grained textual semantics for subsequent fusion.

The image encoder combines convolutional operations with lightweight Transformer modules, following the bottleneck principle (Mehta and Rastegari, 2022). This structure enables efficient local feature extraction while maintaining the ability to

capture global context with minimal computational overhead. As shown in Figure 1, the input image is first passed through a convolutional stem and a set of lightweight MobileNetV2-style bottleneck blocks (Sandler et al., 2018) to extract early features. The resulting feature maps are then processed by three Transformer blocks operating at progressively smaller spatial resolutions (e.g., 14×14 and 7×7). Each block consists of a local convolutional module followed by a stack of Transformer encoder layers — specifically, 2, 4, and 3 layers in the three blocks, respectively, totalling 9 self-attention layers. In each Transformer block, the input feature map is processed by a 3×3 convolution to capture local context, followed by a 1×1 convolution to reduce channel dimensions. The resulting features are flattened into non-overlapping patches and passed through the Transformer encoder for global modelling. The output is then reshaped back into the 2D spatial layout and fused with the original representation using a 1×1 convolution. Finally, the output feature map is flattened across spatial dimensions into a patch-level embedding sequence to produce a three-dimensional representation aligned with the text token embeddings for downstream cross-modal interaction.

The constructed BreezeCLIP achieves a total parameter reduction of approximately 80% (from 151M to 31M), which significantly improves deployment efficiency in resource-constrained environments compared to standard CLIP.

Distillation and Contrastive Learning. A joint optimisation scheme that unifies contrastive learning and knowledge distillation is adopted to strengthen alignment across modalities in a lightweight setting. This enables the student model to learn from both raw image-text pairs and the supervision of a large pretrained teacher model.

We construct image-text pairs by leveraging the scene graphs provided in the GQA dataset. Each image is associated with a structured graph that specifies object categories, spatial relationships (e.g., “a book is on the right of the table”), and attributes (e.g., “the book is red”), as shown in Appendix Figure A1. We convert these graph structures into informative language descriptions to serve as textual counterparts to the original images.

Given image-text pairs, the inputs are encoded using the student model’s image and text encoders, resulting in feature embeddings I_s and T_s . A contrastive InfoNCE loss (Parulekar et al., 2023) is used to bring matched pairs closer and push mis-

matched pairs apart in the shared embedding space:

$$\mathcal{L}_{\text{contrast}} = \frac{1}{\alpha} \left(\text{CE} \left(\frac{I_s T_s^\top}{\tau}, \mathbf{y} \right) + \text{CE} \left(\frac{T_s I_s^\top}{\tau}, \mathbf{y} \right) \right), \quad (1)$$

where τ is a temperature coefficient, \mathbf{y} denotes the index labels of positive pairs and α balances the contributions of the two directional terms.

To further improve the student model’s representations, we perform knowledge distillation from a pretrained CLIP model. Visual and textual embeddings I_t and T_t are extracted from the teacher model and projected into the student embedding space via learnable projection heads. The distillation loss is computed as the mean squared error (MSE) between the L2-normalised student and teacher embeddings:

$$\mathcal{L}_{\text{distill}} = \frac{1}{\beta} (\text{MSE}(I_s, I_t) + \text{MSE}(T_s, T_t)), \quad (2)$$

where β is a normalisation factor.

The total training objective is defined as a weighted combination of the two losses:

$$\mathcal{L}_{\text{total}} = \lambda_1 \cdot \mathcal{L}_{\text{contrast}} + \lambda_2 \cdot \mathcal{L}_{\text{distill}}, \quad (3)$$

where hyperparameters, λ_1 and λ_2 , balance the contributions of the contrastive and distillation losses.

3.2 Multimodal Input Construction

To enable fine-grained visual-language interaction, a Dynamic Gated Cross-Attention module is used. After extracting modality-specific representations from the feature encoders, visual and textual features undergo cross-modal integration to establish contextual grounding for answer generation.

Specifically, the image encoder outputs a latent feature map $F_V \in \mathbb{R}^{B \times C \times H \times W}$, which is spatially flattened into patch-level embeddings $F_I \in \mathbb{R}^{B \times N \times C}$, where $N = H \times W$. Concurrently, the text encoder provides token-wise embeddings $F_T \in \mathbb{R}^{B \times T \times d}$, where d is the embedding dimension, as well as a global semantic representation $\bar{F}_T \in \mathbb{R}^{B \times d}$ derived through average pooling:

$$\bar{F}_T = \frac{1}{T} \sum_{i=1}^T F_T^{(i)}, \quad (4)$$

Visual patch tokens serve as queries, while textual embeddings provide keys and values for attention computation:

$$Q = F_I W^Q, \quad K = F_T W^K, \quad V = F_T W^V, \quad (5)$$

$$\hat{F}_I = \text{Softmax} \left(\frac{QK^\top}{\sqrt{d}} \right) V. \quad (6)$$

To inject global semantic control, the pooled textual vector \bar{F}_T is projected into the same latent space and broadcast across all spatial locations:

$$\bar{F}'_T = \text{Proj}(\bar{F}_T) \in \mathbb{R}^{B \times d}, \quad (7)$$

$$\gamma = \sigma(\text{MLP}([F_I; \bar{F}'_T])) \in [0, 1]^{B \times N \times 1}. \quad (8)$$

The attention-enhanced visual representations \hat{F}_I are adaptively modulated by the gate γ , then integrated with the original features through a gated residual connection:

$$F_{\text{mod}} = F_I + \gamma \odot \hat{F}_I, \quad (9)$$

This intermediate representation is subsequently refined via a feed-forward network and normalisation layer:

$$F_{\text{fused}} = \text{LayerNorm}(F_{\text{mod}} + \text{FFN}(F_{\text{mod}})). \quad (10)$$

To align with the decoder’s token embedding space, the fused visual patches are mapped through a lightweight projection head, yielding pseudo textual tokens:

$$F_{\text{pseudo}} = \text{Adapter}(F_{\text{fused}}) \in \mathbb{R}^{B \times N \times d}. \quad (11)$$

3.3 LLM Decoder

In this work, the LLaMA-3.2-1B model (Grattafiori et al., 2024) is adopted as the language decoder to generate final responses. As a compact yet capable member of the LLaMA series, the 1B variant is pre-trained exclusively on large-scale textual corpora, without exposure to any visual or multimodal data. Its small size enables efficient fine-tuning and inference, which makes it suitable for resource-constrained deployment. More importantly, the absence of any prior visual grounding makes it an ideal testbed to rigorously assess the effectiveness of our visual-text fusion module in aligning visual content for language generation.

Given the pseudo tokens from the image encoder, denoted as $F_{\text{pseudo}} \in \mathbb{R}^{B \times N \times d}$, and the instruction-response text sequence X_{text} , tokenised and embedded as $F_T = \text{Embed}(X_{\text{text}}) \in \mathbb{R}^{B \times T \times d}$, the input to the decoder is constructed by concatenation:

$$F_{\text{llama}} = [F_{\text{pseudo}}; F_T] \in \mathbb{R}^{B \times (N+T) \times d}, \quad (12)$$

Method	LLM	Param.	Res.	GQA	VQA ^{V2}	VisWiz
BLIP-2 (2023)	Vicuna-13B	13.5	224	44.0	65.0	19.6
InstructBLIP (2023)	Vicuna-7B	7.5	224	49.2	–	34.5
InstructBLIP (2023)	Vicuna-13B	13.5	224	49.5	–	33.4
IDEFICS-9B (2023)	LLaMA-7B	9.0	224	38.4	50.9	35.5
IDEFICS-80B (2023)	LLaMA-65B	80.0	224	45.2	60.0	36.0
<i>BcQLM (ours)</i>	<i>Llama-3.2-1B</i>	1.2	224	60.8	71.0	49.5
Qwen-VL (2023)	Qwen-7B	9.6	448	59.3	78.8	35.2
Qwen-VL-Chat (2023)	Qwen-7B	9.6	448	57.5	78.2	38.9
LLaVA-1.5 (2024)	Vicuna-1.5-7B	7.3	336	62.0	78.5	50.0
LLaVA-1.5 (2024)	Vicuna-1.5-13B	13.3	336	63.3	80.0	53.6
VILA-7B (2024)	Llama-2-7B	7.0	336	62.3	79.9	57.8
VILA-13B (2024)	Llama-2-13B	13.0	336	63.3	80.8	60.6
<i>BcQLM (ours)</i>	<i>Llama-3.2-1B</i>	1.2	336	<u>62.4</u>	78.7	56.1

Table 1: Comparison of accuracy (%) across three VQA benchmarks: GQA, VQAv2, and VizWiz. Results for BLIP-2 (Li et al., 2023), InstructBLIP (Dai et al., 2023), IDEFICS (Laurençon et al., 2023), LLaVA-1.5 (Liu et al., 2024), VILA (Lin et al., 2024), and Qwen-VL (Bai et al., 2023) are reported based on their official releases. Our proposed BcQLM achieves competitive performance across all datasets while using only 1.2B parameters.

During training, label supervision is applied only to the response portion of the sequence, with instruction and visual tokens masked out from loss computation. This facilitates alignment of visual features with downstream language tokens via gradient propagation through the decoder:

$$\mathcal{L}_{\text{gen}} = \text{CE}(\text{LLaMA}(F_{\text{llama}}), Y_{\text{response}}). \quad (13)$$

where $\text{CE}(\cdot)$ denotes the cross-entropy loss computed over the generated response tokens Y_{response} .

4 Experiments

4.1 Dataset

To comprehensively evaluate the performance of our model, we conduct experiments on three widely used datasets: GQA, VQAv2, and VizWiz:

- **GQA** (Hudson and Manning, 2019): 113K real-world images from Visual Genome paired with 22.7M programmatically generated questions. Each question is derived from a scene graph and converted into a natural-language query covering comparison, spatial relations, logical inference and other multi-step reasoning types.
- **VQAv2** (Goyal et al., 2017): Built on MSCOCO, comprising approximately 205K images and over 1.1M image–question–answer triplets. Each question is paired with ten

human-provided answers and complementary image–question pairs are included to mitigate language biases.

- **VizWiz** (Gurari et al., 2018): Consists of about 31K images and corresponding free-form questions captured by blind users with mobile phones. Each question has ten crowd-sourced answers; the dataset is characterised by poor image quality, incomplete visual context and diverse real-world query types.

4.2 Settings

We use a pretrained CLIP model as the teacher and BreezeCLIP as the student. Images are resized (224×224 or 336×336), normalised and converted to tensors. Text is tokenised and padded to 77 tokens. Visual and textual features are projected to a shared embedding space (512 or 768 dims). We use contrastive loss with $\tau = 0.07$, $\alpha = \beta = 0.5$, and train for 64 epochs using Adam (lr= 1×10^{-5}), StepLR (step_size=10, $\gamma=0.5$), batch size 32, and gradient clipping (max norm=1.0).

BreezeCLIP visual features are adapted via a VisualAdapter and fused with question embeddings through an 8-head Dynamic Gated Cross-Attention. The fused embeddings are fed into a trainable LLaMA model. Training uses AdamW (lr= 1×10^{-4}), StepLR (step_size=5, $\gamma=0.1$), for 15 epochs, with batch size 32 and gradient clipping.

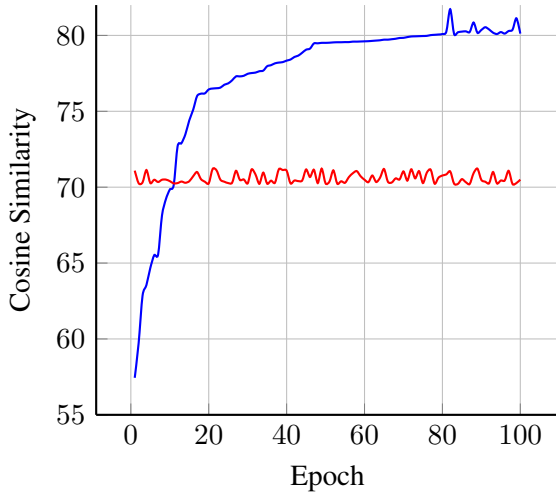


Figure 2: Cosine similarity in BreezeCLIP distillation.

4.3 Evaluation Metrics

During the distillation and contrastive learning stage, we compute the InfoNCE loss on the validation set to measure the convergence of image–text embedding alignment. Additionally, the mean cosine similarity of all positive image–text pairs and the positive–negative similarity gap are reported to evaluate the discrimination capability of the model. For the visual question answering task, we employ VQA accuracy as the sole evaluation metric, quantifying the proportion of generated answers that match the reference answers.

4.4 Quantitative Results

BreezeCLIP is jointly optimised with contrastive and distillation objectives under the supervision of a frozen CLIP ViT-B/32 teacher. During training, the cosine similarity between visual and textual embeddings increases steadily, ultimately exceeding 80%, whereas the teacher model remains around 71% (Figure 2). This demonstrates that despite its compact size, BreezeCLIP achieves stronger alignment in the multimodal embedding space.

To better visualise the feature distribution, we project high-dimensional embeddings into 3D via PCA and show them in Figure 3. The teacher model, although not trained on these datasets, still achieves about 71% positive cosine similarity due to large-scale pretraining and the strong cross-modal representation learning, generalisation ability and robust alignment properties of CLIP. However, its embedding space mixes positive and negative pairs without clear separation. In contrast,




Figure 3: 3D projection of vision-language embeddings after dimensionality reduction. BreezeCLIP (a) exhibits clearer separation between positive and negative pairs compared to the teacher CLIP (b), which shows no separation between positive and negative pairs.

BreezeCLIP clusters positive pairs tightly at the centre and pushes negative pairs outward. This demonstrates that our dual-loss training strategy significantly improves feature discrimination.

The proposed framework is evaluated on three widely used VQA benchmarks (GQA, VQAv2 and VizWiz), using accuracy as the primary evaluation metric. As shown in Table 1, BcQLM achieves an accuracy of 60.8% on GQA, 71.0% on VQAv2 and 49.5% on VizWiz, outperforming most existing methods operating at a resolution of 224×224 and even surpassing several models using higher resolutions. When increasing the input resolution to 336×336, BcQLM improves its performance by achieving 62.4% on GQA as the second best result among published models, 78.7% on VQAv2 and 56.1% on VizWiz as the third best result, bringing its overall performance in line with mainstream large scale vision language models. Notably, our model uses only 1.2B parameters, significantly smaller than existing models, and requires no external data or instruction tuning, demonstrating the effectiveness of our training and encoder design.

4.5 Ablation Study

Five configurations were compared on the validation set (as shown in Table 2): CLIP + LLaMA-3.2-1B (58.8%), BreezeCLIP + LLaMA-3.2-1B (59.2%), Q-GCAM (60.8%), Q-GCAM’ with Token Balance (58.7%) and Q-GCAM’ with Visual Query (55.7%). Token Balance reweights visual features to match textual contributions; Visual Query adds a cross-modal attention layer using visual tokens as queries. These results show that Q-GCAM drives the performance gains, while its Token Balance and Visual Query variants need finer parameter tuning and integration.

In the preceding ablation study, we evaluate the

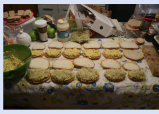




Input Image					
Input Question	Are there any forks in the bowl?	Of which color is the fur?	Is the van to the right of a bus?	Does the ground look bare?	Is there a bus or a window in the scene?
BcQLM(5%)	No, there is no fork in the bowl.	The fur is gray.	No, the van is to the right of a man.	No, the ground is snowy.	Yes, there is a window.
BcQLM(50%)	No, there is some batter in the bowl.	The fur is gray.	No, the van is to the right of people.	No, the ground is snowy.	Yes, there is a window.
BcQLM(100%)	No, there is a spoon in the bowl.	The fur is gray.	No, the van is to the right of the person.	No, the ground is snowy.	Yes, there are windows.
Ground Truth	No, there is a spoon in the bowl.	The fur is white.	No, the van is to the right of the person.	No, the ground is rocky.	Yes, there are windows.

Figure 4: Example BcQLM responses to GQA questions: “Is there a fork in the bowl?”, “What color is the fur?”, “Is the van to the right of a bus?”, “Is there a window in the scene?”, and “Does the ground look bare?” under three different LLaMA-3.2-1B unfreeze ratios (5%, 50%, 100%). Despite unfreezing only 5% of parameters, the model maintains correct semantic judgments; full (100%) unfreezing further improves fine-grained detail recognition, such as correctly identifying a spoon in the bowl.

Model Configuration	Acc. (%)
CLIP + LLaMA	58.8
BreezeCLIP + LLaMA	59.2
BreezeCLIP + Q-GCAM + LLaMA	60.8
BreezeCLIP + Q-GCAM' + LLaMA	58.7
BreezeCLIP + Q-GCAM'' + LLaMA	55.7

Table 2: Comparison of different model configurations on validation accuracy. All LLaMA means LLaMA-3.2-1B, Q-GCAM' indicates the Token Balance strategy, Q-GCAM'' indicates the Visual Query strategy.

impact of varying unfreezing ratios (5%, 50%, and 100%) on BcQLM’s responses using the LLaMA-3.2-1B backbone. Across different settings, the model consistently produces semantically correct answers, even when only 5% of the parameters are unfrozen. Fully unfreezing the model provides slight improvements on fine-grained details, such as object attributes and spatial descriptions. Importantly, every prediction can still be viewed as correct from certain interpretative perspectives, even though its expression may differ from the canonical ground truth. These findings demonstrate BcQLM’s robustness and effectiveness under different unfreezing strategies, with minimal reliance on extensive parameter updates. Example responses are

shown in Figure 4 for qualitative reference.

4.6 Efficiency Analysis

We further evaluate inference efficiency on an NVIDIA RTX 4070 Ti using 20 test samples. BcQLM completes the batch in **2.54 seconds** (127.1 ms per sample), requiring **4989.5 MiB peak GPU memory** and **1.57×10^{11} FLOPs per sample**. In comparison, Qwen2.5-VL-3B (Qwen et al., 2025) takes 5.06 seconds (253.1 ms per sample), uses 7152.6 MiB memory, and costs 1.60×10^{11} FLOPs per sample, while Gemma3-4B (Team et al., 2025) takes 5.48 seconds (273.9 ms per sample), consumes 7432.8 MiB memory, and requires 1.85×10^{11} FLOPs per sample. These results show that BcQLM runs at **twice the speed** of Qwen2.5-VL-3B and Gemma3-4B, with around **30% lower memory usage** and comparable or fewer FLOPs, highlighting its superior efficiency for edge-like deployment scenarios.

5 Conclusion

In this work, we present BcQLM, a lightweight multimodal language model, specifically designed for efficient and effective vision-language understanding under resource-constrained settings. Our

approach is centred on BreezeCLIP, a compact vision-language encoder trained using a dual strategy of contrastive learning and knowledge distillation. It preserves cross-modal alignment while significantly reducing computational overhead. The model further incorporates Q-GCAM, a Q-Gated Cross-Modal Attention module to enable dynamic and context-aware fusion of visual and textual features for adaptive multimodal reasoning.

Our extensive experiments demonstrate that BcQLM achieves strong performance on three widely used VQA benchmarks — GQA (62.4%), VQAv2 (78.7%), and VizWiz (56.1%) — using only 1.2 billion parameters, a fraction of the size of comparable models. These results validate the effectiveness of our lightweight design in balancing computational efficiency and high performance. Beyond its strong quantitative results, BcQLM offers practical advantages for real-world deployment. Its compact architecture enables efficient inference on modest hardware, which makes it suitable for applications in edge computing, assistive technologies and real-time multimodal systems. Our modular framework also allows for flexible adaptation to other multimodal tasks as a versatile solution for resource-limited environments.

Limitations

While our work introduces an efficient and effective multimodal large language model under constrained computational settings, certain areas offer opportunities for further enhancement. First, our model relies on publicly available vision-language datasets, which may not fully capture the complexities and distributions of real-world multimodal scenarios. The limited diversity and quality of these datasets could constrain the model’s generalisation ability, particularly in domains requiring fine-grained or specialised reasoning. Second, the decoder (LLaMA-3.2-1B) employed in our system is pretrained solely on text and remains frozen during fusion training. Although this design choice enhances efficiency and controllability, it may also limit the model’s capacity to adaptively integrate visual semantics during generation. Finally, although our approach has been primarily validated on VQA benchmarks, future work should explore its extension to more dynamic modalities, including video, audio and real-time interactive settings.

References

- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.
- 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. *Preprint*, arXiv:2308.12966.
- Zhendong Chu, Shen Wang, Jian Xie, Tinghui Zhu, Yibo Yan, Jinheng Ye, Aoxiao Zhong, Xuming Hu, Jing Liang, Philip S. Yu, and Qingsong Wen. 2025. Llm agents for education: Advances and applications. *Preprint*, arXiv:2503.11733.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. Instructblip: towards general-purpose vision-language models with instruction tuning. In *Proceedings of the 37th International Conference on Neural Information Processing Systems, NIPS ’23*, Red Hook, NY, USA. Curran Associates Inc.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*.
- Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, and 3 others. 2023. Palm-e: An embodied multimodal language model. *Preprint*, arXiv:2303.03378.
- Akira Fukui, Dong Huk Park, Daylen Yang, Anna Rohrbach, Trevor Darrell, and Marcus Rohrbach.

2016. [Multimodal compact bilinear pooling for visual question answering and visual grounding](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 457–468, Austin, Texas. Association for Computational Linguistics.
- Peng Gao, Shijie Geng, Renrui Zhang, Teli Ma, Rongyao Fang, Yongfeng Zhang, Hongsheng Li, and Yu Qiao. 2023. [Clip-adapter: Better vision-language models with feature adapters](#). *Int. J. Comput. Vision*, 132(2):581–595.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. [Making the v in vqa matter: Elevating the role of image understanding in visual question answering](#). In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6325–6334.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. 2018. [Vizwiz grand challenge: Answering visual questions from blind people](#). In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3608–3617.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Drew A. Hudson and Christopher D. Manning. 2019. [Gqa: A new dataset for real-world visual reasoning and compositional question answering](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. 2018. [Bilinear attention networks](#). In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, page 1571–1581, Red Hook, NY, USA. Curran Associates Inc.
- Wonjae Kim, Bokyoung Son, and Ildoo Kim. 2021. [Vilt: Vision-and-language transformer without convolution or region supervision](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 5583–5594. PMLR.
- Pranath Reddy Kumbam and Kshitij Maruti Vejre. 2024. [Floodlense: A framework for chatgpt-based real-time flood detection](#). *Preprint*, arXiv:2401.15501.
- Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, Siddharth Karamcheti, Alexander M. Rush, Douwe Kiela, Matthieu Cord, and Victor Sanh. 2023. [Obelics: An open web-scale filtered dataset of interleaved image-text documents](#). *Preprint*, arXiv:2306.16527.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024. [Llava-onevision: Easy visual task transfer](#). *Preprint*, arXiv:2408.03326.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. [BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models](#). In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pages 19730–19742. PMLR.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. [Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation](#). In *ICML*.
- Junnan Li, Ramprasaath R. Selvaraju, Akhilesh Deepak Gotmare, Shafiq Joty, Caiming Xiong, and Steven Hoi. 2021. [Align before fuse: Vision and language representation learning with momentum distillation](#). In *NeurIPS*.
- Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. 2024. [Vila: On pre-training for visual language models](#). In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 26689–26699.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024. [Improved baselines with visual instruction tuning](#). In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26286–26296.
- Zheng Liu, Hao Liang, Bozhou Li, Tianyi Bai, Wentao Xiong, Chong Chen, Conghui He, Wentao Zhang, and Bin Cui. 2025. [Synthvlm: High-efficiency and high-quality synthetic data for vision language models](#). *Preprint*, arXiv:2407.20756.
- Sachin Mehta and Mohammad Rastegari. 2022. [Mobilevit: Light-weight, general-purpose, and mobile-friendly vision transformer](#). In *International Conference on Learning Representations*.
- Advait Parulekar, Liam Collins, Karthikeyan Shanmugam, Aryan Mokhtari, and Sanjay Shakkottai. 2023. [Infonce loss provably learns cluster-preserving representations](#). *Preprint*, arXiv:2302.07920.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin

Yang, Jiayi Yang, Jingren Zhou, and 25 others. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.

Alec Radford, Jong Wook Kim, Chris Hallacy, A. Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *ICML*.

Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. [Mobilenetv2: Inverted residuals and linear bottlenecks](#). In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4510–4520.

Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. [MobileBERT: a compact task-agnostic BERT for resource-limited devices](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2158–2170, Online. Association for Computational Linguistics.

Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. [Gemma 3 technical report](#). *Preprint*, arXiv:2503.19786.

Sheng Wang, Zihao Zhao, Xi Ouyang, Qian Wang, and Dinggang Shen. 2023. [Chatcad: Interactive computer-aided diagnosis on medical image using large language models](#). *Preprint*, arXiv:2302.07257.

Xiaojie Wang, Beibei Wang, Yu Wu, Zhaolong Ning, Song Guo, and Fei Richard Yu. 2024. [A survey on trustworthy edge intelligence: From security and reliability to transparency and sustainability](#). *IEEE Communications Surveys & Tutorials*, pages 1–1.

Rui Xiao, Sanghwan Kim, Mariana-Iuliana Georgescu, Zeynep Akata, and Stephan Alaniz. 2025. [Flair: Vlm with fine-grained language-informed image representations](#). *CVPR*.

Pengchuan Zhang, Xiujun Li, Xiaowei Hu, Jianwei Yang, Lei Zhang, Lijuan Wang, Yejin Choi, and Jianfeng Gao. 2021. [Vinvl: Revisiting visual representations in vision-language models](#). In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5575–5584.

A Additional Visualization

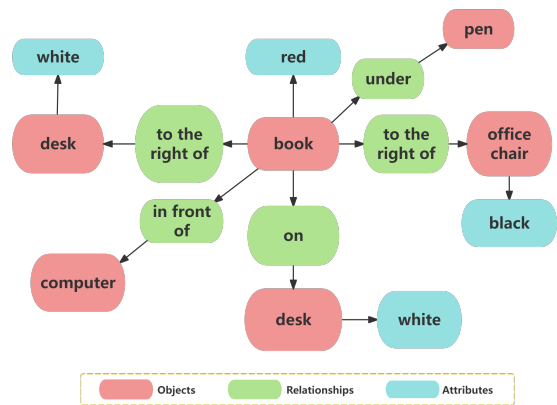


Figure A1: An example scene graph from the GQA dataset, illustrating spatial and attribute relationships among objects, such as a red book on the desk.