

# Generative Multimodal Entity Linking

Senbao Shi, Zhenran Xu, Baotian Hu\*, Min Zhang

Harbin Institute of Technology, Shenzhen

shisenbaohit@gmail.com, xuzhenran@stu.hit.edu.cn, {hubaotian, zhangmin2021}@hit.edu.cn

## Abstract

Multimodal Entity Linking (MEL) is the task of mapping mentions with multimodal contexts to the referent entities from a knowledge base. Existing MEL methods mainly focus on designing complex multimodal interaction mechanisms and require fine-tuning all model parameters, which can be prohibitively costly and difficult to scale in the era of Large Language Models (LLMs). In this work, we propose GEMEL, a Generative Multimodal Entity Linking framework based on LLMs, which directly generates target entity names. We keep the vision and language model frozen and only train a feature mapper to enable cross-modality interactions. To adapt LLMs to the MEL task, we leverage the in-context learning capability of LLMs by retrieving multimodal instances as demonstrations. Extensive experiments show that, with only  $\sim 0.3\%$  of the model parameters fine-tuned, GEMEL achieves state-of-the-art results on two well-established MEL datasets (7.7% accuracy gains on WikiDiverse and 8.8% accuracy gains on WikiMEL). The performance gain stems from mitigating the popularity bias of LLM predictions and disambiguating less common entities effectively. Further analysis verifies the generality and scalability of GEMEL. Our framework is compatible with any off-the-shelf language model, paving the way towards an efficient and general solution for utilizing LLMs in the MEL task. Our code is available at <https://github.com/HITsz-TMG/GEMEL>.

**Keywords:** multimodal entity linking, large language models

## 1. Introduction

**Entity Linking** has attracted increasing attention in the natural language processing community, which aims at linking entity mentions in a document to referent entities in a knowledge base (KB) (Shen et al., 2015). It is a fundamental component in applications such as question answering (De Cao et al., 2019), relation extraction (Lin et al., 2016) and semantic search (Blanco et al., 2015).

Existing EL methods mainly focuses on textual modality and has been proven to be successful for well-formed texts (Wu et al., 2020; Cao et al., 2021). However, with the popularity of multimodal information on social media platforms, more ambiguous mentions appear in the short or coarse text. Due to the vast number of mentions arising from incomplete and inconsistent expressions, the conventional text-only EL methods cannot address cross-modal ambiguity, making it difficult to link these mentions accurately (Adjali et al., 2020). To address this issue, the task of **Multimodal Entity Linking** (MEL) has been proposed, which links mentions with multimodal contexts to their corresponding entities (Moon et al., 2018). In Figure 1, based solely on textual modality, it is difficult to determine whether “Harry Potter” should be linked to the corresponding film series or the novel series. However, using multimodal information allows for correctly associating it with the film series.

Although some recent methods have achieved promising performance for the MEL task (Moon

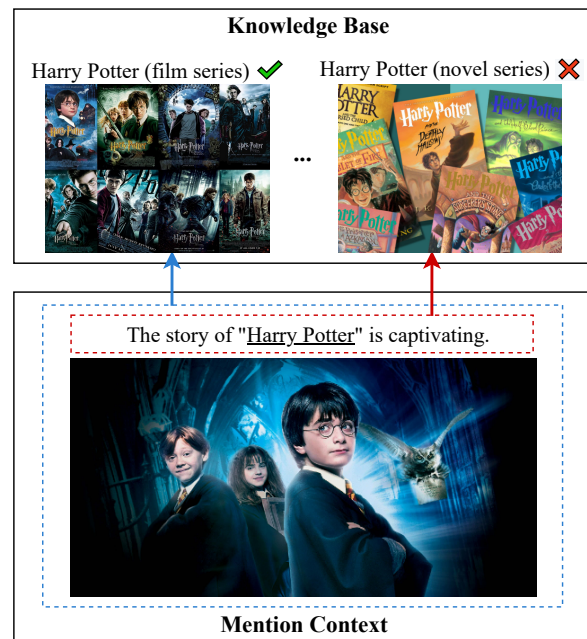


Figure 1: An example of multimodal entity linking with mention underlined in the text. Based solely on text, it is hard to determine whether “Harry Potter” should be linked to the films or the novels. With image as context, “Harry Potter” can be easily linked to the Harry Potter film series.

et al., 2018; Adjali et al., 2020; Wang et al., 2022c,a), challenges still exist. First, these methods often involve a two-stage process of candidate entity retrieval and re-ranking, which can be fur-

\*Corresponding author

ther improved in efficiency. For example, previous works (Wang et al., 2022c; Gan et al., 2021) use various sparse and dense retrievers and then combine candidate entities, followed by a re-ranking process. Second, prior MEL methods (Dongjie and Huang, 2022; Wang et al., 2022a) involve designing complex multimodal interaction mechanisms and training all model parameters, which is prohibitively costly and difficult to scale in the era of large language models (LLMs).

Recently, Vision-Language Models (VLMs) (Wang et al., 2023; Alayrac et al., 2022; Wang et al., 2022b) trained on an enormous amount of image-text data have shown impressive results in various multimodal tasks. However, training a Large Vision-Language Models (LVLM) (e.g., Flamingo (Alayrac et al., 2022)) from scratch is resource-intensive. To alleviate this issue, previous works (Merullo et al., 2023; Eichenberg et al., 2022; Li et al., 2023a) propose that we can construct a LVLM based on the text-only LLM by transforming the visual information into the textual representation space of LLM. Through this approach, LLMs can effectively comprehend visual information and address multimodal tasks. To step forward this direction, adapting LLMs to the MEL task and training models in a parameter-efficient manner have become a promising research direction.

Motivated by the analysis above, we propose a simple yet effective **Generative Multimodal Entity Linking** framework (GEMEL) based on LLMs. Given the multimodal mention contexts, GEMEL can leverage the capabilities of LLMs from large-scale pre-training to directly generate corresponding entity names. We freeze the vision encoder and the language model, and only train a feature mapper to project visual features into a soft prompt for the LLM input. Additionally, we utilize the in-context learning (ICL) capability by constructing multimodal demonstration examples to guide the LLMs to better comprehend the MEL task. Experimental results demonstrate that GEMEL can effectively integrate multimodal information based on LLMs to improve MEL performance, achieving state-of-the-art results on two well-established MEL datasets. Further studies reveal the popularity bias in entity predictions of LLMs, significantly underperforming on rare entities. GEMEL not only excels at common entity but also mitigates the popularity bias of LLMs to further boost performance on the MEL task. Our framework is parameter-efficient and model-agnostic, and it can be transferred to any larger or stronger LLMs in the future.

In summary, the contributions of our work are as follows:

- We propose GEMEL, a simple yet effective framework that utilizes a generative LLM to address the MEL task. To the best of our knowl-

edge, this is the first work to introduce generative methods based on LLMs in the MEL task.

- Extensive experiments show that with only  $\sim 0.3\%$  of the model parameters fine-tuned, GEMEL achieves state-of-the-art results on two well-established MEL datasets (7.7% accuracy gains on WikiDiverse and 8.8% accuracy gains on WikiMEL), exhibiting high parameter efficiency and strong scalability.
- Further studies reveal the popularity bias in LLM predictions, which our framework can effectively mitigate, thereby enhancing overall performance in the MEL task.

## 2. Related Work

### 2.1. Textual EL

Most previous EL methods follow a two-stage “retrieve and re-rank” pipeline (Wu et al., 2020; Barba et al., 2022; Xu et al., 2022; Lai et al., 2022). The initial stage involves retrieving candidate entities from the vast number of entities in the knowledge base. Subsequently, in the second stage, these candidates undergo a re-ranking process to determine the final entity linking results. However, these methods heavily rely on retrieval results during re-ranking, leading to potential error accumulation and performance degradation. GENRE (Cao et al., 2021) introduces a generative approach and uses constrained beam search to directly generate entity names. However, when dealing with short and coarse text, accurately linking entities becomes challenging if based solely on textual modality (Adjali et al., 2020), thus motivating the task of Multimodal EL (Moon et al., 2018).

### 2.2. Multimodal EL

Multimodal EL was initially proposed by Moon et al. (Moon et al., 2018) to address the ambiguous mentions in short social media posts. Although existing MEL methods have made significant progress (Wang et al., 2022c,a; Gan et al., 2021; Dongjie and Huang, 2022; Yang et al., 2023), they still have two limitations that need to be addressed: 1) Prior MEL approaches utilize complex models with many interworking modules (Wang et al., 2022c,a; Gan et al., 2021), which can be further improved in efficiency. For example, Wang et al. (Wang et al., 2022c) combine candidate entity lists from multiple retrieval algorithms and then re-rank; 2) Existing MEL methods mainly focus on improving multimodal fusion based on co-attention mechanisms (Dongjie and Huang, 2022; Wang et al., 2022a), but they require training all model

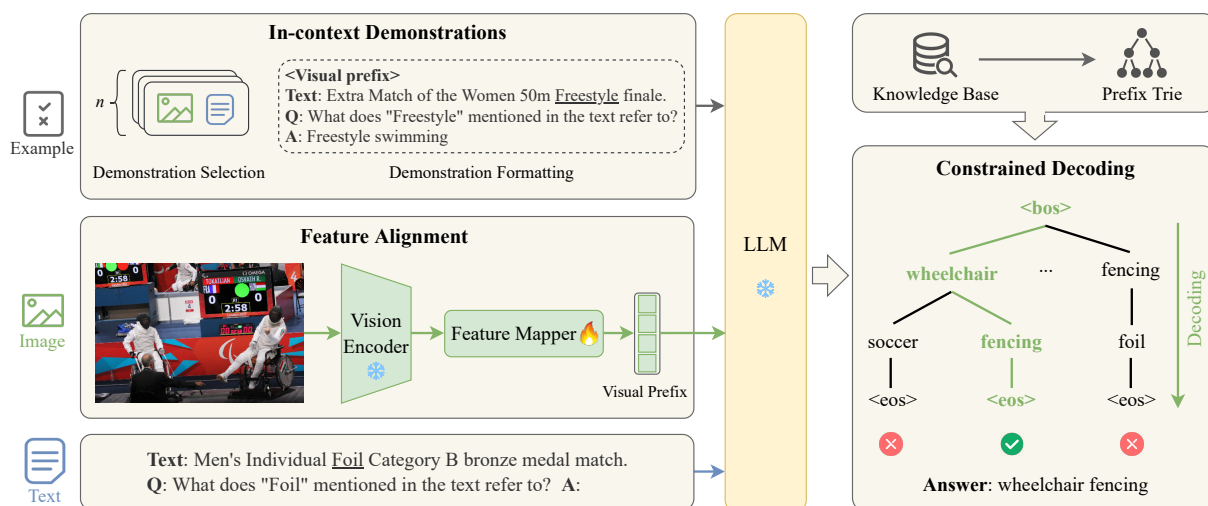


Figure 2: Overview of the GEMEL method. Given the multimodal mention context, GEMEL first uses a feature mapper to transform image features to visual prefix in the textual space. Then, GEMEL leverages the capabilities of LLM to directly generate the target entity name (e.g. “wheelchair fencing”), with  $n$  retrieved multimodal instances as in-context demonstrations. GEMEL applies a constrained decoding strategy to efficiently search the valid entity space. The mention in the text is underlined.

parameters, which becomes difficult to scale in the era of LLMs. To overcome these limitations, we propose a general end-to-end MEL framework that is parameter-efficient and easily scalable.

### 2.3. LLMs for Vision-language Tasks

In recent years, LLMs have demonstrated notable empirical achievements in language understanding (Wei et al., 2022a), generation (Pu and Demberg, 2023), and reasoning (Kojima et al., 2022). These successes have inspired recent research in vision-language tasks. Since LLMs can only perceive text, bridging the gap between natural language and other modalities is necessary.

A feasible approach is to transform visual information into languages with the help of expert models (such as an image captioning model). For example, VideoChat-Text (Li et al., 2023b) enriches the video descriptions with a speech recognition model. However, though using expert models is straightforward, it may result in information loss.

Another effective approach is to introduce a learnable interface between the vision encoder and the LLM to connect information from different modalities. For instance, BLIP-2 (Li et al., 2023a) utilizes a lightweight trainable Q-Former to bridge the modality gap and achieves remarkable performance on various vision-language tasks. Furthermore, some methods (Merullo et al., 2023; Liu et al., 2023; Su et al., 2023) use a projection-based interface to close the modality gap. For example, LLavA (Liu et al., 2023) adopts a simple linear layer to project image features into textual embedding space. Over-

all, previous works demonstrate that it is a potential research direction to utilize frozen LLMs for vision-language tasks.

## 3. Methodology

We introduce our **Generative Multimodal Entity Linking (GEMEL)** framework. As shown in Figure 2, GEMEL takes the multimodal mention context and several in-context demonstrations as input, and directly generates the target entity name. We keep the parameters of the LLM and vision encoder frozen and only train a feature mapper to map image features into the textual space. In Section 3.1, we first present problem formulation. Then, we describe the two components in GEMEL: feature alignment (Section 3.2) and language model generation (Section 3.3).

### 3.1. Problem Formulation

Multimodal Entity Linking (MEL) is the task of mapping mentions with multimodal contexts to the corresponding entities in a knowledge base (KB). Formally, we define  $\mathcal{E}$  as the entity set of the KB, which typically comprises millions of entities. Each mention  $m$  is characterized by the visual context  $V_m$  and textual context  $T_m$ . Here,  $V_m$  represents the associated image for  $m$ , and  $T_m$  represents the textual spans surrounding  $m$ . The task of MEL is to output mention-entity pairs:  $\{(m_i, e_i)\}_{i \in [1, n_m]}$ , where each corresponding entity  $e_i$  is an entry in  $\mathcal{E}$  and  $n_m$  represents the number of mentions. Here we assume each mention has a valid gold entity in

the KB, which is the *in-KB* evaluation problem. We leave the out-of-KB prediction (i.e., *nil* prediction) to future work.

### 3.2. Feature Alignment

Given the inherent incapacity of the LLMs to directly process multimodal information, it becomes imperative to undertake feature alignment procedures on visual data. In our approach for feature alignment, the initial step involves the extraction of image features through utilization of a pre-trained vision encoder. Subsequently, these image features undergo projection into the textual embedding space through the employment of a lightweight feature mapper. Upon this transformation, the resultant features are then introduced into the LLM as a visual prefix, thereby facilitating the LLM’s capacity to process visual information.

**Vision Encoder.** To extract visual features from an input image  $V_m$  corresponding to the mention  $m$ , we utilize a pretrained visual backbone model which produces visual embeddings  $f_\phi(V_m) \in \mathbb{R}^{d_v}$ , where  $d_v$  represents the hidden state size of the vision encoder. The weights of vision encoder  $\phi$  are kept frozen.

**Feature Mapper.** To facilitate cross-modal alignment and fusion, we employ a feature mapper to project visual features into a soft prompt, i.e., a visual prefix for the LLM input. Specifically, we train a feature mapper  $\mathbf{W} \in \mathbb{R}^{d_v \times kd_t}$  to project visual embeddings  $f_\phi(V_m)$  into  $f_\phi(V_m)^T \mathbf{W} \in \mathbb{R}^{kd_t}$ . The result is then reshaped into a visual prefix, which is a sequence of  $k$  embeddings  $\mathbf{v} = \{v_1, v_2, \dots, v_k\}$ , where each embedding shares the same hidden dimensionality  $d_t$  as the text embeddings produced by the LLM for input tokens.

### 3.3. Language Model Generation

In the context of language model generation, the visual prefix derived from the feature alignment module is concatenated with text embeddings and subsequently fed into the LLM. Moreover, we construct a few in-context demonstrations, which serve to enhance the LLM’s understanding of the MEL task. Importantly, this process facilitates the generation of target entity names without necessitating any alterations to the LLM parameters. At test time, we employ a constrained decoding strategy to efficiently search the valid entity space.

**In-context Learning.** To let the LLM better comprehend the MEL task, we leverage its in-context learning (ICL) ability (Brown et al., 2020; Dong et al., 2022), and construct a prompting template with  $n$  demonstration examples from the training set. The demonstration formatting, depicted in Figure 2, includes the image and textual context of mention  $m$ ,

a hand-crafted question (i.e., “What does  $m$  mentioned in the text refer to?”) and the entity name as the answer. For demonstration selection, we take several sparse and dense retrieval methods into consideration:

- **Random selection** will randomly select in-context demonstrations from the training set for each mention.
- **BM25**<sup>1</sup> is one of the most advanced sparse retrieval methods. We use all the mentions in the training set as the corpus and retrieve demonstrations based on the mention.
- **SimCSE**<sup>2</sup> (Gao et al., 2021) is a dense retrieval method for semantic matching. For a pair of mentions, we take the cosine similarity of the mention embeddings as the relevance score.

For  $n$  demonstrations and a new instance  $q$ , we sequentially concatenate the visual prefix  $\mathbf{v}$  and text embedding  $\mathbf{t}$  to obtain the LLM input  $\mathbf{x}$ :

$$\mathbf{x} = [\mathbf{v}_1; \mathbf{t}_1; \dots; \mathbf{v}_n; \mathbf{t}_n; \mathbf{v}_q; \mathbf{t}_q] \quad (1)$$

**Large Language Model.** GEMEL takes an autoregressive LLM  $p_\theta$ , which was originally trained with the maximum likelihood objective on text-only data, and keeps its parameters  $\theta$  frozen. With the input embeddings  $\mathbf{x}$  and  $N$  tokens of the entity name as the target output  $\mathbf{y} = \{y_1, y_2, \dots, y_N\}$ , we can express the teacher forcing training objective as follows:

$$\mathcal{L}_{teacher} = - \sum_{i=1}^N \log p_\theta(y_i | \mathbf{x}; \mathbf{y}_{<i}) \quad (2)$$

**Constrained Decoding.** During the testing phase, if the LLM is allowed to choose any word from its vocabulary at each decoding step, it may generate output strings that are not valid identifiers. To address this issue, we exploit Constrained Beam Search (Sutskever et al., 2014; Cao et al., 2021), a tractable decoding strategy to efficiently search the valid entity space. We define our constrain in terms of a prefix trie  $\mathcal{T}$  where nodes are annotated with tokens from the LLM vocabulary. For each node  $t \in \mathcal{T}$ , its children represent all the allowed continuations from the prefix defined traversing  $\mathcal{T}$  from the root to  $t$ . We tokenize all entity names in the knowledge base to construct a prefix tree, ensuring that the content generated by the LLM consists of valid entity names, thereby matching specific entities in the knowledge base.

<sup>1</sup>We adopt the implementation from [https://github.com/dorianbrown/rank\\_bm25](https://github.com/dorianbrown/rank_bm25).

<sup>2</sup>We utilize the supervised RoBERTa-large checkpoint from <https://github.com/princeton-nlp/SimCSE>.



	WikiDiverse	WikiMEL
# Image-text Pairs	7824	22136
# Mentions	16327	25846
# Text Length (avg.)	10.2	8.2
# Mentions (avg.)	2.1	1.2

Table 1: Statistics of WikiDiverse and WikiMEL, two multimodal entity linking datasets.

## 4. Experiments

### 4.1. Experimental Setup

We elaborate the experimental protocols from the following four aspects: *Datasets*, *Baselines*, *Evaluation Metric* and the *Implementation Details*.

**Datasets.** To assess the capability of GEMEL on the MEL task, we conduct experiments on two MEL datasets, WikiDiverse (Wang et al., 2022c) and WikiMEL (Wang et al., 2022a). WikiDiverse is a human-annotated MEL dataset with diversified contextual topics and entity types from Wikinews. WikiMEL is a large human-verified MEL dataset extracted from Wikidata and Wikipedia. Both Wikidiverse and WikiMEL datasets have been split into training, validation, and test sets, with ratios of 8:1:1 and 7:1:2, respectively, and our experimental setting follows the partition. The statistics of datasets are summarized in Table 1.

**Baselines.** Following previous works (Wang et al., 2022a; Yang et al., 2023; Wang et al., 2022c), we compare our method to recent state-of-the-art methods, categorized as follows: (1) text-only methods relying solely on textual features, and (2) text + vision methods utilizing both textual and visual features. Specifically, we consider:

- **BERT** (text-only) (Devlin et al., 2019) is used as a text encoder to capture textual features and then calculate the relevance score.
- **BLINK** (text-only) (Wu et al., 2020) adopts a bi-encoder for candidate entity retrieval, followed by a cross-encoder for re-ranking.
- **GENRE** (text-only) (Cao et al., 2021) is the first system that retrieves entities by generating their names in an autoregressive fashion.
- **GPT-3.5-Turbo-0613** (text-only) is a powerful LLM developed by OpenAI. We utilize `GPT-3.5-Turbo-0613` (hereafter referred to as `GPT-3.5`) to generate entity names directly and employ the same ICL configuration as in our framework.
- **DZMNED** (text + vision) (Moon et al., 2018) uses a concatenated multimodal attention

mechanism to combine visual, textual, and character features of mentions and entities.

- **JMEL** (text + vision) (Adjali et al., 2020) employs fully connected layers to map the visual and textual features into an implicit space.
- **GHMFC** (text + vision) (Wang et al., 2022a) uses gated hierarchical multimodal fusion and contrastive training to facilitate cross-modality interactions and reduce noise of each modality.
- **LXMERT** (text + vision) (Wang et al., 2022c) is a bi-encoder framework based on the multimodal encoder LXMERT (Tan and Bansal, 2019) for calculating the relevance scores between mentions and entities.
- **MMEL** (text + vision) (Yang et al., 2023) proposes a joint learning framework to solve the multi-mention entity linking task in the multimodal scenario.

**Evaluation Metric.** Following previous works (Moon et al., 2018; Wang et al., 2022a; Yang et al., 2023), we use Top-1 accuracy as the evaluation metric.

**Implementation Details.** Our GEMEL framework is implemented with PyTorch (Paszke et al., 2019). We employ Llama-2-7B (Touvron et al., 2023b) and CLIP ViT-L/14 (Radford et al., 2021) as our default LLM and vision encoder, respectively, unless otherwise stated. For feature mapping, we employ a linear layer (Merullo et al., 2023), leaving the exploration of more complex feature mappers for future research. We adopt SimCSE (Gao et al., 2021) to retrieve a total of  $n = 16$  demonstration examples, sorted in ascending order based on their relevance to the mention. We search the visual prefix length  $k$  among [1, 2, 4, 8] and find  $k = 4$  performs the best based on the development set. We set the beam size to 5 during inference. Our models are trained with a batch size of 1 and 16 steps of gradient accumulation for 5 epochs on a single A100 GPU. We utilize the AdamW optimizer (Loshchilov and Hutter, 2019) with a learning rate of  $1e-6$  and a warmup ratio of 10%.

### 4.2. Main Results

Table 2 presents the model performances on two MEL datasets. According to the experimental results, we can see that: First, GEMEL surpasses all other approaches and achieves state-of-the-art performance on both MEL datasets, with a 7.7% improvement on WikiDiverse (78.6%  $\rightarrow$  86.3%) and a 8.8% improvement on WikiMEL (73.8%  $\rightarrow$  82.6%), showing the effectiveness of our framework. This indicates that, by fine-tuning a feature mapper ( $\sim 0.3\%$  of model parameters), GEMEL

Modality	Model	Top-1 Accuracy (%)	
		WikiDiverse	WikiMEL
Text	BERT (Devlin et al., 2019)	69.6	31.7
	BLINK (Wu et al., 2020)	70.9	30.8
	GENRE (Cao et al., 2021)	78.0*	60.1*
	GPT-3.5-Turbo-0613	72.7	73.8
Text + Vision	JMEL (Adjali et al., 2020)	38.4	31.3
	DZMNED (Moon et al., 2018)	70.8	30.9
	GHMFC (Wang et al., 2022a)	62.7*	43.6
	LXMERT (Wang et al., 2022c)	78.6	-
	MMEL (Yang et al., 2023)	-	71.5
Text + Vision	GEMEL (ours)	<b>86.3</b>	<b>82.6</b>

Table 2: Model performance on the test set. **Bold** denotes the best results. “\*” means our implementation with official Github repositories. “-” means not reported in the cited paper.

Model	Top-1 Accuracy (%)	
	WikiDiverse	WikiMEL
GEMEL	86.3	82.6
w/o Visual Information	79.5	74.2
w/o In-context Learning	80.2	75.2

Table 3: Ablation results of GEMEL.

enables the frozen LLM to comprehend visual information effectively and efficiently, and then leverage it to enhance MEL performance. Second, methods based on LLMs (namely GPT-3.5 and GEMEL) demonstrate powerful performance in both textual modality and multimodality. In the text modality, GPT-3.5 can match or even exceed the performance of previous multimodal methods. We reckon there are two main reasons for this: 1) Textual modality still plays a dominant role in the MEL task, while visual modality primarily serves as supplementary information; 2) LLMs pretrained on large-scale datasets can capture extensive language patterns, context, and knowledge, leading to outstanding performance in common entity prediction (see Section 5.3 for details).

Table 3 shows the ablation study results. First, removing visual information significantly impairs GEMEL’s performance, which indicates the importance of visual information when the textual information is short and insufficient (see cases in Section 5.4). Second, eliminating the demonstrations in prompts leads to a performance decrease of 6.1% and 7.4% on WikiDiverse and WikiMEL respectively. This suggests that providing a few in-context demonstrations facilitates the LLM in recognizing and comprehending the MEL task (see Section 5.2 for details).

LLM	Parameters	Top-1 Accuracy (%)	
		WikiDiverse	WikiMEL
OPT	6.7B	82.7	75.5
Llama	7B	85.8	82.2
Llama-2	7B	86.3	82.6

Table 4: Results of different LLMs.

Method	Top-1 Accuracy (%)	
	WikiDiverse	WikiMEL
ResNet-101	85.3	82.1
BEIT-large	84.7	81.0
CLIP ViT-L/14	86.3	82.6

Table 5: Results of different vision encoders.

## 5. Analysis

In this section, we will analyze the performance of our framework from four aspects. First, we will explore the generality and scalability of our framework. Second, we will investigate the influence of different demonstration selection methods in our framework. Subsequently, we identify the popularity bias in LLM predictions (i.e., significantly under-performing in tail entities), which our framework effectively mitigates. Finally, we conduct case study and limitation analysis.

### 5.1. Generality and Scalability

To test the generality of GEMEL across different LLMs, we conduct experiments on various LLMs including OPT (Zhang et al., 2022), Llama (Touvron et al., 2023a), and Llama-2 (Touvron et al., 2023b). As shown in Table 4, our GEMEL framework is generally effective for the widely-used LLMs.

Different vision encoders may affect the model

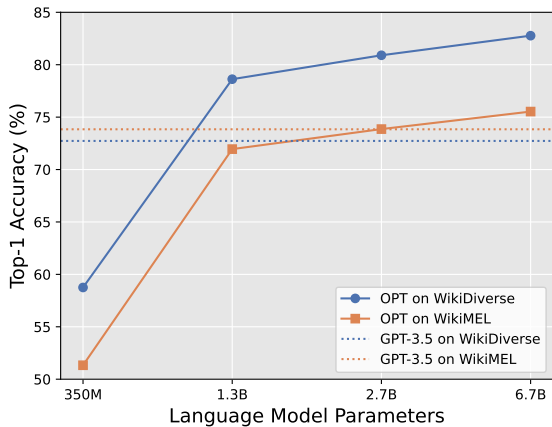


Figure 3: Results of scaling up parameters of OPT. As language models continue to scale up, GEMEL consistently exhibits enhanced performance in the MEL task.

Method	Top-1 Accuracy (%)	
	WikiDiverse	WikiMEL
No-Retrieval	80.2	75.2
Random Selection	83.1	78.3
BM25	85.3	82.5
SimCSE	<b>86.3</b>	<b>82.6</b>

Table 6: Results of different demonstration selection methods. **Bold** indicates the best performance.

performance. We compare three widely-used types of vision encoders, ResNet-101 (He et al., 2016), BEiT-large (Bao et al., 2022), and CLIP ViT-L/14 image encoder (Radford et al., 2021). As shown in Table 5, our framework performs well with various visual encoders, with CLIP achieving the best results. Therefore, we use CLIP by default in our framework.

As shown in Figure 3, we employ OPT models of varying scales (i.e., 350M, 1.3B, 2.7B, 6.7B) to assess the influence of language model scale on GEMEL performance. We do not perform the ICL prompting for the OPT 350M model due to its insufficient scale for demonstrating the ICL ability. It is evident that with the language model scaling up, GEMEL consistently improves its performance on the MEL task. Our framework even outperforms GPT-3.5 on small-scale language models, as demonstrated by the results of OPT-1.3B on the WikiDiverse dataset. This indicates that our framework is highly effective and model-agnostic, meaning it can be applied to larger or more powerful LLMs that may be released in the future, thereby further enhancing performance. Given resource constraints, our investigation is limited to models up to 6.7B in scale, leaving the exploration of larger language models for future work.

## 5.2. Demonstration Selection

To investigate the influence of demonstration selection on the performance of GEMEL, we conduct experiments using several sparse and dense retrieval methods: random selection, BM25, and SimCSE (Gao et al., 2021). The results of demonstration selection are shown in Table 6, where we also present the result of no-retrieval method (i.e., without ICL). The experimental findings indicate the following: 1) Whether employing random selection, BM25, or SimCSE, all of these methods surpass the approach of no-retrieval. This indicates that incorporating in-context demonstrations facilitates LLMs in recognizing the format of the MEL task. 2) Methods utilizing similarity retrieval outperform random selection. We believe that demonstrations retrieved based on similarity are more likely to include similar mentions and entity candidates, thus learning through similar demonstrations and enhancing MEL performance; 3) The retrieval approach utilizing SimCSE outperforms the BM25 method. We reckon that employing dense retrieval methods, such as SimCSE, enhances the comprehension of semantic meaning in mentions compared to sparse methods like BM25. This results in the retrieval of more relevant and higher quality demonstrations.

## 5.3. Popularity Bias of LLMs

To explore why methods based on LLMs can exhibit remarkable performance in both textual modality and multi-modality, we conduct a statistical analysis of prediction outcomes using LLM-based approaches. Following previous work (Chen et al., 2021), we tallied the occurrences of each entity in Wikipedia<sup>3</sup>. Entities that appear less than 5% of the total count are considered as *tail entities*, while the rest are regarded as *common entities*. We conduct statistical analysis on the accuracy of GPT-3.5, GEMEL (w/o Visual Information), and GEMEL in terms of common entity and tail entity predictions. The results are shown in Table 7.

Based on the statistical results, we can observe the following: 1) The LLM-based approaches demonstrate impressive performance in the textual modality, primarily attributed to LLM’s exceptional performance in predicting common entity after large-scale pre-training. For instance, both GPT-3.5 and GEMEL achieve accuracy rates exceeding 70% in predicting common entity; 2) In terms of predicting tail entity, methods based on LLMs exhibit a popularity bias, i.e., significantly under-performing on less common entities. Taking GPT-3.5’s result

<sup>3</sup>We adopt the result from GENRE (Cao et al., 2021) repository: [https://dl.fbaipublicfiles.com/GENRE/mention2wikidataID\\_with\\_titles\\_label\\_alias\\_redirect.pkl](https://dl.fbaipublicfiles.com/GENRE/mention2wikidataID_with_titles_label_alias_redirect.pkl).

Model	WikiDiverse		WikiMEL	
	Common Entity	Tail Entity	Common Entity	Tail Entity
GPT-3.5-Turbo-0613	72.0	37.5	74.4	65.5
GEMEL (w/o Visual Information)	80.8	40.0	74.4	68.8
GEMEL	87.4	57.5	82.7	75.2

Table 7: Accuracy (%) of LLM-based methods on common entity and tail entity.


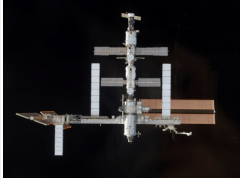


Case	1	2	3	4
Visual Information				
Textual Information	The official seal of the <u>House of Representatives</u>	The ISS from <u>Atlantis</u> on June 10, 2007.	<u>Hawthorn</u> have won their first premiership in 17 years.	<u>Black Widow</u> and Captain America.
Text-only Prediction	House of Representatives	Atlantis, Florida	Hawthorn	Scarlett Johansson
GEMEL Prediction	<u>United States House of Representatives</u>	<u>Space Shuttle Atlantis</u>	<u>Hawthorn Football Club</u>	<u>Black Widow (Marvel Comics)</u>
Golden Entity	United States House of Representatives	Space Shuttle Atlantis	Hawthorn Football Club	Black Widow (Natasha Romanova)

Figure 4: Case study. For underlined mentions, green and red text mean successful and failed predictions, respectively. The text-only prediction results are obtained by GEMEL without utilizing visual information.

on WikiDiverse as an example, the prediction accuracy for common entity reaches 72.0%, whereas for tail entity, it is only 37.5%; 3) Our framework not only ensures significant performance in predicting common entities but also demonstrates substantial improvements in tail entity prediction. For instance, on the WikiDiverse dataset, compared to GPT-3.5, GEMEL exhibits notable enhancement in common entity prediction and a significant increase of 20.0% (37.5% → 57.5%) in tail entity prediction. This indicates the effectiveness of our multimodal framework in mitigating the bias of LLMs towards tail entity prediction, thereby enhancing overall performance in the MEL task.

#### 5.4. Case Study

In Figure 4, we compare prediction results of GEMEL in textual modality and multimodality.

Cases 1, 2, and 3 are examples where predictions are incorrect in the text-only modality but correct in the multimodality. It can be observed that when the information provided by the text is insufficient, visual modality can complement information and help eliminate ambiguity. Taking Case 1 as an example, in situations with only textual modal-

ity, the model struggles to determine the region to which “House of Representative” belongs. However, when visual information is introduced, GEMEL effectively utilizes visual information to accurately link the mention to a specific entity.

Case 4 predicts incorrectly in both textual modality and multimodality. It requires linking to the specific role of “Black Widow” based on multimodal information, which necessitates making judgments considering details such as attire. This inspires us to explore more fine-grained multimodal information in future research.

## 6. Conclusion and Future Work

We propose GEMEL, a simple yet effective generative multimodal entity linking framework based on LLMs, which leverages the capabilities of LLMs to directly generate target entity names. Experimental results demonstrate that GEMEL outperforms state-of-the-art methods on two MEL datasets, exhibiting high parameter efficiency and strong scalability. Further studies reveal the existence of bias in LLMs predictions for tail entity, which our framework can effectively mitigate, thereby enhancing overall performance in the MEL task. Moreover,



our framework is model-agnostic, enabling its application to larger or more powerful LLMs in the future. Further research can explore how to mitigate bias in tail entity prediction for LLMs and extend GEMEL to more modalities (such as video, speech, etc.).

## Acknowledgments

We would like to thank Yunxin Li, Ziyang Chen, Qian Yang, Xinshuo Hu, and Yulin Chen for the valuable comments and useful suggestions. This work is supported by grants: Natural Science Foundation of China (No. 62376067).

## References

- Omar Adjali, Romaric Besançon, Olivier Ferret, Hervé Le Borgne, and Brigitte Grau. 2020. [Multi-modal entity linking for tweets](#). In *Advances in Information Retrieval - 42nd European Conference on IR Research, ECIR 2020, Lisbon, Portugal, April 14-17, 2020, Proceedings, Part I*, volume 12035 of *Lecture Notes in Computer Science*, pages 463–478. Springer.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736.
- Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. 2022. [Beit: BERT pre-training of image transformers](#). In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Edoardo Barba, Luigi Procopio, and Roberto Navigli. 2022. [ExtEnD: Extractive entity disambiguation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2478–2488, Dublin, Ireland. Association for Computational Linguistics.
- Roi Blanco, Giuseppe Ottaviano, and Edgar Meij. 2015. [Fast and space-efficient entity linking for queries](#). In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining, WSDM 2015, Shanghai, China, February 2-6, 2015*, pages 179–188. ACM.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Nicola De Cao, Gautier Izacard, Sebastian Riedel, and Fabio Petroni. 2021. [Autoregressive entity retrieval](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Anthony Chen, Pallavi Gudipati, Shayne Longpre, Xiao Ling, and Sameer Singh. 2021. [Evaluating entity disambiguation and the role of popularity in retrieval-based NLP](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4472–4485, Online. Association for Computational Linguistics.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. 2019. [Question answering by reasoning across documents with graph convolutional networks](#). 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 2306–2317, Minneapolis, Minnesota. Association for Computational Linguistics.
- 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, and Zhifang Sui. 2022. A survey for in-context learning. *arXiv preprint arXiv:2301.00234*.
- Zhang Dongjie and Longtao Huang. 2022. [Multi-modal knowledge learning for named entity disambiguation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 3160–3169, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

- Constantin Eichenberg, Sidney Black, Samuel Weinbach, Letitia Parcalabescu, and Anette Frank. 2022. [MAGMA – multimodal augmentation of generative models through adapter-based finetuning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2416–2428, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jingru Gan, Jinchang Luo, Haiwei Wang, Shuhui Wang, Wei He, and Qingming Huang. 2021. [Multimodal entity linking: A new dataset and a baseline](#). In *Proceedings of the 29th ACM International Conference on Multimedia, MM '21*, page 993–1001, New York, NY, USA. Association for Computing Machinery.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [Simcse: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 6894–6910. Association for Computational Linguistics.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. [Deep residual learning for image recognition](#). In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27-30, 2016*, pages 770–778. IEEE Computer Society.
- Takeshi Kojima, Shixiang (Shane) Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 22199–22213. Curran Associates, Inc.
- Tuan Lai, Heng Ji, and ChengXiang Zhai. 2022. [Improving candidate retrieval with entity profile generation for Wikidata entity linking](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3696–3711, Dublin, Ireland. Association for Computational Linguistics.
- Guillaume Lample, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer. 2016. [Neural architectures for named entity recognition](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 260–270, San Diego, California. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. 2023a. [BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models](#). In *International Conference on Machine Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of Machine Learning Research*, pages 19730–19742. PMLR.
- KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2023b. [Videochat: Chat-centric video understanding](#). *arXiv preprint arXiv:2305.06355*.
- Yankai Lin, Shiqi Shen, Zhiyuan Liu, Huanbo Luan, and Maosong Sun. 2016. [Neural relation extraction with selective attention over instances](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2124–2133, Berlin, Germany. Association for Computational Linguistics.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. [Visual instruction tuning](#). In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*.
- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2022. [What makes good in-context examples for GPT-3? In Proceedings of Deep Learning Inside Out \(DeeLIO 2022\): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures](#), pages 100–114, Dublin, Ireland and Online. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Jack Merullo, Louis Castricato, Carsten Eickhoff, and Ellie Pavlick. 2023. [Linearly mapping from image to text space](#). In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

- Seungwhan Moon, Leonardo Neves, and Vitor Carvalho. 2018. [Multimodal named entity disambiguation for noisy social media posts](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2000–2008, Melbourne, Australia. Association for Computational Linguistics.
- Thien Huu Nguyen, Nicolas Fauceglia, Mariano Rodriguez Muro, Oktie Hassanzadeh, Alfio Massimiliano Gliozzo, and Mohammad Sadoghi. 2016. [Joint learning of local and global features for entity linking via neural networks](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2310–2320, Osaka, Japan. The COLING 2016 Organizing Committee.
- OpenAI. 2023. [GPT-4 technical report](#). *CoRR*, abs/2303.08774.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Dongqi Pu and Vera Demberg. 2023. [ChatGPT vs human-authored text: Insights into controllable text summarization and sentence style transfer](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 4: Student Research Workshop)*, pages 1–18, Toronto, Canada. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya 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 *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Wei Shen, Jianyong Wang, and Jiawei Han. 2015. [Entity linking with a knowledge base: Issues, techniques, and solutions](#). *IEEE Transactions on Knowledge and Data Engineering*, 27(2):443–460.
- Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. 2023. Pandagpt: One model to instruction-follow them all. *arXiv preprint arXiv:2305.16355*.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. *Advances in neural information processing systems*, 27.
- Hao Tan and Mohit Bansal. 2019. [LXMERT: learning cross-modality encoder representations from transformers](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 5099–5110. Association for Computational Linguistics.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [Llama: Open and efficient foundation language models](#). *CoRR*, abs/2302.13971.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutvi Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Rangan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *CoRR*, abs/2307.09288.
- Maria Tsimpoukelli, Jacob Menick, Serkan Cabi, S. M. Ali Eslami, Oriol Vinyals, and Felix Hill. 2021. [Multimodal few-shot learning with frozen language models](#). In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pages 200–212.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,



- Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Peng Wang, Jiangheng Wu, and Xiaohang Chen. 2022a. [Multimodal entity linking with gated hierarchical fusion and contrastive training](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '22*, page 938–948, New York, NY, USA. Association for Computing Machinery.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. 2022b. [Opa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework](#). In *International Conference on Machine Learning*, pages 23318–23340. PMLR.
- Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. 2023. [Image as a foreign language: Beit pretraining for vision and vision-language tasks](#). In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19175–19186.
- Xuwu Wang, Junfeng Tian, Min Gui, Zhixu Li, Rui Wang, Ming Yan, Lihan Chen, and Yanghua Xiao. 2022c. [WikiDiverse: A multimodal entity linking dataset with diversified contextual topics and entity types](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4785–4797, Dublin, Ireland. Association for Computational Linguistics.
- Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, Ed H. Chi, Tatsunori Hashimoto, Oriol Vinyals, Percy Liang, Jeff Dean, and William Fedus. 2022a. [Emergent abilities of large language models](#). *Transactions on Machine Learning Research*. Survey Certification.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022b. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Ledell Wu, Fabio Petroni, Martin Josifoski, Sebastian Riedel, and Luke Zettlemoyer. 2020. [Scalable zero-shot entity linking with dense entity retrieval](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 6397–6407. Association for Computational Linguistics.
- Zhenran Xu, Yulin Chen, Senbao Shi, and Baotian Hu. 2022. [Enhancing entity linking with contextualized entity embeddings](#). In *Natural Language Processing and Chinese Computing - 11th CCF International Conference, NLPCC 2022, Guilin, China, September 24-25, 2022, Proceedings, Part II*, volume 13552 of *Lecture Notes in Computer Science*, pages 228–239. Springer.
- Chengmei Yang, Bowei He, Yimeng Wu, Chao Xing, Lianghua He, and Chen Ma. 2023. [MMEL: A joint learning framework for multi-mention entity linking](#). In *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, volume 216 of *Proceedings of Machine Learning Research*, pages 2411–2421. PMLR.
- Klim Zaporozhets, Johannes Deleu, Yiwei Jiang, Thomas Demeester, and Chris Develder. 2022. [Towards consistent document-level entity linking: Joint models for entity linking and coreference resolution](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 778–784, Dublin, Ireland. Association for Computational Linguistics.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona T. Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer. 2022. [OPT: open pre-trained transformer language models](#). *CoRR*, abs/2205.01068.