

MoEMoE: Question Guided Dense and Scalable Sparse Mixture-of-Expert for Multi-source Multi-modal Answering

Vinay Kumar Verma * †

Private Brands - Discovery, Amazon
vkvermaa@amazon.com

Happy Mittal †

CMT Systems, Amazon
mithappy@amazon.com

Shreyas Sunil Kulkarni *

International Machine Learning, Amazon
kulkshre@amazon.com

Deepak Gupta

International Machine Learning, Amazon
dgupt@amazon.com

Abstract

Question Answering (QA) and Visual Question Answering (VQA) are well-studied problems in the language and vision domain. One challenging scenario involves multiple sources of information, each of a different modality, where the answer to the question may exist in one or more sources. This scenario contains richer information but is highly complex to handle. In this work, we formulate a novel question-answer generation (QAG) framework in an environment containing multi-source, multimodal information. The answer may belong to any or all sources; therefore, selecting the most prominent answer source or an optimal combination of all sources for a given question is challenging. To address this issue, we propose a question-guided attention mechanism that learns attention across multiple sources and decodes this information for robust and unbiased answer generation. To learn attention within each source, we introduce an explicit alignment between questions and various information sources, which facilitates identifying the most pertinent parts of the source information relative to the question. Scalability in handling diverse questions poses a challenge. We address this by extending our model to a sparse mixture-of-experts (sparse-MoE) framework, enabling it to handle thousands of question types. Experiments on T5 and Flan-T5 using three datasets demonstrate the model’s efficacy, supported by ablation studies.

1 Introduction

The field of question-answer generation (QAG) (Touvron et al., 2023; Jiang et al., 2023) and visual question-answer generation (VQAG) (Li et al., 2022) holds significant promise with extensive applications across various domains. Recent advancements in large-scale language

¹Equal contribution.

²This work was done while author was in International Machine Learning team.




Figure 1: Example of the multi-modal and multi-source attribute extraction using the proposed question answering mechanism.

models (Jiang et al., 2023; Taori et al., 2023) and vision models (Zhang et al., 2023; Li et al., 2022; Verma et al., 2023) have demonstrated notable progress. However, current models are often constrained to QAG tasks that use a single source of information, generating answers solely from either language or visual signals. In practical applications, handling multiple sources of information is crucial, as answer signals may exist in any or all sources. For example, when reading a paper or article, relying solely on textual content may be insufficient, requiring references to images for a more comprehensive understanding. Similarly, in E-commerce platforms, questions related to attributes such as pattern, fabric, or material can be answered using diverse sources, including images, product descriptions, or external references. Figure 1 illustrates this scenario: questions about attributes like pattern color can be inferred from images, while fabric and material are extracted from text descriptions. Notably, certain attributes, such as pattern, may be present in both sources of information.

Recently, various models (Li et al., 2022; Workshop et al., 2022; Almazrouei et al., 2023) have emerged for answer generation, leveraging single-source information from either images or text. To handle multi-source information, these models often rely on separate models for different sources, integrating their outputs through post-processing. Novel frameworks, such as PAM (Lin et al., 2021) and MXT (Khandelwal et al., 2023), have intro-

duced multi-source, multi-modal generative approaches, showing promising results in attribute-related question answering. However, significant challenges remain in developing efficient mechanisms for training and integrating models to handle diverse sources of information.

Existing approaches face several limitations that hinder their effectiveness in answer generation tasks. Firstly, reliance on textual data introduces language bias, potentially leading to skewed attribute generation. Additionally, these models often neglect crucial visual information contained in product images, relying primarily on textual descriptions, which undermines the benefits of multi-source and multi-modal data (Verma et al., 2024). Effective answer generation also requires selectively attending to the most relevant source, focusing on key visual or textual information within that source. Furthermore, handling a diverse range of questions with a single model poses scalability challenges, necessitating expert models tailored to specific question types. Unfortunately, models such as MXT (Khandelwal et al., 2023) and PAM (Lin et al., 2021) fail to address these limitations.

The proposed model addresses the aforementioned limitations by incorporating a question-guided attention (QGA) mechanism and a sparse mixture-of-experts (MoE) model. The QGA mechanism enables the model to autonomously discern attention patterns across multiple sources in scenarios involving diverse information streams. These attention patterns are tailored to the specific posed question. When the answer relies on visual information, the model focuses its attention on visual embeddings. Conversely, when the answer is within the textual context, the model assigns higher weights to textual information. In cases where the answer is derived from all available sources, the model distributes attention appropriately across each source. While cross-modal attention aids in aligning different modalities, it is insufficient for acquiring robust attention patterns within a single source. To address this, we introduce separate embeddings for the question, context, and image, aligning question-image and question-context pairs by maximizing their correlation. This alignment process allows the model to learn precise attention patterns within individual sources based on the given question. Given the diverse nature of the questions, a single model struggles to handle all question types effectively. To address this, We incorporate an MoE strategy into our model,

allowing experts to specialize in different question types. Experiments on a large-scale multi-modal dataset show state-of-the-art performance in attribute-based answer generation. Ablation studies analyze the contribution of each model component.

2 Related Work

Extensive work has been conducted on attribute answer extraction, which can be broadly categorized as *extractive*, *predictive*, and *generative*. Extractive models tag each word in a description using Named Entity Recognition (NER) and extract answers based on these tags. Recent works such as OpenTag (Zheng et al., 2018), LATEX-numeric (Mehta et al., 2021), and MQMRC (Shrimal et al., 2022) leverage NER for answer extraction. While effective for certain categories, these models face limitations in predicting novel entities, and defining entity classes remains challenging. Furthermore, NER-based approaches rely solely on unimodal data, ignoring the richer context available in multi-modal sources such as text and images.

Predictive models form another popular category, where answers are predicted from predefined classes using classification models. These approaches accept unimodal or multimodal data with a question and predict attributes from a fixed set. CMA-CLIP (Liu et al., 2021a) is a recent multimodal approach for attribute prediction. However, these models are limited to predefined attributes and cannot perform zero-shot inference. Given the vast diversity and continuous growth of data, defining a fixed answer set is impractical, and managing large classifier sizes is challenging.

Generative models offer a more flexible solution by generating attributes rather than predicting or extracting them. These models take a question and unimodal or multimodal information as input. AVGPT (Roy et al., 2021) generates attribute answers using text data, while PAM (Lin et al., 2021) and MXT (Khandelwal et al., 2023) introduce multimodal generative frameworks. PAM and MXT are closely related to our approach, both employing generative models in multimodal settings. However, MXT uses two image encoders (ResNet152 (He et al., 2016) and Xception (Chollet, 2017)), making image encoding computationally expensive, and relies on a joint encoder for questions and context. This design prevents direct interaction between the question and the image, limiting the model’s ability to focus on relevant image regions. Additionally, MXT uses a cross-modal




Figure 2: The proposed model architecture consists of two T5 encoders for processing the question and context, along with one image encoder. The question and context are aligned using the *Question Context Alignment Loss* while the question and image are aligned through the *Question Image Alignment Loss*

mechanism that restricts the question’s ability to attend to the most relevant source.

Proposed model addresses these limitations by using a single image encoder and learning patch-wise attention, enabling efficient and precise focus on image regions. Furthermore, question-guided attention facilitates attending to the most relevant mode across all sources of information. We also incorporate the MoE (Shazeer et al., 2017), a recent advancement combining specialized expert models for specific tasks or modalities. MoE has demonstrated significant performance improvements in decoder-only architectures, such as Mixtral (AI, 2023) and MoE-LLAVA (Lin et al., 2024), compared to their non-MoE counterparts.

3 Problem Setting

The proposed model solves the QAG task, unlike standard VQA or QA tasks, our approach incorporates multi-source information, where the answer to a given question may originate from any of the available sources. We define the dataset as $\mathcal{D} = \{q_i, c_i, i_i\}_{i=1}^N$, comprising N samples, each represented by a triplet of question q_i , context c_i , and image i_i . Here, q_i denotes the question posed for attribute generation, c_i represents the context, including the question, product type (PT), product description, and bullet points, while i_i corresponds to the associated image.

4 Proposed Model

To obtain a robust and highly generalizable model, an approach is needed that can automatically attend to various sources in a multi-modal information scenario. To address this, our approach employs three encoders with unshared parameters for context, image, and question. We have developed a question-guided attention mechanism to automatically learn weights for different data sources. The

following section provides a detailed discussion of the proposed model and its components.

4.1 Source Information Embedding

Let us consider a context (c_i) and question (q_i), where $c_i, q_i \in \mathbb{R}^{k \times d}$. The context is encoded using the T5 text encoder model with parameters θ_c and θ_q . Here, c_i includes product descriptions, bullet points, titles, and other relevant information. The T5 architecture is based on the transformer model (Vaswani et al., 2017) and employs self-attention and multi-head attention (MHA). The encoded embeddings of the context and question are defined as follows:

$$C = T5_{\theta_c}(c_i), \quad Q = T5_{\theta_q}(q_i) : \quad Q, C \in \mathbb{R}^{k \times d} \quad (1)$$

The image is encoded using the SwinV2 (Liu et al., 2021b) vision transformer model. Let i_i denote the image, where $i_i \in \mathbb{R}^{3 \times 256 \times 256}$, and let S represent the Swin model with parameter θ_s . The patch embedding from the model is obtained as:

$$i'_o = S\theta_s(i_i) \quad : \quad q_o \in \mathbb{R}^{k' \times d} \quad (2)$$

$$\mathcal{I} = \text{repeat}(i'_o, \text{int}(k/k')) \quad (3)$$

Here, we return the patch embedding rather than the final layer logits. The operation $\mathcal{I} = \text{repeat}(i'_o, \text{int}(k/k'))$ ensures that the image embedding matches the dimensions of the question and context, i.e., dimension k .

4.2 Question Guided Attention (QGA)

The answer to a question may be derived from one or more sources of information. To address this, we developed a QGA mechanism for handling multiple sources. This generic approach applies to any number of sources and can be viewed as a dense MoE, where question-guided attention acts as gating by attending to all sources via a weighted combination. Let $Q \in \mathbb{R}^{k \times d}$ represent the question embedding.

We transform q_o into a c -dimensional embedding using a fully connected (FC) layer, where $c = 2$ (representing two sources of information). This operation is given by:

$$q_a = FC_{\theta_f}(Q), \quad q_a \in \mathbb{R}^{k \times c} \quad (4)$$

These c -dimensional values for each k are used to assign weights to the various sources. We then learn a joint embedding e_i as follows:

$$e_i = \alpha * \mathcal{I} + \beta * \mathcal{C}, \quad e_i \in \mathbb{R}^{k \times d} \quad (5)$$

Here, $\alpha = q_a[:, 0]$ and $\beta = q_a[:, 1]$ are k -dimensional vectors. Rather than learning a scalar weight for each source, we learn token-specific weights, allowing for finer adjustments compared to a single weight per source. The joint embedding for the i^{th} sample, e_i , is passed to the decoder. Since it is guided by the question, if the answer exists in the context, the model learns a higher weight for the context. If it exists in the image, the image weight is higher. However, for a solution in both sources, the model learns a balanced weight value between sources.

4.3 Sources and Question Alignment

In the previous section, the model attends to information from various sources guided by the question but does not learn to focus on relevant information within the source data itself. Here, we align the question with the source data, enabling the model to attend to the most pertinent parts of the source. This alignment is performed for both the image and context relative to the question. The embeddings obtained in Equations 4 and 5 are projected to a single vector of dimension k using a linear transformation and aligned by maximizing cosine similarity. Let q_p , c_p , and i_p denote the projected embeddings. The alignment losses between the question and sources are defined as:

$$\mathcal{L}_{QCA} = |1 - (q_p \cdot c_p) / (|q_p|_2 |c_p|_2)|, \quad (6)$$

$$\mathcal{L}_{QIA} = |1 - (q_p \cdot i_p) / (|q_p|_2 |i_p|_2)|, \quad (7)$$

where \mathcal{L}_{QCA} and \mathcal{L}_{QIA} represent the alignment losses for context and image, respectively. This alignment mechanism improves model performance by focusing on the most relevant parts of the source information.

4.4 Sparse MoE

To handle the diverse set of question and obtained a highly scalable model we leverages sparse MoE

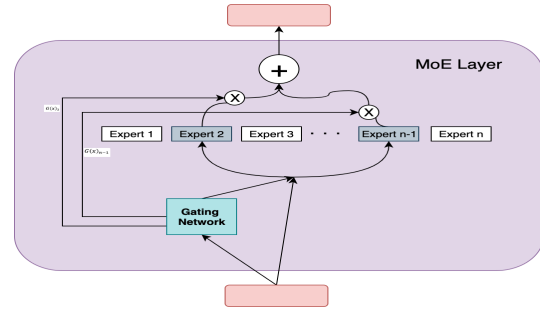


Figure 3: Illustration of working of the Mixture of Experts (MoE) Layer (Shazeer et al., 2017)

model (Figure-3), where different expert can handle the various type of questions using a single model. In an MoE framework, we have a set of experts $\{f_1, \dots, f_n\}$, each taking the same input x and producing outputs $f_1(x), \dots, f_n(x)$, respectively. Additionally, there is a gating function w that takes x as input and produces a vector of weights $(w(x)_1, \dots, w(x)_n)$. The gating network is defined by $w(x) = \text{softmax}(\text{top}_k(Wx + \text{noise}))$. Given an input x , the MoE produces a single combined output by aggregating the outputs of the experts $f_1(x), \dots, f_n(x)$ according to the gating weights $w(x)_1, \dots, w(x)_n$. At the each layer we choose only $\text{top} - k$ expert which produce the sparsity to the model and saves the significant computations. Load balancing is a key issue in the MoE model, to overcome the same we use the load balancing loss. Let n denote the number of experts, and for a given batch of queries $\{x_1, x_2, \dots, x_T\}$, the **auxiliary loss** for the batch is defined as: $\mathcal{L}_{aux} = n \sum_{i=1}^n f_i P_i$ Here, $f_i = \frac{1}{T} \#(\text{queries sent to expert } i)$ represents the fraction of times where expert i is ranked highest, and $P_i = \frac{1}{T} \sum_{j=1}^T w_i(x_j)$ denotes the fraction of weight assigned to expert i , where $w_i(x_j)$ is the weight assigned by the gating mechanism to expert i for query x_j .

4.5 Joint Objective

Let g is the ground truth token and \hat{g} is the generated token, the decoder loss over the generated token is calculated as follows: $\mathcal{L}_{\theta_d}(q_i, i_i, c_i) = \text{CrossEntropy}(\hat{g}, g)$. The complete objective over the decoder and encoder is given as:

$$\begin{aligned} \mathcal{L}_{\theta_q, \theta_i, \theta_c}(q_i, i_i, c_i) &= \mathcal{L}_{\theta_d}(q_i, i_i, c_i) + \mathcal{L}_{QCA} \\ &+ \mathcal{L}_{QIA} + \lambda \mathcal{L}_{aux} \end{aligned} \quad (8)$$

The model is jointly optimized with respect to parameter $\Theta = [\theta_q, \theta_i, \theta_c, \theta_d, \theta_f]$, where θ_q, θ_i , and θ_c are the encoder parameters, θ_f is the fully con-

PT	#Top Attr.	CMA-CLIP	NER-MQMRC	MXT	MoE-MoE
Kurta	K=5	60.69	54.53	76.86	76.55
	K=10	56.67	49.97	66.86	76.93
	K=15	46.49	44.68	57.91	60.31
Shirt	K=5	79.60	71.26	87.89	88.87
	K=10	70.47	52.01	76.99	78.17
	K=15	56.81	45.09	63.60	69.86

Table 1: Results ($Recall@90$) on the 30PT dataset for Kurta and Shirt product types. Results shown for top K attributes ($K = 5, 10, 15$).

Attribute	CMA-CLIP	NER-MQMRC	KNN	MXT	MoE-MoE
Color Map	48.26	26.54	45.95	34.48	49.61
Dress Style	20.34	20.97	13.27	23.79	20.23
Item Length	66.39	47.13	63.08	65.57	69.92
Neck	30.58	13.09	33.67	31.90	34.81
Pattern	14.48	11.61	23.37	24.83	41.62
Season	67.93	16.45	65.37	73.10	69.43
Sleeve	61.68	35.37	44.71	54.38	65.77
Average	44.23	24.45	41.34	44.01	50.19

Table 2: Attribute-wise results ($Recall@90$) on the CMA-CLIP dataset for Dress product type. MoE-MoE outperforms MXT for most attributes, showing significant improvement on average.

ected layer parameter for question embedding projection, and θ_d is the decoder parameter.

5 Experiment and Results

This section, briefly discusses the datasets, baselines and the results obtained using the proposed model.

5.1 Data Description and Base Model

We utilize the *30PT dataset* introduced by MXT (Khandelwal et al., 2023), comprising 30 selected product types (PTs) and 38 distinct attributes sourced from an online platform. The CMA-CLIP (Liu et al., 2021a) paper employed a different dataset with approximately 2.2 million samples for training and 300 samples for validation and testing per attribute. To extend our experiments to a larger scale, we collected a more extensive dataset from the online platform, referred to as OHLSL. This dataset encompasses data for the OHL (other hardlines) and SL (softlines) categories, consisting of 20 million samples for 318 and 145 product types, respectively. The details descriptions about data, baselines and implementations are provided in the supplementary material.

5.2 Results

The result over the three standard datasets are discussed below.

30PT dataset The 30 PT dataset utilized in our study is a comprehensive dataset containing data from various marketplaces and diverse PTs. Our

Attr.	#Prod	CMA-CLIP	NER-MQMRC	MXT	MoE-MoE
age range	27.7k	97.67	13.33	99.03	99.35
department	27.0k	98.39	87.92	98.09	98.57
care inst.	23.3k	36.59	24.62	46.04	48.73
neck	22.2k	52.74	48.01	68.99	74.47
color	21.2k	84.04	74.79	86.03	87.65
design	19.4k	24.97	–	32.69	35.41
occas.	17.3k	19.93	29.67	50.58	52.63
pattern	13.7k	25.83	–	31.92	32.61
season	12.9k	5.07	0.19	33.20	27.15
fit	8.0k	94.66	41.59	95.61	95.72
closure	7.7k	5.05	–	9.33	18.48
collect.	7.3k	0.00	–	30.02	46.87
sleeve	5.7k	60.53	47.99	75.63	80.75

Table 3: $Recall@90P\%$ on Kurta PT. MoE-MoE shows superior performance on visual attributes (neck, color, design), reducing bias towards textual descriptions.




Figure 4: Results on the OHLSL dataset over the Flan-T5 architecture. We report the average Accuracy and $Recall@90$ metric for the OHL and SL category for all the attribute.

trained model underwent evaluation on two PTs, namely, *Kurta* and *Shirt*, encompassing 16 and 19 attributes, respectively. A detailed breakdown of attribute information is available in the supplementary section. Attributes in our evaluation are associated with either visual information or product descriptions. We employed $Recall@90$ (recall with precision ≥ 90) as our evaluation metric for the top k attributes, where $k = 5, 10, 15$. The results, as presented in Table-1, unveil a notable improvement in our proposed approach compared to the recent work MXT (Khandelwal et al., 2023). Specifically, our method, MoEMoE, exhibits an absolute improvement of 6.26% and 2.4% over the top 15 attributes for the Shirt and Kurta datasets, respectively. Our analysis indicates that the majority of the improvement over the MXT model stems from attributes related to visual information. In terms of product description-related attributes, both MXT and MoEMoE yield competitive results.

CMA-CLIP dataset The dataset employed in this study aligns with the one utilized in the CMA-CLIP paper (Liu et al., 2021a). Training was conducted using this standardized dataset, and subsequent inference focused on the "dress" category, comprising nine distinct attributes outlined





Figure 5: The figure shows the ablation over the various component of the proposed model. We can observe that without question guidance (WoQG) or without alignment (WoAL) the model performance significantly drops. Also, single encoder (S-Enc) shows degraded result.

in Table-2. Notably, the proposed model showcases superior performance in six out of the nine attributes when compared to its most competitive counterpart, MXT. Specifically, MoEMoE demonstrates an average absolute improvement of 6.18%. In appendix we provide further discussions regarding current model challenges.

OHL SL dataset OHL SL constitutes a large-scale dataset, with each of the OHL and SL categories containing 20 million samples, while the test set encompasses 1 million samples. The results for the OHL SL dataset are evaluated using the Flan-T5 architecture. In Figure-4, the MoEMoE results are compared with those of the Flan-T5 architecture. Notably, we observe an absolute performance improvement of 2.5% and 5.8% for accuracy and recall@90, respectively, over the Flan-T5 architecture in the OHL category. Similarly, for the SL category, we note an absolute improvement of 2.4% and 3.8% over the same architecture.

6 Ablations

We conducted the extensive ablation over the various proposed components and Figure-5 shows the results for the same. Notably, we observed that the presence of question-guided attention and alignment loss had a substantial impact on the model’s performance. In the absence of question guidance (WoQG), the performance dropped from 59.9 to 58.1 for Kurta PTs and from 69.4 to 68.1 for Shirt PTs. When leveraging the joint encoder (S-Enc), where the question and context are merged, the application of question-guided attention was not feasible, leading to a significant performance drop and the lowest results for both Kurta and Shirt product types. The alignment loss emerged as a crucial factor in directing attention within the source information, enabling the model to focus on the most relevant parts of the image or context. Incorporating the alignment loss further enhanced the model’s performance, raising it from 59.4 and 68.8 to 59.9 and 69.4 for the Kurta and Shirt PTs,

Table 4: Results on the Softlines Dataset (1500 PT-attribute test set)

Model	Acc.	R@90
MXT	63.94	53.52
QGA (Question Guided Attention)	66.04	56.69
QGA Enc-Dec MoE Full Training	62.45	54.19
QGA Enc-Dec MoE Odd*	62.14	52.96
QGA Enc-Dec MoE Even*	62.30	53.18
QGA Encoder MoE Full Training	52.33	41.70
QGA Decoder MoE Full Training	63.81	55.08
QGA Decoder Last MoE	64.29	55.70
QGA Decoder Last-2 MoE	63.67	55.07
QGA Decoder Even MoE*	66.29	57.13
QGA Decoder Odd MoE†	64.79	55.50
MoEMoE (QGA Dec. Odd MoE, Expert Training)	66.57	57.03

*Expert Training Only †MoE Frozen, Backbone Training

respectively. The MoE model further helps to improve the model’s performance, while maintaining the model’s complexity.

6.1 MoE Ablations

We conducted extensive experiments over various settings discussed in the Section-5. To the best of our knowledge, there is no existing literature that has conducted experiments for the encoder-decoder architecture. Most recent works on MoE (AI, 2023) (Lin et al., 2024) focus on the decoder only architecture. In our experiment we have tried to explore all the experimental scenarios for the encoder-decoder architecture. We measure the results on the Softlines test dataset across 1500 PT-attributes and showcase the results in Table 4. This is a challenging dataset and has a huge, diverse output space across all the product-types. In our experiments, we investigate the application of the Mixture of Experts (MoE) architecture within the QGA Model over the previously discussed scenarios and our key observations are as follows:

MoE in Decoder: Applying the MoE architecture exclusively to the decoder layers of the model yields superior performance compared to incorporating it in the encoder layers or across the entire model. This finding suggests that the MoE mechanism is particularly effective in leveraging specialized experts during the output generation phase.

Table-4 shows detailed results over the encoder-decoder architecture. The addition of the MoE layer to the full network degrades the model performance and is unable to outperform the base architecture. However, adding the MoE layer to the decoder layer only helps improve the model accuracy and recall@90 by 0.53% and 0.34% absolute gain, respectively. However adding the MoE to the encoder-decoder layer degrades the model accuracy and recall@90 by 3.59% and 2.50% respectively in the absolute value. Similarly, adding the MoE to the encoder layer only shows the worst performance and the decrease in the baseline accuracy by 13.71% absolute value. We also observe that training the whole model along with the experts slightly degrades the model performance, however training only the MoE layer helps and outperforms the other baselines. Therefore we can conclude that adding the MoE layer to the decoder only and training the MoE expert only, while freezing the basemodel parameters shows the highest improvement and no other setting works as well. In the future it will be interesting to explore how the internal MoE experts are selected if there are there any intrinsic patterns in the question, context and data that helps to select the MoE expert. In the future we will explore the same.

Layer Distribution: We observe that the choice of applying MoE to even or odd decoder layers does not significantly impact the model’s performance, indicating a degree of flexibility in the layer-wise distribution of experts.

Training Strategy: The optimal training strategy involves selectively training only the expert modules and the routing network responsible for assigning inputs to experts, while keeping the remaining model parameters frozen. This focused training approach outperforms the conventional end-to-end training of the entire model, including the MoE components. Interestingly, our experiments reveal that fully training the entire model, encompassing the MoE components (experts and routing network) alongside the rest of the model parameters, tends to degrade the overall performance. This observation highlights the potential challenges of jointly optimizing the MoE architecture and the base model in an end-to-end fashion.

6.2 Auxiliary Loss Ablations

We conducted the ablation for the MoE loss, the results are shown in the Table 5. The ablations are conducted over the best model obtained in the

Table 5: Impact of Auxiliary Loss Weight on Model Accuracy

Model Type	Accuracy
Enc-Dec (wt=0.01)	40.34%
Enc-Dec (wt=0.1)	48.60%
Decoder Only (wt=0.01)	62.73%
Decoder Only (wt=0.1)	66.57%
Decoder Only (wt=0.5)	62.95%

Table-4.

We observe that while Enc-Dec MoE with different weights shows degraded results, the decoder-only model demonstrates significant improvement. The MoE loss weight tuning further enhances performance, with $w = 0.1$ outperforming other baselines. Too low a weight causes the model to ignore the MoE component, while too high a weight overly prioritizes the MoE loss at the expense of the base model’s learning. Thus, the weight must be carefully balanced to enable effective learning of both components.

However, it is important to note that these observations are derived from experiments conducted on a specific task, model architecture, and dataset. The optimal training strategies and deployment of the MoE architecture may vary depending on the problem domain, model characteristics, and data properties.

7 Conclusions

In this work, we introduce MoEMoE, a robust model designed for question answering from multi-source, multi-modal information. Our approach leverages automatic attention learning across diverse information sources, facilitating the identification of the most reliable source for robust answer generation. The proposed question-guided attention mechanism employs a dense-MoE architecture combined with alignment loss and sparse-MoE training in the intermediate layer, which significantly enhances the model’s ability to extract robust features in a scalable manner. The MoE-MoE model achieves state-of-the-art results compared to recent baselines. The proposed attention mechanism, operating both between and within multiple sources, is versatile and applicable to various contexts. By incorporating alignment loss between question-context and question-image pairs, the model effectively explores attention within each source, enabling it to focus on the most pertinent parts of the image or context based on the given question. Extensive experiments on a large-scale dataset, coupled with ablation studies, validate the efficacy of our approach.

References

- Mistral AI. 2023. *Mixtral*. <https://mistral.ai/news/mixtral-of-experts/>.
- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Heslow, Julien Launay, Quentin Malartic, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance.
- François Chollet. 2017. Xception: Deep learning with depth-wise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258.
- Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition supplementary materials. In *IEEE conference on computer vision and pattern recognition*, pages 770–778.
- Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Anant Khandelwal, Happy Mittal, Shreyas Sunil Kulkarni, and Deepak Gupta. 2023. Large scale generative multi-modal attribute extraction for e-commerce attributes. In *To be appear in 61th Annual Meeting of the Association for Computational Linguistics Industry Track*.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. 2022. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR.
- Bin Lin, Zhenyu Tang, Yang Ye, Jiayi Cui, Bin Zhu, Peng Jin, Jinfa Huang, Junwu Zhang, Munan Ning, and Li Yuan. 2024. *Moe-llava: Mixture of experts for large vision-language models*. *Preprint*, arXiv:2401.15947.
- Rongmei Lin, Xiang He, Jie Feng, Nasser Zalmout, Yan Liang, Li Xiong, and Xin Dong. 2021. Pam: Understanding product images in cross product category attribute extraction. *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*.
- Huidong Liu, Shaoyuan Xu, Jinmiao Fu, Yang Liu, Ning Xie, Chien Wang, Bryan Wang, and Yi Sun. 2021a. Cma-clip: Cross-modality attention clip for image-text classification. *ArXiv*, abs/2112.03562.
- Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021b. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022.
- Kartik Mehta, Ioana Oprea, and Nikhil Rasiwasia. 2021. Latex-numeric: Language agnostic text attribute extraction for numeric attributes. In *NAACL*.
- Kalyani Roy, Pawan Goyal, and Manish Pandey. 2021. Attribute value generation from product title using language models. In *Proceedings of The 4th Workshop on e-Commerce and NLP*, pages 13–17.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *Preprint*, arXiv:1701.06538.
- Anubhav Shrivastava, Avi Rajesh Jain, Kartik Mehta, and Promod Yenigalla. 2022. Ner-mqmc: Formulating named entity recognition as multi question machine reading comprehension. *ArXiv*, abs/2205.05904.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. 2023. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models*. <https://crfm.stanford.edu/2023/03/13/alpaca.html>, 3(6):7.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Vinay Verma, Dween Sanny, Abhishek Singh, and Deepak Gupta. 2024. Cod: Coherent detection of entities from images with multiple modalities. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 8015–8024.
- Vinay K Verma, Dween Rabi Sanny, Shreyas Sunil Kulkarni, Prateek Sircar, Abhishek Singh, and Deepak Gupta. 2023. Skill: Skipping color and label landscape: self supervised design representations for products in e-commerce. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3503–3507.
- BigScience Workshop, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, et al. 2022. Bloom: A 176b-parameter open-access multi-lingual language model. *arXiv preprint arXiv:2211.05100*.
- Pan Zhang, Xiaoyi Dong Bin Wang, Yuhang Cao, Chao Xu, Linke Ouyang, Zhiyuan Zhao, Shuangrui Ding, Songyang Zhang, Haodong Duan, Hang Yan, et al. 2023. Internlm-xcomposer: A vision-language large model for advanced text-image comprehension and composition. *arXiv preprint arXiv:2309.15112*.
- Guineng Zheng, Subhabrata Mukherjee, Xin Dong, and Feifei Li. 2018. Opentag: Open attribute value extraction from product profiles. *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.